

WORKING PAPER NO: 575

**More Heat than Light: Census-scale Evidence for the
Relationship between Ethnic Diversity and Economic
Development as a Statistical Artifact**

Naveen Bharathi

*Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 5600 76
naveen.bharathi@iimb.ac.in*

Deepak Malghan

*Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 5600 76
dmalghan@iimb.ac.in*

Andaleeb Rahman

*Cornell University
Ithaca NY 14853, USA
ar687@cornell.edu*

Year of Publication – August 2018

More Heat than Light: Census-scale Evidence for the Relationship between Ethnic Diversity and Economic Development as a Statistical Artifact*

By NAVEEN BHARATHI AND DEEPAK MALGHAN AND ANDALEEB RAHMAN[†]

The association between diversity and development – both negative and positive – has been empirically tested for a limited set of diversity variables despite its centrality to the political economy discourse. Using a unique census-scale micro dataset from rural India containing detailed caste, religion, language, and landholding data ($n \approx 13.25$ million households) in combination with administrative data on human development, satellite measurements of luminosity as proxy for sub-national economic development, we show that an association between social heterogeneity and economic development is tenuous at best, and is likely an artifact of geographic, political, and ethnic units of analysis. We develop a cogent framework to jointly account for these ‘units of analysis’ effects – in particular by introducing the MEUP or the Modifiable Ethnic Unit Problem as the counterpart of MAUP (Modifiable Areal Unit Problem) in spatial econometrics. We use seventeen different diversity metrics across multiple combinations of ethnic and geographic aggregations to empirically validate this framework, including the first ever census-scale enumeration and coding of elementary Indian caste categories (jatis) since 1931.

JEL: O12; Z12; Z13

Keywords: Ethnic-geographic Aggregation; Modifiable Ethnic Unit Problem (MEUP); Caste; Night Lights; Ethnic Inequality; India

* Draft of August 6, 2018. Download latest version [HERE](#).

[†] Bharathi: Indian Institute of Management Bangalore, Bannerghatta Road, Bengaluru 560076 INDIA, naveen.bharathi@iimb.ac.in. Malghan: Indian Institute of Management Bangalore, Bannerghatta Road, Bengaluru 560 076 INDIA, dmalghan@iimb.ac.in. Rahman: Cornell University, Ithaca NY 14853 USA, ar687@cornell.edu. We thank Arnab Basu, James Fenske, Ravi Kanbur, Sharun Mukand, Imran Rasul, Matthias Vom Hau, and Andreas Wimmer for useful discussions and comments on earlier drafts. Conversations with Abdul Aziz, Prithvi Datta Chandra Shobi, Satish Deshpande, Chandan Gowda, Rajalaxmi Kamath, M.S.Sriram, Anand Teltumbde, and M. Vijaybhaskar were crucial in validating our understanding of the political sociology of historical and contemporary caste groups (the usual disclaimers apply). Devvrat Vegad assisted with the processing of DMSP-OLS data. Bharathi acknowledges a seed grant from IIM Bangalore that funded archival research and fieldwork used to harmonize social groups coding. Part of the research reported here was carried out when Malghan was affiliated with Princeton University. He thanks Princeton Institute for International and Regional Studies for funding support.

I. Introduction

The relationship between social heterogeneity and economic outcomes is one of the central themes punctuating political economy of development. The dominant empirical literature has found a negative association – if not a robust statistical correlation – between ethnic diversity and a range of political economy outcomes including economic performance (Easterly and Levine, 1997); public goods provisioning (Alesina, Baqir and Easterly, 1999; Miguel and Gugerty, 2005); quality of governance (La Porta et al., 1999; Alesina and Zhuravskaya, 2011); civil war or strife (Collier and Hoeffler, 1998; Collier, 2004; Montalvo and Reynal-Querol, 2005*b*); and social trust (Alesina and La Ferrara, 2002). This apparent empirical success has even spawned attempts to delineate causal pathways linking social heterogeneity and negative development outcomes – especially public goods provisioning (Habyarimana et al., 2007; Dahlberg, Edmark and Lundqvist, 2012). While the multitude of evidence for the “diversity debit hypothesis” (Gerring et al., 2015) has dominated the literature, a competing “diversity dividend” (Gisselquist, Leiderer and Niño Zarazúa, 2016) argument has also been advanced both empirically and theoretically – a complex diversified economy can achieve significant productivity gains by harnessing the skill complementarities available in a heterogeneous society (Alesina and La Ferrara, 2005).

The sovereign nation state has been the principal geographic and political unit of analysis in the extant empirical literature investigating the diversity debit hypothesis. The potential empirical biases stemming from the “methodological nationalism” of cross-national regressions are well-documented but data limitations have meant that the sovereign state has continued to rule the empirical roost (Wimmer and Glick Schiller, 2002; Kanbur, Rajaram and Varshney, 2011). Limited evidence from sub-national settings in developing countries have generally not supported the diversity debit hypothesis. The nature of governance regimes, as well as political processes at the subnational level are different from those at the national level, confounding the potential effect of diversity and leading to greater empirical ambiguity (Glennster, Miguel and Rothenberg, 2013; Gerring et al., 2015). Indeed, in modern urban centers supporting a complex economy, diversity has a positive effect on both wages and productivity (Ottaviano and Peri, 2006).

With few notable exceptions, almost all of the empirical work linking diversity and economic development has relied on a narrow set of ethnicity, language, and religion variables. Historically minded political scientists and sociologists have questioned the use of ethnic divisions as an independent variable in econometric models testing the diversity debit hypothesis, especially when the sovereign nation state is the unit of analysis. Ethnic boundaries are not exogenous to the extent that the historical processes that created them are also the ones that resulted in extant national boundaries. Further, these macro-historical processes of

nation-state formation have a direct bearing on contemporary state-capacity that is in turn is reflected in current levels of economic development.¹ In the context of caste in India – our primary empirical context – administrative classifications of caste like census designated categories, rather than elementary ascriptive categories have been used (Banerjee and Somanathan, 2007). In this paper, we investigate the extent to which empirical results are robust across ethnic units of analysis.

The central contribution of this paper is to show how the relationship between diversity and economic development is likely an artifact of *where* diversity is measured, *how* diversity is measured, and *what* diversity is measured. We show that the *where*, the *how*, and the *what* are jointly determined in an indivisible ethnic-geographic continuum. Using a large census-scale micro dataset, we are able to empirically test the diversity - development relationship across the universe of values assumed by various diversity metrics.

Our first contribution is to show how geographic and ethnic units of analysis operate together, and account for much of the instability in diversity - development empirical models. We use two subnational units of analysis (village, $n \approx 27,000$; sub-district, $n = 175$) and three different aggregations of caste, our principal ethnic diversity variable. The six resulting ethnic-geographic composites are all aggregated from a common household micro dataset ($n \approx 13.25$ million households). We develop the Modifiable Ethnic Unit Problem (MEUP) as the ethnic analogue of the Modifiable Areal Unit Problem (MAUP) used to characterize geographic aggregation, and also delineate the intersections between MAUP and MEUP. We develop a formal framework for testing theories (including diversity-development hypotheses) in the indivisible ethnic-geographic continuum.

Second, we demonstrate how diversity - development models are sensitive to particular diversity metrics that are used. With few notable exceptions, the literature has relied on ethnolinguistic fractionalization (ELF) as a measure of social heterogeneity. Being a population-share metric, a fractionalization metric does not always capture the most “politically relevant ethnic groups” (Posner, 2004) that contribute to saliency of ethnic tensions. At each level of ethnic and geographic aggregation we construct a polarization metric (Esteban and Ray, 1994; Montalvo and Reynal-Querol, 2005*b,a*, 2008). In addition to a discrete polarization metric that assumes equal ethnic distance between different groups, we also adapt the original formulation developed by Esteban and Ray (1994) using differences in landholding between different groups as the ethnic distance between groups. We also show how fractionalization and polarization metrics are impacted

¹See for example, Singh and vom Hau (2016) and other essays that follow it in the special section of *Comparative Political Studies*. Also see Chandra (2006) for a detailed discussion on why the empirical literature must pay greater attention to construction of ethnic categories.

differentially by MAUP and MEUP. Our rich dataset allows use the full range of values in the theoretical fractionalization-polarization map that has hitherto not been possible.

Third, we test the diversity - development relationship using not only fractionalization and polarization metrics but also using ethnic inequality. We construct a standard entropic inequality measure at each level of geographic aggregation (mean log deviation of household landholding) and decompose this inequality into within-group and between-group components, with the latter representing ethnic inequality (Alesina, Michalopoulos and Papaioannou, 2016). We demonstrate how the relationship between ethnic inequality and development is also impacted by ethnic and geographic units of analysis. We show that our argument about the relationship between diversity and development being a statistical artifact is further strengthened by explicitly accounting for ethnic inequality.

Fourth, this is the first study to use a household level micro dataset containing elementary caste categories (*jati*) at census-scale. The census-scale survey used in constructing our analytic dataset returned more than 1600 caste names that we coded into ≈ 700 unique locally endogamous ascriptive caste categories. Beyond allowing us to study MEUP, our caste coding represents the first such census-scale attempt since the colonial decennial census of 1931. Our rich micro-data allows us to interact caste with other diversity variables such as religion, language, and even economic class. Thus, we are able to account for all major social cleavages that define agrarian Indian society.

The remainder of this paper is organized as follows. In the Section II, we introduce the “units of analysis problem” in the ethnic-geographic continuum using a general taxonomic map for quantitative diversity and ethnic inequality metrics. In Section III, we describe how caste is the most important social cleavage in agrarian India, and delineate the pathways through which it influences economic outcomes. We also discuss the political economy of caste aggregations (both ethnic and geographic) in contemporary India. In Section IV, we detail the construction of our analytic dataset that is centered on the first census-scale enumeration and coding of elementary caste categories in India since 1931. In Section V, we present our empirical models, and discuss principal regression results. In Section VI, we develop the Modifiable Ethnic Unit Problem (MEUP) as an analogue of the Modifiable Areal Unit Problem (MAUP) in spatial econometrics to account for our empirical results. We also illustrate the intersections between MAUP and MEUP by formally defining the ethnic-geographic continuum, and delineating a framework for validation of theories in the ethnic-geographic space. We conclude in Section VII with a brief discussion of the key implications of our findings.

II. The Units of Analysis Problem

There are three implicit questions underlying the construction of diversity metrics that the extant diversity-development literature has largely ignored — where is diversity measured? How is diversity measured? What diversity is measured? Nearly all available evidence for the diversity debit hypothesis comes from cross-national datasets using a limited set of diversity variables.² However, there are no innate theoretical reasons for a particular geographic scale, or a particular set of diversity variables to circumscribe the diversity debit hypothesis. Varying preferences across social groups, or ethnic strife in a socially heterogeneous society — the two principal mechanisms underlying the diversity debit hypothesis — are not specific to any particular geographic scale or particular forms of diversity. If the underlying theory is general enough, it must be subject to empirical tests across varying geographic scales and differing heterogeneity variables. In Figure 1, we symbolically describe how the construction of any diversity metric must account for geographic aggregation, salient heterogeneity axes, ethnic aggregation, as well as the actual measurement of diversity. Using examples from rural India — the empirical context of this paper — we discuss why a diversity metric must be sensitive to these factors. A particularly important contribution of this paper is to show why the ethnic unit of analysis — a topic that has completely escaped extant literature — is at least as important as the geographic unit of analysis. Indeed, we show that ethnic and geographic aggregations are theoretically and empirically entangled together in an ethnic-geographic continuum.

A. Geographic Unit of Analysis

Much of the support for diversity debit hypothesis comes from cross-national datasets. Data unavailability has been the most significant barrier in transcending the limitations of “methodological nationalism” that has plagued the literature on diversity and development (Wimmer and Glick Schiller, 2002; Kanbur, Rajaram and Varshney, 2011). The seminal works that pioneered the empirical study of diversity-development relationship derived their analytical datasets from the *Atlas Narodov Mira* (Atlas of the Peoples of the World) — published in 1964 using data collected by Soviet ethnographers in the 1960s (Easterly and Levine, 1997; Alesina et al., 2003; Fearon and Laitin, 2003). The *Atlas Narodov* data enumerates the homelands of over nine-hundred ethnic groups using 1964 national boundaries. This original dataset has since been geo-coded at higher resolutions (Weidmann, Rød and Cederman, 2010), and has formed the basis for more recent works revisiting the ethnic-diversity and development conundrum (Alesina, Michalopoulos and Papaioannou, 2016; Montalvo and Reynal-Querol, 2017). Ethnic data collated by *Ethnologue* — the fifteenth edition of which contains over seven thousand five-hundred language-country groups mapped to national boundaries in 2000 — has

²The best-known exception is the seminal paper of Alesina, Baqir and Easterly (1999) that documents a negative relationship between ethnic diversity and public goods for US cities.

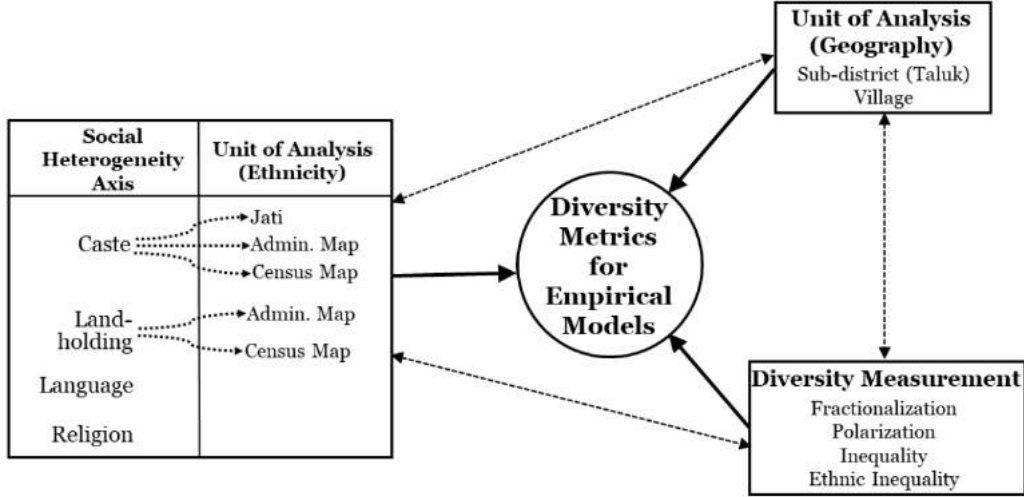


Figure 1. : *Ethnic-geographic Aggregation, and Diversity Measurement*

served as yet another global ethnic dataset (Gordon Jr, 2005). Religion data from L'Étatdes Religions Dans le Monde (ET), World Christian Encyclopedia, and the Encyclopedia Britannica datasets have also been used in the literature (Montalvo and Reynal-Querol, 2005a; Akdede, 2010). Most recently, the Spatially Interpolated Data on Ethnicity (SIDE) dataset collates information on ethno-linguistic, religious, and ethno-religious settlement patterns across forty-seven LMICs – low and middle income countries (Müller-Crepon and Hunziker, 2018).

Despite evidence that the diversity debit hypothesis likely breaks down at sub-national scales, there is no consensus on why the relationship between diversity and development might be sensitive to geographic unit of analysis (Ottaviano and Peri, 2006; Gerring et al., 2015; Singh, 2015b; Gisselquist, Leiderer and Niño Zarazúa, 2016). Several explanations explicating the nature of subnational governance and political economy processes at the subnational level have been offered. Typical subnational units are geographically and politically nested – for example, towns are embedded in counties that are themselves contained in a state or province. The spatial distribution of ethnic groups across nested geographies rather than simple intra-unit diversity accounts for the empirical instability in the relationship between diversity and development (Tajima, Samphantharak and Ostwald, 2018; Bharathi et al., 2018).

Potential endogeneity in the relationship between development outcomes and ethnic segregation is particularly difficult to tease out at sub-national levels. While we discuss the longer-term co-evolution of ethnic categories and development outcomes in detail below, another source of plausible endogeneity lies in

relative ease of subnational migration. Unlike restrictions on cross-border migration across international boundaries, movements within a country – to regions with better economic opportunities – are typically unfettered. Subnational migrations can lead to “optimal sorting” so that economically productive individuals are more likely to be found in socially diverse geographic units (Damm, 2009; Gerring et al., 2015; Gisselquist, Leiderer and Niño Zarazúa, 2016). While a theoretical possibility, the empirically observed levels and patterns of internal migration are not always congruent with “optimal sorting” predictions. In rural India for example – the empirical context of analysis in this paper – there is scant support for optimal sorting. Migrations from rural India are largely driven by marriages within endogamous castes, or are seasonal in nature.³

At smaller subnational scales, the coordination problems arising out of varying ethnic preferences are more easily ameliorated such that benefits of diversity dominate. However, not all subnational scales are created equal – at some scales, diversity debit rather than diversity credit can be more pronounced (Gerring et al., 2015). Cognate evidence from around the world (Ostrom, 1990), including from rural India (Thapliyal, Mukherji and Malghan, 2017), suggests that informal non-market and non-state mechanisms that are unlikely to work at national scales can develop at subnational scales.

Determining the saliency of any one (or more) of the plausible pathways implicating spatial scales as a determinant of the diversity-development relationship is an empirical exercise. Irrespective of the particular pathway(s) through which the geographic unit of analysis is salient, there is no sound theoretical or empirical argument beyond data limitations for ignoring the saliency of spatial scales. This is precisely what we capture in Figure 1 that shows how we account for the geographic units of analysis effect in our own data. We aggregate a common household dataset with identity variables of interest into two political-geographic aggregates – villages and sub-districts. We test the diversity-development association at both these geographic units of analysis.

The diversity-development literature using Indian data has traditionally relied on district-level data (Banerjee and Somanathan, 2007). The typical Indian district is a large aggregation,⁴ and the subnational units within a district are heterogeneous. Table 1, along with Figure 2 and Figure 3 present the preliminary picture of the extent of this subnational heterogeneity in our dataset — a theme that we will develop in detail when we discuss our empirical results. Figure 2 shows the intra-district heterogeneity in fractionalization indices computed at

³*cf.* Section VI for further discussion on migration patterns in rural India. For a comprehensive bibliography of Indian migration patterns, see Tumble (2013).

⁴The average size of the thirty districts in the Indian state of Karnataka that we study is 2.2 million people per district.

Table 1—: Geographic Aggregation, and Diversity Correlations

	Fractionalization	Polarization
Jati	0.336	0.035
Admin Map	0.273	-0.053
Census Map	0.464	0.419

Note: The table reports the correlations (Pearson) between respective diversity metrics measured at village, and sub-district scales.

the village-level, and Figure 3 represents variation for polarization indices.⁵ Table 1 shows the correlations (or lack thereof) between diversity metrics computed at village and sub-district levels. The low correlations are indicative of heterogeneous villages within a sub-district. The ecological inference problem associated with spatial aggregations is well-established across disciplines (Robinson, 1950; Openshaw, 1984; Fotheringham and Wong, 1991; Freedman, 1999; King, 2013). In this paper, we investigate how (if) this modifiable areal unit problem (MAUP) clouds the extant empirical base of the diversity debit hypothesis. Our census-scale household-level micro dataset allows a particularly “clean” test of the impact of MAUP as the analytic dataset at both spatial units of analysis (village and the sub-district) are constructed from the common household-level micro dataset.

B. Ethnic Unit of Analysis

The widespread use of cross-national datasets in empirical development-diversity models has also meant that extant evidence for the diversity-debit hypothesis rests on a narrow set of aggregate ethnic categories. The *Atlas Narodov Mira* aggregates the world population into approximately 1,600 ethnic groups, and the empirical literature has relied on a subset of these groups. For example, Fearon (2003) uses 822 ethno-religious classifications. The essentialist premise of such data that has dominated the political economy discourse has been criticized by by constructivists (Chandra, 2006). The ideal ethnic category dataset will contain self-reported ascriptive ethnic markers from individuals residing in a given geographical territory. In the absence of such ideal data (especially for historical categories), researchers have coded ethnicity using secondary sources. Such enumerations, however, are subjective and not rooted in a cogent theory of ethnic categorization with the result that ethnic group assignment is contingent on specific secondary sources consulted (Marquardt and Herrera, 2015).

Individuals have multiple group identities – ethnic, religious, and linguistic etc. The political saliency of these identities is neither time-invariant nor space-invariant. Thus, a significant challenge for any research is determining the subset

⁵cf. Section IV for details about how these diversity indices were constructed.

of heterogeneity axes that are relevant to given empirical context. For example, an individual who identifies herself as a “Bengali,” or “Tamil” in interactions within India might self-identify as an “Indian” or even a “South Asian” in a foreign land. The same individual might use her religious or caste affiliation as her primary identity within a Tamil or Bengali linguistic group. Global datasets widely used in the diversity-development literature are susceptible to an erroneous delineation of the most important societal dividing lines (Marquardt and Herrera, 2015). For example, the *Atlas Narodov Mira* groups Jews into sub-groups like *Ashkenazi*, *Bukharan*, etc., based upon linguistic and cultural divides, but *Tatar Muslims* and *Christians* are considered as a unified group. The arbitrary definitions of group boundaries can potentially contaminate empirical models that use diversity metrics derived from ethnic group shares.

If ethnic boundaries vary across space and time, the historical geography of cultural heritage, or discrimination determine which of the ethnic boundaries are salient at a given time and place (Wimmer, 2008*b,a*). The saliency of particular ethnic boundaries are also determined by the spatial scale of interest. For example, of the many ethnic groups that are salient locally in several African countries, national-level political economy is dominated by select aggregate ethnic coalitions or major individual groups (Scarritt and Mozaffar, 1999). New ethno-political identities are often activated in response to major political upheavals, and these newly forged identities are a result of both fusion and fission of historical groups – a process that is central to making and remaking contemporary caste boundaries in India as it is to forging of African ethnic identities. Further, subnational majorities are not automatically also always politically relevant at the national level, and country-wide aggregate ethnic classifications can obfuscate the actual inter-group dynamics driving conflicts (Marquardt and Herrera, 2015).

Static country-wide ethnic classifications often ignore assimilation, amalgamation, or further sub-divisions among ethnic-linguistic-religious categories – especially the actual salience of groups from the perspective of inter-group conflicts, or heterogeneous preferences at the heart of the diversity debit hypothesis. Thus not all ethnic divisions based on cultural divisions, and enumerated in anthropological accounts, are politically relevant in the context of diversity-development associations. We show in this paper that not accounting for the political saliency of ethnic divisions is at the heart of why ethnic scales have been neglected in empirical diversity-development models.

SUBNATIONAL ETHNIC BOUNDARIES IN RURAL INDIA

In Figure 1, we list four different social heterogeneity axes that we use in this paper – caste, landholding, language, and religion. Caste is the most important social cleavage in agrarian India. The elementary endogamous and hereditary

caste groups are not the only caste identity that is politically salient in India. The colonial state as well as the successor post-colonial state of independent India have both used aggregate caste categories for administrative reasons such as census enumeration and affirmative action. Over a period of time, these state created identities are assimilated so that they gain wide currency. In some instances, administrative interventions such as the colonial census enumeration of castes can even make and remake elementary endogamous caste categories.⁶

Empirical diversity-development models have hitherto ignored ethnic aggregation despite the fact that diversity metrics such as the fractionalization index (ELF) are sensitive to aggregation effects – in that a diversity rank order of a spatial unit is sensitive to the level of ethnic aggregation used in computing diversity. Ethnic aggregation is a “many-to-one” function, $f : \mathbb{N} \rightarrow \mathbb{N}$, whose inverse is not defined. As a specific example of f , consider the mapping of an elementary caste group or *jati*, $J_i \in \{J_1, \dots, J_n\}$ to the three-fold census categorization of social groups that is the staple of empirical work using Indian caste data⁷

$$(1) \quad \begin{aligned} \text{CensusMap} : J_i &\mapsto f(J_i) \ni \\ &f(J_i) \in \{SC, ST, OTH\} \end{aligned}$$

A diversity metric (for example, a fractionalization metric) for a given spatial unit can be evaluated on either elementary n ethnic groups ($\{J_1, \dots, J_n\}$), or $k \leq n$ aggregate groups ($\{f(J_1), \dots, f(J_n)\}$).

Let X be a set of distinct and labelled spatial units (for example, all the villages, or sub-districts in our dataset):

$$(2) \quad X = \{x_1, \dots, x_n\}$$

A diversity metric like the fractionalization index, evaluated for all the spatial units in X is a distinct partially ordered set (or, poset) at any given ethnic scale. Formally, we represent these posets for elementary ethnic groups and some arbitrary aggregation of elementary groups as:

$$(3a) \quad P_0 = (X, \preceq_0); \text{ Diversity computed on elementary ethnic groups}$$

$$(3b) \quad P_1 = (X, \preceq_1); \text{ Diversity computed on aggregate ethnic groups}$$

The extant empirical literature on diversity and development has glossed over the fact that P_0 and P_1 are not isomorphic. Like any posets, P_0 and P_1 are best

⁶*Cf.* Section III for an overview of when different aggregations of caste are politically salient. For a more general and historically-grounded *longue dur e* introduction to the evolution of the institution of caste in India, see Banerjee-Dube (2008) and Guha (2013).

⁷We code over 700 elementary *jati* categories in our empirical models used in this paper.

conceptualized as directed graphs. Graphs $G_0 = G(P_0, \preceq_0)$ and $G_1 = G(P_0, \preceq_1)$ have the spatial units (countries, villages, sub-districts, etc.) in P_0 or P_1 as the respective vertex-sets, and edge-sets defined by relations \preceq_0 and \preceq_1 respectively.

In Figure 4, we illustrate how ethnic aggregation alters the ordinal rank-ordering of spatial units using Hasse diagram snippets corresponding to fractionalization computed using elementary caste categories or *jati* groups (G_0); and two ethnic aggregations (G_1 and say, G_2). These Hasse Diagrams are derived from the actual dataset used in this paper.⁸ The three Hasse snippets⁹ in Figure 4 show the rank-ordering of these sub-district posets obtained by computing the fractionalization index (ELF) at three ethnic scales of caste – *jati*, administrative mapping of *jatis*, and the census mapping of *jatis*. The three Hasse diagram snippets in the figure represent three different partially ordered sets of the kind represented in equation 3. We show in this paper that this change in ordinal rank-ordering is not randomly distributed in space. We show why the problems with ethnic aggregation are compounded when the spatial structure of such aggregation is also taken into account.

While we have presented ethnic aggregation of elementary caste categories as a technical problem for econometric models of diversity and development, the real substantive question is however the political saliency (or lack thereof) of any particular ethnic scale. Among other things, this saliency is at least partly contingent on the dependent variable of interest. A direct corollary of dependent variable determining the choice of ethnic scale is that the choice of ethnic scale is not divorced from the choice of geographic scale – a theme that we will explore in much detail in this paper. For example, when studying collective action problems surrounding village-level public goods where the village community has substantive agency, elementary *jati* cleavages are likely to be the most salient ethnic scale. However, administrative caste categories might assume saliency if state provisioning of public goods were the primary outcome variable of interest.

While the ethnic-geographic aggregation of caste is the primary theoretical focus of this paper, we also benchmark our empirical results using the two most widely used diversity axes in the literature – language and religion. In our empirical setting, different caste groups speak a common language – at least at local scales. Thus, linguistic differences are not likely to be the most salient axis and indeed there is some evidence that a unifying language can under certain circumstances even promote subnational solidarity that can in turn drive positive development outcomes in societies with other social cleavages (Singh, 2015a).

⁸For the algorithm used to construct the Hasse Diagram representation of our three posets, see Weisstein, Eric W. “Hasse Diagram.” From MathWorld—A Wolfram Web Resource. <http://mathworld.wolfram.com/HasseDiagram.html>. (accessed, July 10, 2018).

⁹The snippets show six randomly selected sub-districts out of a total of 175 in the corresponding complete Hasse diagrams.

Religious differences are however a significant source of conflict and strife in our empirical context.¹⁰ Our empirical models also account for economic diversity and inequality using household-level landholding data. In computing ethnic inequality, Figure 1 shows how we are once again confronted with the question of selecting the most salient ethnic aggregation of caste.

C. *The Ethnic-Geographic Continuum, and Diversity Metrics*

The general framework we introduced in Figure 1 includes the diversity measurement framework besides geographic and ethnic scales. Empirical models of diversity-development have largely relied on a single diversity metric – the ethno-linguistic fractionalization index or ELF (Roeder, 2001). Besides fractionalization as a measure of diversity, empirical models have also included a measure of polarization (Esteban and Ray, 1994) as a more theoretically appealing measure of ethnic conflict (Montalvo and Reynal-Querol, 2005*b,a*, 2008; Esteban, Mayoral and Ray, 2012*a,b*; Gören, 2014). While fractionalization is merely a descriptive measure of ethnic diversity, polarization also measures the underlying conflict-potential that is implicit in characterizations of ethnic divisions.

In general, the choice of diversity metric is driven by the dependent variable of interest. If ethnic fractionalization negatively impacts economic growth or public goods provisioning, polarization is theoretically more appealing in empirical models of ethnic conflicts. A small number of influential empirical studies have included both fractionalization and polarization metrics in empirical models that shed some light on how fractionalization and polarization interact with each other. Fractionalization is implicated when the conflict is over a “public” prize such as political power but polarization better explains a “private” prize like government subsidy, infrastructure resources that can be privatized, or even war-loot (Esteban, Mayoral and Ray, 2012*a*). Polarization explains how consequential ethnic divisions are “activated” (Chandra and Wilkinson, 2008) when a majority ethnic group is faced with a significantly-sized minority group. Mere presence of primordial differences does not always lead to ethnic divisions that matter (Esteban, Mayoral and Ray, 2012*b*; Ray and Esteban, 2017). There is empirical evidence from the diversity-debit genre that polarization is a better predictor predictor of the negative association between diversity and economic development (as measured by economic output, investment, or government consumption (Montalvo and Reynal-Querol, 2003, 2005*b*). More recently, based on a cross-national study of 100 countries, Gören (2014) showed that the fractionalization and polarization channels are qualitatively different in how diversity impacts economic development. If ethnic fractionalization has a direct impact on economic growth, ethnic polarization’s impact on economic growth is mediated

¹⁰ Cf. Section IV for a detailed account of religion as a significant diversity axis in our empirical context.

through civil wars, political instability, and other manifestations of strife.

While the choice of diversity metric has largely been driven by empirical contingencies, our framework in Figure 1 shows that this choice is related to geographic and ethnic units of analysis. Figure 1 shows how the choice of ethnic and geographic units of analysis as well as the choice of diversity metric are jointly determined. The figure shows that questions of *what* diversity is measured, *where* is diversity measured, and *how* diversity is measured are all inextricably linked with each other. Our principal argument in this paper is that the relationship between diversity and development is embedded in an ethnic-geographic continuum and the most relevant diversity metric is determined by *both* ethnic and geographic scales. Continuing with ethnic and geographic aggregation of caste discussed above, the choice of diversity metric (fractionalization or polarization) is determined by specific ethnic-geography of caste. For example, it is fair to hypothesize that ethnic fractionalization is a more relevant metric at larger geographic aggregations (for a given ethnic scale), and that local conflicts are best captured by a polarization metric. While we discuss the specific metrics relevant to the our empirical context (agrarian India) in Section V below, we only make the preliminary observation here that the questions of *how* diversity must be measured cannot be answered without reference to ethnic and geographic units of analysis.

DIVERSITY, OR ETHNIC INEQUALITY?

The empirical literature linking diversity and development has developed independent of the older concern about the relationship between inequality and development (Alesina, Michalopoulos and Papaioannou, 2016). Economic inequality across ethnic groups exacerbates the potential for intergroup conflict (Ray and Esteban, 2017). For example, Houle and Bodea (2017) show that between 1960 and 2005, ethnic inequality increased the likelihood of a authoritarian coup in 32 sub-Saharan countries in Africa. The likely channel for this linkage is the process by which ethnic inequality contributes to consolidation and hardening of group preferences. This paper builds on the recent empirical focus on ethnic *distance* rather than mere ethnic diversity (Baldwin and Huber, 2010; Alesina, Michalopoulos and Papaioannou, 2016).

There is no theoretical reason why ethnic inequality and ethnic diversity measurements must coincide. To the extent that ethnic inequality and diversity are both related to development outcomes, empirical models must test the relative strengths (or lack thereof) of both inequality and diversity channels. Figure 5 records the relationship between ethnic diversity and ethnic inequality in our principal dataset. Ethnic inequality as well as diversity is measured at two different ethnic as well as geographic aggregations of caste. Ethnic inequality is

simply the between-group component of the landholding inequality at village or sub-district (*taluka*) levels. In each of the four panels, the points are color-coded by district.¹¹ Each of the panels rank spatial units (villages or sub-districts) by respective fractionalization and ethnic inequality measures. The top two panels shows the diversity-inequality ranking map for $\approx 27,000$ villages in our dataset and the bottom panel is from 175 sub-districts (*talukas*). The left panels at each of the spatial scales uses the eight-fold administrative aggregation of caste categories, and the right panels use the three-fold census categorization of caste. The figure provides empirical proof for why ethnic inequality and ethnic diversity are distinct concepts and models of diversity and development must control of ethnic inequality.

III. Political Economy of Caste Aggregations

The Indian caste system is one of the of the most stable and entrenched institutions of hierarchy and social stratification in history (Ghurye, 1969; Jodhka, 2012). Caste is the principal source of social heterogeneity in India’s agrarian society – the empirical context of this paper. An individual’s caste marker is also closely correlated with her station on the class hierarchy (Zacharias and Vakula-bharanam, 2011; Iversen et al., 2014). The defining feature of caste, relevant to delineating the relationship between social heterogeneity and economic outcomes is complete and total social closure. A caste ridden society restricts every group, or groups of castes to pre-ordained socio-economic opportunities that serves as the canvas for intergroup and intra-group competition (Weber, 1968). Institutional scholars have argued that “caste rules prevent the free operation of markets; and these rules have led to inefficiency and stagnation as well as inequality” (Olson, 1982).

An individual’s caste identity can take the form of an endogamous *jati* marker, or the *varna* marker, with both being hereditary, and defined at birth. The census-scale data that we use (detailed in the next section) in our analysis enumerated over 1600 distinct *jati* groups in rural Karnataka. The lower bound for an all-India estimate for number endogamous *jati* categories is well over three-thousand. The *jati* hierarchy maps on to a five-fold *varna* hierarchy with the *Brahmins* or the priestly class at the top and erstwhile untouchables at the bottom. The *varna* hierarchy is abstract, theoretical, and maps *jati* groups to specific occupations. This pre-modern mapping between *jati* and *varna* was substantively modified during the colonial period and has served as the basis for administrative classification of various *jati* groups in contemporary India. While the *varna* hierarchy has a pan-India structure, *jati* operates with a sub-regional character with fluid and ambiguous hierarchy, power and status (Srinivas, 1965; Dumont, 1980). At the

¹¹There are 30 districts in our dataset. Cf. Section IV for detailed data description.

local level, *jati* rather than *varna* is the principal operational axis of stratification and heterogeneity (Srinivas, 1996; Gupta, 1991). Even while *jatis* are thought to be ontologically stable objects, they have been subject to a continuous metamorphosis in both space and time. Colonization, encounter with modern political economy, ensuing de-ritualization, and democratization have brought about great changes in the hierarchical nature of the *jati* structure (Bayly, 2001; Dirks, 2011). Changing nature of *jati* structure constantly puts *jati* groups in conflict with each other in their quest for political and socio-economic power where numerical and economic strength of a *jati* group is a significant predictor of its economic fortunes (Iversen et al., 2014). For example, in the late nineteenth century, and continuing into early twentieth century, Mysore state (part of the present Indian state of Karnataka that we study here), many *jatis* forged themselves together under a single umbrella identity to increase their numerical strength to claim a share of political power. Conversely, *jatis* have also fragmented into sub-*jatis* to reinforce their identity in representative politics. Hence, *jati* identities both divide and unite people in their pursuit of economic prosperity and power. This process of “fission” and “fusion” is an important driver of fragmentation and polarization of Indian society. (Hardgrave Jr, 1968).

Historically, caste structure has helped define economic relations in India’s agrarian society. Besides constraining occupational mobility, caste structure stratifies rural India into proprietors and tenants. The erstwhile untouchable groups at the bottom of caste hierarchy have historically been excluded from tenancy and relegated to labouring on farm lands (Otsuka, Chuma and Hayami, 1993). Social distance between caste groups is congruent with economic distance and caste fragmentation influences economic contracting (Anderson, 2011). Democratic politics has been the most significant driver of change in the nature of relationship between castes groups so that several land owning middle peasantry have made rapid political progress if not always accompanied by corresponding economic gains. Caste has been implicated in the ability of various groups to seize opportunities accorded by a rapidly urbanizing society embedded in a global economy. Caste in a modern economy can function independent of its traditional ritual role. For example, Damodaran (2009) in his pioneering study of new Indian capitalists has documented how dominant landowning castes have benefited from rapid growth of the urban economy. The low status of castes such as Dalits (erstwhile untouchables), artisan groups, small and marginal peasants in the traditional socio-economic hierarchy has been a significant hurdle preventing their entry into new economic sectors where labor markets continue to discriminate along caste lines (Sheth, 1999; Munshi and Rosenzweig, 2006).

A central contribution of our paper is to empirically account for the fact that the political saliency of caste as a marker of social heterogeneity is not static in time, or in space. We show how specifying caste at different levels of ethnic aggre-

gation in diversity - development models helps uncover the variation of impacts of caste across varying sub-national geographic aggregations. We present robust evidence for why caste is a variable that is embedded in the ethno-geographic continuum.

IV. Data

In order to test the relationship between diversity and development at multiple ethnic and geographic units of analysis, we combine data from various sources including satellite “night lights,” village-level administrative data on human development, census data on village-level provision of public goods and centrally, independent India’s first census-scale enumeration and coding of detailed endogamous caste group (*jati*) information. In this section, we describe the construction of our analytic dataset, and also present key descriptive statistics for both dependent and independent variables used in our models.

A. Karnataka: Study Area Description

The geographic site of our empirical investigation is the south Indian state of Karnataka – a state whose population is comparable to that of France.¹² With an area of over 191,000 square kilometers, it is the seventh largest state in India. The left panel in Figure 6 shows the geographic location of Karnataka within India. The contemporary state of Karnataka was formed in 1956 (as part of a wave that reorganized Indian states along linguistic lines) by integration of five different colonial-era territories, shown in the right-hand-side panel of Figure 6. The historical administrative map of Karnataka also shows thirty contemporary districts – the highest sub-state political and administrative unit. Present day Karnataka has Kannada¹³ speaking regions from two British presidencies (Madras and Bombay), and two nominally independent states under British suzerainty (the “princely states” of Hyderabad, Mysore), and the district of Coorg that was administered by the British resident commissioner in the present day capital of Karnataka, Bangalore. The historical differences in colonial administrative and land tenure regimes continue to influence the political economy of development in contemporary Karnataka. The northern part of the state and especially districts from erstwhile Hyderabad princely state continue to lag behind the rest of the state on almost every indicator of social and human development.¹⁴

¹²The census-scale household survey (2015) that is at the heart of our empirical analysis enumerated 5.99 million residents and the 2011 decennial census returned a population figure of 6.11 million residents (Chandramouli, 2011). France had a population of 6.28 million in 2010.

¹³Kannada is the principal language of the state, and the etymological root for the state name.

¹⁴For example, see Nanjundappa et al. (2002); UNDP (2005). Also see Figure 7 and Table 2 below for a discussion about the historical geography of our dependent variables. Hyderabad Karnataka (the districts which were integrated from the erstwhile Hyderabad princely state) was accorded special status under Article 371 of Indian constitution in 2012 to better address these developmental disparities.

Like other states in India, Karnataka is also administratively divided into districts and sub-districts (called *talukas* in Karnataka but also referred to as *tehsils* in other parts of the country). In 2015, there were 30 districts and 176 sub-districts.¹⁵ Beyond the diverse administrative and colonial history, Karnataka is also among the most physically and ecologically diverse states in the country. The state consists of three principal geographic divisions that are consequential for development outcomes – coastal region, hilly region, and plains. The state also counts ten different agro-climatic zones within its boundaries. This agro-climatic diversity accounts for some of the regional disparity in observed patterns of rural development. About 65% of Karnataka’s land area is under cultivation and about 77% of the total area of the state is classified as arid or semi-arid (Ramachandra, Kamakshi and Shruthi, 2004).

B. ‘Night Lights’ Luminosity

There is now considerable evidence that satellite images of night light intensity is a good proxy of economic development (Ebener et al., 2005; Chen and Nordhaus, 2011; Henderson, Storeygard and Weil, 2012). In the absence of reliable sub-national income data, we use “night lights” as our primary dependent variable in our regression models. For the rural and agrarian context that forms our empirical backdrop, economic activity outside the formal sector accounts for a large fraction of the aggregate and not reflected in sub-national GDP accounts even when such accounts are available. At any rate, sub-national GDP estimates are not reliably computed each year and a reliable time series is difficult, if not impossible to construct – at least at village and sub-district levels that are our geographic units of analysis.

We use DMSP-OLS Night-time Lights Time Series produced by NGDC-NOAA (Version-4). We construct our data using averaged multiple-satellite stable cloud-free monthly composites with 30 arc-second resolution.¹⁶ We obtained geospatial boundaries of census-designated villages in Karnataka from the Karnataka State Remote Sensing Applications Centre (KSRSAC; Census of India 2011 village boundary revision) and computed aggregate luminosity for each village from 2001 to 2013. For sub-districts (*talukas*), we simply aggregated village level data so that urban centers are excluded. We used village population numbers from decennial census of 2001 and 2011 to obtain population numbers for non-census years with the assumption of uniform annual population growth between 2001 and 2013 – to obtain per-capita luminosity for all years.

Computation of per-capita luminosity at the village level is now well-established

¹⁵Of these 176 sub-districts, the Bengaluru-Urban sub-district that contains the capital and the largest city in the state is not part of our dataset as it contains no rural settlements and is fully urban.

¹⁶A detailed description of how the composites were constructed is available from NOAA – https://www.ngdc.noaa.gov/eog/gcv4_readme.txt (accessed April 26th 2018).

(Roychowdhury et al., 2012; Min et al., 2013). However, using such luminosity measures as indicators of economic performance can potentially be contested — night lights intensity might reflect village electrification as a proxy for state provisioning of public goods rather than economic output (Paik, 2013). There is also evidence that rural electrification does not always lead to short-term economic gains (Burlig and Preonas, 2016). However, in our specific empirical context — rural Karnataka — these concerns do not hold. First, more than 90% of all villages in our dataset were already electrified by 2001 — the base year in our analysis. Further, night-lights intensity in rural India is a reflection of hours of electricity available rather than mere grid connectivity (Chakravorty, Pelli and Ural Marchand, 2014); and there is a positive spillover from reliable and longer hours of electricity supply on economic growth, with higher non-farm enterprise income being the primary channel (Rao, 2013; Chakravorty, Pelli and Ural Marchand, 2014; Van de Walle et al., 2017). Empirically, we find significant variation in the number of hours of electricity across our sample of villages (the coefficient of variation is 0.45).

We use night lights intensity to construct two dependent variables at each geographic level (village and the sub-district) — per-capita luminosity growth between 2001 and 2013, and mean of per-capita luminosity for all years between 2001 and 2013. The first two columns of Table 2 contain elementary descriptive statistics for our two night light variables. The actual cardinal values for mean per-capita luminosity are not relevant as the distribution — luminosity is coded as a pixel value in NOAA’s composite `tiff` images of night lights. Figure 7 shows the geographic variation of these variables. For 1256 villages, luminosity growth variable is not defined as these villages were not electrified in the base year.

Table 2—: *Village-level Dependent Variables, Descriptive Statistics*

	Mean PC. Lum. (2001-13)	Per. PC. Lum. Gr. (2001-13)	HDI (2015)
Min.	0	-100	0.37
1st Qu.	0.01	10.81	34.28
Median	0.02	34.5	42.01
Mean	0.03	39.62	40.76
3rd Qu.	0.04	62.75	48.13
Max.	0.2	458.92	75.29

Note: The data in this table is from 25,732 villages. For the geographic spread of these variables refer to Figure 7.

C. Village-level Human Development Index

The third dependent variable that we use in addition to the two night lights based variables described above is the village-level human development index (HDI). In 2015, the government of Karnataka initiated a unique exercise – a first in India – to compute HDI for each one of the over 27,000 villages in the state. Following the United Nation Development Program (UNDP) methodology (UNDP, 2010) village-level HDI used information about standard of living as well as status of health and education. The standard of living component of HDI was proxied as the percentage of households in a village with access to modern cooking fuel, toilets, safe drinking water, electricity, *pucca* or durable housing, share of non-farm worker. The data from the census-scale survey (described below) was used to construct this standard of living component. Education and health components were constructed using this survey data in combination with administrative data from various government sources including the Planning Department, *Zilla Panchayats* or the district-level local government, Backward Classes welfare Department, Department of Rural Development and Panchayati Raj, Department of Health and Family Welfare, and the Department of Women and Child Development (Shivashankar and Prasad, 2015). We digitized the village-level summary tables in Shivashankar and Prasad (2015) and matched $\approx 27,000$ villages with the 2011 national census village codes – representing over 98% of all villages.¹⁷ We calculated sub-district level HDI as a census 2011 population weighted average of all villages in the sub-district.

We present elementary descriptive statistics for HDI in the last column of Table 2, and Figure 7 shows the geographic variation of HDI. The geographic spread of HDI shown in the figure is consistent with regional patterns of development described earlier in this section, and with the geographical distribution of our dependent variables constructed from night lights. The coastal districts have the highest HDI and north eastern districts are the least developed. As with night lights, we have shown quartiles rather than actual values to better contextualize our results from quantile regressions.

D. Census-scale Primary Survey Data on Caste

Even while caste is the principal axis of heterogeneity in agrarian India, systematic large- n data on elementary caste groups of contemporary India is scarce. The decennial census data categorizes the population into three broad aggregate categories – Scheduled Castes (SC), Scheduled Tribes (ST), and a residual others (OTH). The last time Indian census data collected detailed information on

¹⁷The HDI data in Shivashankar and Prasad (2015) is presented without the census village codes as it is a collation of a decentralized exercise coordinated by the local governments at each one of the 175 sub-districts in the state.

elementary *jati* categories was in 1931.¹⁸ Nationally representative surveys such as the National Sample Surveys (NSS), or the National Family Health Surveys (NFHS) that inform much of the large-scale empirical studies of India also collect aggregate caste categories with other backward castes (OBCs) added to the three-fold census division. While the India Human Development Survey (IHDS) collected *jati* information, the data has not been coded or standardized (Desai and Vanneman, 2005, 2011). The village-level rosters prepared as part of IHDS surveys collected caste information through key informants in a village rather than through a village census and the *jati* information is available for only the largest caste groups in a village.

The lack of detailed *jati*-level caste data has also constrained the diversity-development literature that uses Indian data. India-focussed studies have relied on the 1931 census data at the district level. For example, Banerjee and Somanathan (2007), the pioneering study in India, uses the 1931 census data mapped to electoral constituencies that are roughly (but not wholly) congruent with district boundaries.¹⁹ In this paper, we construct the first fully-coded census-scale *jati* data since 1931, and the first ever census-scale micro dataset to combine such detailed *jati*-level caste information with landholding data at the household-level.

In 2015, the Government of Karnataka conducted a census-scale survey (henceforth, GOKS15) that collected detailed household level information on caste, language, religion, landholding, and individual educational attainment. The raw data in GOKS15 contained over 2500 caste names in Kannada (the regional language in Karnataka) from over thirteen million households ($n = 13,255,421$). After eliminating orthographic duplicates and transliteration into English, we were left with 1641 *jati* names. We used several published historical, literary, anthropological, and administrative, accounts to code these 1641 caste names into 717 distinct locally endogamous groups (*inter alia*, Anantha Krishna Iyer and Nanjundayya, 1928; Thurston and Rangachari, 1975; Enthoven, 1990; Singh, 2002). We also used the colonial census records (1871-1931) from Mysore, Bombay and Madras as well as the administrative list of Scheduled Castes and Scheduled Tribes in developing our *jati* codebook. The final iteration of the codebook was vetted by social anthropologists of rural Karnataka including members of the independent advisory committee that advised GOKS15.²⁰

¹⁸Owing to the Second World War and the imminent end of colonial rule, the caste data from the 1941 census was not tabulated. Post colonial independent India has not collected caste data in any of its decennial census exercises starting 1951, except in 2011 when a companion census operation known as the Socio-Economic Caste Census (SECC 2011) was conducted. However, SECC 2011 *jati*-level data has not been coded and made public.

¹⁹Also see Banerjee, Iyer and Somanathan (2005).

²⁰We alone are responsible for any remaining errors. A full-length paper describing our *jati* coding will accompany public release of our codebook (expected, c. 2020). The empirical results reported here are robust to any reasonable alternate coding strategy. Cf. Section VI for these robustness tests.

The central objective of our coding exercise was to standardize *jati* names across the state and account for synonyms that arise partly from distinct dialects of the Kannada language spoken in different parts of the state. For example, *Agasa* and *Madiwala* are synonyms for the ‘washerman’ caste. For all religions except Hindus (the majority religion practiced by 78.7% of households in our dataset) and nomadic groups that practice Islam (accounting for less than 0.5% of all Muslim households in the dataset), we collapsed individual *jati* groups into a monolithic aggregate category. This coding of *jati* groups defines the **ModJati** or ‘modified *jati*’ variable in our final analysis dataset.

After the *jati* names from the raw data were coded into 717 **ModJati** groups, we mapped these groups to the eight-fold administrative categories used by the Government of Karnataka for state-level affirmative action and local government political quotas. The **ModMap** variable in our dataset represents this mapping between **ModJati** and the eight-fold administrative categories. Two of these eight categories – “SC,” and “ST” – are congruent with the corresponding categories in the three-fold census division. The remaining named categories correspond to administrative classification based on social and economic status of various caste groups. Every state in India has a state-specific backward classes classification in addition to an aggregate list of backward classes (administratively known as OBC or other backward classes) maintained by the federal government. While the central (federal) government classification is used for affirmative action in federal government jobs and institutions of higher education funded by the federal government, state-level classification is also used for determining quotas for various elected positions in the the local government in addition to being used for state-level affirmative action. Local political economy factors determine the exact nature of classification of backward castes at the state level. Thus while states like Karnataka and Andhra Pradesh follow a five-fold classification of backward social groups, states like Bihar (two-fold) and Tamil Nadu (three-fold) have fewer.²¹ At the state-level, these administratively constructed aggregate categories have wide-ranging currency and large sections of the society identifies themselves by these administrative categories – especially in the secular public realm even as endogamous *jati* categories are salient in personal matters.

We present the summary of **ModJati**–**ModMap** mapping in Table 3. Karnataka classifies socially and economically backward social groups into five administrative categories – “I,” “II A,” “II B,” “III A,” and “III B.” Category I includes all nomadic and semi-nomadic castes which are beyond the pale of the traditional caste hierarchy and caste groups that were historically nearly as persecuted as the erstwhile untouchable castes classified as Scheduled Castes (SC). This group

²¹Proceedings of the *Rajya Sabha*, written statement by Minister of State for Social Justice and Empowerment, 14 August, 2014.

also includes Scheduled Caste converts to Christianity who cannot access affirmative action quotas reserved by the federal government. Category II A consists of traditional occupational castes such as *Kuruba* (shepherd), *Agasa* (washer-men), *Devanga* (weaver), and *Kumbara* (potter) castes. Category II B includes all Muslims including nomadic groups that practice Islam. Landholding communities including “dominant castes” (Srinivas, 1959, 1994) constitute Category III. The subdivision of this category into III A and III B categories reflects both the regional patterns of distribution of dominant castes as well as the political economy of affirmative action quotas. Category III A consists mainly of sub-castes of *Vokkaligas*, *Reddys*, and *Kodavas* while III B classification includes sub-castes of *Lingayats*, *Marathas*, *Jains*, *Bunts*. The landed caste groups in III A and III B continue to be underrepresented in modern sectors of the economy even when they are relatively prosperous within the agrarian economy.

The final aggregation of ModMap administrative categories into the three-fold census of India taxonomy (SC, ST, and Others) was trivially accomplished by subsuming all categories except “SC” and “ST” groups into the residual category, “Others.” This final mapping also helps us benchmark our *jati* coding exercise with the 2011 census data (the latest decennial census). At the village-level, the correlation between fractionalization computed using our mapping and the census-level data is very high (Pearson coefficient ≈ 0.95). The census aggregation has been the most widely used caste variable in the Indian empirical context, and we use this specification of caste as a benchmark in all our models. The mapping of elementary *jati* into administrative and census categories is at the heart of one of the principal results from this paper — the impact of ethnic aggregation (MEUP) on econometric models of diversity and development.

E. Language and Religion Data

In addition to *jati*-level caste data, GOKS15 also contains detailed data on language and religion — the two commonly used heterogeneity axes in the diversity-development literature. We coded the raw data into six major religious categories, reported in Table 4. Our coding of the GOKS raw data follows the taxonomy used by the decennial census in India with two modifications for rural Karnataka — Sikhism and Zoroastrianism were merged into the residual “Others” category, and we collapsed all tribal religions that are not denominations of major organized religions into a composite “Tribal religion” category. Additionally, while the survey included a separate code for “Atheism,” we subsumed atheist households into the residual “Others” category. The distribution of religions affiliation as reported by GOKS15 is consistent with the national census data from 2011. GOKS15 also collected language data that was coded at source – we directly use the language coding in the original data (summarized in Table 5). Eleven different languages are spoken by at least 0.5 million people in rural Karnataka and

Table 3—: Summary of Caste Groups in Rural Karnataka

ModMap	Description	Households (millions)	ModJati, Species Count
ST	Scheduled Tribes	0.896	55
SC	Scheduled Castes	2.381	105
I	Nomadic and semi-nomadic castes	1.073	175
II A	Traditional occupational castes (eg., shepherds and washermen)	2.791	192
II B	Muslims (including nomadic groups that practice Islam)	1.431	9
III A	Landholding castes	1.865	60
III B	Landholding castes	2.085	27
OTH	All other castes	0.723	94

Note: Caste information is not available for 10,307 households ($\approx 0.05\%$ of households in our dataset).

Table 4—: Distribution of Religions in Rural Karnataka, GOKS15

Religion	Prop. of HH
Hindusim	78.73%
Islam	10.77%
Tribal	6.75%
Christianity	1.93%
Jainism	0.68%
Buddhism	0.016%
Others	0.0085%
<i>Not Reported</i>	1.12%

the top ten languages all count at least a million native speakers each.

While the primary diversity marker analyzed in this paper is caste, religion and language identities as additional diversity axes serve as important robustness checks for our main results. Unlike with caste, language and religion are not subject to the modifiable ethnic unit problem (MEUP) that we introduce in this paper. Thus, our empirical models using language and religion help empirically illustrate how MEUP combines with the familiar spatial aggregation problem (MAUP, or the modifiable areal unit problem) by providing a MAUP-only benchmark. As the most ubiquitous diversity variables in the literature, language and religion models also help us compare our results with extant evidence for the relationship between diversity and development. Even in a rural agrarian context where caste is the most important heterogeneity axis, anecdotal evidence suggests that ethnic strife is concentrated in regions with greater religious polarization and greater economic inequality along religious lines. This is especially true in coastal districts of the state (also the most prosperous) that have seen sustained mobilization of the majority Hindu population in the region (Assadi, 1999, 2002; Sayeed, 2016).

F. Landholding Data

In the context of agrarian India, the most important productive asset is the agricultural landholding of a household. GOKS15 is the first ever census-scale dataset that combines detailed landholding data with information on caste, language, and religion enabling the development of a high-resolution snapshot of ethnic inequality. Table 6 summarizes state-wide landholding patterns by administrative caste categories. The table reports landholding for only rural residents and does not include rural agricultural land owned by urban residents. The statewide aggregates presented in the table can mask significant regional variations in ethnic landholding patterns. In Figure 8, we present district-level distribution of agri-

Table 5—: Distribution of First Language in Rural Karnataka, GOKS15

Language	Prop. of HH
Kannada	68.70%
Urdu	9.39%
Telugu	5.31%
Marathi	3.19%
Tamil	3.12%
Tulu	2.65%
Hindi	1.64%
Konkani	1.47%
Malayalam	0.91%
Byari	0.65%
Kodava	0.25%
Arebhashe	0.16%
English	0.04%
Others	1.63%
<i>Not Reported</i>	0.89%

Table 6—: Agricultural Landholding by Administrative Caste Categories (ModMap)

ModMap	Landless Proportion	Median Acres	Mean Acres	Land Share
ST	51.21%	0.00	1.66	7.11%
SC	64.53%	0.00	1.01	11.49%
I	55.33%	0.00	1.53	7.82%
IIA	54.84%	0.00	1.72	22.87%
IIB	82.84%	0.00	0.59	4.05%
IIIA	48.78%	0.25	2.05	18.27%
IIIB	46.56%	0.50	2.57	25.56%
OTH	81.96%	0.00	0.82	2.84%

Note: Data from GOKS15 ($n = 13,255,421$ rural households). Agricultural land owned by urban households not included here.

cultural landholding across the thirty districts in Karnataka. Taken together, Table 6 and Figure 8 underscore the limitations of using large census aggregates for caste, or relying on historical data from the 1931 census. India’s agrarian society has changed significantly if not undergone a metamorphosis in the over eight intervening decades. Landholding across caste groups have become more fragmented as seen from the median and mean landholding sizes in Table 6. Over 80% of households who belong to the residual “others” category are landless despite being the most socially and educationally advanced social group – indicating the extent to which these groups, and especially landholders among these groups are no longer rural residents. The dominant landowning castes (categories III A and III B) own nearly 45% of all agricultural land held by rural residents of Karnataka. Unlike in many other parts of India, SC and ST groups (the most marginalized of social groups) have historically owned agricultural lands and limited land reforms in the 1970s have also contributed to these groups collectively owning just under a fifth of agricultural lands (Kohli, 1982, 1987; Srinivas and Panini, 1984). As seen from Table 3, each of the categories in the administrative caste taxonomy includes several elementary *jatis* so that there is great diversity in landholding patterns within any given administrative category. Within the SC category, there are caste groups that have historically held small lands and other groups that have been castigated to servicing the dominant farming castes.²²

The debate on diversity-development has largely developed independent of the inequality-development literature (Alesina, Michalopoulos and Papaioannou, 2016). Caste continues to be the principal cleavage in India’s agrarian society partly because an individual’s caste identity is strongly associated with her economic prospects (Anderson, 2011; Zacharias and Vakulabharanam, 2011; Iversen et al., 2014; Munshi, 2014). Local power relations in India’s agrarian society are often mediated by landholding patterns and a social group that controls this central agrarian asset can be locally dominant even when their ritual status on the caste hierarchy is not congruent with their economic status (Srinivas, 1959). In this paper, we exploit the availability of household landholding data along with detailed caste information to study how the political economy of development is mediated through both diversity and inequality channels. Besides explicitly computing ethnic inequality as the between-group component of landholding inequality, we also use landholding data to construct more meaningful diversity metrics that do not implicitly treat the ethnic distance as a constant between any two groups. In particular, we use landholding data to construct a polarization metric that uses ethnic landholding in a village as a proxy of localized ethnic distance.

²²Similar historical trajectories have been documented in other parts of India including in Uttar Pradesh, the largest state in India (Rawat, 2011).

G. Population Census, and other Administrative Data

We merged the village-level national census data from 2001 and 2011 with GOKS data. The Census Village Directory data contain information on the availability of basic public infrastructure in the village such as the distance to nearest town, number of hours of electricity, paved roads and national highway, the provision of education and health facilities among others. Our empirical models use these variables as controls (we create a PCA-based asset index of village-level public goods). The three category caste groups classification from the Primary Census Abstract have been used to create diversity measures as a complement to the more detailed GOKS data. Similarly, we used the sub-district level population from the census to create diversity metrics and other demographic indicators. The census population numbers were also used to compute per-capita luminosity as detailed above.

In addition to census data, we obtained detailed sub-district level meteorological data with information on annual rainfall, administrative history during the colonial period and agro-climatic classifications – all variables that we use as controls in our sub-district level regressions.

V. Empirical Models

We use the discussion in Section II to develop empirical models that are able to test the diversity-development relationship at various levels of ethnic and geographic aggregations using multiple ethnic diversity and ethnic inequality metrics. We construct models that simultaneously account for spatial scales (G), ethnic scales for caste (E), and measurement frameworks (M). We estimate our models at two subnational spatial scales — village, and sub-district (*taluka*); three aggregation levels of caste — elementary *Jati*, administrative map, and census map; and five different diversity measurement frameworks — fractionalization, discrete polarization, polarization, inequality, and ethnic inequality. Equation 4 records the framework that we use in developing our empirical models.

$$(4) \quad \left\{ \begin{array}{l} \underbrace{G = \{\text{Village, Sub-district}\}}_{\text{Geographic Aggregation}} \\ \underbrace{E = \{\text{ModJati, ModMap, Census2011}\}}_{\text{Ethnic Aggregation}} \\ \underbrace{M = \{\text{FRA, POL, ERPOL, MLD, EIQ}\}}_{\text{Measurement Framework}} \end{array} \right.$$

Our base family of OLS models is described in Equation 5.

$$(5) \quad Y_i^{(G)} = \alpha^{(G)} + \underbrace{\vec{\beta}^{(G)} \circ \vec{C}_i^{(G,E,M)}}_{\text{Caste}} + \underbrace{\vec{\gamma}^{(G)} \circ \vec{D}_i^{(G,M)}}_{\text{Lang, Religion, Land}} + \underbrace{\vec{\theta}^{(G)} \circ \vec{P}_i^{(G)}}_{\text{Controls}} + \epsilon_i^{(G)}$$

Each of our three dependent variables (mean per-capita luminosity, per-capita luminosity growth, and human development index) are modelled at village and sub-district levels. The caste diversity vector, \vec{C}_i , is aggregated in ethnic and geographic dimensions from a common household-level micro dataset with elementary caste (*jati*) information. The language, religion, and land class diversity variables do not have a ethnic aggregation dimension. The control vector \vec{P}_i at the village-level includes availability of public goods, share of irrigated land, literacy rates, distance of the village to the nearest town, presence of a highway, and hours of electricity during summers and winter days. For public goods provisioning, we create a PCA (principal component analysis) based cardinal index using ordinal information about presence or absence of a vector of public goods including primary and secondary schools, health facilities, water and drainage facilities, transport services, and banking infrastructure. We further control for the sub-district level fixed effects using dummy variables. In the sub-district level regressions, we control for literacy rate, agro-ecological zones, and colonial-era administrative histories. For growth regressions at both levels of geographic aggregation, we control for the intensity of nightlight in 2001 to account for any potential base effect. Before we proceed with a discussion of our regression models, we define the five different measurement frameworks (M in Equation 4) used in our models.

A. Diversity Measurement Frameworks

FRACTIONALIZATION

The ethnolinguistic fractionalization index (ELF) has been the most commonly used metric to measure diversity. First constructed using data from the *Atlas Narodov Mira* (Taylor and Hudson, 1972), the fractionalization index is related to the well-known Herfindahl-Hirschman Index (Hirschman, 1964). The fractionalization index, defined in Equation 6 is simply the Herfindahl-Hirschman Index subtracted from unity.

$$(6) \quad FRA_i^{(G,E)} = 1 - \left(\sum_{k \in i|_E} \pi_k^2 \right)$$

In Equation 6, π_k is the population share of subgroup k – the proportion of people who belong to a particular caste, religion, language, etc. in geographic unit i for a given diversity axis, E . The fractionalization index, FRA_i measures

the probability that two randomly selected individuals from i are drawn from two distinct subgroups in E . Given its relationship to the Herfindhal-Hirschman index, FRA_i is an increasing function of the number of subgroups of interest. In case of ethnic aggregation of caste that we study in this paper, for a given geographic aggregation:

$$(7) \quad FRA_i |_{E=Jati} \geq FRA_i |_{E=Admin-Map} \geq FRA_i |_{E=Census-Map}$$

Beyond this dependence of the fractionalization index on the number of subgroups, it also does not discriminate “politically relevant ethnic groups” (Posner, 2004) that are relevant for diversity-development relationships and those that are not. We address these shortcomings by complementing the fractionalization index with polarization and ethnic inequality measures.

POLARIZATION METRICS

The fractionalization index is a species-diversity measure that does not account for the distribution of relative sizes of heterogeneous social groups. However, tension and strife between social groups generally tend to increase when the share of the minority group increases (Horowitz, 1985). In their seminal paper, Esteban and Ray (1994) introduced an axiomatic framework to model such subgroup tensions using a family of polarization metrics. In our main empirical models, we include a discrete polarization metric (Montalvo and Reynal-Querol, 2005b, 2008) in addition to the fractionalization index. The discrete polarization metric is defined in Equation 8.

$$(8) \quad POL_i^{(G,E)} = 1 - \left(\sum_{k \in i|_E} \left\{ \left(\frac{0.5 - \pi_k}{0.5} \right)^2 \pi_k \right\} \right)$$

In Equation 8, π_k is once again the population share of group k in geographic unit i for some given diversity axis. Unlike $FRACT_i$ which is an increasing function of the number of distinct social groups, POL_i is non-monotonic with respect to number of groups and attains its maxima with two equal sized groups — consistent with the polarization axioms of Esteban and Ray (1994). Thus, the fixed relationship under ethnic aggregation that characterizes fractionalization index (Equation 7) breaks down for POL_i :

$$(9) \quad POL_i |_{E=Jati} \gtrless POL_i |_{E=Admin-Map} \gtrless POL_i |_{E=Census-Map}$$

Figure 9 plots this relationship for village-level discrete polarization along caste lines at two different ethnic scales – elementary *jati* and the three-fold census classification. We show in this paper that this lack of a predictable relationship between polarization at different levels of ethnic aggregation is a significant econo-

metric challenge for empirical models of diversity-development that has hitherto not been addressed adequately.

The discrete polarization metric in Equation 8 makes the simplifying assumption that the “distance” between any two subgroups is identical. This assumption is empirically untenable in the context of agrarian India – especially in the context of hierarchical caste as the diversity axis. In Equation 10a, we reproduce the original formulation of polarization (Esteban and Ray, 1994, eq-3).

$$(10a) \quad P_i^{(G,E)}(\pi, y) = K \sum_{j \in i|_E} \sum_{k \in i|_E, k \neq j} \pi_j^{1+\alpha} \pi_k |y_j - y_k|$$

$$(10b) \quad |y_j - y_k| = \mathbf{L}_{jk} = f\left(\vec{L}_H|_{H \in j}, \vec{L}_H|_{H \in k}\right)$$

$$(10c) \quad ERPOL_i^{(G,E)} = 4 \sum_{j \in i|_E} \sum_{k \in i|_E, k \neq j} \pi_j^2 \pi_k \mathbf{L}_{jk}$$

In Equation 10a, π is once again the population shares of respective groups j and k ; y is the population characteristic of interest (income for example); and $\alpha \in (0, 1.6]$, $K > 0$ are arbitrary constants. With these bounds for constants K and α , Equation 10a is the only specification that satisfies the polarization axioms of Esteban and Ray (1994). We use household agricultural landholding – the most important endowment in an agrarian regime – to account for the distance term in Equation 10a, $|y_j - y_k|$. Equation 10b defines this distance (\mathbf{L}_{jk}) as a function of vector of households (\vec{L}_H) belonging to social groups of interest, j and k . In our main regressions, we operationalize equation 10b by setting f to be the difference in average landholding of each social group in a village. Using $K = 4$, $\alpha = 1$ as parameter values in equation 10a – the constant values used in the discrete polarization metric – we define our operational polarization metric, $ERPOL_i$ in equation 10c.

For all three diversity metrics introduced here – fractionalization, discrete polarization, and polarization, we also use language and religion as diversity axes in addition to three different aggregations of caste.

INEQUALITY, AND ETHNIC INEQUALITY

There has been little intersection between inequality-development and diversity-development literatures (Alesina, Michalopoulos and Papaioannou, 2016). In this paper, we offer a corrective by using the detailed household landholding data described above to explore the joint effects of diversity and inequality. In order for

Table 7—: Regression Model Taxonomy

	Sub-District Regressions	Village Regressions
Dependent Variable		
Per-capita Luminosity Growth (2001-11)	Suite of 11 Models	Suite of 11 Models
Mean Per-capita Luminosity (2001-11)	Suite of 11 Models	Suite of 11 Models
HDI, 2015	Suite of 11 Models	Suite of 11 Models

our ethnic inequality measure to be “path independent” (Foster and Shneyerov, 2000; Shorrocks and Wan, 2005), we use the additively decomposable mean log deviation (MLD) as the metric of choice.

$$(11a) \quad MLD_i^{(G)} = \frac{1}{N_i} \sum_{H \in i} \ln \frac{\bar{L}_{H \in i}}{L_H}$$

$$(11b) \quad N_i = \sum_{H \in i} H$$

The MLD metric defined in 11 uses standard notation such that $\bar{L}_{H \in i}$ is the mean landholding in spatial unit i and N_i is the number of households in i .

The MLD metric is perfectly sub-group decomposable and this decomposition is evaluated in equation 12.

$$(12) \quad MLD_i^{(G,E)} = \underbrace{\sum_{k \in i|_E} \pi_k \cdot MLD_{ik}}_{\text{Within Group}} + \underbrace{\sum_{k \in i|_E} \pi_k \cdot \ln \left(\frac{\bar{L}_{H \in i}}{\bar{L}_{H \in i, k}} \right)}_{\substack{\text{Between Groups} \\ \text{Ethnic Inequality } (EIQ_i^{(G,E)})}}$$

Our regression models use the “between groups” component of the decomposition, as ethnic inequality (EIQ_i). We compute ethnic inequality for two different aggregations of caste as well as for religion.

B. Main Regression Models and Results

Using the base OLS framework of equation 5, we develop and test sixty-six principal models across different combinations of dependent variables, diversity axes, ethnic-geographic aggregation, and measurement frameworks. Table 7 summarizes the taxonomic structure of our OLS models. The first model in our common suite of 11 models is a ‘kitchen-sink’ model with all 18 diversity metrics (across different diversity axes and measurement frameworks) as independent variables. The 10 other models test the diversity-development relationship for individual diversity axes (and at different aggregation levels for caste) – landholding, lan-

guage, religion, and caste. Our main results are not altered by models that use two or three diversity axes in a model instead of the “one, or all” models that are part of the 11-regression suite described in the table.²³ In addition to OLS, we used right-hand-side specifications from equation 5 in quantile regression models to estimate the 25%, 50%, and 75% quantiles for each of our three dependent variables. Across, OLS and quantile regressions models, we find no stable evidence for neither the diversity-debit hypothesis, nor the diversity-dividend hypothesis. Taken together, we find support for the argument that both negative and positive associations between diversity and development are mere statistical artifacts.

OLS RESULTS

The first dependent variable that we model is the per-capita luminosity growth between 2001 and 2013. The village-level model results are presented in Table 12, and the sub-district level models are presented in Table 13. The regression results are neither stable across model specifications that use different aggregations of caste (our primary diversity variable), nor across the two spatial scales. While at the village-level we see broad support for the diversity-debit hypothesis, the negative association largely disappears at the sub-district level. Even at the village-level, ethnic inequality is positively associated with growth in per-capita luminosity – an effect that is not statistically significant at the sub-district level. The effect of ethnic aggregation is seen in how the coefficient on fractionalization goes from negative and significant (elementary *jati* categories) to positive and insignificant (administrative mapping of caste).

We also use the DMSP-derived “night lights” data to construct another dependent variable – average per-capita luminosity between 2001 and 2013. Table 14 and Table 15 record the regression results with this changed dependent variable specification but with an unaltered right-hand-side. The coefficients on these models largely mirror the growth regressions. The ethnic aggregation effect is even more pronounced here so that the coefficient on caste fractionalization goes from negative and significant to positive and significant when the ethnic scale changes from elementary *jati* categories to administrative categories at the village level. The regressions coefficients are once again not robust to a change in geographic scale.

Our final dependent variable is the Human Development Index (HDI) computed at the village level. The regression results with HDI as the dependent variable are presented in Table 16 and Table 17 – village and sub-district regressions respectively. The evidence for diversity-debit that we find in several models at the

²³In the interest of parsimonious exposition, we do not report these additional results in the main paper. Replication code available for these additional models upon request.

village level regressions with nights lights as our dependent variable largely disappears and we instead find limited evidence for diversity-dividend hypothesis. Geographic scale continues to drive the results so that we find limited evidence for diversity-debit at the sub-district level.

The central “result” from our OLS models is that there is no systematic relationship between diversity and development. Both the direction and strength of any such association is contingent on particular ethnic and geographic scales. We are unable to rule out the possibility that both diversity-debit and diversity-credit relationships are spurious statistical artifacts. The instability of coefficient signs is summarized in Table 18. The first three columns of the table are from village-level regressions, and the last three columns are from sub-district regressions. We report the sign (and if the relevant coefficient is significant) on fractionalization, polarization, and ethnic inequality terms in respective regressions discussed above. The only stable association seen from the table is a negative relationship between caste polarization (at all three ethnic aggregations, and at both geographic aggregations) and development. However, this negative association is not statistically significant at the sub-district level. Similarly, the positive association between religious, linguistic diversity and development is not statistically robust across dependent variable specification or geographic scales.

QUANTILE REGRESSION RESULTS

In addition to estimating the conditional means using the OLS framework described in equation 5, we also used quantile regressions to estimate the 25%, 50%, and 75% quantiles for our two night-lights based dependent variables. We used the exact same right-hand-side specification described in equation 5, including respective controls at the village level. Our quantile regressions help test the contention that the relationship between diversity and development is not monotonic with diversity-debit prevalent at lower levels of development (Alesina and La Ferrara, 2005). The results of our quantile regressions are summarized in Table 19. In the interest of expository clarity, we present only village-level regressions in the summary so as to retain the focus on coefficient stability (or lack thereof) across the three quantiles. However, our quantile regressions, like the OLS counterparts are also not robust across geographic aggregation.²⁴

As seen from Table 19, our data does not support the hypothesis that the diversity-debit dominates at lower levels of development and diversity credit is more pronounced at higher levels of development. For a given quantile, the regression coefficients are not robust across different ethnic aggregations of caste. The estimated coefficients are unstable – in sign as well as statistical significance

²⁴The replication code for generating detailed quantile regression tables akin the OLS tables available with the authors.

– regardless of the choice of diversity metrics, or even the levels of inequality between various social or ethnic groups. The impact of ethnic aggregation of caste is clearly seen in Table 19. When the three-fold census category aggregation of caste as the diversity axis, the diversity debit relationship is stable when mean per-capita luminosity is the dependent variable. However, this relationship is not robust to alternate specifications of the dependent variable – thus diversity debit relationship vanishes when the dependent variable is changed to growth in per-capita luminosity. Census caste categories are large aggregates that mask more variation than they reveal, leading to misleading statistically significant associations.

ROBUSTNESS OF DIVERSITY-DEVELOPMENT MODELS

Table 8 summarizes the instability observed in diversity-development association across OLS and Quantile regressions. The entries in table are “*ceteris paribus*” in the sense that each column of the table reports coefficient stability on a single dimension. The table reports a coefficient as being stable (represented by a “✓” in the table) if at least 80% of the relevant models have the same significant coefficient sign, or are insignificant in 80% of the models. When one of these

Table 8—: Robustness of Diversity-Development Models

	Dependent Variable	Ethnic Agg.	Geographic Agg.	Measurement Framework	Quantiles
<i>Caste</i>	✗	✗	✗	✗	✗
<i>Language</i>	✗	NA	✗	✗	✗
<i>Religion</i>	✓	NA	✗	✗	✗
<i>Landholding</i>	✗	NA	✗	✗	✗

Note: See main text for details

very modest requirements for robustness is not met, Table 8 uses a “✗” to indicate coefficient instability, across four different diversity axes that we use in our models. This table distills the essence of regression instability described in Table 18 (OLS instability) and Table 19 (Quantile Regression instability). As seen from the table, even with a modest bar for robustness, the association between diversity and development is unstable across all the five dimensions – dependent variable specification, ethnic aggregation, geographic aggregation, measurement framework, and conditional quantiles estimations.

The coefficient on caste, our principal diversity axis is unstable across all dimensions. For example, a change in dependent variable specification at given levels of geographic aggregation, ethnic scale and the diversity metric is enough to change the sign on caste. Similarly, the coefficient on landholding is unstable across all dimensions. The coefficient on religion is unstable on all dimensions but narrowly meets our modest criteria so that it is “stable” across different specifi-

cations of the dependent variable for given ethnic-geographic scale, and diversity metric. Our quantile regressions are similarly unstable so that we are unable to detect either a monotonic relationship between development and diversity across all quantiles, or confirm the non-monotonicity hypothesis (with diversity debit at lower levels of development, and diversity credit at higher levels of development). Thus, Table 8 shows why both diversity-debit as well as diversity-credit hypotheses are likely statistical artefacts. In the next section, we provide a diagnosis, and develop a framework to account for the regression instability that we observe.

VI. MEUP and MAUP

VALIDATING THEORIES IN THE ETHNIC-GEOGRAPHIC CONTINUUM

We formally define the ethnic-geographic continuum to account for the modifiable ethnic unit problem (MEUP) and its intersection with the spatial MAUP introduced in Section II. A formal definition helps underscore the point that ethnic aggregation has the same logical structure as the more familiar geographic aggregation. The geographic space contains multiple political and administrative aggregations – villages, sub-districts, districts, states, etc. Consider a geographic space \mathbf{G} with n individuals (or households, or whatever is the elementary unit in an empirical analysis).

$$(13) \quad \mathbf{G} = \{g^0, \dots, g^i, \dots\}$$

Equation 13 shows that are an arbitrary number of geographic aggregates that can be constructed by aggregating n elementary units into some aggregation $g^{i \neq 0}$. Equation 13 defines the geographic space \mathbf{G} as a collection of all possible such aggregations including g^0 that is trivially the collection of n elementary units – formally represented by Equation 14.

$$(14) \quad \mathbf{G} = \left(\begin{array}{c} \underbrace{g^0 = \{g_1^0, \dots, g_n^0\}}_{n \text{ Elementary Units}} \\ \vdots \\ \underbrace{g^{i \neq 0} = \{g_1^i, \dots, g_{k < n}^i\}}_{k < n \text{ Aggregate Units}} \\ \vdots \end{array} \right)$$

The modifiable areal unit problem (MAUP) stems from the fact that the choice of i – the geographic unit of analysis – is arbitrary, and that empirical models are

sensitive to this choice of i . Aggregation changes the number of observations used in empirical analysis. A household level analysis will use n observations while at some arbitrary aggregation, i , the n households are aggregated into $k < n$ units.

Before we define the ethnic-geographic continuum, we need a formal definition of the ethnic space, \mathbf{E} , as a counterpart of \mathbf{G} . For some ethnic variable – caste for example – the n households in Equation 14 can return $m \leq n$ elementary self-identified caste groups. We can now define the ethnic space \mathbf{E} as:

$$(15) \quad \mathbf{E} = \left(\begin{array}{c} \underbrace{e^0 = \{e_1^0, \dots, e_m^0\}}_{m \leq n \text{ Elementary Ethnic Groups}} \\ \vdots \\ \underbrace{e^{j \neq 0} = \{e_1^j, \dots, e_{l < m}^j\}}_{l < m \text{ Aggregate Ethnic Groups}} \\ \vdots \end{array} \right)$$

The structural similarity between how we have defined \mathbf{E} and \mathbf{G} makes it clear that the choice of a particular ethnic aggregation, j in empirical analysis is as arbitrary as the choice of geographic aggregation, i . An empirical model with elementary ethnic units will use $m \leq n$ observations; and at the arbitrary ethnic aggregation j , these m groups are aggregated into $l < m$ groups.

Equations 14 and 15 show \mathbf{G} and \mathbf{E} as having an arbitrary number of geographic and ethnic aggregations respectively. For the purposes of studying the impact of ethnic and geographic aggregations on empirical models, it is only fair to assume that both \mathbf{G} and \mathbf{E} have a countably finite number of aggregations, N_G and N_E respectively. The ethnic-geographic continuum, \mathbf{C} is simply a set of all the $(N_G \cdot N_E)$ 2-tuples from the Cartesian product of \mathbf{G} and \mathbf{E} .

$$(16) \quad \mathbf{C} = \mathbf{G} \times \mathbf{E} = \{(x, y) | x \in \mathbf{G}, y \in \mathbf{E}\}$$

A theory or a hypothesis (for example, the diversity debit hypothesis) is aggregation independent only if it can be validated for all the ordered pairs in \mathbf{C} – with each ordered pair representing different combinations of ethnic and geographic aggregation. Conceivably, a hypothesis that is not sustained for all combinations of ethnic *and* geographic aggregations might however be valid at each geographic (ethnic) aggregation but not at all ethnic (geographic) aggregations. Table 9 presents a formal taxonomy for possible theory categories in the ethnic-geographic continuum based on the extent of their validity. A theory belongs to

Table 9—: *Theory Taxonomy for the Ethnic-Geographic Continuum*

Theory	Valid across G	Valid across E
${}^G\mathbb{T}_E$	✓	✓
${}^G\mathbb{T}$	✓	✗
\mathbb{T}_E	✗	✓
\mathbb{T}	✗	✗

${}^G\mathbb{T}_E$ if it is universally valid at all levels of geographic and ethnic aggregations. A theory belongs to \mathbb{T} if it is neither valid at all geographic aggregations, nor at all ethnic aggregations. Finally, ${}^G\mathbb{T}$ (\mathbb{T}_E) represents theories that are valid across all geographic (ethnic) aggregations.²⁵

One of the goals of empirical research is – or, at any rate should be – to classify theories like the diversity-debit theory, or the diversity-credit theory into one of the four categories of the taxonomy in Table 9. Our empirical results show that both diversity-debit and diversity-credit are at best \mathbb{T} -theories – valid only at certain specific geographic and ethnic aggregations. The diversity-development theory can be formally written as set of competing hypotheses in the ethnic-geographic continuum, **C**.

$$(17) \quad \left. \begin{array}{l} \text{Null Hypothesis, } H_0 : \frac{\partial Y}{\partial(M(X))} = 0 \\ \text{Diversity Debit, } H_1 : \frac{\partial Y}{\partial(M(X))} < 0 \\ \text{Diversity Credit, } H_2 : \frac{\partial Y}{\partial(M(X))} > 0 \end{array} \right\} \ni M(X) \in \mathbf{C}, Y \in \mathbf{G}$$

In Equation 17 Y is the development outcome of interest, and $M(X)$ is a measure of ethnic diversity – M is a diversity metric like fractionalization, polarization, etc. (the empirical specification can include more than one diversity metric, like the models in Section V).

We represent the diversity development theory from Equation 17 as a Theory

²⁵Trivially, any theory that is ${}^G\mathbb{T}_E$ is also both ${}^G\mathbb{T}$ and \mathbb{T}_E .

Validity Matrix (\mathbf{T}) in equation 18 below.

$$(18) \quad \mathbf{T} = \begin{matrix} & & e^0 & \dots & e^j & \dots \\ & g^0 & \begin{pmatrix} {}^0T_0 & \dots & {}^1T_j & \dots \end{pmatrix} \\ & \vdots & \vdots & \ddots & \vdots & \ddots \\ & g^i & \begin{pmatrix} {}^iT_1 & \dots & {}^iT_j & \dots \end{pmatrix} \\ & \vdots & \vdots & \ddots & \vdots & \ddots \end{matrix}$$

The generic entry in the theory validity matrix, ${}^iT_j \in \mathbf{T}$ represents the relationship between diversity and development at geographic aggregation $g^i \in \mathbf{G}$, and the ethnic aggregation $e^j \in \mathbf{E}$.

$$(19) \quad {}^iT_j \in \mathbf{T}|_{i,j=0,\dots} = \begin{cases} 0, & \forall (g_i, e_j) \in \mathbf{C} \text{ where } H_0 \text{ cannot be rejected} \\ -1, & \forall (g_i, e_j) \in \mathbf{C} \text{ where } H_1 \text{ (diversity debit) is sustained} \\ 1, & \forall (g_i, e_j) \in \mathbf{C} \text{ where } H_2 \text{ (diversity credit) is sustained} \end{cases}$$

If there are N_G geographic aggregations and N_E ethnic aggregations, the theory validity matrix has the dimensions of $(N_G \times N_E)$; and there are $3^{(N_G \cdot N_E)}$ such matrices (for each entry in the matrix can take on any of the three values in Equation 19).

As an illustration of the family of theory validity matrices, we consider a hypothetical case with $N_G = N_E = 3$. Of the possible $3^9 (\approx 20,000)$ (3×3) theory validity matrices that can be constructed in this case, we illustrate 12 such matrices in Table 10. The table shows how there are only three diversity validity matrices (irrespective of the number of ethnic or geographic aggregations) that correspond to a theory that is valid across all levels of ethnic and geographic aggregations. The top row of the table contains the three matrices in our example with $N_G = N_E = 3$. A theory that posits a monotonic relationship between diversity and development can be a ${}^G\mathbb{T}_E$ theory in three different ways – each corresponding to one of the three possibilities listed in equations (17) and (19). A diversity-debit theory, diversity-credit theory, or a even a null relationship can all be ${}^G\mathbb{T}_E$ theories that are valid at every combination of ethnic and geographic aggregations.

If ${}^G\mathbb{T}_E$ represents the most robust category of theories in the ethnic-geographic continuum, the next two rows of Table 10 illustrate the theory validity matrices that correspond to ${}^G\mathbb{T}$ theories and \mathbb{T}_E theories respectively. These theories are valid at either all of the geographic aggregations (for at least one level of ethnic aggregation), or are valid at all ethnic aggregations (at least at one geographic scale). From an empirical perspective, ${}^G\mathbb{T}$ and \mathbb{T}_E theories are more salient. In-

Table 10—: *Theory Validity Matrices: Examples*

Diversity Credit, ${}^G\mathbb{T}_{\mathbb{E}}$ $\mathbf{T} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$	Diversity Debit, ${}^G\mathbb{T}_{\mathbb{E}}$ $\mathbf{T} = \begin{pmatrix} -1 & -1 & -1 \\ -1 & -1 & -1 \\ -1 & -1 & -1 \end{pmatrix}$	Null Sustained, ${}^G\mathbb{T}_{\mathbb{E}}$ $\mathbf{T} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$
Diversity Credit, ${}^G\mathbb{T}$ $\mathbf{T} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}$	Diversity Debit, ${}^G\mathbb{T}$ $\mathbf{T} = \begin{pmatrix} -1 & 1 & 0 \\ -1 & 0 & -1 \\ -1 & -1 & 1 \end{pmatrix}$	Null Sustained, ${}^G\mathbb{T}$ $\mathbf{T} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix}$
Diversity Credit, $\mathbb{T}_{\mathbb{E}}$ $\mathbf{T} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix}$	Diversity Debit, $\mathbb{T}_{\mathbb{E}}$ $\mathbf{T} = \begin{pmatrix} -1 & -1 & -1 \\ 1 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix}$	Null Sustained, $\mathbb{T}_{\mathbb{E}}$ $\mathbf{T} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix}$
Statistical Artifact, $\cancel{\mathbb{T}}$ $\mathbf{T} = \begin{pmatrix} -1 & 1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix}$	Statistical Artifact, $\cancel{\mathbb{T}}$ $\mathbf{T} = \begin{pmatrix} 1 & 0 & -1 \\ 0 & -1 & 1 \\ -1 & 1 & 0 \end{pmatrix}$	Statistical Artifact, $\cancel{\mathbb{T}}$ $\mathbf{T} = \begin{pmatrix} 0 & -1 & 1 \\ -1 & 1 & 0 \\ 1 & 0 & -1 \end{pmatrix}$

deed, much of the extant evidence for both diversity-debit and diversity-credit hypotheses is presented such theories even if only by extrapolation. When a diversity-development theory is neither a ${}^G\mathbb{T}$ -theory nor a $\mathbb{T}_{\mathbb{E}}$ -theory, the ‘null’ explanation is a likely one of spurious association between diversity and development at particular geographic or ethnic scales. The last row of Table 10 shows three theory validity matrices that are consistent with the interpretation of the relationship between diversity and development as no more than an artifact. The interpretation of the association between diversity and development when the empirical results correspond to theories consistent with ${}^G\mathbb{T}$ and/or $\mathbb{T}_{\mathbb{E}}$ is non-trivial. Any non-monotonic association between diversity and development (the key feature of any ${}^G\mathbb{T}$ and/or $\mathbb{T}_{\mathbb{E}}$ theories) has to be consistent with political, historical, and sociological explanation (Gerring et al., 2015; Singh and vom Hau, 2016). In the absence of such arguments that can be sustained, it is prudent to interpret even results consistent with ${}^G\mathbb{T}$ or $\mathbb{T}_{\mathbb{E}}$ theories as statistical artifacts. If theoretically, the diversity-development linkage is not circumscribed by ethnic or geographic scales, a scale-specific empirical finding must be interpreted with skepticism.

THE DIVERSITY CONTEXT

Implicit in our taxonomy of theories in Table 9 is the assumption that even when a theory is tested at specific levels of geographic or ethnic aggregation, the empirical sample used for testing the theory is representative of those aggregations. For example, if the diversity-debit theory is being tested geographic and ethnic aggregations $(g^i, e^j) \in \mathbf{C}$, we would ideally test the theory across the universe of values in g^i and e^j . From Equation 14 and Equation 15, there are k geographic units corresponding to geographic aggregation level i , and l ethnic groups at ethnic aggregation level j . A census-scale testing of a theory (like we have done in this paper) is not always feasible and the established strategy in empirical work is to test the theory on a sample of representative units from geographic aggregation g^i . In the ethnic-geographic continuum, this strategy is fraught for two reasons. First, and somewhat trivially, for the sampled units from i be ‘representative,’ it must also be representative of the ethnic groups in j . In theory, there might not exist a stratified random sampling strategy – the workhorse sampling strategy in empirical research – that is able to achieve sufficient statistical power.

Second, and more crucially, even when a stratified random sample is able to achieve ethnic representation with statistical power, achieving representation on the universe of values assumed by M (the diversity metric of choice, *cf.* equation 17) is not guaranteed. For example, consider an empirical model testing the diversity debit hypothesis that includes fractionalization and discrete polarization as diversity metrics. If the model was being tested at a geographic aggregation (g^i) corresponding to a district, it is entirely possible that an arbitrarily large sample or even a census-scale sample might not span the range of values taken by both fractionalization and polarization. A similar argument can be made about how even a census-scale sample at a particular ethnic aggregation (e^j) might not circumscribe the full range of theoretical values that a diversity metric M can assume. If $(\mathbf{M} \in \mathbb{R})$ is the range of all values that a metric M can take, we define the ideal theory, \mathbb{T}^* as:

$$(20) \quad \mathbb{T}^* \ni \begin{cases} \mathbb{T}^* \in \mathbf{G}\mathbb{T}_{\mathbf{E}} \\ \mathbb{T}^* \text{ is valid for all values in } \mathbf{M} \end{cases}$$

In Figure 10, we illustrate how the hypotheses in Equation 17 straddle across multiple geographic and ethnic aggregations, and diversity metrics. The figure shows the multiple diversity contexts faced by a household embedded in the ethnic-geographic continuum. There are many ‘diversity forcing functions’ faced by a household, each corresponding to multiple combinations of ethnic and geographic aggregations that are possible. Depending on the context (for example, public good provisioning or economic development), one, more, or even none of the ethnic and geographic aggregations are politically salient.

For any given level of ethnic-geographic aggregation, one or more diversity metrics (M) might be politically relevant. Three different diversity models that we have discussed in this paper (fractionalization, discrete polarization, and polarization) are represented in the figure. The diversity numbers in the matrices of this figure are for a single randomly selected rural household from the dataset used in this paper ($n \approx 13.25$ million). Figure 10 describes ethnic aggregation of caste in three different geographic contexts — village, sub-district, and district. **ModJati** is the elementary ethnic group that is aggregated into administrative categories (**ModMap**) or into one of the three census categories (**CensusMap**). Similarly, the village is the elementary geographic category that can be aggregated into a sub-district (*Taluk*), or a district. All households in a given village have identical diversity contexts. In general, if there are N_E levels of ethnic aggregations, N_G levels of geographic aggregations, and N_M different diversity metrics, there are $(N_E \cdot N_G \cdot N_M)$ possible diversity indices — with all, some, or none of them being politically salient depending on the context.

The theory validity matrix (**M**) that we introduced in Equation 18 must be suitably modified to account for multiple diversity measurement frameworks, and to be consistent with the ideal theory (\mathbb{T}^*) introduced in Equation 20. A \mathbb{T}^* theory must not only be consistent across all levels of ethnic and geographic aggregation but also across multiple diversity metrics, and the range of values assumed by these metrics. We now need a rank-3 tensor-like object to represent the equivalent theory validity information, and Table 10 will need an additional row to account for \mathbb{T}^* (all the cells of the table are now rank-3 tensors rather than a matrix). Our empirical results in Section V show that besides geographic and ethnic units of analysis, the relationship between diversity and development is also sensitive to particular diversity measurement framework used. In our empirical models we have used four different metrics – fractionalization, discrete polarization, polarization, and ethnic inequality – to show why neither diversity debit not diversity credit is a \mathbb{T}^* theory.

ETHNIC-GEOGRAPHIC CONTINUUM AND THE FRA_i - POL_i MAP

We use the two of the four diversity metrics that we have used in our empirical models – fractionalization (FRA_i) and discrete polarization (POL_i) – to illustrate why a diversity metric’s empirical range is limited by particular ethnic and geographic aggregations. In order to test if a diversity-development hypothesis is a \mathbb{T}^* theory, empirical models must be able to capture the range of possible theoretical values that can be assumed by the diversity metric being used to measure diversity. Both fractionalization and discrete polarization have identical theoretical bounds — $FRA_i, POL_i \in [0, 1] \forall i$. While it might be possible to span this range individually for both fractionalization and polarization at sev-

eral specific ethnic-geographic aggregations in \mathbf{C} , the true import of the diversity context framework we have introduced is the question of joint distribution – in the present example, the bivariate distribution of fractionalization and discrete polarization.

The theoretical relationship between fractionalization and discrete polarization is non-monotonic and has also been observed empirically (Montalvo and Reynal-Querol, 2003, 2005a). At low levels of polarization, the relationship between fractionalization and polarization is linear – with a positive correlation at low levels of fractionalization and a negative correlation at higher level of fractionalization.²⁶ At high levels of polarization, the relationship between fractionalization and polarization is indeterminate. The non-linear relationship between fractionalization and polarization is also what accounts for the non-monotonic relationship between polarization and conflict – polarization is a relevant predictor of ethnic conflict only with ethnic fractionalization above a certain threshold level (Bleaney and Dimico, 2017).

In Figure 11, we present the relationship between fractionalization and polarization at different levels of ethnic and geographic aggregations. The top panels (A–C) present the fractionalization polarization map at the village level ($n = 26,890$) and the bottom panels (D–F) present this map at the sub-district level ($n = 175$). At both these geographic aggregations, we present the fractionalization–polarization map at three different ethnic aggregations of caste — elementary *jati* with 717 subgroups, administrative mapping of caste with 8 subgroups, and finally the three-fold census category mapping of caste. As described in Section IV, both the village-level dataset and the sub-district level dataset are constructed from a common household-level micro dataset with ≈ 13.25 million households. Figure 11 illustrates how both ethnic aggregation and geographic aggregation impact that joint distribution of fractionalization and polarization. With elementary *jati* categories at the village-level, the empirical fractionalization–polarization map covers the full range of the theoretically expected joint distribution. At the village level, using the three-fold census categories (the mainstay of much empirical work around caste in India) we see the non-monotonic relationship between fractionalization and polarization reduced to a linear map. Even at the intermediate level of ethnic aggregation, not surprisingly, we lose high-fractionalization portion of the map.²⁷

The transformation of the fractionalization–polarization map is even more dramatic when we move from village to the sub-district level (the bottom three

²⁶In this section, while we use “polarization” as a short-hand for discrete polarization (in the interest of parsimonious exposition), we always refer to the discrete polarization metric (POL_i) and not true polarization ($ERPOL_i$).

²⁷For a discussion on the relationship of fractionalization and discrete polarization metric to number of subgroups, cf. equations(7),(9), and the accompanying discussion.

panels in Figure 11). The standard unit of subnational empirical work in India is a district (in our dataset, the 175 sub-districts are aggregated into 30 districts), and cannot account for the full range of theoretical values that a diversity metric can assume. The comparison of the fractionalization-polarization map between village and the sub-district for elementary *jati* categories is instructive and points to why the diversity-debit, or diversity-credit hypothesis being a unit of analysis artifact cannot be ruled out. Even with 717 self-reported elementary caste groups, the geographic unit of analysis will determine if a diversity-development theory is being tested using the map in panel-A, or alternatively, panel-D in Figure 11. A diversity-development theory cannot be confirmed as a T^* theory at the sub-district level even if such a theory actually existed.

To underscore the fact that the metamorphosis of the fractionalization polarization map is a product of both geographic and ethnic aggregation, we present similar maps for religion, language, and economic class (measured in terms of census-defined landholding categories) in Figure 12. Unlike caste, we do not have to account for ethnic aggregation and can empirically study the ‘uncontaminated’ impact of geographic aggregation. As seen from Figure 12, changing the unit of analysis from village to the sub-district does not change the fractionalization-polarization map as starkly as in the case of caste – as seen by LOESS fits across religion, language, and land class. Most significantly, the non-monotonic relationship between fractionalization and polarization in the case of language and land class are preserved when the geographic unit of analysis changes from the village to the sub-district.

Beyond the technical explanations for why the relationship between diversity and development might be no more than a statistical artifact (a \mathcal{F} in terms of the framework we have developed here), what are the political economy channels that make both diversity debit and diversity credit theories contingent on ethnic and geographic scales? Particular patterns of ethnic geography (Hodler, Valsecchi and Vesperoni, 2017; Ezcurra and Rodríguez-Pose, 2017), ethnic competition (Schaub, 2017), politicization of ethnicity (Lieberman and Singh, 2012; Singh and vom Hau, 2014), and power differential between ethnic core and the periphery (Green, Forthcoming) have all been implicated. One of the central criticisms of “methodological nationalism” underlying the bulk of empirical evidence for the diversity debit hypothesis is that it is blind not only to questions of aggregation but also the historical geography of such aggregation (Wimmer and Glick Schiller, 2002; Wimmer, 2015).

In the context of the fractionalization-polarization relationship that we have been working with historical geography is important to uncover why polarization is a better predictor of national-level conflicts but masks sub-national ethnic variation that a fractionalization metric can capture. A plausible explanation is

that ethnically marginalized groups are spatially concentrated and this leads to pockets of higher polarization (Bleaney and Dimico, 2017). Segregated regions further enable the marginalized group to more cohesively organize protests and rebellions (Fearon and Laitin, 1996). Spatial segregation negatively affects governance, lowers trust, and leads to spatial inequality that feed further conflicts and lower levels of development outcomes (Alesina and Zhuravskaya, 2011; Ezcurra and Rodríguez-Pose, 2017). Conflicts are often a result of “local ethnic configurations” determined at various sub-national levels (Cunningham and Weidmann, 2010). The subnational historical geography of how concessions to, and accommodation of marginalized ethnic groups – arising out of political and social tensions around ethnicity – is redressed has implications for any diversity-development theory (Singh and vom Hau, 2014). For example, sectarian violence in Northern Ireland is explained by ethnic segregation that inhibits social contact and social capital formation (Balcells, Daniels and Escribà-Folch, 2016; Brown, Mccord and Zachary, 2017).

FRA_i-POL_i QUADRANT SUBSAMPLE REGRESSIONS

The centrality of ethnic geography and history in accounting for spatial patterns of the relationship between fractionalization and polarization is evident in our dataset. The fractionalization-polarization map in Figure 11 contains various spatial units (villages or sub-districts) are not randomly distributed in space. In Figure 13 we have divided the nearly 27,000 villages in our dataset into four quadrants based on village-level fractionalization (FRA_i) and discrete polarization (POL_i) map with these metrics computed for elementary caste groups (*jatis*). These quadrants are defined by $FRA_i = 0.5$ and $POL_i = 0.5$ lines in panel-A of Figure 11. As seen from Figure 13, these quadrants are spatially segregated. Areas with high fractionalization and high polarization dominate the map accounting for little under 60% of all villages in our dataset. The smallest quadrant (representing villages with low fractionalization and low polarization) also contains over 3000 villages – a sample size that is larger than any used in extant empirical studies of the relationship between diversity and development. We separately ran our OLS model described in Equation 5 on four different subsamples corresponding to the four quadrants described in Figure 13.

The results from the subsample “quadrant regressions” are summarized in Table 20. As all the regression models are run at the village-level, we are able to isolate the effect of ethnic aggregation. Thus, these regressions are not able to test if diversity debit or diversity credit are $\mathbb{G}\mathbb{T}$ -theories; instead we test if either of the two competing hypotheses about the relationship between diversity and development qualify as a $\mathbb{T}\mathbb{E}$ -theory. The individual subsample regressions also help test various hypotheses about the combined effect of fractionalization and polarization on development outcomes – as outlined above. As seen from the table,

these subsample regressions are as unstable as the full-sample OLS regressions that we discussed in Section V. The regression models are not robust to different specifications of the dependent variable. Coefficients on all three diversity metrics (fractionalization, discrete polarization, and ethnic inequality) are not robust to dependent variable specification even within a quadrant, and for a given diversity axis.

The findings from quadrant regressions add further ballast to our finding that extant support for both diversity debit and diversity credit theories are likely spurious statistical artifacts – \mathcal{X} theories that appear to hold at particular points on the ethnic-geographic continuum. The subsample regressions corresponding to fractionalization-polarization quadrants most starkly demonstrate the need to test diversity-development theories across the ethnic-geographic space and across the universe of values that a diversity metric can theoretically assume. Given the geographic segregation of the fractionalization-polarization quadrants as seen on the map in Figure 13, it is possible to construct a large enough spatially contiguous sample of villages that will support either diversity-debit, or the diversity-credit hypotheses.²⁸

ETHNIC GEOGRAPHY OF MEUP AND MAUP

In Section II, we introduced how the Modifiable Ethnic Unit Problem (MEUP) using non-isomorphic partially ordered sets (posets). Here, we present empirical examples from the dataset used in our empirical models, and also explore the spatial structure (ethnic geography) of ranking non-isomorphic posets. We begin by observing that the simple univariate distribution of diversity metric is sensitive to ethnic and geographic aggregation. In Figure 14, we present the density plots for fractionalization at three different levels of ethnic aggregation of caste at two spatial scales – village and sub-district. The resulting six distributions are shown for both fractionalization as well as discrete polarization. The kernel density plots in Figure 14 readily illustrate the problem of both the modifiable ethnic unit problem (MEUP) as well as the more traditionally recognized modifiable areal unit problem (MAUP). For both fractionalization and polarization, there is very little overlap between the distribution for census aggregation (the most commonly used aggregation of caste in empirical work) and the other two aggregations of caste. The effect of geographic aggregation is best seen from the fact that while village distributions are bimodal, sub-district distributions are unimodal. At the *taluka* level (bottom panel in the figure), we note how the three different ethnic aggregations of caste fractionalization are essentially different distributions, and minimally overlap for polarization. The ethnic-geographic aggregation of caste

²⁸One such exercise we carried out was to run our OLS models on 200-village spatially contiguous sample from within the NW quadrant and demonstrate apparent support for the diversity credit hypothesis. A similar exercise in the NE quadrant showed strong support for the diversity debit hypothesis.

presented in Figure 14 immediately suggests that any support for the diversity deficit (credit) hypothesis is likely no more than a statistical artefact of particular ethnic or geographic aggregations used in the empirical model.

The ecological fallacy that manifests in the form of MEUP or MAUP is a direct byproduct of using large spatial aggregates such as the nation state. The MAUP problem is documented in the diversity-development literature (Mateos, Singleton and Longley, 2009; Gershman and Rivera, 2018; Montalvo and Reynal-Querol, 2017; Gerring et al., 2015). At what level of geographic aggregation must the diversity development relationship be tested? Limited studies that have tried to overcome the MAUP problem do not find unambiguous support for either the diversity debit hypothesis, or the diversity credit hypothesis. For example, Montalvo and Reynal-Querol (2017) use grid-level data to test the diversity-development relationship at varying spatial scale. At finer grid resolutions, they find support for the diversity credit hypothesis (positive association between diversity and economic growth), while diversity debit dominates at larger aggregations (Montalvo and Reynal-Querol, 2017). In another study, McDoom and Gisselquist (2015) report that when a discrete polarization metric is constructed at multiple spatial scales, polarization scores decrease at lower levels of spatial (administrative region) aggregation.

The Modifiable Ethnic Unit Problem that we have introduced here is particularly serious because an ethnic variable like caste is embedded in the ethnic-geographic continuum. Here, we present evidence for how MAUP and MEUP combine in the ethnic-geographic space. In Figure 15, we show the spatial variation in caste fractionalization at both the village-level (top panel) as well as the sub-district level (bottom panel). A common scale is used to represent the six diversity indices across three ethnic aggregations and two spatial aggregations. The figure directly shows how the particular spatial structure of the ethnic aggregation problem can potentially influence the sign and magnitude of coefficients in the regression models used to test the diversity debit conjecture. The figure once again shows why any theory linking diversity and development must be tested at multiple ethnic-geographic aggregations. Figure 16 presents similar data as Figure 15 for discrete polarization (POL_i). Once again, all six maps are drawn to a common scale (that is also shared with maps in Figure 15). Once again, spatial structure of how MEUP and MAUP jointly operate is starkly represented.

The intersections between MEUP and MAUP also impacts how aggregation differentially impacts ethnic groups. In Table 11, we report the correlation coefficients (Pearson) for fractionalization and discrete polarization computed at village and district levels ($n = 13, 255, 421$). As seen from the table, there is significant variation in how geographic aggregation impacts particular ethnic groups.

Table 11—: *Village-District Diversity Correlations by Ethnic Groups*

Caste Group	Fractionalization	Polarization
I	0.14	0.07
II A	0.17	0.06
II B	0.04	0.13
III A	0.41	-0.09
III B	0.13	0.03
OTH	0.25	0.28
SC	0.12	0.08
ST	0.04	0.04

Note: The table reports correlation coefficient (Pearson) for fractionalization and polarization computed at village and district levels. $n = 13,255,421$ rural households.

An even starker evidence for the need to consider the spatial structure of ethnic aggregation is presented in Figure-17. In this figure we show how the ordinal ranking of villages in our data set changes dramatically based on the level of ethnic aggregation at which fractionalization and polarization are computed. This figure shows the spatial structure of non-isomorphic posets that we introduced in Section II. In the left panel, we computed change in ordinal ranking going from fractionalization index calculated using elementary *jati* categories to the corresponding index calculated using the three-fold census categorization of caste. To underscore the point that average change in ranking masks the spatial structure of ethnic aggregation, we have also shown the density plot of the ranking difference variable as an inset that is symmetrically distributed. The right panel repeats this exercise for discrete polarization.

In Figure 18, we present the spatial structure of non-isomorphic posets at the sub-district level. Once again we see that changes in ranking while symmetrical on the average are spatially heterogeneous. An explanation of the spatial structure of non-isomorphic posets shown in Figure 17 and Figure 18, requires an exploration of the micro-histories of different regions in our study area to understand the spatial variation in ‘species diversity’ (or number of subgroups) that is the principal driver of the observed patterns of ordinal rank difference. Historical analysis is an exercise beyond the scope of the present paper but our results underscore the need to integrate such analysis into accounts of how ethnic diversity impacts development.

VII. Conclusion

We have shown conclusively that the relationship between diversity and development is likely an artifact of *where* diversity is measured, *how* diversity is measured, and *what* diversity is measured. In particular, we have shown that the *where*, the *how*, and the *what* are jointly determined in an ethnic-geographic con-

tinuum – an object that we formally defined. We also developed a framework to test for validity of a theory in the ethnic-geographic continuum. The taxonomy of theories in the ethnic-geographic continuum is general and applicable to problems beyond diversity-development models. Our taxonomic structure for theories also contributes to recent interest in accounting for null, or non-significant empirical results in empirical models. We show how the “difference between ‘significant’ and ‘not significant’ is not itself statistically significant” in the ethnic-geographic continuum (Gelman and Stern, 2006). Indeed, we showed how a ‘null’ empirical result can be G_{TE} -theory (one that is valid across the universe of ethnic-geographic combinations) in the same way as a ‘statistically significant’ result. Our null results are informative because they update extant priors on the relationship between diversity and development (Abadie, 2018).

As a second key theoretical contribution, this paper has introduced a rigorous framework that accounts for ethnic aggregations as a counterpart of spatial aggregation. The Modifiable Ethnic Unit Problem (MEUP), in conjunction with the spatial aggregation derived MAUP (Modifiable Areal Unit Problem) explains the observed empirical instability in models connecting ethnic diversity and development. Taken together, the MEUP framework that we developed and the formal taxonomy of theories in the ethnic-geographic continuum helps account for the empirical observations about instability in the direction of the relationship between diversity and development.

To the best of our knowledge, this paper presents the most comprehensive multi-scale test of the empirical relationship between ethnic diversity and development. Our village-level regressions ($n \approx 27,000$) is the largest number of observations that have been used to test the diversity-development relationship. Our unique census-scale micro dataset allowed aggregating household-level self-reported ethnic group information. This rich dataset allowed us to cover the universe of values that a diversity metric can theoretically assume.

The empirical models presented in this paper tested the diversity-development relationship using seventeen different diversity metrics across different levels of ethnic and geographic aggregation. In addition to fractionalization and discrete polarization that have been the staple of empirical literature, we used household landholding data to construct a true polarization metric (Esteban and Ray, 1994) as well as account for ethnic inequality (Alesina, Michalopoulos and Papaioannou, 2016). We are able to present robust evidence for why the relationship between diversity and development is likely no more than a statistical artifact.

Finally, for the first time since 1931, this paper uses census-scale elementary caste *jati* data and makes a broader contribution to empirical literature about caste in India.

REFERENCES

- Abadie, Alberto.** 2018. "Statistical Non-Significance in Empirical Economics." MIT Economics Department Working Paper.
- Akdede, Sacit Hadi.** 2010. "Do More Ethnically and Religiously Diverse Countries have lower Democratization?" *Economic Letters*, 106(2): 101–104.
- Alesina, Alberto, and Ekaterina Zhuravskaya.** 2011. "Segregation and the Quality of Government in a Cross Section of Countries." *The American Economic Review*, 101(5): 1872–1911.
- Alesina, Alberto, and Eliana La Ferrara.** 2002. "Who Trusts Others?" *Journal of Public Economics*, 85(2): 207–234.
- Alesina, Alberto, and Eliana La Ferrara.** 2005. "Ethnic Diversity and Economic Performance." *Journal of Economic Literature*, 43(3): 762–800.
- Alesina, Alberto, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat, and Romain Wacziarg.** 2003. "Fractionalization." *Journal of Economic Growth*, 8(2): 155–194.
- Alesina, Alberto, Reza Baqir, and William Easterly.** 1999. "Public Goods and Ethnic Divisions." *The Quarterly Journal of Economics*, 114(4): 1243–1284.
- Alesina, Alberto, Stelios Michalopoulos, and Elias Papaioannou.** 2016. "Ethnic Inequality." *Journal of Political Economy*, 124(2): 428–488.
- Anantha Krishna Iyer, L. Krishna, and H. V Nanjundayya.** 1928. *The Mysore Tribes and Castes*. University of Mysore Press.
- Anderson, Siwan.** 2011. "Caste as an Impediment to Trade." *American Economic Journal: Applied Economics*, 3(1): 239–263.
- Assadi, Muzaffar.** 1999. "Karnataka: Communal Violence in Coastal Belt." *Economic and Political Weekly*, 34(08).
- Assadi, Muzaffar.** 2002. "Karnataka: Hindutva Policies in Coastal Region." *Economic and Political Weekly*, 37(23).
- Balcells, Laia, Lesley-Ann Daniels, and Abel Escribà-Folch.** 2016. "The Determinants of Low-intensity Intergroup Violence." *Journal of Peace Research*, 53(1): 33–48.
- Baldwin, Kate, and John D. Huber.** 2010. "Economic versus Cultural Differences: Forms of Ethnic Diversity and Public Goods Provision." *The American Political Science Review*, 104(4): 644–662.

- Banerjee, Abhijit, and Rohini Somanathan.** 2007. "The Political Economy of Public Goods: Some Evidence from India." *Journal of Development Economics*, 82(2): 287–314.
- Banerjee, Abhijit, Lakshmi Iyer, and Rohini Somanathan.** 2005. "History, Social divisions, and Public Goods in Rural India." *Journal of the European Economic Association*, 3(2-3): 639–647.
- Banerjee-Dube, Ishita,** ed. 2008. *Caste in History*. Oxford University Press.
- Bayly, Susan.** 2001. *Caste, Society and Politics in India from the Eighteenth Century to the Modern Age*. Cambridge University Press.
- Bharathi, Naveen, Deepak Malghan, Sumit Mishra, and Andaleeb Rahman.** 2018. "Spatial Segregation, Multi-scale Diversity, and Public Goods." *Charles H. Dyson School of Applied Economics and Management at Cornell University Working Paper Series*, AEM-WP2018(7).
- Bleaney, Micheal, and Arcangelo Dimico.** 2017. "Ethnic Diversity and Conflict." *Journal of Institutional Economics*, 13(02): 357–378.
- Brown, Joseph, Gordon C Mccord, and Paul Zachary.** 2017. "Sunday, Bloody Sunday: Evidence from Northern Ireland for the Effect of Ethnic Diversity on Violence." *mimeo*.
- Burlig, Fiona, and Louis Preonas.** 2016. *Out of the Darkness and Into the Light? Development Effects of Rural Electrification in India*. Energy Institute at Haas, University of California, Berkeley.
- Chakravorty, Ujjayant, Martino Pelli, and Beyza Ural Marchand.** 2014. "Does the Quality of Electricity Matter? Evidence from Rural India." *Journal of Economic Behavior & Organization*, 107: 228–247.
- Chandra, Kanchan.** 2006. "What Is Ethnic Identity and Does It Matter?" *Annual Review of Political Science*, 9(1): 397–424.
- Chandra, Kanchan, and Steven Wilkinson.** 2008. "Measuring the Effect of Ethnicity." *Comparative Political Studies*, 41(4-5): 515–563.
- Chandramouli, C.** 2011. *Census of India 2011*. Registrar General of India, Government of India.
- Chen, Xi, and William D Nordhaus.** 2011. "Using Luminosity Data as a Proxy for Economic Statistics." *Proceedings of the National Academy of Sciences*, 108(21): 8589–8594.
- Collier, Paul.** 2004. "Greed and Grievance in Civil War." *Oxford Economic Papers*, 56(4): 563–595.

- Collier, Paul, and Anke Hoeffler.** 1998. "On Economic Causes of Civil War." *Oxford Economic Papers*, 50(4): 563–573.
- Cunningham, Kathleen Gallagher, and Nils B. Weidmann.** 2010. "Shared Space: Ethnic Groups, State Accommodation, and Localized Conflict." *International Studies Quarterly*, 54: 1035–1054.
- Dahlberg, Matz, Karin Edmark, and Heléne Lundqvist.** 2012. "Ethnic Diversity and Preferences for Redistribution." *Journal of Political Economy*, 120(1): 41–76.
- Damm, Anna Piil.** 2009. "Ethnic Enclaves and Immigrant Labor Market Outcomes: Quasi-experimental Evidence." *Journal of Labor Economics*, 27(2): 281–314.
- Damodaran, Harish.** 2009. *India's New Capitalists: Caste, Business, and Industry in a Modern Nation*. Permanent Black.
- Desai, Sonalde, and Reeve Vanneman.** 2005. *India Human Development Survey (IHDS)*. NCAER, India and Inter-university Consortium for Political and Social Research, MI, USA.
- Desai, Sonalde, and Reeve Vanneman.** 2011. *India Human Development Survey-II (IHDS-II)*. NCAER, India and Inter-university Consortium for Political and Social Research, MI, USA.
- Dirks, Nicholas B.** 2011. *Castes of Mind: Colonialism and the Making of Modern India*. Princeton University Press.
- Dumont, Louis.** 1980. *Homo Hierarchicus: The Caste System and its Implications*. University of Chicago Press.
- Easterly, William, and Ross Levine.** 1997. "Africa's Growth Tragedy: Policies and Ethnic Divisions." *Quarterly Journal of Economics*, 112(4): 1203–1250.
- Ebener, Steeve, Christopher Murray, Ajay Tandon, and Christopher C Elvidge.** 2005. "From Wealth to Health: Modelling the Distribution of Income per capita at the Sub-national Level using Night-time Light Imagery." *International Journal of Health Geographics*, 4(1): 5.
- Enthoven, Reginald E.** 1990. *The Tribes and Castes of Bombay (Vol. I)*. Asian Educational Services.
- Esteban, Joan, and Debraj Ray.** 1994. "On the Measurement of Polarization." *Econometrica*, 62(4): 819–851.
- Esteban, Joan, Laura Mayoral, and Debraj Ray.** 2012a. "Ethnicity and Conflict: An Empirical Study." *American Economic Review*, 102(4): 1310–1342.

- Esteban, Joan, Laura Mayoral, and Debraj Ray.** 2012*b*. “Ethnicity and Conflict: Theory and Facts.” *Science*, 336(6083): 858–865.
- Ezcurra, Roberto, and Andrés Rodríguez-Pose.** 2017. “Does Ethnic Segregation Matter for Spatial Inequality?” *Journal of Economic Geography*, 17(6): 1149–1178.
- Fearon, James D.** 2003. “Ethnic and Cultural Diversity by Country.” *Journal of Economic Growth*, 8(2): 195–222.
- Fearon, James D, and David D Laitin.** 1996. “Explaining Interethnic Cooperation.” *The American Political Science Review*, 90(4): 715–735.
- Fearon, James D, and David D Laitin.** 2003. “Ethnicity, Insurgency, and Civil War.” *American Political Science Review*, 97(1): 75–90.
- Foster, James E, and Artyom A Shneyerov.** 2000. “Path Independent Inequality Measures.” *Journal of Economic Theory*, 91(2): 199–222.
- Fotheringham, A Stewart, and David WS Wong.** 1991. “The Modifiable Areal Unit Problem in Multivariate Statistical Analysis.” *Environment and Planning A*, 23(7): 1025–1044.
- Freedman, David A.** 1999. “Ecological Inference and the Ecological Fallacy.” *International Encyclopedia of the social & Behavioral sciences*, 6: 4027–4030.
- Gelman, Andrew, and Hal Stern.** 2006. “The Difference between “Significant” and “Not Significant” is not Itself Statistically Significant.” *The American Statistician*, 60(4): 328–331.
- Gerring, John, Strom C. Thacker, Yuan Lu, and Wei Huang.** 2015. “Does Diversity Impair Human Development? A Multi-level Test of the Diversity Debit Hypothesis.” *World Development*, 66: 166–188.
- Gershman, Boris, and Diego Rivera.** 2018. “Subnational Diversity in Sub-Saharan Africa: Insights from a New Dataset.” *Journal of Development Economics*, 133: 231–263.
- Ghurye, Govind Sadashiv.** 1969. *Caste and Race in India*. Popular Prakashan.
- Gisselquist, Rachel M., Stefan Leiderer, and Miguel Niño Zarazúa.** 2016. “Ethnic Heterogeneity and Public Goods Provision in Zambia: Evidence of a Subnational “Diversity Dividend”.” *World Development*, 78: 308–323.
- Glennerster, Rachel, Edward Miguel, and Alexander D Rothenberg.** 2013. “Collective Action in Diverse Sierra Leone Communities.” *The Economic Journal*, 123(568): 285–316.
- Gordon Jr, Raymond G.** 2005. *Ethnologue, Ethnologue: Languages of the World*. SIL International.

- Gören, Erkan.** 2014. "How Ethnic Diversity Affects Economic Growth." *World Development*, 59: 275–297.
- Green, Elliott D.** Forthcoming. "Ethnicity, National Identity and the State: Evidence from Sub-Saharan Africa." *British Journal of Political Science*.
- Guha, Sumit.** 2013. *Beyond Caste: Identity and Power in South Asia, Past and Present*. Brill.
- Gupta, Dipankar.** 1991. *Social Stratification*. Oxford University Press.
- Habyarimana, James, Macartan Humphreys, Daniel N Posner, and Jeremy M Weinstein.** 2007. "Why Does Ethnic Diversity Undermine Public Goods Provision?" *American Political Science Review*, 101(4): 709–725.
- Hardgrave Jr, Robert L.** 1968. "Caste: Fission and Fusion." *Economic and Political Weekly*, 1065–1070.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil.** 2012. "Measuring Economic Growth from Outer Space." *American Economic Review*, 102(2): 994–1028.
- Hirschman, Albert O.** 1964. "The Paternity of an Index." *The American Economic Review*, 54(5): 761–762.
- Hodler, Roland, Michele Valsecchi, and Alberto Vesperoni.** 2017. "Ethnic Geography: Measurement and Evidence." *CESifo Working Paper Series 6720*.
- Horowitz, Donald L.** 1985. *Ethnic Groups in Conflict*. University of California Press.
- Houle, Christian, and Cristina Bodea.** 2017. "Ethnic Inequality and Coups in sub-Saharan Africa." *Journal of Peace Research*, 54(3): 382–396.
- Iversen, Vegard, Adriaan Kalwij, Arjan Verschoor, and Amaresh Dubey.** 2014. "Caste Dominance and Economic Performance in Rural India." *Economic Development and Cultural Change*, 62(3): 423–457.
- Jodhka, Surinder.** 2012. *Caste*. Oxford University Press.
- Kanbur, Ravi, Prem Kumar Rajaram, and Ashutosh Varshney.** 2011. "Ethnic Diversity and Ethnic Strife. An Interdisciplinary Perspective." *World Development*, 39(2): 147–158.
- King, Gary.** 2013. *A solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton University Press.
- Kohli, Atul.** 1982. "The State and Agrarian Policy." *The Journal of Commonwealth & Comparative Politics*, 20(3): 309–328.

- Kohli, Atul.** 1987. *The State and Poverty in India: The Politics of Reform.* Cambridge Univ. Press.
- La Porta, Rafael Raphael, Florencio Lopez-de Silanes, Andrei Shleifer, Robert W Vishny, and Robert Vishney.** 1999. "The Quality of Government." *Journal of Law, Economics, and Organization*, 15(1): 222–279.
- Lieberman, Evan S., and Prerna Singh.** 2012. "Conceptualizing and Measuring Ethnic Politics: An Institutional Complement to Demographic, Behavioral, and Cognitive Approaches." *Studies in Comparative International Development*, 47(3): 255–286.
- Marquardt, Kyle L, and Yoshiko M Herrera.** 2015. "Ethnicity as a Variable: An Assessment of Measures and Data Sets of Ethnicity and Related Identities." *Social Science Quarterly*, 96(3): 689–716.
- Mateos, Pablo, Alex Singleton, and Paul Longley.** 2009. "Uncertainty in the Analysis of Ethnicity Classifications: Issues of Extent and Aggregation of Ethnic Groups." *Journal of Ethnic and Migration Studies*, 35(9): 1437–1460.
- McDoom, Omar Shahabudin, and Rachel M. Gisselquist.** 2015. "The Measurement of Ethnic and Religious Divisions: Spatial, Temporal, and Categorical Dimensions with Evidence from Mindanao, the Philippines." *Social Indicators Research*, 129(2): 863–891.
- Miguel, Edward, and Mary Kay Gugerty.** 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics*, 89(11–12): 2325–2368.
- Min, Brian, Kwawu Mensan Gaba, Ousmane Fall Sarr, and Alasane Agalassou.** 2013. "Detection of Rural Electrification in Africa using DMSP-OLS Night Lights Imagery." *International Journal of Remote Sensing*, 34(22): 8118–8141.
- Montalvo, Jose G., and Marta Reynal-Querol.** 2003. "Religious Polarization and Economic Development." *Economics Letters*, 80(2): 201–210.
- Montalvo, Jose G., and Marta Reynal-Querol.** 2005a. "Ethnic Diversity and Economic Development." *Journal of Development Economics*, 76(2): 293–323.
- Montalvo, Jose G., and Marta Reynal-Querol.** 2005b. "Ethnic Polarization, Potential Conflict, and Civil Wars." *American Economic Review*, 95(3): 796–816.
- Montalvo, Jose G., and Marta Reynal-Querol.** 2008. "Discrete Polarisation with an Application to the Determinants of Genocides." *Economic Journal*, 118(533): 1835–1865.

- Montalvo, Jose G., and Marta Reynal-Querol.** 2017. "Ethnic Diversity and Growth: Revisiting the Evidence." Barcelona Graduate School of Economics Working Papers 992.
- Müller-Crepon, Carl, and Philipp Hunziker.** 2018. "New Spatial Data on Ethnicity: Introducing SIDE." *Journal of Peace Research*, 0022343318764254.
- Munshi, Kaivan.** 2014. "Community Networks and the Process of Development." *Journal of Economic Perspectives*, 28(4): 49–76.
- Munshi, Kaivan, and Mark Rosenzweig.** 2006. "Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy." *The American Economic Review*, 96(4): 1225–1252.
- Nanjundappa, DM, A Aziz, B Sheshadri, G Kadekodi, and MJM Rao.** 2002. *Report of the High Power Committee for Redressal of Regional Imbalances in Karnataka.* Government of Karnataka.
- Olson, Mancur.** 1982. *The Rise and Decline of Nations: Economic Growth, Stagnation, and Social Rigidities.* Yale University Press.
- Openshaw, Stan.** 1984. "Ecological Fallacies and the Analysis of Areal Census Data." *Environment and Planning A*, 16(1): 17–31.
- Ostrom, Elinor.** 1990. *Governing the Commons.* Cambridge: Cambridge University Press.
- Otsuka, Keijiro, Hiroyuki Chuma, and Yujiro Hayami.** 1993. "Permanent Labour and Land Tenancy Contracts in Agrarian Economies: An Integrated Analysis." *Economica*, 57–77.
- Ottaviano, Gianmarco I.P., and Giovanni Peri.** 2006. "The Economic Value of Cultural Diversity: Evidence from US cities." *Journal of Economic Geography*, 6(1): 9–44.
- Paik, Christopher.** 2013. "Income or Politics: A Study of Satellite Streetlight Imagery Application in Pakistan." *mimeo*, Accessed: 2018-07-10.
- Posner, Daniel N.** 2004. "Measuring Ethnic Fractionalization in Africa." *American Journal of Political Science*, 48(4): 849–863.
- Ramachandra, TV, G Kamakshi, and BV Shruthi.** 2004. "Bioresource status in Karnataka." *Renewable and Sustainable Energy Reviews*, 8(1): 1–47.
- Rao, Narasimha D.** 2013. "Does (Better) Electricity Supply Increase Household Enterprise Income in India?" *Energy Policy*, 57: 532–541.
- Rawat, Ramnarayan S.** 2011. *Reconsidering Untouchability: Chamars and Dalit History in North India.* Indiana University Press.

- Ray, Debraj, and Joan Esteban.** 2017. "Conflict and Development." *Annual Review of Economics*, 9: 263–293.
- Robinson, W. S.** 1950. "Ecological Correlations and the Behavior of Individuals." *American Sociological Review*, 15(3): 351–357.
- Roeder, Philip G.** 2001. "Ethnolinguistic Fractionalization (ELF) Indices, 1961 and 1985." <http://web.archive.org/web/20080207010024/http://www.808multimedia.com/winnt/kernel.htm>, Accessed: 2018-07-10.
- Roychowdhury, Koel, Simon Jones, Karin Reinke, and Colin Arrowsmith.** 2012. "Night-Time Lights and Levels of Development: A Study Using DMSP-OLS Night Time Images At the Sub-National Level." *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXIX-B8: 93–98.
- Sayeed, Vikar Ahmed.** 2016. "Communal Cauldron." *Frontline*, 33(19).
- Scarritt, James R, and Shaheen Mozaffar.** 1999. "The Specification of Ethnic Cleavages and Ethnopolitical Groups for the Analysis of Democratic Competition in Contemporary Africa." *Nationalism and Ethnic Politics*, 5(1): 82–117.
- Schaub, Max.** 2017. "Second-order Ethnic Diversity: The Spatial Pattern of Diversity, Competition and Cooperation in Africa." *Political Geography*, 59: 103–116.
- Sheth, D. L.** 1999. "Secularisation of Caste and Making of New Middle Class." *Economic and Political Weekly*, 34(34/35): 2502–2510.
- Shivashankar, P., and G.S. Ganesh Prasad.** 2015. *Human Development Performance of Gram Panchayats in Karnataka 2015*. Mysuru: Abdul Nazir Sab State Institute of Rural Development and Panchayat Raj, Government of Karnataka.
- Shorrocks, Anthony, and Guanghua Wan.** 2005. "Spatial Decomposition of Inequality." *Journal of Economic Geography*, 5(1): 59–81.
- Singh, Kumar Suresh.** 2002. *People of India: An Introduction*. Anthropological Survey of India.
- Singh, Prerna.** 2015a. *How Solidarity Works for Welfare*. Cambridge University Press.
- Singh, Prerna.** 2015b. "Subnationalism and Social Development: A Comparative Analysis of Indian States." *World Politics*, 67(03): 506–562.

- Singh, Prerna, and Matthias vom Hau.** 2014. "Ethnicity, State Capacity and Development: Reconsidering Causal Connections." In *The Politics of Inclusive Development: Interrogating the Evidence.*, ed. Sam Hickey, Kunal Sen and Badru Bukenya, 231–256. Oxford University Press.
- Singh, Prerna, and Matthias vom Hau.** 2016. "Ethnicity in Time: Politics, History, and the Relationship between Ethnic Diversity and Public Goods Provision." *Comparative Political Studies*, 49(10): 1303–1340.
- Srinivas, Mysore Narasimhachar.** 1959. "The Dominant Caste in Rampura." *American Anthropologist*, 61(1): 1–16.
- Srinivas, Mysore Narasimhachar.** 1965. *Religion and Society among the Coorgs of South India.* New York:Asia Pub. House.
- Srinivas, Mysore Narasimhachar.** 1994. *The Dominant Caste and Other Essays.* Oxford University Press.
- Srinivas, Mysore Narasimhachar.** 1996. *Caste: Its Twentieth Century Avatar.* Viking.
- Srinivas, Mysore Narasimhachar, and M. N. Panini.** 1984. "Politics and Society in Karnataka." *Economic and Political Weekly*, 19(2): 69–75.
- Tajima, Yuhki, Krislert Samphantharak, and Kai Ostwald.** 2018. "Ethnic Segregation and Public Goods: Evidence from Indonesia." *American Political Science Review*, 112(3): 1–17.
- Taylor, Charles L, and Michael C Hudson.** 1972. *World handbook of social and political indicators.* ICSPR.
- Thapliyal, Sneha, Arnab Mukherji, and Deepak Malghan.** 2017. "Inequality and Erosion of Commons: Causal Evidence for Differing Roles of Scale of Governance Hierarchies from India." *IIMI Working Paper Series*, IIMI-WP 17.
- Thurston, Edgar, and K Rangachari.** 1975. *Castes and Tribes of Southern India.* Delhi:Cosmo.
- Tumbe, Chinmay.** 2013. "India Migration Bibliography." *eSocialSciences*, 5311: 1–183.
- UNDP.** 2005. *Karnataka Human Development Report 2005.* Government of Karnataka, Planning and Statistics Department.
- UNDP.** 2010. *Human Development Report 2010: The Real Wealth of Nations: Pathways to Human Development.* United Nations Development Program and Palgrave Macmillan.

- Van de Walle, Dominique, Martin Ravallion, Vibhuti Mendiratta, and Gayatri Koolwal.** 2017. "Long-term Gains from Electrification in Rural India." *The World Bank Economic Review*, 31(2): 385–411.
- Weber, Max.** 1968. *Economy and Society: An Outline of Interpretive Sociology*. New York:Bedminster Press.
- Weidmann, Nils B, Jan Ketil Rød, and Lars-Erik Cederman.** 2010. "Representing Ethnic Groups in Space: A New Dataset." *Journal of Peace Research*, 47(4): 491–499.
- Wimmer, Andreas.** 2008a. "Elementary Strategies of Ethnic Boundary Making." *Ethnic and Racial Studies*, 31(6): 1025–1055.
- Wimmer, Andreas.** 2008b. "The Making and Unmaking of Ethnic Boundaries: A Multilevel Process Theory." *American Journal of Sociology*, 113(4): 970–1022.
- Wimmer, Andreas.** 2015. "Is Diversity Detrimental? Ethnic Fractionalization, Public Goods Provision, and the Historical Legacies of Stateness." *Comparative Political Studies*, 49(11): 1407–1445.
- Wimmer, Andreas, and Nina Glick Schiller.** 2002. "Methodological Nationalism and Beyond: Nation-state Building, Migration and the Social Sciences." *Global Networks*, 2(4): 301–334.
- Zacharias, Ajit, and Vamsi Vakulabharanam.** 2011. "Caste Stratification and Wealth Inequality in India." *World Development*, 39(10): 1820–1833.

Table 12—: *Growth in Per-capita Luminosity and Social Heterogeneity Indicators (Village level models)*

	Growth in Mean Per-capita Luminosity, 2001-2013										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ModJatiFRAC	20.72 (18.22)	-28.64** (11.59)									
ModJatiPOL	14.73 (13.71)	-54.01*** (10.96)									
ModMapFRAC	19.68 (22.11)		20.10 (16.86)	-32.33*** (12.30)							
ModMapPOL	-36.67** (17.12)		-63.63*** (12.53)								
ERPolar	-0.48 (1.27)			-4.68*** (1.61)							-5.76*** (1.58)
ReligionFRAC	-32.93 (47.44)				22.55 (63.88)						
ReligionPOL	20.26 (25.95)				-8.57 (35.09)						
LanguageFRAC	-16.56 (29.30)					9.55 (39.00)					
LanguagePOLAR	13.68 (16.61)					-0.38 (22.74)					
Census2011FRAC	-11.55 (28.28)						-70.95*** (10.44)				
Census2011POL	-14.99 (16.23)							-46.70*** (6.13)			
LandClassFRAC	-38.02** (17.40)								-27.81 (17.22)		
LandClassPOL	15.17 (14.54)								-145.64*** (17.68)		
LandMLD	-3.11*** (0.84)									-3.25*** (0.62)	
LandModMap_BG	4.27 (3.02)									7.86*** (2.89)	
LandReligion_BG	-5.51 (4.09)									-4.23 (4.00)	
LandLanguage_BG	-5.78** (2.79)									-5.62** (2.66)	
LandCensus_BG	15.08*** (3.80)									13.21*** (3.76)	
Constant	183.56*** (22.23)	175.43*** (27.06)	191.62*** (27.30)	173.59*** (26.98)	132.42*** (26.65)	125.51*** (26.52)	118.45*** (26.25)	124.48*** (26.29)	349.99*** (29.73)	93.96*** (18.60)	114.19*** (26.30)
R-squared	0.046	0.023	0.024	0.023	0.020	0.020	0.021	0.021	0.030	0.040	0.019
N	25572	25633	25633	25633	25633	25633	25633	25633	25633	25572	25633

Standard errors are reported in parentheses (to two decimal places). * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Per-capita luminosity is computed using 2001 and 2011 census population for each village with linear extrapolation (see main text for more details). All models included appropriate ‘species count’ variable(s) indicating number of distinct social categories (for example, count of distinct *jaitis* in the village), and are run with sub-district (taluka) fixed effects that among other things, accounts for variation in rainfall, colonial administrative history, and agro-ecological classification. Additionally, all models also include following village-level controls from 2011 Census of India data: a PCA-based ‘asset index’ of village-level public goods (computed using data for availability of primary school, secondary school, health care services, treated water, sanitation, bus service, and financial institutions); irrigated land as share of total land; literacy rates; distance to nearest town in kilometers; electricity provision; and number of hours of electricity during summer and winter months.

Table 13—: *Growth in Per-capita Luminosity and Social Heterogeneity Indicators (Sub-district Models)*

	Growth in Mean Per-capita Luminosity, 2001-2013										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ModJatiFRAC	18.96 (171.48)	-61.02 (76.64)									
ModJatiPOL	-19.12 (78.91)	13.51 (52.33)									
ModMapFRAC	1.95 (153.39)		-36.88 (64.66)	-100.50*** (37.64)							
ModMapPOL	135.12 (114.25)		72.97 (67.90)								
ERPolar	-2.02 (13.21)			-7.20 (7.12)							-3.88 (7.06)
ReligionFRAC	-69.82 (107.97)				-6.85 (96.85)						
ReligionPOL	69.71 (79.61)				-36.83 (62.71)						
LanguageFRAC	82.95* (44.21)					29.61 (42.00)					
LanguagePOLAR	-88.11*** (33.38)					-43.92 (32.87)					
Census2011FRAC	-157.73* (92.58)						-99.56*** (21.75)				
Census2011POL	43.47 (58.03)							-59.27*** (16.60)			
LandClassFRAC	22.89 (45.23)								31.21 (24.83)		
LandClassPOL	106.35** (48.39)								73.70** (33.57)		
LandMLD	-1.77 (2.80)									-0.37 (1.55)	
LandModMap_BG	36.51 (82.48)									15.77 (75.01)	
LandReligion_BG	-72.74 (90.12)									46.96 (80.06)	
LandCensus_BG	-23.84 (80.66)									-121.85 (81.75)	
Constant	-3.98 (119.08)	145.15 (97.57)	75.72 (103.26)	179.82*** (54.35)	117.03*** (38.05)	72.24 (47.07)	185.54*** (36.08)	161.70*** (35.98)	3.03 (49.38)	111.42*** (34.27)	102.31*** (33.17)
R-squared	0.313	0.198	0.154	0.153	0.221	0.145	0.227	0.190	0.157	0.137	0.126
N	175	175	175	175	175	175	175	175	175	175	175

Standard errors are reported in parentheses (to two decimal places). * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Per-capita luminosity is computed using 2001 and 2011 census population for rural areas of each sub-district with linear extrapolation (see main text for more details). All models included appropriate ‘species count’ variable(s) indicating number of distinct social categories (for example, count of distinct *jaitis* in the sub-district). Additionally, all models also include controls for colonial administrative history (ten dummies); agro-ecological zone (six dummies), and literacy rates.

Table 14—: *Per-capita Luminosity and Social Heterogeneity Indicators (Village level models)*

	Mean Per-capita Luminosity, 2001-2013										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ModJatiFRAC	0.05** (0.02)	-0.10*** (0.01)									
ModJatiPOL	0.03* (0.02)	-0.10*** (0.01)									
ModMapFRAC	0.07** (0.03)		0.12*** (0.02)	-0.03** (0.01)							
ModMapPOL	-0.07*** (0.02)		-0.16*** (0.01)								
ERPolar	0.02*** (0.00)			0.00 (0.00)							-0.00 (0.00)
ReligionFRAC	0.08 (0.06)				0.26*** (0.07)						
ReligionPOL	-0.06* (0.03)				-0.13*** (0.04)						
LanguageFRAC	-0.15*** (0.04)					0.17*** (0.04)					
LanguagePOLAR	0.11*** (0.02)					-0.05** (0.03)					
Census2011FRAC	0.11*** (0.04)						-0.19*** (0.01)				
Census2011POL	-0.11*** (0.02)							-0.13*** (0.01)			
LandClassFRAC	-0.37*** (0.02)								-0.18*** (0.02)		
LandClassPOL	0.08*** (0.02)								-0.30*** (0.02)		
LandMLD	-0.02*** (0.00)									-0.01*** (0.00)	
LandModMap_BG	0.04*** (0.00)									0.06*** (0.00)	
LandReligion_BG	-0.03*** (0.01)									-0.02*** (0.01)	
LandLanguage_BG	0.02*** (0.00)									0.02*** (0.00)	
LandCensus_BG	-0.03*** (0.00)									-0.04*** (0.00)	
Constant	0.65*** (0.03)	0.33*** (0.03)	0.41*** (0.03)	0.35*** (0.03)	0.24*** (0.03)	0.23*** (0.03)	0.21*** (0.03)	0.22*** (0.03)	0.89*** (0.03)	0.17*** (0.02)	0.18*** (0.03)
R-squared	0.109	0.034	0.046	0.042	0.023	0.024	0.028	0.030	0.109	0.040	0.018
N	26821	26889	26889	26889	26889	26889	26889	26889	26889	26821	26889

Standard errors are reported in parentheses (to two decimal places). * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Per-capita luminosity is computed using 2001 and 2011 census population for each village with linear extrapolation (see main text for more details). All models included appropriate ‘species count’ variable(s) indicating number of distinct social categories (for example, count of distinct *jaitis* in the village), and are run with sub-district (taluka) fixed effects that among other things, accounts for variation in rainfall, colonial administrative history, and agro-ecological classification. Additionally, all models also include following village-level controls from 2011 Census of India data: a PCA-based ‘asset index’ of village-level public goods (computed using data for availability of primary school, secondary school, health care services, treated water, sanitation, bus service, and financial institutions); irrigated land as share of total land; literacy rates; distance to nearest town in kilometers; electricity provision; and number of hours of electricity during summer and winter months.

Table 15—: *Per-capita Luminosity and Social Heterogeneity Indicators (Sub-district Models)*

	Mean Per-capita Luminosity, 2001-2013											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	
ModJatiFRAC	-0.13*** (0.04)	-0.05*** (0.02)										
ModJatiPOL	-0.05** (0.02)	-0.03** (0.01)										
ModMapFRAC	0.06 (0.04)		-0.04** (0.02)	-0.02* (0.01)								
ModMapPOL	0.07** (0.03)		-0.03 (0.02)									
ERPolar	-0.00 (0.00)			-0.00** (0.00)							-0.00* (0.00)	
ReligionFRAC	-0.02 (0.03)				0.01 (0.03)							
ReligionPOL	0.03* (0.02)				-0.01 (0.02)							
LanguageFRAC	-0.01 (0.01)					0.02 (0.01)						
LanguagePOLAR	0.00 (0.01)					-0.00 (0.01)						
Census2011FRAC	-0.08*** (0.02)						-0.01 (0.01)					
Census2011POL	0.06*** (0.01)							0.00 (0.00)				
LandClassFRAC	-0.01 (0.01)								-0.03*** (0.01)			
LandClassPOL	-0.03** (0.01)								-0.01 (0.01)			
LandMLD	0.00 (0.00)									0.00*** (0.00)		
LandModMap_BG	-0.04** (0.02)									-0.04** (0.02)		
LandReligion_BG	0.03 (0.02)									0.02 (0.02)		
LandCensus_BG	0.05** (0.02)									0.03 (0.02)		
Constant	0.06** (0.03)	0.07*** (0.03)	0.05 (0.03)	0.01 (0.02)	-0.01 (0.01)	-0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.04*** (0.01)	-0.00 (0.01)	0.01 (0.01)
R-squared	0.693	0.613	0.525	0.530	0.582	0.539	0.517	0.513	0.594	0.584	0.524	
N	175	175	175	175	175	175	175	175	175	175	175	

Standard errors are reported in parentheses (to two decimal places). * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Per-capita luminosity is computed using 2001 and 2011 census population for rural areas of each sub-district with linear extrapolation (see main text for more details). All models included appropriate ‘species count’ variable(s) indicating number of distinct social categories (for example, count of distinct *jatis* in the sub-district). Additionally, all models also include controls for colonial administrative history (ten dummies); agro-ecological zone (six dummies), and literacy rates.

Table 16—: *Human Development Index (HDI) and Social Heterogeneity Indicators (Village-level Models)*

	Dependent Variable, Human Development Index (HDI)										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ModJatiFRAC	0.29 (0.81)	3.29*** (0.37)									
ModJatiPOL	-2.30*** (0.61)	-0.19 (0.35)									
ModMapFRAC	-1.14 (0.99)		0.22 (0.54)	-0.03 (0.39)							
ModMapPOL	2.64*** (0.76)		0.03 (0.40)								
ERPolar	0.05 (0.06)			0.17*** (0.05)							0.19*** (0.05)
ReligionFRAC	2.63 (2.13)				2.60 (2.04)						
ReligionPOL	-1.30 (1.17)				-2.02* (1.12)						
LanguageFRAC	1.81 (1.30)					-0.15 (1.23)					
LanguagePOLAR	-0.56 (0.74)					0.39 (0.72)					
Census2011FRAC	-4.43*** (1.27)						2.46*** (0.34)				
Census2011POL	1.75** (0.73)							1.88*** (0.20)			
LandClassFRAC	2.97*** (0.77)								-3.36*** (0.55)		
LandClassPOL	0.63 (0.65)								2.70*** (0.57)		
LandMLD	0.12*** (0.04)									0.31*** (0.03)	
LandModMap_BG	0.02 (0.13)									-0.39*** (0.13)	
LandReligion_BG	0.12 (0.18)									0.10 (0.18)	
LandLanguage_BG	-0.01 (0.12)									0.30** (0.12)	
LandCensus_BG	-0.01 (0.17)									-0.20 (0.17)	
Constant	37.06*** (0.99)	45.01*** (0.87)	43.21*** (0.87)	43.02*** (0.86)	45.82*** (0.86)	45.77*** (0.85)	48.12*** (0.85)	47.81*** (0.85)	40.34*** (0.95)	47.55*** (0.85)	48.26*** (0.85)
R-squared	0.372	0.356	0.360	0.360	0.349	0.355	0.342	0.343	0.359	0.346	0.341
N	25090	25141	25141	25141	25141	25141	25141	25141	25141	25090	25141

Standard errors are reported in parentheses (to two decimal places). * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

The data for dependent variable, HDI at village level, is from [Shivashankar and Prasad \(2015\)](#). All models included appropriate ‘species count’ variable(s) indicating number of distinct social categories (for example, count of distinct *jaitis* in the village), and are run with sub-district (taluka) fixed effects that among other things, accounts for variation in rainfall, colonial administrative history, and agro-ecological classification. Additionally, all models also include following village-level controls from 2011 Census of India data: a PCA-based ‘asset index’ of village-level public goods (computed using data for availability of primary school, secondary school, health care services, treated water, sanitation, bus service, and financial institutions); irrigated land as share of total land; literacy rates; distance to nearest town in kilometers; electricity provision; and number of hours of electricity during summer and winter months.

Table 17—: *Human Development Index (HDI) and Social Heterogeneity Indicators (Sub-district Models)*

	Dependent Variable, Human Development Index (HDI)										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ModJatiFRAC	26.08 (23.08)	3.84 (10.79)									
ModJatiPOL	6.66 (10.71)	0.59 (7.39)									
ModMapFRAC	-29.61 (20.99)		-9.80 (10.29)	4.08 (5.88)							
ModMapPOL	-1.55 (15.56)		-23.00** (10.78)								
ERPolar	-2.03 (1.81)			-4.41*** (1.11)							-4.46*** (1.07)
ReligionFRAC	6.68 (14.79)				17.01 (15.51)						
ReligionPOL	3.64 (10.84)				-7.62 (10.04)						
LanguageFRAC	4.84 (6.05)					12.07* (6.46)					
LanguagePOLAR	-6.42 (4.57)					-5.13 (5.11)					
Census2011FRAC	-20.23* (12.12)						-1.70 (3.73)				
Census2011POL	5.49 (7.54)							-1.55 (2.78)			
LandClassFRAC	-13.40** (6.14)								-27.75*** (2.90)		
LandClassPOL	-2.18 (6.30)								4.29 (4.24)		
LandMLD	0.50 (0.38)									1.49*** (0.19)	
LandModMap_BG	12.99 (11.11)									3.06 (9.92)	
LandReligion_BG	-19.51 (12.26)									-12.54 (10.69)	
LandCensus_BG	6.44 (10.78)									1.13 (10.81)	
Constant	40.43** (15.79)	25.49* (13.78)	53.62*** (16.52)	29.38*** (8.51)	18.81*** (6.07)	15.76** (7.29)	31.84*** (6.19)	31.98*** (6.03)	48.47*** (6.01)	25.00*** (4.56)	30.23*** (5.09)
R-squared	0.791	0.730	0.641	0.665	0.677	0.666	0.631	0.632	0.768	0.751	0.667
N	175	175	175	175	175	175	175	175	175	175	175

Standard errors are reported in parentheses (to two decimal places). * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

The data for dependent variable, HDI at village level, is from [Shivashankar and Prasad \(2015\)](#). Sub-district (*taluka*) HDI numbers are obtained as weighted averages of village HDI values (using census 2011 village population as weights). All models included appropriate ‘species count’ variable(s) indicating number of distinct social categories (for example, count of distinct *jaitis* in the sub-district). Additionally, all models also include controls for colonial administrative history (ten dummies); agro-ecological zone (six dummies), and literacy rates.

Table 18—: *Ethnic-geographic Aggregation and Diversity-Development Model Instability (OLS)*

	VILLAGE			SUB-DISTRICT		
	FRAC	POLAR	ETHNIC-INEQ	FRAC	POLAR	ETHNIC-INEQ
<i>Mean Per-capita Luminosity (2001-13)</i>						
Jati	-	-	NA	-	-	NA
Admin Map	+	-	+	-	-	-
Census Map	-	-	-	-	+	+
Religion	+	-	-	+	-	+
Language	+	-	+	+	-	NA
<i>Per-capita Luminosity growth (2001-13)</i>						
Jati	-	-	NA	-	+	NA
Admin Map	+	-	+	-	+	+
Census Map	-	-	+	-	-	-
Religion	+	-	-	+	-	+
Language	+	-	-	+	-	NA
<i>HDI, 2015</i>						
Jati	+	-	NA	+	+	NA
Admin Map	+	+	-	-	-	+
Census Map	+	+	-	-	-	+
Religion	+	-	+	+	-	-
Language	-	+	+	+	-	NA

The table presents a summary of the various OLS models that we have used to test the diversity deficit hypothesis at varying ethnic-geographic aggregation levels. The results for the three primary dependent variables (mean per-capita luminosity, growth in per-capita luminosity, and human development index) are presented for village level ($n \approx 27,000$) and sub-district level ($n = 175$) regressions. As seen from the table, the evidence for diversity deficit hypothesis is an unit of analysis artefact – with coefficient signs unstable across both ethnic and geographic aggregations, and across the three different dependent variables used here as indicator of economic, social, and human development. The table records the sign of coefficients on various social heterogeneity axes (as measured by fractionalization, polarization, or ethnic inequality) in respective regression models. Coefficients that are statistically significant ($p \leq 0.05$) are in colored font. See main text for further discussion.

Table 19—: *Quantile Regression Models: Summary*

	FRAC	POLAR	ETHNIC-INEQ
<i>Mean Per-capita Luminosity (2001-13)</i>			
Jati	$+,-,-$	$-,-,-$	NA
Admin Map	$+,+,+$	$-,-,-$	$+,+,+$
Census Map	$-,-,-$	$-,-,-$	$-,-,-$
Religion	$-,+,+$	$+,+,-$	$-,-,-$
Language	$+,+,+$	$-,-,-$	$+,+,+$
<i>Per-capita Luminosity Growth (2001-13)</i>			
Jati	$+,+,-$	$+,-,-$	NA
Admin Map	$-,-,+$	$+,+,-$	$-,-,+$
Census Map	$+,+,-$	$+,+,-$	$+,+,+$
Religion	$-,-,+$	$-,-,-$	$-,-,-$
Language	$-,-,-$	$+,+,+$	$-,-,-$

The table presents a summary of the various quantile regression models that we have used to test the diversity deficit hypothesis at varying ethnic-geographic aggregation levels. The results for the two primary dependent variables (mean per-capita luminosity, and growth in per-capita luminosity) are presented for village level ($n \approx 27,000$) quantile regressions. Each cell reports coefficient sign and significance for 25%, 50% and 75% quantile regressions. As seen from the table, the evidence for diversity deficit hypothesis is an unit of analysis artefact – with coefficient signs unstable across ethnic aggregation, and across two different dependent variables. The table records the sign of coefficients on various social heterogeneity axes (as measured by fractionalization, polarization, or ethnic inequality) in respective regression models. Coefficients that are statistically significant ($p \leq 0.05$) are in colored font. Thus, the regression results are not stable across different quantiles. See main text for further discussion.

Table 20—: *Regression Models by Fractionalization-Polarization Quadrants*

	FRAC	POLAR	ETHNIC-INEQ
<i>Mean Per-capita Luminosity (2001-13)</i>			
Jati	+,-,-,+	+,+,-,-	NA
Admin Map	+,-,+ ,+	-,+,-,-	+,+ ,+,-
Census Map	-,-,-,-	-,-,-,-	-,-,+,-
Religion	+,-,+ ,+	-,+,-,-	-,-,+,-
Language	+,+ ,+ ,+	+,-,-,+	+,+ ,+,-
<i>Per-capita Luminosity Growth (2001-13)</i>			
Jati	-,-,-,+	+,+ ,+,-	NA
Admin Map	+,-,+ ,+	-,+ ,+,-	+,+ ,+,-
Census Map	-,-,-,-	-,-,-,-	+,+ ,+ ,+
Religion	-,-,+,-	+,+ ,+,-	-,-,-,-
Language	-,-,-,+	+,+ ,+,-	-,-,-,-
<i>HDI, 2015</i>			
Jati	-,+,-,-	-,-,-,-	NA
Admin Map	+,-,-,-	+,-,-,+	-,-,-,-
Census Map	-,-,-,+	-,-,-,+	-,-,+ ,+
Religion	+,+ ,+,-	-,-,-,-	+,-,+ ,-
Language	+,+ ,+,-	-,-,+,-	+,+ ,+ ,+

This table presents a summary of the regression models for each of the quadrant along the fractionalization-polarization distribution as represented in Figure-13. Each cell reports coefficient sign and significance for NE ($FRAC \geq 0.5; POLAR \geq 0.5$), SE ($FRAC \geq 0.5; POLAR < 0.5$), SW ($FRAC < 0.5; POLAR < 0.5$), and NW ($FRAC < 0.5; POLAR \geq 0.5$). The results for our primary dependent variables (mean per-capita luminosity, growth in per-capita luminosity and HDI) are presented for village level ($n \approx 27,000$). As seen from the table, the evidence for diversity deficit hypothesis is an unit of analysis artefact – with coefficient signs unstable across the four quadrant, and across the three different dependent variables. The table records the sign of coefficients on various social heterogeneity axes (as measured by fractionalization, polarization, or ethnic inequality) in respective regression models. Coefficients that are statistically significant ($p \leq 0.05$) are in colored font.

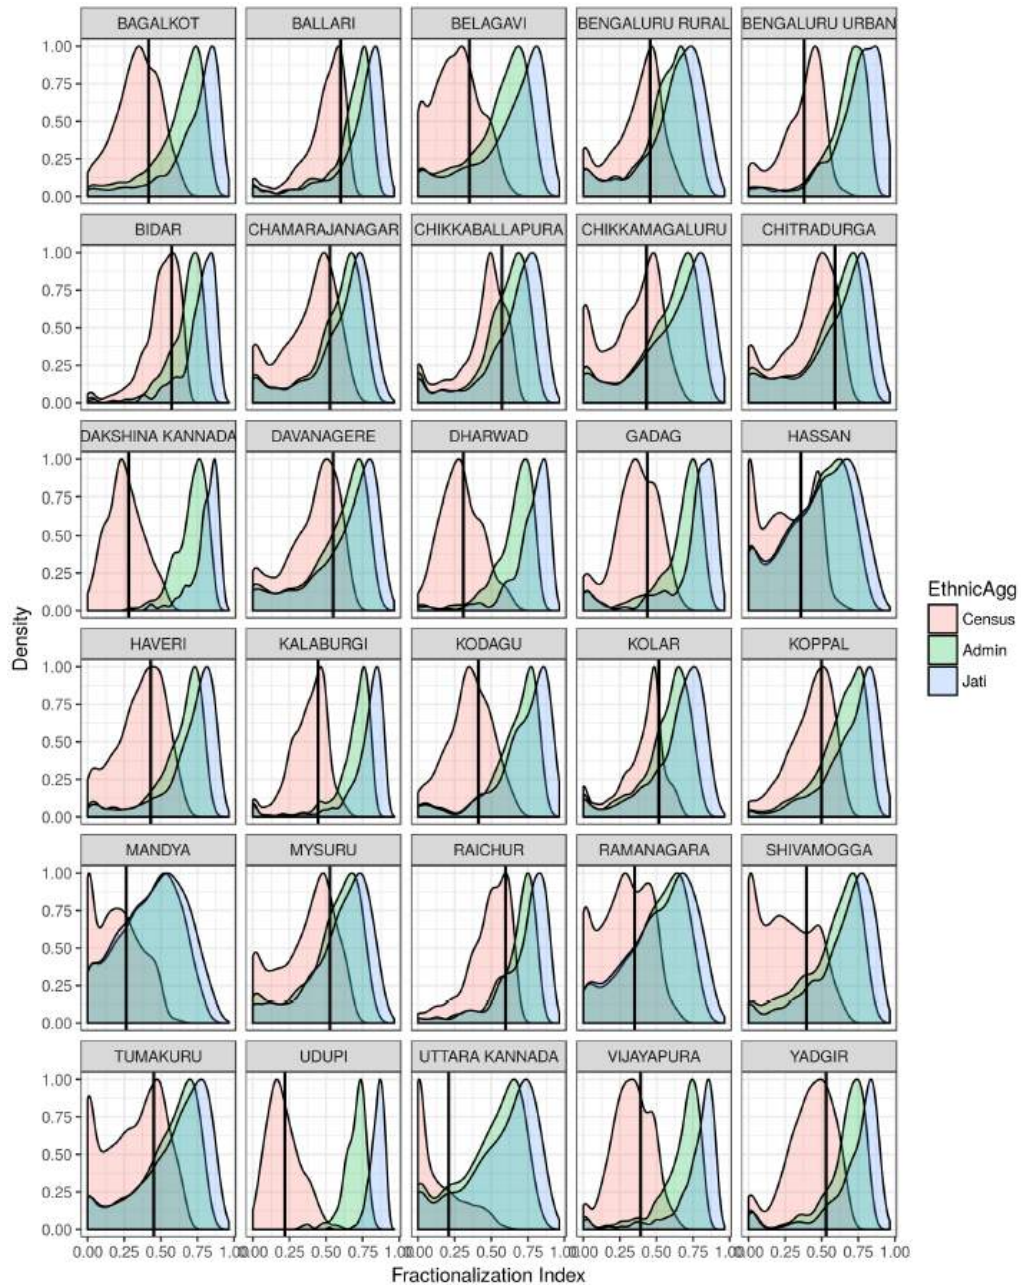


Figure 2. : *Distribution of Village-level Fractionalization by District*

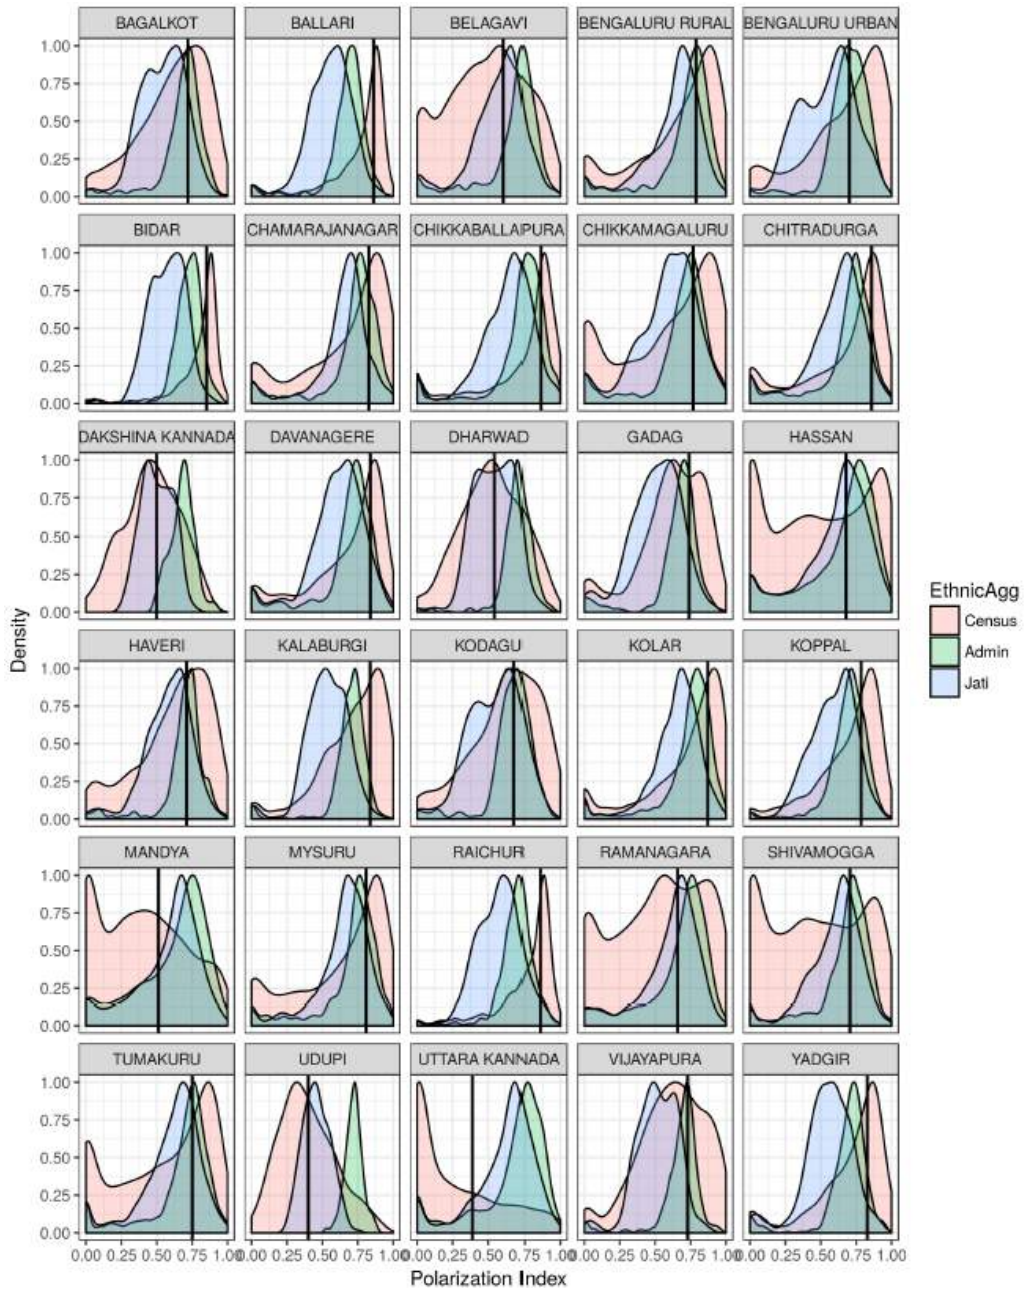


Figure 3. : Distribution of Village-level Polarization by District

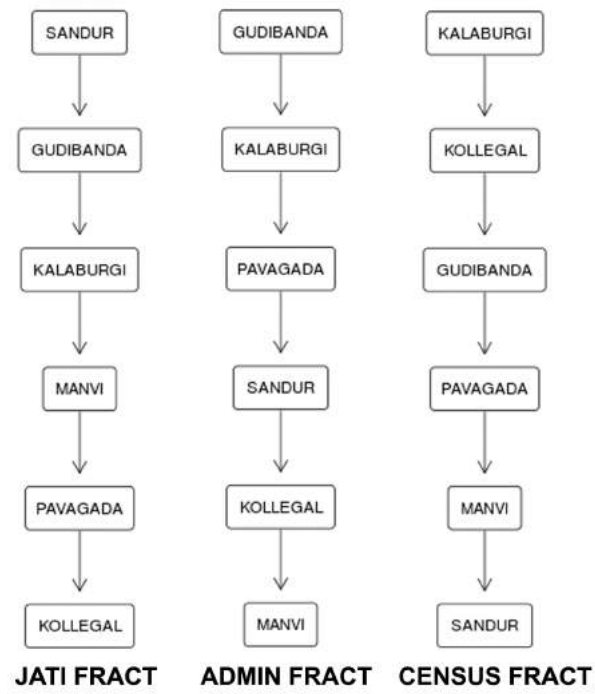


Figure 4. : *Hasse Diagrams, Sub-district Fractionalization*

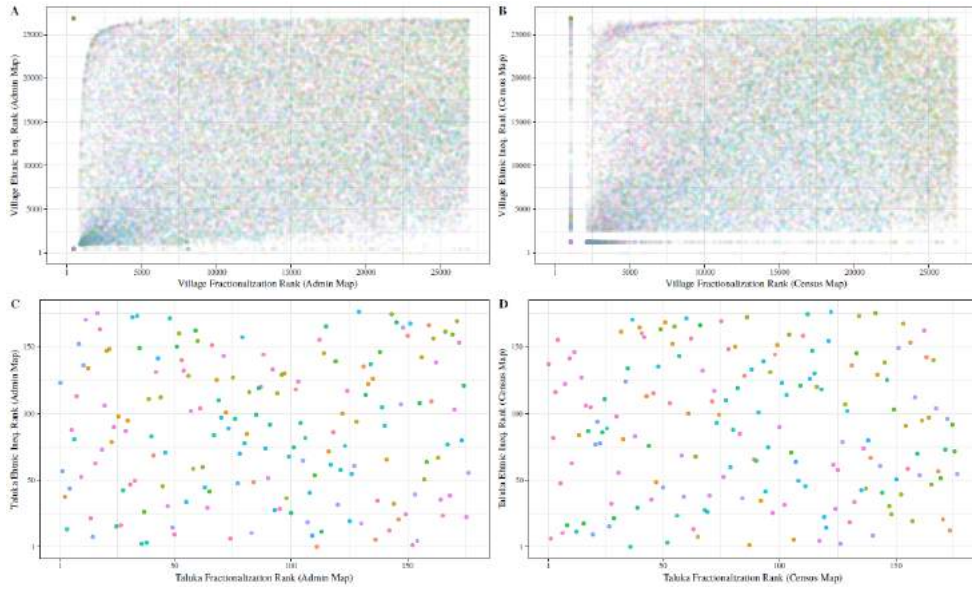


Figure 5. : Diversity and Ethnic Inequality Ranks




Figure 6. : Karnataka: Geographic Location, and Administrative History

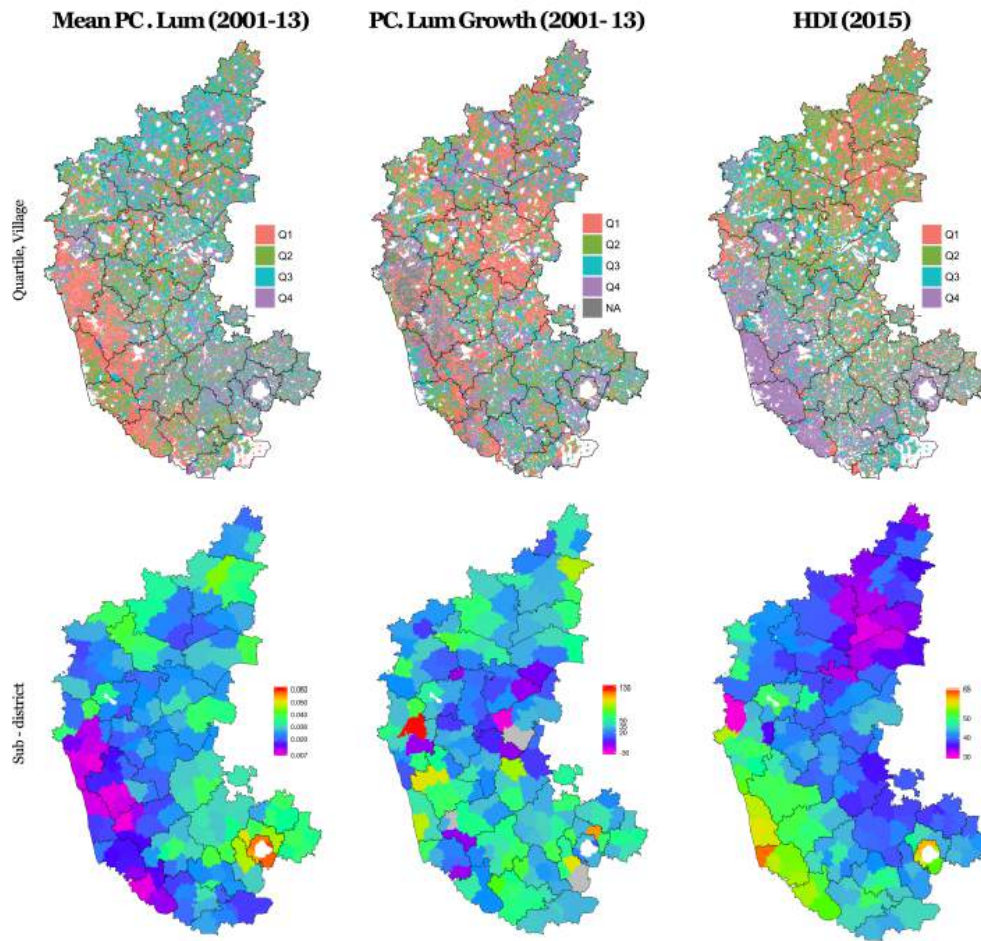


Figure 7. : *The Geographic Distribution of Dependent Variables. The top panel shows the quartile-rank for each of the three dependent variables used in our regression models at the village level. At the village level we have shown quartiles rather than actual variable values to better illustrate results from our quantile regressions at the village level. The bottom panel shows the geographic spread of actual variable values at the sub-district (taluka) level. For descriptive statistics, refer to Table 2*

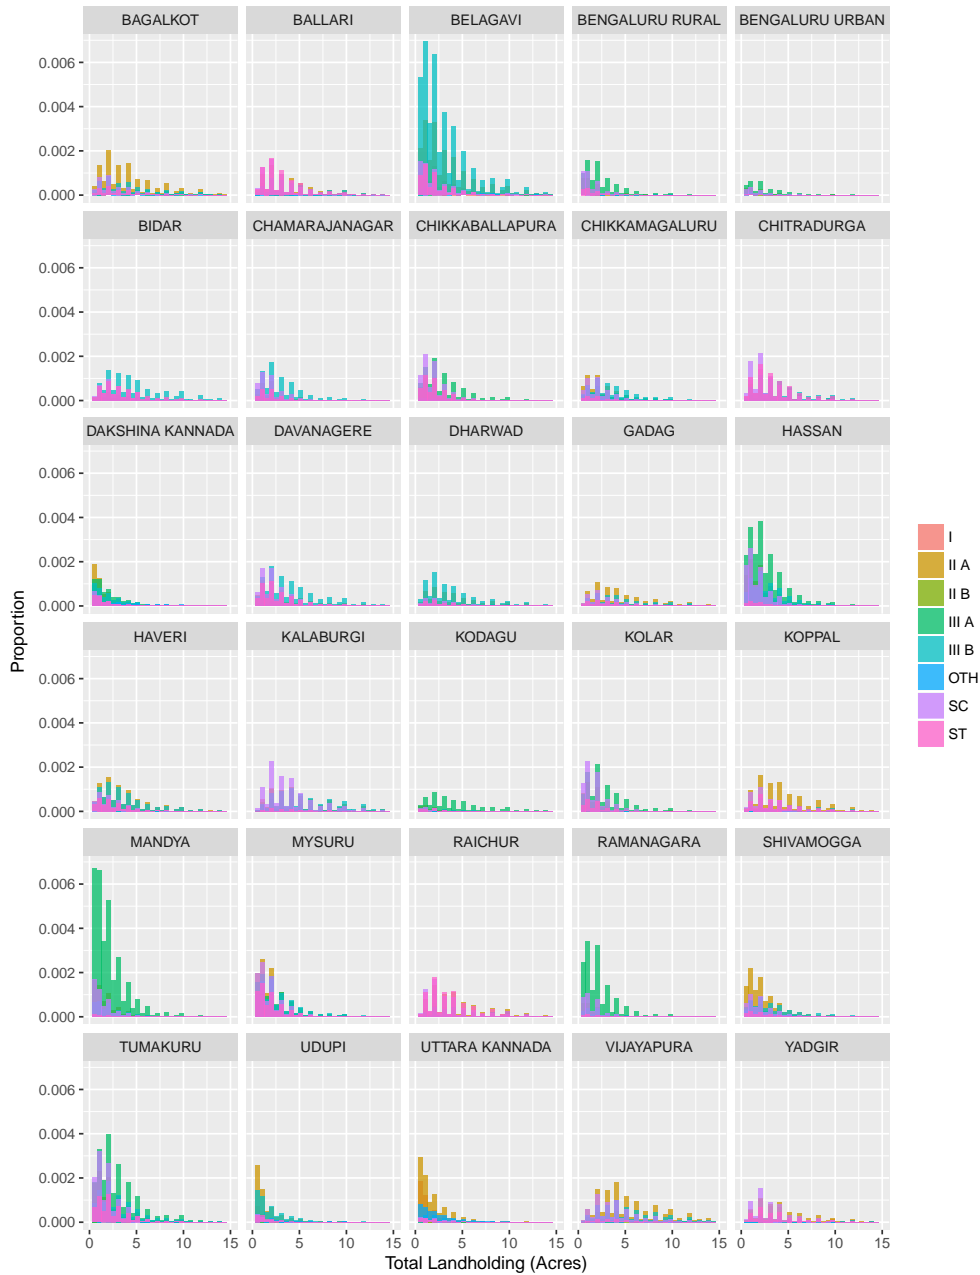


Figure 8. : *Landholding by Administrative Caste Categories. Only rural landed households with landholding below fifteen acres in 30 districts of Karnataka $n = 4,638,107$ are included here. About 40% of rural households are landless. Approximately 2% of rural households have landholding in excess of 15 acres and are also excluded here.*




Figure 9. : *Polarization Metric and Ethnic Aggregation. Village-level data ($n = 26,890$)*




Figure 10. : *Ethnic-Geographic Aggregation, and Multiple Diversity Contexts*

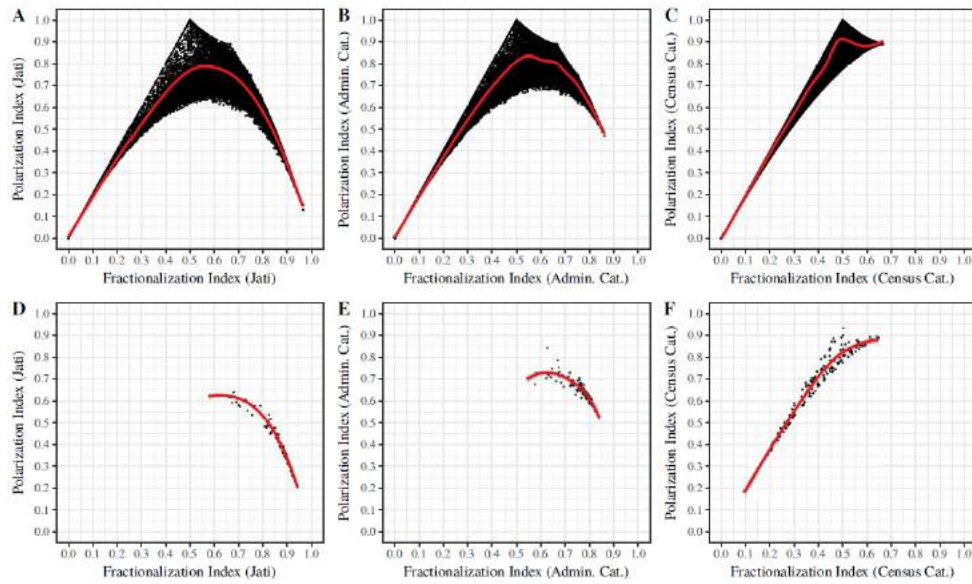


Figure 11. : *Fractionalization, Polarization, and the Ethnic-geographic Aggregation of Caste.*

Note: The top panel is from village-level data ($n = 26,890$), and the bottom panel is from sub-district data (*talukas*, $n = 175$). Both village and sub-district level data aggregated from a common household-level dataset – a census of all rural households in Karnataka ($n = 13,255,421$). Besides scatter points, all the six charts also show the Locally Weighted Scatter-plot (LOESS) fitted smoothing curve along with the 95% confidence-band. See main text for more explanation.

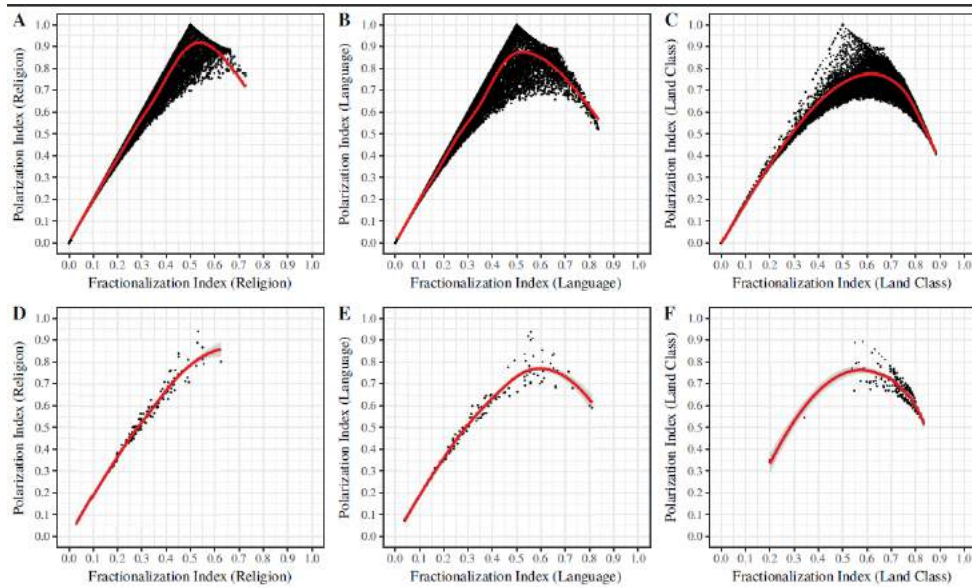


Figure 12. : *Fractionalization, Polarization, and Geographic Aggregation of Religion, Language, and Economic class*

Note: The top panel is from village-level data ($n = 26,890$), and the bottom panel is from sub-district data (*talukas*, $n = 175$). Both village and sub-district level data aggregated from a common household-level dataset – a census of all rural households in Karnataka ($n = 13,255,421$). Besides scatter points, all the six charts also show the Locally Weighted Scatter-plot (LOESS) fitted smoothing curve along with the 95% confidence-band.

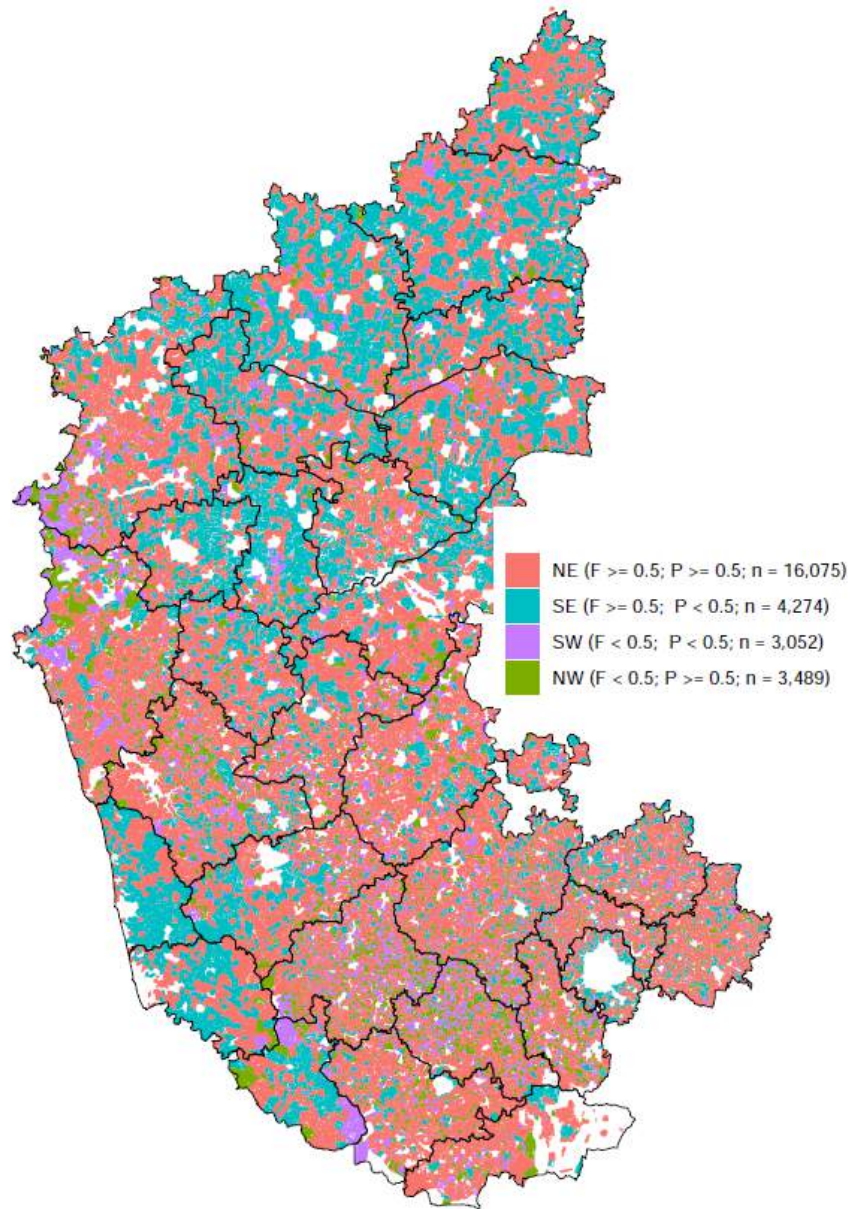


Figure 13. : *Villages in Karnataka by (Jati) Fractionalization-Polarization Quadrant*

Note: Census-designated urban areas are not in our data and are shown in white on the map. Missing data from villages ($< 0.1\%$ of all inhabited villages) are also shown in white. Data from $n = 26,890$ villages. District boundaries are shown.

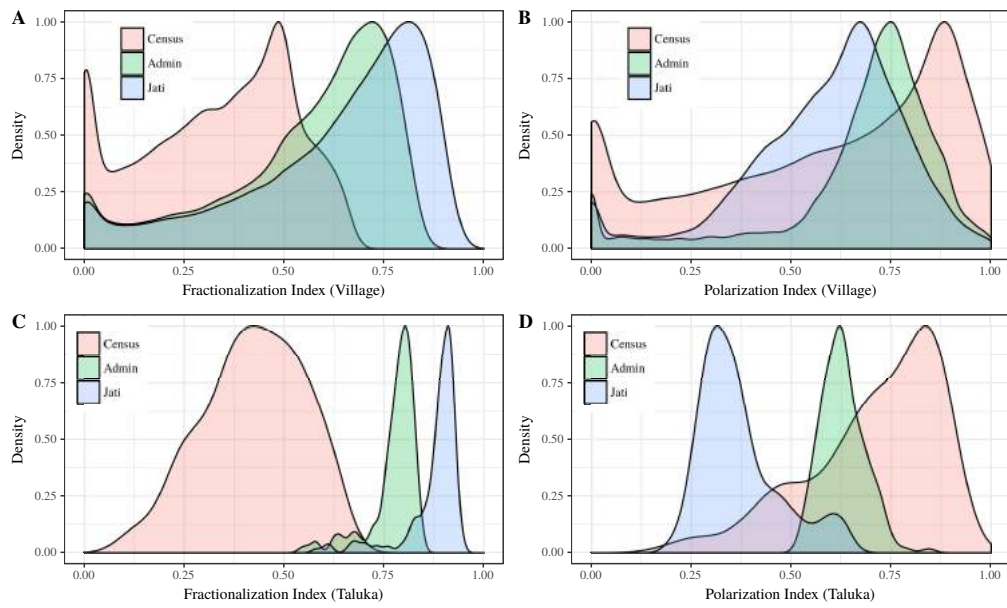


Figure 14. : *Ethnic-geographic Aggregation of Caste: Density Plots*

Note: The top panel is from village-level data ($n = 26,890$), and the bottom panel is from sub-district data (*talukas*, $n = 175$). Both village and sub-district level data aggregated from a common household-level dataset – a census of all rural households in Karnataka ($n = 13,255,421$).

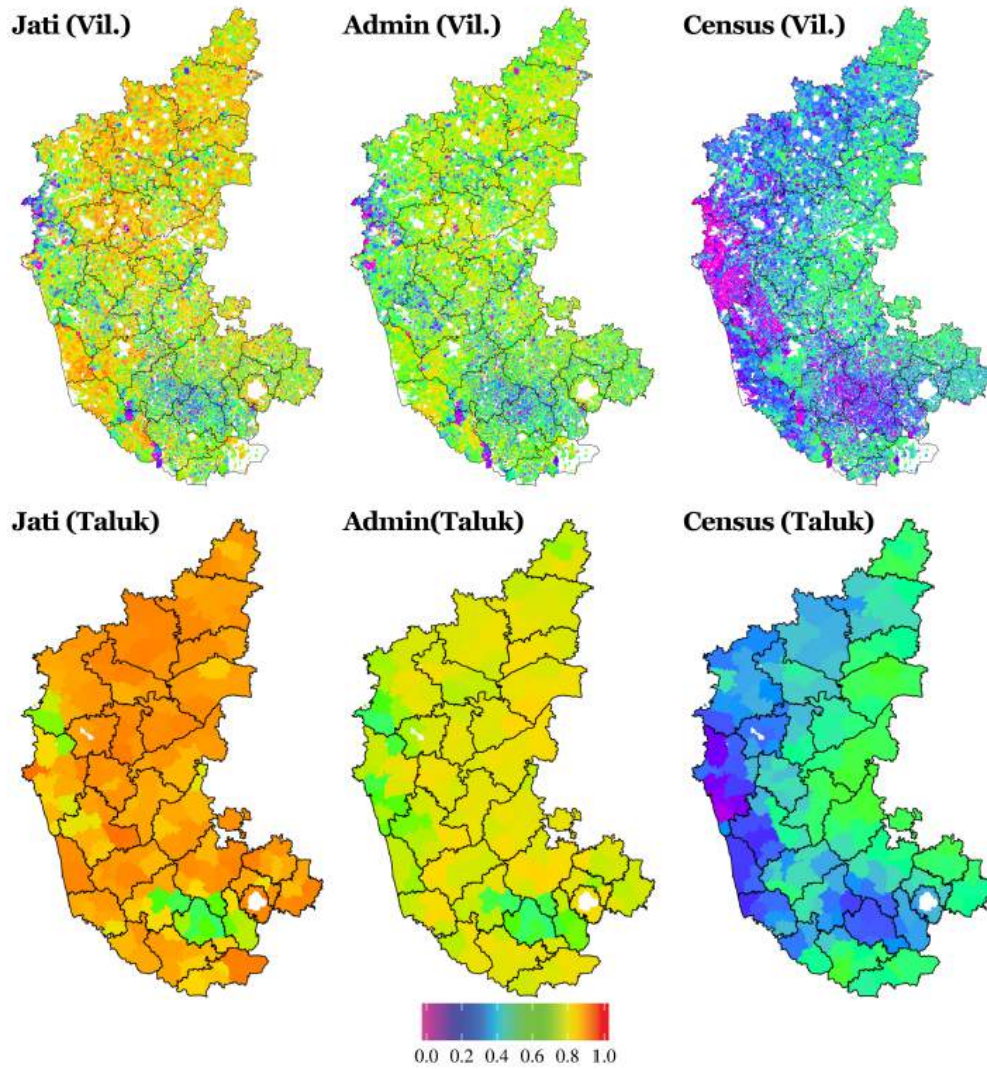


Figure 15. : *Ethnic-Geographic Aggregation and Fractionalization*

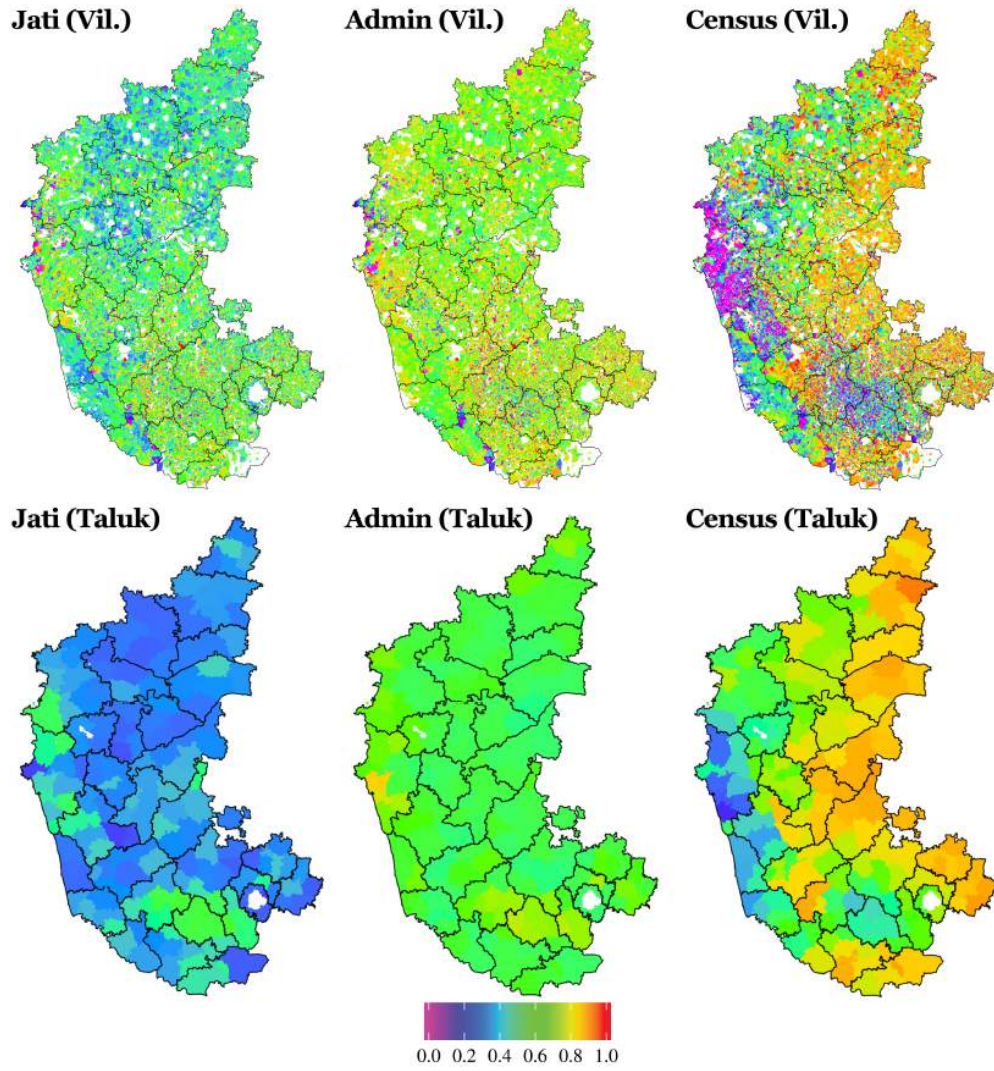


Figure 16. : *Ethnic-Geographic Aggregation and Polarization*

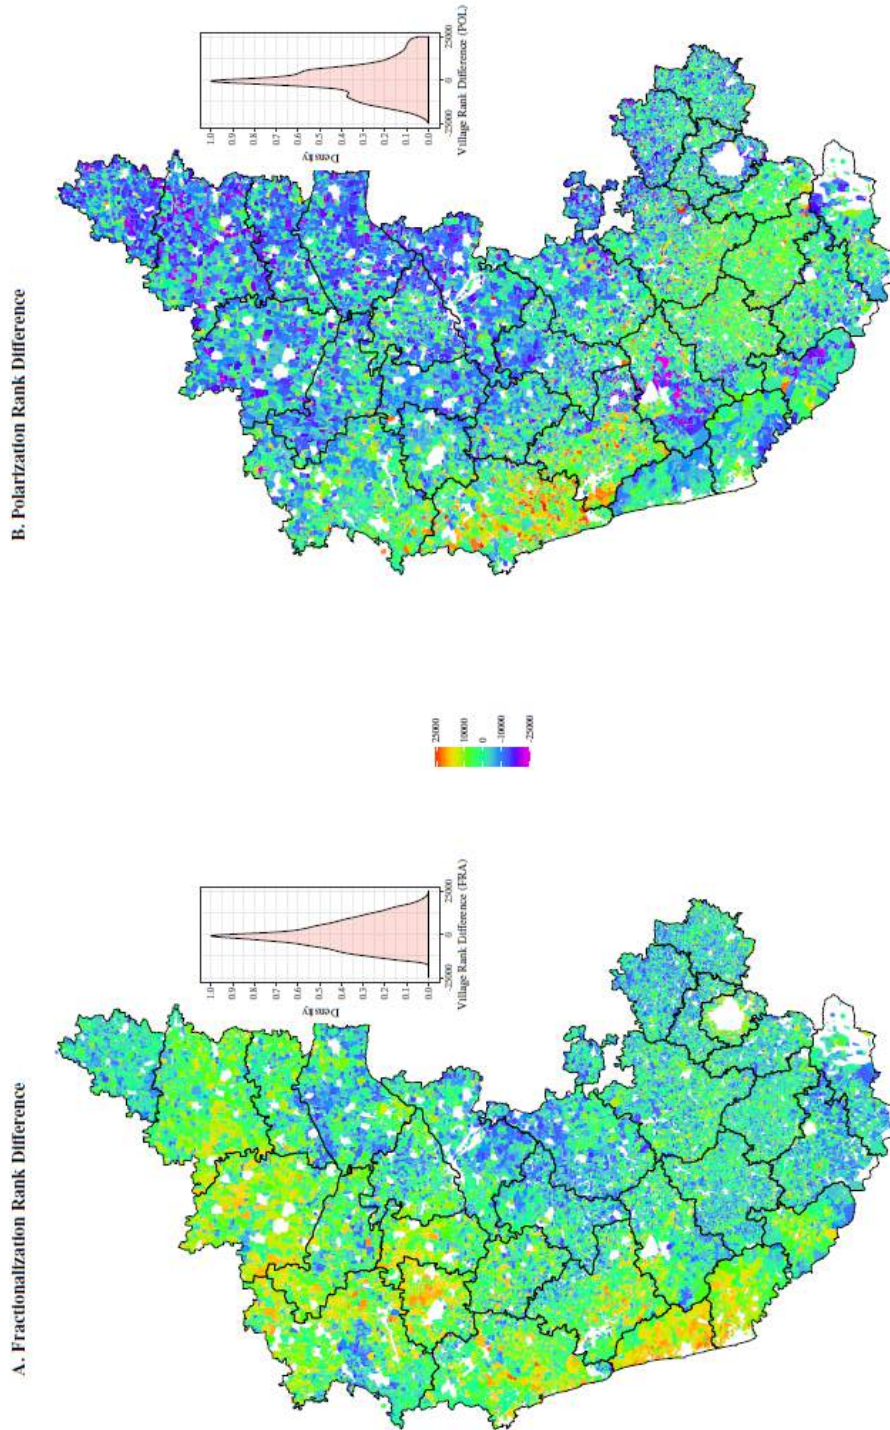
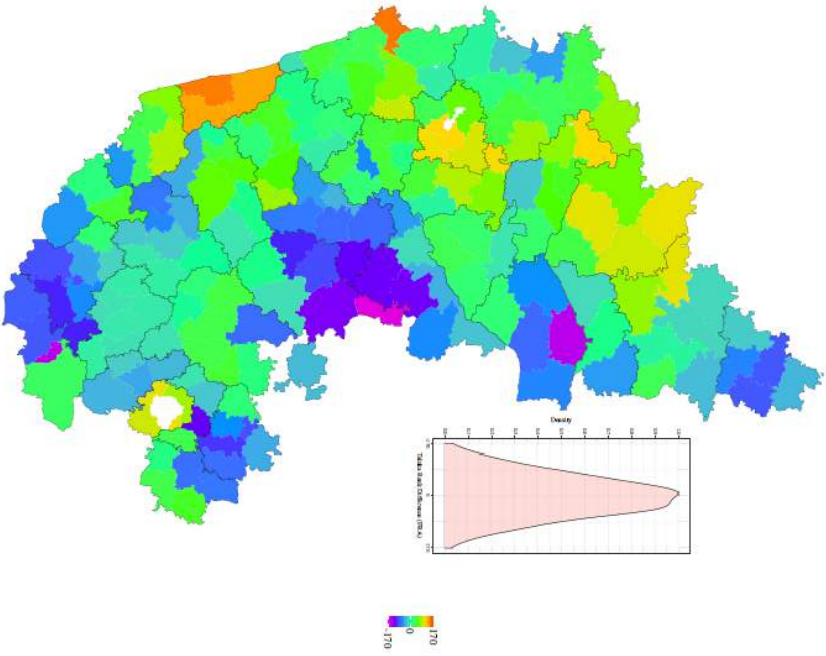


Figure 17. : The Spatial Structure of the Modifiable Ethnic Unit Problem. The left panel represents the difference in ranking of villages ($n = 26, 890$) when ranked by jati fractionalization and when ranked using the three-fold census categories. The density plot (shown as an inset) of this difference (ModJatiFRAC – Census2011FRAC) hides the uneven spatial spread of ranking-difference. The right panel uses corresponding polarization metrics. District boundaries are shown on both maps. See text for more explanation.

A. Fractionalization Rank Difference



B. Polarization Rank Difference




Figure 18. : *The Spatial Structure of the Modifiable Ethnic Unit Problem. The left panel represents the difference in ranking of sub-districts ($n = 175$) when ranked by jati fractionalization and when ranked using the three-fold census categories. The density plot (shown as an inset) of this difference (ModJatiFRAC – Census2011FRAC) hides the uneven spatial spread of ranking-difference. The right panel uses corresponding polarization metrics. District boundaries are shown on both maps. See text for more explanation.*