Causal Inference with Spatially Disaggregated Data: Some Potentials and Limits

Castner, Johannes
jac2130@columbia.edu

Dumas, Marion
mmd2172@columbia.edu

Gočev, Petr
petr.gocev@yale.edu

May 2012

Abstract

In studies of civil strife, the ecological fallacy seems to befall all large-n studies and thus there has been a big push, by several researchers, in recent years to gather disaggregated, spatially explicit data. An interesting example is the recently assembled geo-reference Ethno-Power Relations dataset (Wimmer et al.), and used by Cederman et al. (2011) to argue that political exclusion and group-level economic inequality lead groups to organize to rebel against the political status quo. While we think there is much potential in this new approach to measuring the determinants of conflict, we find that the resulting data can not be analysed in conventional ways, if the estimation of causal effects is the goal. The reason is that analysis of sub-national entities may bring about other dangers: the violation of the Stable Unit Treatment Value Assumption (SUTVA). To be specific, one "treated" group's enemy could hardly be its control. We get around this problem by changing the causal effect of interest and by carefully re-aggregating the lower level data so as to preserve its most salient information. Restricting our analysis to groups that are excluded from power, we find some tentative evidence that such groups are less likely to engage in conflict if they are more spatially integrated with other groups.

1 Introduction

In recent years, several authors have pointed to the pervasive problems of ecological fallacy pervading the quantitative empirical literature on civil conflict (Buhaug and Lujala, 2006; Gleditsch and Ward, 2001). Indeed, salient variables are usually summarized by the polity-level vector of their means, but such means may describe no one. Key sub-national interactions and dependence structures are left out. Thus, there has been a push for spatially disaggregated data. The PRIO team has put together an impressive set of geo-referenced data for both conflict and the explanatory variables commonly invoked in the literature, conveniently integrated through a common 0.5'x0.5' grid of the world. Overall, this trend towards greater availability of spatially disaggregated yet systematic data seems to be a promising turn for the empirical literature and better still, for creatively complementing case material (Fearon and Laitin, 2008). However, we contend that this new data brings forth important methodological challenges that are overlooked in analyses so far. To illustrate, we reanalyze the data and model recently contributed by Cederman et al. (2011), in which they report that political and economic inequalities between ethnic groups lead to civil conflict. We do this as a way to address methodological issues arising throughout the recent literature that uses disaggregated, sub-national data to evaluate causes of conflict onset and to probe the potential and limits of these new datasets.

Questions concerning the role of inequality and ethnicity in conflict are arguably those for which arguments calling for spatial disaggregation of the description of countries cry the loudest. Indeed, country-level analyses have largely dismissed ethnicity as a relevant factor in civil conflict. But they summarized the distribution of ethnicities in a polity by a single measure of diversity, the fractionalization index (e.g. Fearon and Laitin, 2003). Inequality as summarized by the country GINI of household-level incomes, has similarly been dismissed (Esteban and Ray, 2011). However, several authors have argued that neither diversity as measured by the ELF nor inequality aggregated by the GINI captures the tensions arising from the covariation between ethnic groups and the distribution of power and resources, or what some authors have termed "horizontal inequalities" (Stewart, 2008) as they overlap - in space! - with ethnic identities. In the context of this debate and the push for spatially disaggregated data, Wimmer, Cederman and Min (2009) assembled the Ethno-Power Relations dataset (EPR) via a large-scale consultation of area experts. It encodes access to executive power of politically organized ethnic groups from 1946 to 2005 for nearly all countries in the world and in addition, provides data in the form of a shapefile on the geographic location of these groups. The aim was to describe how ethnicities relate to the state and open the door to large-N analysis of the role and functioning of ethnic politics in influencing political and economic outcomes.

In this paper, we show that applying familiar statistical models[1] to sub-national entities amounts to pooling units of observation that are in fact interdependent, or drawing counterfactual comparisons that are logically inconsistent or that do not address the question at hand. We explain how, by opening our choices over the relevant scale at which to define the unit of observation, spatially disaggregated data brings forth an important tension between SUTVA and ignorability, two fundamental assumptions of causal inference. We argue that for most questions relating to the causes and consequences of conflicts, at least those involving the state, carrying inference with sub-national entities as units of observation dramatically violates the assumption of non-interference. In this case, spatially disaggregated data may need to be re-aggregated at the polity level, albeit in ways that capture the covariance structures it helped reveal and deemed causally important for political instability and violence. We advocate for new theoretical work to build better indices reflecting features of the sub-national data, as have in the past the work on polarization by Reynal-Queyrol and Montalvo (2008) and others. We also suggest that the new wave of spatially disaggregated data may lend itself to a fruitful new wave of descriptive inference, as well as to theory testing other than by causal effect hypotheses.

The rest of the paper is organized as follows: Section II presents the data, Section III lays out the methodological challenges to be addressed and shows why the original question is unanswerable. Section IV delimits answerable causal questions and Section V presents the results of our reanalysis.

[1] In Cederman et al. (2011), the model was a logit specification for time series binary cross sectional data as specified in Beck, Katz and Tucker (1998) and widely used in cross-country regressions since then.

2 Spatially disaggregated data: an overview of new data sets

The data set we examine and assembled by Cederman et al. (2011) codes the access to power of "politically relevant ethnic groups" from 1946 to 2005. Specifically, it distinguishes between nine power status categories: monopoly group, dominant group, group sharing power as junior or senior partner, excluded from politics with further distinction between discriminated groups, powerless groups or groups enjoying some degree of regional autonomy. These are *de facto* and not *de jure* power relations. A group is considered ethnic if it constitutes a community with a "subjectively experience sense of commonality based on a belief in common ancestry and shared culture". An ethnic category is then considered politically relevant "if at least one significant political actor claims to represent the interests of that group in the national political arena, or if members of an ethnic category are systematically and intentionally discriminated against in the domain of public politics" (quotes from Wimmer, Cederman and Min, 2009). It is important to note that since the data records groups on the basis of their engagement in politics, it does not represent the whole universe of ethnic groups and thus cannot be used to understand what causes groups to engage in ethnic politics, only perhaps what causes a group to engage in conflict conditional on it already being politically organized[2]. In this data set, politically relevant groups have a start and end date and their political status may change from year to year. Thus, as argued by several authors (e.g. Posner, 2004), ethnic groups - as they relate to the political system - form and come undone over time: they hardly are fixed entities.

The ethno-power relations (EPR) data is combined with the Armed Conflict Dataset (ACD, by Gleditsch et al. 2002) and its supplementary information on rebel organizations and their affiliations, the Non-State Actors dataset (NSA). The NSA was complemented by Cunningham et al. (2009) who expanded the coding of government-insurgent dyads back to 1945. NSA includes information on attributes of rebel groups for 404 separate conflict

[2]While the definitions given by the authors are quite clear, it is hard to know what politically relevant, as well as *excluded,dominant* etc... mean in practice. In the Appendix, I compare the subset of African groups with the groups in Ethnologue, as well as Posner's selection of "ethnic groups" deemed relevant for competition over macroeconomic policy ($PREG_{macroeconomic}$, exposed in Posner, 2004). I also select 10 groups at random and link the data to secondary qualitative case material

dyads, including the groups it draws on for recruitment and whom it claims to represent. The ACD defines an armed conflict as "a contested incompatibility that concerns government or territory or both where the use of armed force between two parties, at least one of which is the government of a state, results in at least 25 battle-related deaths." It includes information on location of operations and the name of the opposition organization. It is important to note that when compared to other data sets of organized violence, the ACD represents a small subsample. The spatially disaggregated economic data comes from the dataset put together by Nordhaus and Chen (2006), with regional estimates of wealth at a 1'x 1' resolution for 1990, 1995, 2000 and 2005.

3 Challenges of inference on the causes of civil conflict using sub-national data

3.1 The requirements of causal inference reveal a fundamental tradeoff between cross-polity and within-polity analysis

Isolated from their theoretical background, the hypotheses motivating many statistical analyses in the "large-n" literature on civil conflict could be interpreted non-causally; they could be interpreted as purely descriptive propositions or as theoretically predicted statistical patterns without making reference to causal effects of specific variables. However, authors writing in the large-n tradition typically motivate their hypotheses as proposition of salient causal mechanisms and thus they often give causal interpretation to their findings. Cederman, Weidmann and Gleditsch (2011), for example, set out to test whether political and economic inequalities between ethnic groups *lead them to mobilize* into violent collective action, through the formation of grievances and group identification - mechanisms they draw from social psychology. In light of the qualitative theory on which such studies are based, they should thus be interpreted as causal and accordingly the validity of their inferences should be judged in light of the requirements for valid causal inference[3].

There are four main requirements for causal inference to be valid. In examining them in the context of the causal relationship between group-level political economic inequalities and civil conflict, we unveil some of the implicit assumptions made in prior analyses. In doing so, we show that having spatially disaggregated measures to describe a society does not imply that we can straightforwardly transfer our statistical models from the polity level to sub-national entities. In the discussion that follows, we consider the causal effect of a variable B on an outcome variable Y and a set of covariates Z. In referring to the substantive questions asked by Cederman et al. (2011), B is in some instances political exclusion (as opposed to inclusion) and in others it is being far off the average wealth of the country (either positively or negatively). Following the standard notation in the causal inference literature, we call $Y(b = 1|Z)$ the potential outcome of unit i for a treatment variable B given Z and call $Y(b = 0|Z)$ the potential outcome for the value of B that is considered to be the control.

3.1.1 Requirement 1: No un-modeled interactions between units of observation (SUTVA)

The technical term for this assumption in the statistics literature is the Stable Unit Treatment Value Assumption (SUTVA for short). It requires that the causal effect of a variable B on the outcome Y for each unit of observation be unaffected by the value of B for other units and further that the value of B for i be unaffected by the other units. In the analysis of the effects of horizontal political and economic inequalities on conflict using groups as units of observation, SUTVA translates into an assumption that group A's access to power affects the decision of its members to instigate conflict independently of the access to power by other groups in the polity, and - on the economic side - that group A's share of wealth pacifies or enrages its members independently of any other groups' share (but of course, a moment's reflection should convince the reader that this is impossible, as any one group's share of the country's wealth is a direct function of the shares of all other groups in the country). The slightest presence of strategic interactions in this process

[3] We refer here to the Neyman-Rubin potential outcomes framework, supplemented by the non-parametric structural models framework by Judea Pearl, as exposed in Morgan and Winship (2007), and Glynn and Quinn (2008).

makes the SUTVA assumption untenable[4].

Cognizant of the fact that groups within a polity interact, Cederman et. al (2011) include two polity-level controls: the size of the group relative to those in power and the number of excluded groups. However, this does not fix the problem: keeping these variables fixed leads us - implicitly or explicitly - to compare the behaviors of excluded and included groups in polities taking on similar values of Z, perhaps within the very same polity (this is in anticipation of our discussion of the geographic dimension of matching below). This approach does not alleviate the fact that the "control" group(s) to which the excluded groups are compared respond to the current status of the excluded group: the control group is hardly a control, but rather a full participant in the unfolding drama.

For some questions and some settings, it is possible to model interactions within the causal inference framework (see Chapter 8 of Gerber and Green, 2012). To return to our running example, assume that the propensity of a group to violently contest the status quo depends on that group's political or economic status ($b = 1$ versus $b = 0$) and some feature of the distribution of political and economic status of other groups in the same polity, which we call $\mathbf{s}$. Then the true potential outcome under the treatment is $Y(1, \mathbf{s}|Z)$. Here the treatment is an *interaction* of the group's characteristic with that of the other groups that are present in the same polity. Assuming for argument's sake that there are two types of distributions s and s^*, we could make comparisons between $Y(1, \mathbf{s}|Z)$, $Y(0, \mathbf{s}|Z)$, $Y(1, \mathbf{s}^*|Z)$ and $Y(0, \mathbf{s}^*|Z)$. Depending on the question at hand, two logical inconsistencies may arise. First, if i's treatment determines what distribution s it can be associated with, then some comparisons may not be possible. For example, while it is possible to find an included group in both inclusive and non-inclusive polities, one can only find excluded groups in countries that in fact do channel power towards a non-inclusive set of groups. Thus, one has to be very careful in defining the question and pick counterfactuals accordingly. Second, if i's treatment depends on j's treatment while they both influence each others' outcomes, then it is not possible to keep s fixed while varying b and vice versa. In this case, it

[4]SUTVA is distinct from the assumption of stochastic independence, which is violated when units are influenced by common unmodeled factors. Unlike SUTVA, stochastic dependence can be corrected for by directly modeling spatial and temporal auto-dependence (for example, using a spatial error lag model).

becomes impossible to formulate a causal effect question at the group-level. Figure 1 reformulates the argument using a game matrix. It is interesting to realize that more comparisons are logically possible if we redefine the causal variable(s) as being features of the full distributions $S = \{1, s\}$, $S' = \{0, s\}$, $S'' = \{0, s^*\}$ and $S''' = \{1, s^*\}$, in other words, polity-level characteristics (for example comparing a polity marked by equal resource distribution to one marked by great between-group inequality). In all cases, if there is interference (or rather interdependence!) between groups, we are lead to redefine the causal effect at the level of the polity, not of the group. Rebellion may be manifested regionally, but in most cases, at least if the state is involved, it is a local manifestation of a process unfolding in the political system as a whole, which probably should be the unit of analysis for most causal questions about civil conflicts[5] - albeit informed by the structural features characterizing the polity as revealed by disaggregated data. We will come back to this issue to carve out some questions emanating from the theoretical literature that can be answered in the causal inference framework and with the type of data discussed here.

3.1.2 Ignorability of treatment conditional on controls

The choice of the set of control variables to be included is surely the most important modeling step in the estimation of a causal effect. We say that a causal effect is identified by conditioning on a set of control variables, Z, if and only if all back-door paths in the directed acyclic graph between the causal variable, B and the outcome variable Y, are blocked after conditioning on Z (Pearl, 2000, Glynn and Quinn 2008). For a more thorough explanation of Judea Perl's "back-door criterion" see Winship and Morgan (2007).

However, typically, the choice of Z is made by considering which other variables have been found significantly associated with the outcome in prior papers. The back-door criterion clarifies what can go wrong with such a strategy: in addition to including variables that are unnecessary to answer the question at hand[6], so-called collider variables may be included that will cause spurious associations between B and Y. In particular, Morgan and

[5]This argument needs a more precise mathematical formulation to make clear what exact conditions limit the possibility of causal inferences over interacting units

[6]which increases model dependence, or alternatively reduce the set of valid comparisons, and narrow the scope of the local average treatment effect.

Winship (2007) note that the common practice of using proxies of suspected unobserved omitted variables as controls will fail because these proxies are usually colliders (a collider is a variable which is causaly effected by two or more other variables; the dangers of including colliders and their "descendent" as adjustment variables are also made clear in Glynn and Quinn, 2008). One example that is often found in the literature is the use of GDP as a proxy for state capacity and everything else about state success that seems related to both conflict and its hypothesized causes. Multiple factor summary variables such as the Polity Score will probably cause a similar problem in many cases.

One justification that is often cited for a kitchen sink strategy (to add as many pre-treatment variables as possible to an analysis) is to provide a robustness check of the estimated effect. Because the task of thinking about each possible pre-treatment factor and how it fits in a causal graph is daunting if not impossible, it would seem that a sensible thing to do would be to introduce fixed effects (for regions and/or time periods). However, apart from the fact that such dummy variables may be collider variables and thus should *not* be included in Z for proper identification, the more humbling realization is that the more similar the units are and close to satisfying ignorability, the more they are likely to be interdependent and thus violate SUTVA. Herein lies the fundamental tradeoff between SUTVA and ignorability, manifesting itself in our problem as a difficult choice between cross-country and within-country comparisons (should we compare exclusive or excluding *political systems* to inclusive ones or should we instead compare excluded and included *groups*?).

3.1.3 No post-treatment bias

This condition demands that Z contain only variables that are causally prior to B. Indeed, controlling for something occurring after B will subsume part or all of the effect of B on Z since it may be a result of B. For example, it does not seem possible to both estimate the effect of economic and political inequality on conflict at the same time: economic inequality may cause political inequality in the next period and vice versa, while both may directly contribute to a group's desire to contest the status quo. Both may also result from prior conflicts. As in the case of no-interference, many systems of interest to political scientists are replete with mutual causations, ruling out asking certain questions within the causal inference framework or forcing very careful and explicit construction of structural models, which then serve as a common ground for defending the plausibility of causal interpretations of patterns found in the data.

3.1.4 No undue extrapolation

What we mean by no undue extrapolation is that the functional form used in the systematic component of a parametric statistical model (e.g. linear) must be well established theoretically or - in the absence of theoretical justification for one form versus the other - the inference must make minimal assumptions about the functional form, for example by using non-parametric methods such as matching. Otherwise, the results will rely on extrapolation in the area that lies outside of the common support, which can be highly biased if the functional form is wrong (see King and Zheng, 2000). As an illustration, we have found the causal effects of horizontal inequalities on conflict, reported by Cederman et al. (2011), to be highly dependent on the specification of the functional form: they implicitly rely on comparisons with counterfactuals that are not in the data and that are therefore constructed by means of extrapolation via the logit form, which itself is not well justified for event processes (Alt et al. 2001). Figure 1 shows this graphically: it replicates Figure 5 in Cederman et al. (2011), showing the variation in the causal effect on violence as the economic status of a group varies and for excluded and included groups. We added a measure of uncertainty (simulating both estimation and fundamental uncertainty of the causal effect) and coded values on this curve the estimation of which relies on extrapolation: these are the points which lie outside of the convex hull - or common support of the data - (see King and Zheng, 2000). We see that the causal effect reported by the authors is both very tenuous once all sources of uncertainty are taken into account and relies on extrapolations that are dependent on the logit form being correct. To summarize the issue of model dependence, while a model is essential to make inference from observational data (as the earlier points should have made clear), this model need not be parametric and in fact it probably shouldn't be unless our theory is precise enough to pick out and defend a particular parametric choice. Given the state of the theory concerning the problem that concerns us here, we rely on non-parametric inference (via causal graphs and matching) for the remainder of this paper.

Although one can argue that the four requirements are too stringent for most questions of interest to social scientists, including the study of group char-

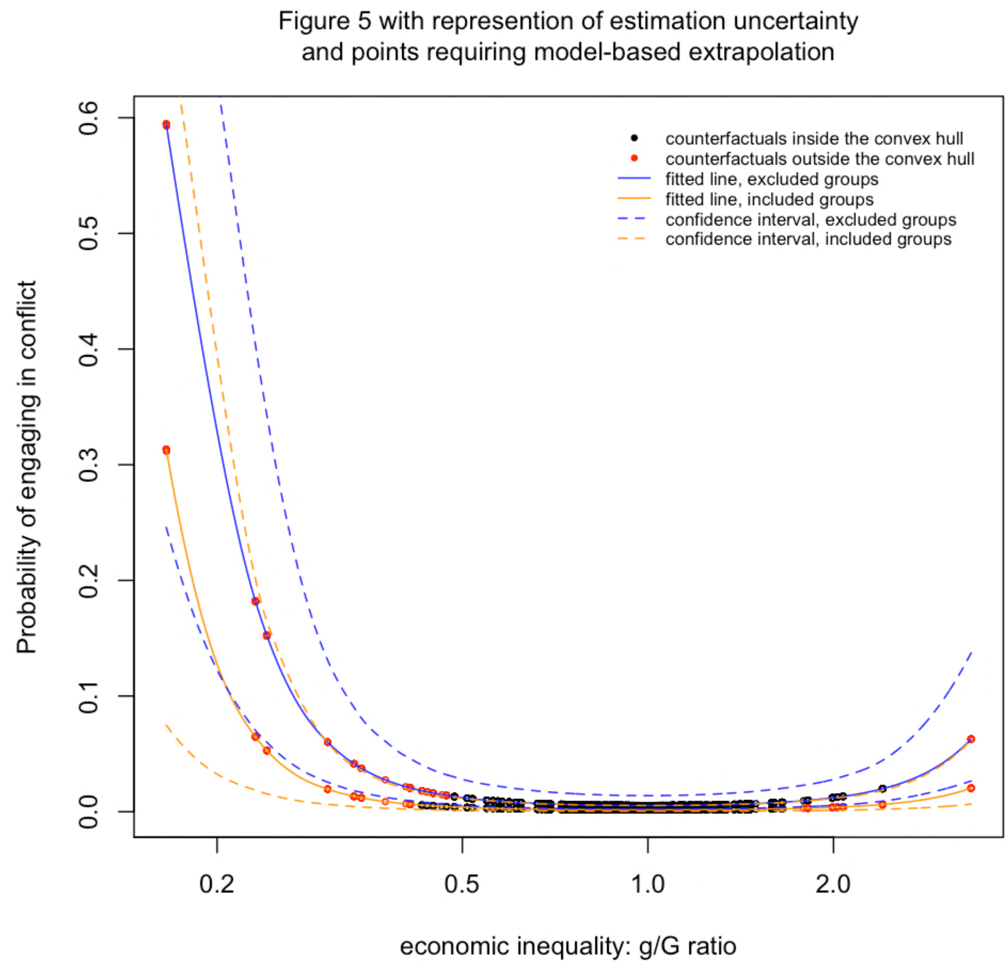

Figure 1: Model dependence in the estimation of the causal effect of political exclusion and economic inequality on groups' engagement in conflict, where the causal effect is estimated keeping all covariates at their median value. g/G is the measure of inequality used in Cederman et al. (2011): the ratio of a group's wealth to that of the country average. Saying that counterfactuals are outside of the convex hull means that they are not in the support of the data and thus rely on extrapolations.

acteristics' role in political instability, we contend that they are extremely useful for clarifying our thinking about which questions are answerable given the state of our theory and given the data that we have at hand at any point in time. In addition, it helps us make reasoned choices concerning the analytical lens with which to approach a social phenomenon: can we reliably learn something of interest by thinking in terms of causal effects? by asking descriptive questions? or by engaging in theory testing? In what follows, we will try to better circumscribe what type of causal effect questions linking group-level political-economic inequalities to civil conflict can be answered given the stringent requirements of causal inference. We will illustrate our claims by asking what the causal effect of ethnic homogeneity is on politically excluded groups' propensity to contest the political status quo. We will lay out the assumptions needed for this inference and show how the results change if we refuse to maintain SUTVA at the group level and accordingly shift the inference to the cross-polity level. Remaining at the polity level, we will also ask what effect certain features of the economic distribution of wealth within and across groups in a polity have on political stability.

4 Explicit definition of counterfactuals

Matching is primarily conceived as a technique to reduce model dependence in making inferences from observational data, especially inferences of causal effects (Ho et al. 2007). Given the model dependence identified in recent analyses, we thus turn to matching. However, matching is not just a technique to neutralize parametric assumptions: the no-interaction assumption and the explicit reasoning about counterfactual that matching forces upon us both help us delimit the set of meaningful causal effect questions one can ask about the political-economic determinants of civil conflicts.

4.1 Group-level and polity-level variables

Some forms of matching such as Coarsened Exact Matching, or the construction of synthetic counterfactuals in quantitative case studies (Abadie and Gardeazabal, 2003) force us to explicitly define which units are to be compared: can we compare an excluded group with an included group to measure the effect of exclusion on civil conflict? Does it matter that the included group is in the same polity as the excluded group? If it is in a different polity, does it matter that the control is in a fully inclusive polity on another continent? Or is it more meaningful to compare power-sharing polities to polities that exclude some to the detriment of others? We have argued that in making these types of choices, there is an inherent tradeoff between SUTVA and ignorability. This tradeoff seems particularly salient in sub-national studies of civil conflict. For example, in constructing a synthetic control for the Basque Country to study the effect of the insurgency on the local economy, Abadie and Gardeazabal (2003) built a synthetic counterfactual Basque region by finding the optimal set of weights to be put on multivariate distribution of covariates of other Spanish regions to best mimic the determinants of growth in the Basque country in the counterfactual situation of no conflict. But in doing so, the regions most similar to the Basque country were chosen, which were Madrid and Catalonia, regions that probably are both affecting and being affected by the insurgency in the Basque country. Could the author have chosen regions outside of Spain? or would that allow too many unobserved factors to change[7]? Thus, although we started our analysis with sub-national political entities, the no-interference assumption lead us to argue that cross-country comparisons may in many cases be the most natural when formulating questions about civil conflict as causal effect questions, at least those conflicts involving the state (as opposed to inter-communal violence)!

4.2 Mapping counterfactual matches

To illustrate our argument graphically, we ran the Coarsened Exact Matching algorithm on one year of Cederman et al. (2011)'s original data and using the same variables as they did in their model (which includes both group-level and polity-level covariates). Specifically, this procedures picks out excluded groups and included groups that take on the same values on a coarsened version of the covariates for each combination of values of the coarsened covariates. We then retrieved all the so-called subclasses thus formed and picked ten of them at random and mapped them, as shown in Figure **??**. What we can see is that some subclasses are highly clustered in space (even

[7] A way of perhaps bounding the effect and quantifying this tradeoff in particular applications would be to draw a distribution of synthetic counterfactuals by varying a parameter giving more or less weight to various fixed effects or distance from the studied case. This was a suggestion made by Kara Ross Camarena. Work may soon start to assess the potential of and develop this solution.

in the same country) while others are dispersed in space. This means that some comparisons are made on a very local scale and some on a much more global scale. This makes sense since the matches were made both on polity-level and group-level covariates (including variables such GDP/capita which clearly will create regional regrouping). Had we included region dummies, even more subclasses would have been closely clustered in space. Hence, we see that matching proceeds oblivious of space, unless we pay close attention (regression also proceeds oblivious of space in making comparisons and it is harder to uncover the comparisons actually being made). This can be a problem if either SUTVA and/or stochastic dependence are a concern. In the latter case, we may need to include a spatial variable to prevent matches that are too close (assuming the causal question still makes sense when doing that). In the former case, we can still model the spatial dependence in the model run after having matched units.

Ten Randomly Selected Matching Strata

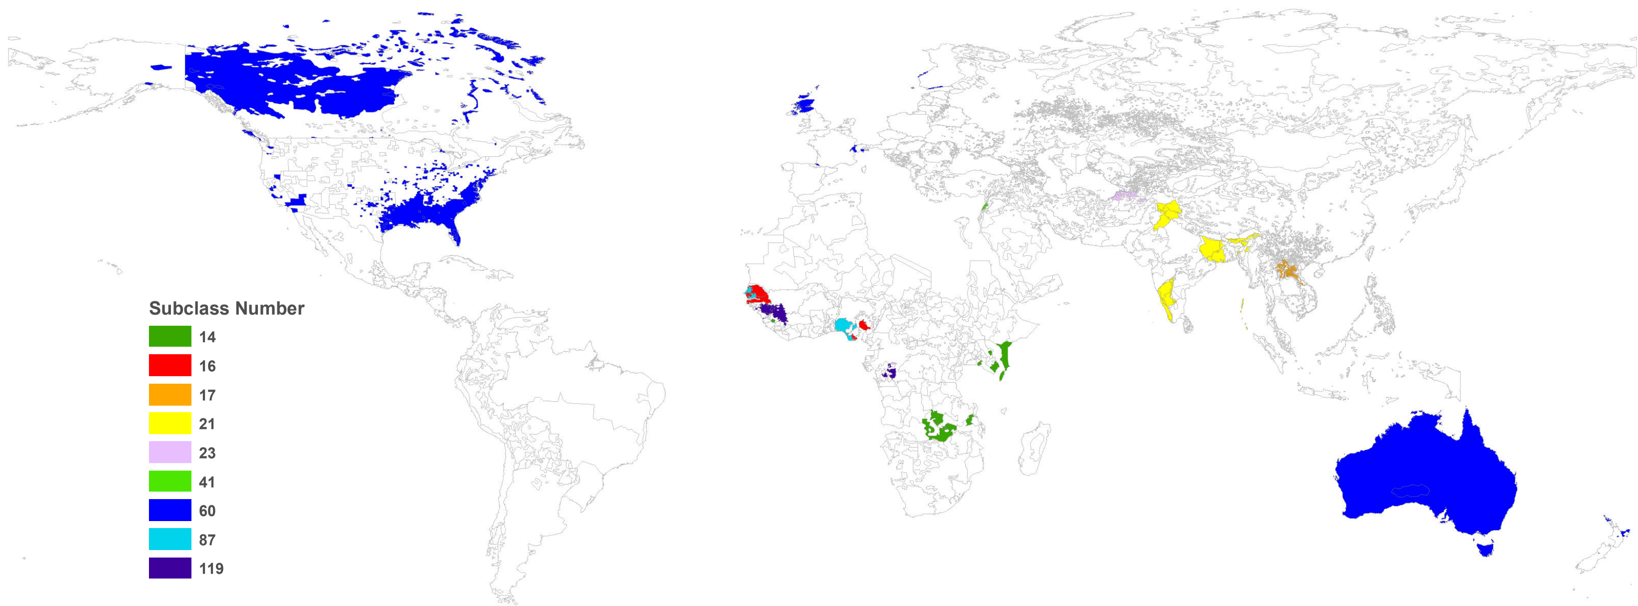

4.3 Group-level logically consistent questions

Warning against inference over sub-national units of observation does not mean that disaggregated data is useless. First, some causal effect questions may be more severely prone to violating ignorability than SUTVA and thus justify a within-country, group-level question. Second, and most importantly, characteristics of the joint distribution formed by the disaggregated data may be key. Imagine that we have at hand a large number of variables thought to bear some relationship to the prevalence of conflict in a political system. One way to rephrase our core argument is to say that for many causal questions about political stability, one would want to compare one distribution to another rather than sub-parts of the distribution to each other. Up to now, this multivariate distribution has usually been described by the vector of means. But this is obviously not the only way it could be described. For example, consider two non-inclusive polities each with two competing groups and suppose the analyst is convinced that economic inequality is an important factor influencing the conflictual dynamics between these groups. It surely must matter whether the politically dominant is also the economically advantaged group or whether instead the excluded group is endowed with the major share of the country's resources. This is the type of information that one gains with more disaggregated data and that one ought to capture by appropriate statistics of the joint distribution. If informed by case material and theory we judge spatial structure to be important, we must find ways to make use of disaggregated data to appropriately reflect salient structural features in the polity level variables. We will now explore these two ways in which spatially disaggregated data can be used by re-analyzing the EPR data with new questions.

4.3.1 Theoretically informed questions on ethnic group characteristics and political stability

Several theories imply or propose group-level characteristics that affect the propensity of an excluded group to form an organized opposition to engage in contentions politics, which can escalate into violence. We can distinguish between factors causing sub-national tensions, factors preventing these tensions from being resolved through the political process and factors that facilitate violent collective action. Tensions may arise because of diverging interests in key economic domains, for example the allocation of land rights, redistribution through transfers and public goods and access to employment (Bates and Yackovlev, 2002). They may also be cultural (alienation of religious freedom or exclusion of languages). Quite clearly, the distribution of resources affects the prospects of a reconfiguration of power relations for all parties involved and thus also lead to claims between ethno-regional groups (Acemoglu and Robinson, 2002). Spatial features may make these tensions more likely or exacerbate them: ecological gradients may force the proximity of communities engaged in different types of agrarian production, distance from the center of power, while a center and periphery configuration may lead to unequal provision of public goods or diverging preferences (La Ferrara and Bates, 2001). This latter factor, interacted with degree of market integration and size of the polity, has been predicted to affect groups' choices to gain territorial independence (Alesina and Spolaore). Implicitly, these theories assume that salient ethnic groups are also spatially organized, forming ethnically homogeneous regional entities. Independently of other factors, this latter factor has been found in some context to affect the organizational capacity of a group to mount a movement (Lewis, 2010) - a facilitating factor for rebellion rather than a cause of tension. The degree to which ethnic groups form homogeneous regional entities can at least to some extent be measured with the Geo-EPR dataset: indeed groups may overlap (see Figure 3). We thus ask: does the degree to which a politically organized but excluded ethnic group forms an ethnically homogeneous regional entity within the polity affect the likelihood that it will rebel against the political status quo?

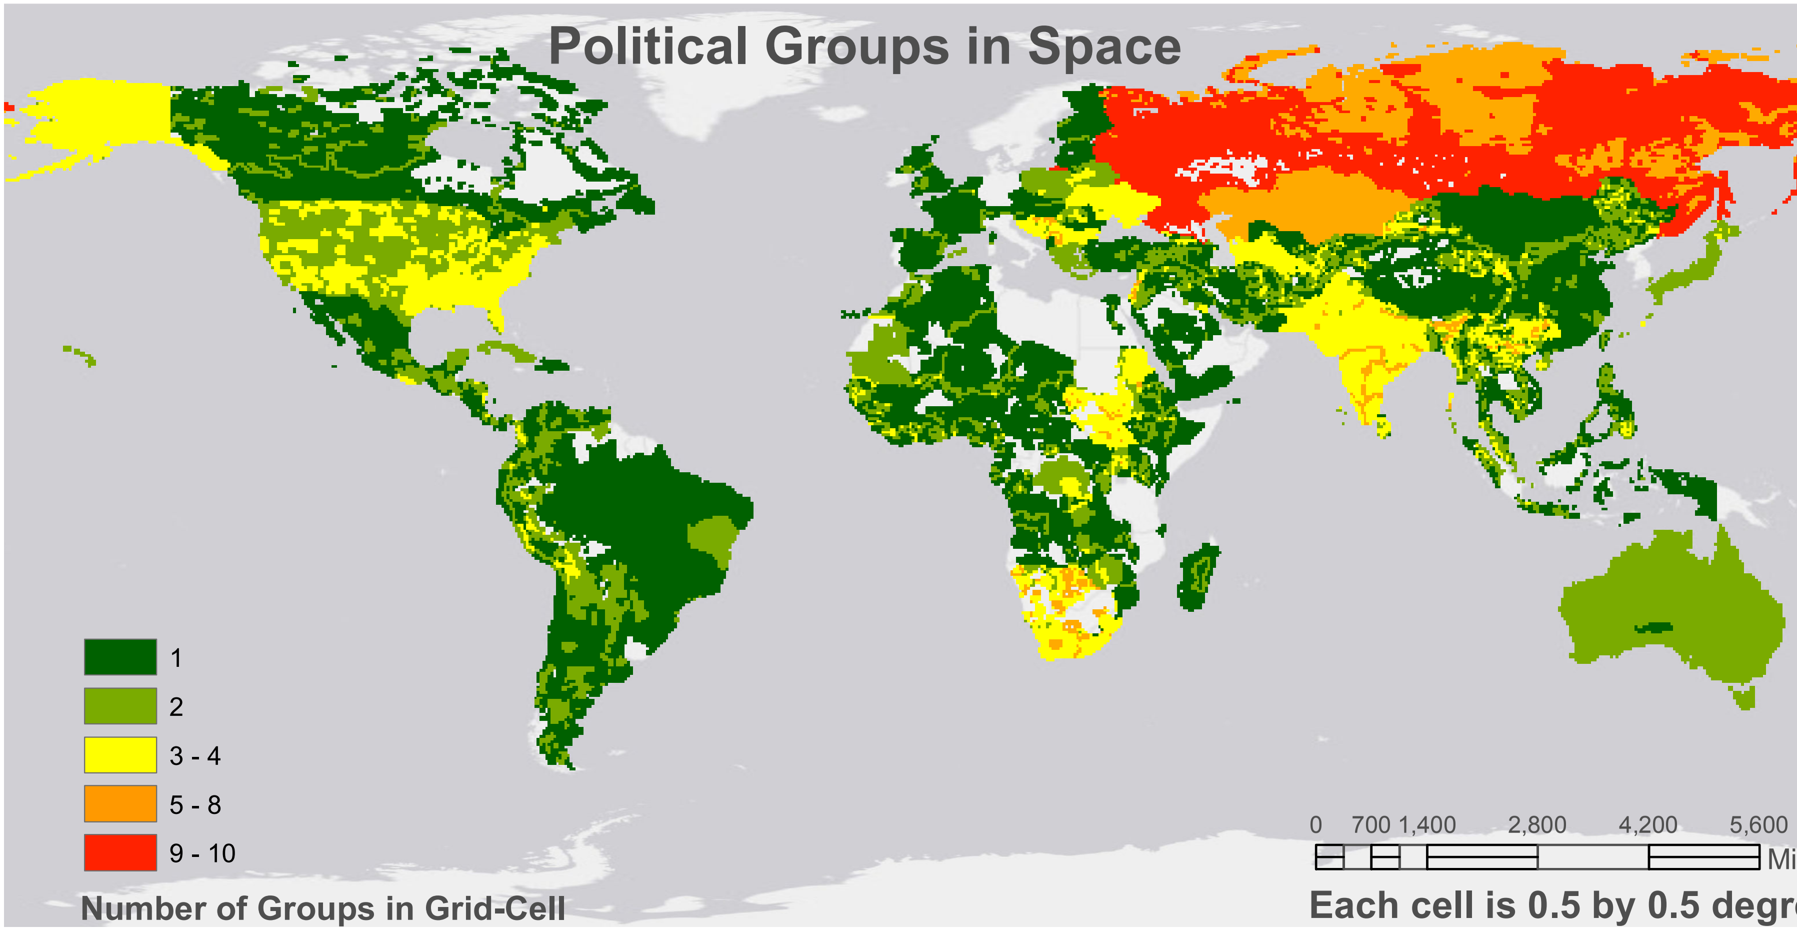

Figure 3: A map of the number of political organized ethnic group in each cell of the 0.5’ by 0.5’ grid. This map does not represent diversity: most of sub-Saharan African is coded as having only one group because in each region delineated by a cell, the EPR datasets indicates that there is only one group present. In contrast India has three to four groups almost everywhere, while Russia’s whole territory seems made up of overlapping groups. Note that some of the overlaps are artefacts of the rasterized version of the geographic dataset: cells that span two groups are coded as including both groups even if those do not overlap. This must be corrected for in order to give the reader the right perception.

We contend that unlike asking whether the political economic status of a group affects its chances of instigating conflict, the question just posed allow counterfactual comparisons that seem to reasonably withstand SUTVA and ignorability. Looking first at within polity comparisons, let us consider two excluded groups in the same polity with different degrees of regional separation from other groups (let us call the treatment spatial segregation, by which we mean a high level of ethnic homogeneity within the home region of an ethnic group). It seems at least plausible that the spatial segregation of one group does not affect the way in which the other excluded groups are affected by their own degree of spatial segregation. We do not know what causes spatial segregation in the first place (most likely a complex life history of the group) and given that this variable is static or changing very sporadically for most groups in our dataset, choosing pre-treatment variables and judging the ignorability assumption are difficult tasks. Ideally, we would include variables coding political history before the period at which a certain pattern of spatial organization crystallized, but such variables would be very idiosyncratic and unlikely to be systematic confounders of the process across the full population of cases. We recognize that this argument might fail if ethnic spatial segregation is systematically linked to the expansionary dynamics of agricultural systems (as described in Bates and Yackovlev, 2002), which could also independently affect ethnic tensions. However, it is a process probably too fine to have been measured by the area experts' geographical referencing of ethnic-political groups in this survey, even though the referencing changes to some degree over the time period. Amongst the observables at hand (including going to the PRIO database to include more covariates - see Data section below), distance to the capital and other "remoteness factors" (forested and mountainous land), as well as a lag of the number of groups relative to the territory and population seemed to be most sensible without inducing post-treatment bias.

Theories that emphasize the role of the state (Azam and Mesnard, 2003 or Bates, 2008) make clear that the factors described above are necessary but not sufficient conditions for conflict: the state must also be incapable of meeting the demands and the threat of collective action. Thus, the outcome variable in our causal question is *propensity to engage in collective action to challenge the status quo if claims are not satisfied by other means* and should be measured not by armed conflicts only, but by instances of protests etc... Unfortunately we do not yet have systematic data on collective action at the group level and so must for now carry out a test of the importance of this factor with group-level conflict data.

4.3.2 Polity-level questions

If we are not willing to maintain the assumption that the spatial segregation of one excluded group affects the spatial segregation of and its effect on other excluded groups, we cannot choose groups as units of analysis and must shift the question to the polity-level. We thus ask whether polities where excluded groups are less spatially integrated are more prone to violence compared to those where excluded groups are more integrated. The data produced to answer this question allows to reassess a related and important hypothesized effect of the demographic balance of groups: that near dominance by one group leads to fear of permanent exclusion for others (Bates and Yackovlev, 2002), thus creating a "danger zone" for ethnic politics.

Turning back to the original question of the effect of horizontal inequalities on conflict that opened this paper, we reformulate them by asking about the effect of different *distributions* of power and wealth *between groups* on civil conflict: 1) Do political configurations involving a greater share of the population excluded from power have increased risks of civil conflict[8]? 2) Do political configurations in which excluded groups also enjoy a lesser degree of welfare have increased risk of conflict, (measured by differentials in GDP/capita as well as by differences in the natural productivity of the land)? We are not convinced that these are the most key features of the distribution of power and wealth along ethnic lines that actually matter (and it is hard to know before having a more complete assessment of what the EPR dataset

[8]We note that we do not ask whether exclusion per se breeds conflict, as have others (Wimmer, Cerderman and Min (2009)) because after explicitly considering the counterfactual, the question loses some of its substantive meaning. Indeed, recall that we are studying conflicts classified as ethnic and as involving a non-state actor against the state on the one hand and ethnic groups' access to power on the other. Almost automatically by definition of the outcome, we are going to find that *ethnic rebellions* are waged by and large by excluded ethnic groups rather than included ethnic groups: descriptive statistics are sufficient to establish this fact. We could extend the outcome to conflicts other than those involving non-state and state actors divided across ethnic lines to other types of conflict, but that would artificially loosen the quasi-tautological status of the initial question since we know that different causal mechanisms are at play in different types of conflicts. We thus focus on degrees of exclusion. The type of tautology just described seems frequent in the literature: for example, it is quite obvious that groups near a frontier have a higher likelihood of being engaged in a secessionist war compared to groups embedded in the territory, yet such question has been asked of this kind of data (Buhaug et al., 2006).

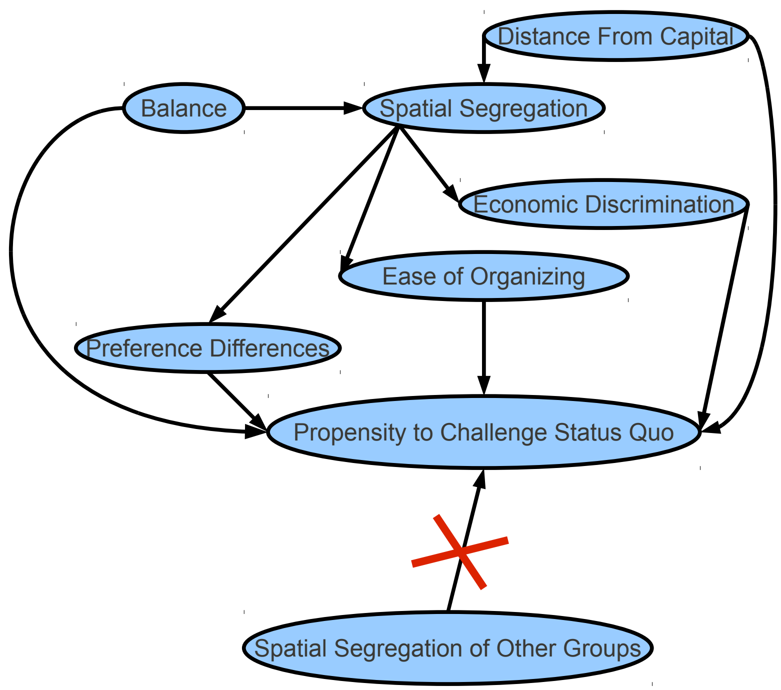

Figure 4: The causal graph underlying our inference. "Remoteness factors" and the demographic balance of groups included as plausible pre-treatment and confounding covariates. Note that the arrow with a cross means that we are assuming no interference in the effect of the treatment by other groups' treatment.

really measures, see forthcoming Appendix on this point), but we propose these questions as illustrations of the idea that one can re-aggregate spatially disaggregated data in different ways to capture different hypothesized mechanisms. Although the causal effect now seems more logically defined, making ignorability assumptions seems to be a daunting task, especially without a longer panel to construct lags of the wealth distribution. This is a limit of the data at hand, and our results are thus descriptive.

5 Re-analysis of the Ethno-Power Relations dataset

5.1 Defining spatio-temporal units of observation

The first step in our re-analysis has been to transform the original data-set to code the variables of interest at the degree of aggregation we deemed appropriate. Turning first to the temporal dimension, in the context of matching, it seemed infeasible to distinguish units of observation over time on the basis of years since years are not independent of each other and it is unclear how to control for that dependence within a matching framework. We thus re-aggregated our units in time by defining political periods instead. We defined a political period as a span of time defined by a given political status quo: any change in the power status of a group or any change in regime (to and from democracy, autocracy and anocracy) was coded as a change in period, which made the new temporal unit of observation, giving us group-political-periods as our units of observation for the first set of questions and country-political-periods as our units in the second set.

Turning to spatial aggregation, we went back to the most disaggregated data available (the 0.5'x0.5' gridded dataset made available by PRIO). We coded our causal variables - spatial segregation - for each grid cell (0.5' x0.5') in which an ethnic group was coded as present in the Geo-EPR dataset with two different measures ranging from 0 to 1. The first defined segregation as the percent of territory in which a group is present that is shared with an included group (segregation is zero if all territory of an excluded group is shared with an included group). This measure varies over time as the political status of groups change. The second was defined as:

$$S_{i,k} = \frac{\sum_j 1 - \frac{\#G_j}{\#G_k}}{\sum_j 1 - \frac{1}{\#G_k}}$$

where j are cells, G_x stands for number of groups in spatial unit x, i is a specific group and k is a country[9]. Thus, if a group shares all of its space with all other groups present in the country, it has a segregation coefficient of 0. In reverse, if the full extent of a group's territorial base is occupied by that group alone, it is coded as 0 (this corresponds to darker blue cells in Figure 3)[10]. For the group level questions, a population weighted mean of cell values across cells in which the group is present was computed for each group in the dataset for the period of time for which population data was available. Otherwise it was weighted by area. Note that since some cells include several groups, we had to make the assumption that a cell's population was evenly distributed across groups present in the cell: despite spatial disaggregation, the measures remain very coarse in part because of the spatial overlaps between entities. The PRIO dataset also provided us with disaggregated data that allowed us to construct measures of groups' remoteness or centrality. Thus we built a measure of a group's distance from the center of power relative to the maximum distance in the polity, as well as the average mountain and forest extent of a group's territory.

We then proceeded to construct aggregated measures at the polity level: for most variables this was simply population weighted means of the variables of interest (share of the population excluded from political power, average spatial segregation of excluded groups, average spatial segregation of all politically organized groups etc...). For the economic variable, we used the Nordhaus and Chen (2006) G-Econ data as integrated within the PRIO gridded data and we computed the ratio of a group's population weighted mean GDP/capita to that of the country as a whole[11]. Unfortunately, this data is only available for 1990 (with updated values until 2005 that seem

[9]To be precise, since a group's territorial basis may not fully cover each cell that records its presence, we correct for the percent of the cell area coded as being occupied by that group.

[10]At this stage, this choice of index is still quite ad-hoc.

[11]We also computed the difference between the between-group and within-group GINI coefficients of each country-political-period unit but this measure turns out to be too noisy given the spatial resolution of the economic data.

to mostly integrate changes in population rather than underlying changes in value added).

As the number of covariates available in Cederman et al. (2011)'s initial data set was too small to improve our analysis and create a strong imputation model to remedy the problem of missing data, and as the ACD data on conflict is limited to certain forms of conflict, we also constructed an expanded dataset for sub-Saharan Africa, drawing mostly from the Harvard Africa Research Program. We added the new data set on events of unrest and violence, controls for urbanization and education. It also allowed us to add more precise institutional controls (electoral competition occurring during the executive selection project) than the Polity codings of democracy and anocracy originally available. From the FAO's *Farming Systems and Poverty* we also obtained a map of the major environmental constraints for agricultural productivity and coded groups according to whether the majority of their home territory falls in a poorly or highly productive region (Dixon et al., 2001). The aim was to obtain some other and more prior measure of horizontal inequalities. Given the coarseness of the G-Econ data on the spatial distribution of value added by Nordhaus and Chen (2006)[12], a further step would be to integrate household wealth from the Demographic and Health Surveys (many countries include ethnic identification in the survey).

5.2 Results

We find suggestive evidence - based on the global comparison - that when excluded groups are more regionally homogeneous, the probability that they organize to challenge the status quo (and that this escalates in political violence) roughly raises by 50% the probability of this otherwise very rare event (precisely, the probability increases by about 0.06 from a base rate 0.04). As shown in Figure 5, this result is associated with a great deal of uncertainty, which is natural given the relative coarseness of the measurement and the fact that this group characteristic is only one of many factors that can influence

[12] A drawback of the G-Econ data, for our purposes (it is a strength in other applications), is that the wealth is defined for a unit of space (about 100 km by 100 km cell at the equator), whereas salient aspects of the wealth distribution across ethnic groups may include the wealth embedded in a social network spanning both the rural home territory and the cities. Wealth data for households would thus be a better basis for creating group-level aggregate measures and measures features of the polity-level distribution. Measures of inequality are extremely sensitive to the level of aggregation.

a group's behavior in response to its political exclusion. We estimated a rare events logit after first matching on demographic balance of group relative to those in power, number of groups in the country, distance from capital, presence of mountainous terrain and region dummies. The causal effect remained positive and large with some variation in the mean estimated values ranging from 0.4 to 0.6 depending on the model (a table of coefficients and standard errors can be found in the Appendix). The result was not robust to inclusion of lags of regime type, conflict history and lags in political status of the group. However, since these lags are relative to the political period forming the unit of analysis, they may still be post-treatment relative to the measure of spatial segregation, which changes for some group in the period from 1946 to 2005 but remains stable for a large number of groups. The results were robust to changing the measure of spatial segregation and remained for a range of the cutoffs used to dichotomize the treatment. Nevertheless, given the continuous measure and the arbitrariness of choosing this cutoff, it would be better to conduct multichotomous matching. In its place, Figure 6 shows the estimated causal effect of the continuous treatment after matching on its dichotomous version. Finally, we expected that spatial segregation would only make a difference when interacted with other variables, such as regime type and distance from center of power, but subsetting on these variables reduced the statistical power of the data too much to report anything about interactions.

As explained earlier, we scaled up the analysis of the effect of the spatial segregation of excluded groups to the polity level. The conclusion is maintained: amongst polities with excluded ethnic groups, those whose excluded population form ethnically homogenous regions are more likely to see groups take up arms (Figure 7). Turning to our polity-level variables capturing features of the distribution of power and wealth across countries, we find much noisier results. As expected, increased share of population excluded is associated with greater probability of ACD-type conflicts over the period 1946-2005 in country-political-periods having at least one politically organized group excluded from power. However, we are unable to find evidence for interactions between economic disadvantage and political exclusion with the data at hand. The number of cases is reduced to our spatial aggregation and the smaller panel for which spatially disaggregated economic data is available (1990-2005). Many fewer matches can be found at this scale of analysis for the explanatory variables that capture the idea of horizontal inequalities across ethnic groups. This we consider to be an important point:

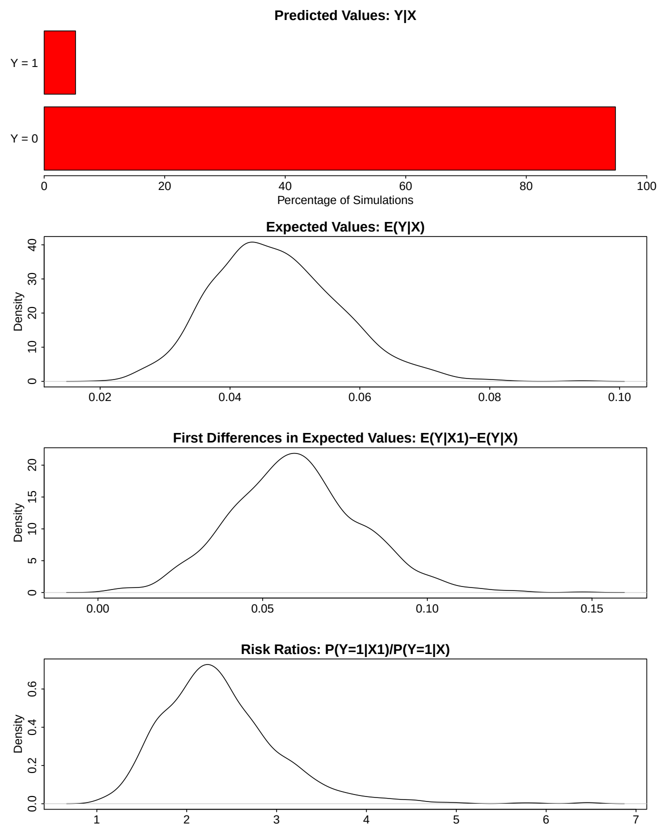

Figure 5: Our results for the effect of spatial segregation on excluded group engagement in conflict. The x-axis is our measure of regional spatial segregation ranging from 0 to 1. The covariates are held at their median values and include region dummies, distance from capital, demographic balance and mountainous terrain.

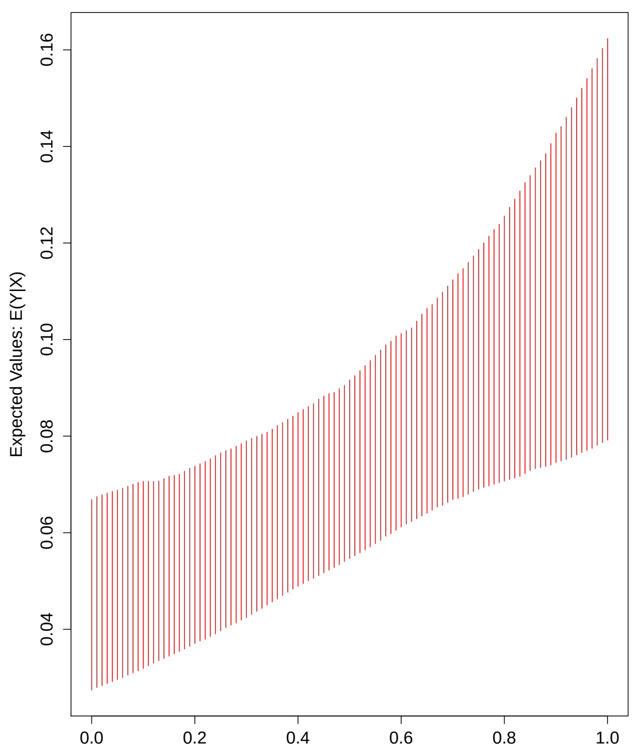

Figure 6: The continuous treatment variable.

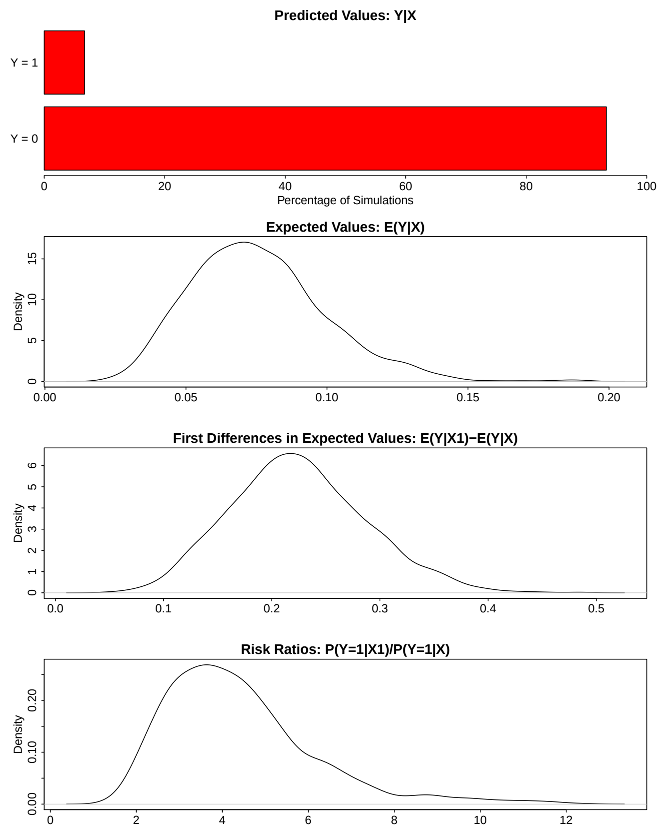

Figure 7: The results when shifting the analysis to the country level.

by constructing indices representing social structure and important aspects of the co-variation between wealth, power and the spatial organization of organized groups, we may find that the set of comparable polities is too small for large-N analyses. Yet, this data may be fruitfully used to enrich small-N case comparisons.

6 Conclusion

In the past years, a number of authors have pointed out the ecological fallacies that aggregated studies of conflict fall prey to. They argue, rightly, that it is senseless to use an average measure of mountainous terrain or ethnic diversity for the whole polity, given that conflicts rarely involve the whole country and that no relevant group may actually be characterized by the country-level average values of the explanatory variable. Yet, having spatially disaggregated measures to describe a society does not imply that we can straightforwardly transfer our statistical models from the polity level to sub-national entities. In particular, the assumption that sub-national entities do not interfere with each other (SUTVA) will be untenable for most mechanisms of interest to political scientists. We thus argue that in most cases, the disaggregated data will have to be re-aggregated at the polity level in ways that capture key structural properties of the polity and that allow researchers to search for adequate controls in the population of polities. In this paper, then, we have explored the kinds of causal questions which may be answered with disaggregated data on ethnic groups and civil conflics. While there are some causal questions that can be meaningfully addressed, we find that the requirements for valid causal inference often are very stringent and that it may be fruitful and insighful to conduct some descriptive analysis first. As space and time have been largely neglected in political theory for some time, most theories concerning ethnic politics or other forms of identity-based group divisions in politics and conflict do not yet include spatial parameters. Thus, describing any regularities in the data concerning the spatial configurations of groups in conflict-prone versus peaceful polities and in polities marked by different degrees of political-economic discriminations may be a very important first step on the way to formulating questions that can build on these new and exciting data.

7 References

Abadie, A. and J. Gardeazabal (2003). "The Economic Costs of Conflict: A Case Study of the Basque Country." American Economic Review 93(1): 113-132.

Acemoglu, D and Robinson, J. 2002. A Theory of Political Transitions. *American Economic Review*, Vol 91: 938-963.

Azam, J-P and Mesnard, A. 2003. Civil War and the Social Contract. *Public Choice*, Vol 115: 455-475.

Alt et al. 2001. Aggregation Among Binary, Count, and Duration Models: Estimating the Same Quantities from Different Levels of Data. *Political Analysis*, Vol 9:21-44.

Bates, R. and Yackovlev, I. 2002. Ethnicity, capital formation adn conflict: evidence from Africa. In: Grootaert. C and van Bastelaer, T. 2002. *The Role of Social Capital in Development, An Empirical Assessment.* Cambridge University Press.

Bates, R. 2008. *When Things Fell Apart: State Failure in Late-Century Africa.* Cambridge University Press.

Buhaug H. and Lujala, P. 2003. Accounting for scale: Measuring geography in quantitative studies of civil war. *Political Geography*, Vol 25: 399-418.

Buhaug, H. and Rod, J.K.2006. Local determinants of African civil wars, 1970-2001, *Political Geography*, Vol 25: 315-335.

Cederman, L.K, Weidmann, N.B., Gleditsch, K.S. 2011. Horizontal Inequalities and Ethnonationalist Civil War: A Global Comparison. *American Political Science Review*, Vol. 105: 478-495.

Cunningham, D. E., Gleditsch, K.G. and Salehyan, I. 2009. It Takes Two: A Dyadic Analysis of Civil War Duration and Outcome. *Journal of Conflict Resolution*, 53 (4): 570-97. .

Dixon, J, Gulliver, A and Gibbon, D. 2001. Farming Systems and Poverty. FAO and World Bank.

Fearon, J and Laitin, D. 2003. Ethnicity, Insurgency and Civil War. *American Political Science Review*, Vol. 97: 75-90.

Fearon and Laitin, 2008. Integrating Qualitative and Quantitative Methods. Oxford Handbook of Political Methodology. Oxford University Press.

Gates, S. 2002. Recruitment and Allegiances: the micro-foundations of rebellion. *Journal of Conflict Resolution*, Vol 46: 111-130.

Gleditsch, K. S., and Ward, M. D. 2001. Measuring space: A minimum distance database. *Journal of Peace Research*, Vol 38:749-768. .

Gleditsch, N. P., Wallensteen, P., Eriksson, M., Sollenberg, M. and Strand, H. 2002. Armed Conflict 19462001: A New Dataset. *Journal of Peace Research*, Vol 39(5): 615-637. .

Glynn and Quinn (2008). Choosing an Identifying Set of Matching or Conditioning Variables. Working Paper Harvard.

Ho, D., Imai, K., King, S., Stuart, E. 2007. Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference *Political Analysis*, Vol.15:199-236.

La Ferrara, E. and Bates, R. 2001. Political competition in Weak States: Rent Seeking, Violence and Identity Politics. *Economics and Politics*, Vol 13: 159-84.

Lewis, J. 2010. Rethinking Civil War Onset: How Nascent Rebel Groups Become Viable. Harvard Working Paper.

Morgan, S.L. and Winship. C. 2007. Counterfactuals and Causal Inference: Methods and Principles for Social Research. Cambridge University Press.

Nordhaus, W., Azam, Q., Corderi, D., Hood, K., Victor, N.M., Mohammed, M.,Miltner, A. and Weiss, J. The G-Econ Database on Gridded Output: Methods and Data,Yale University, May 12, 2006, available at gecon.yale.edu. .

Stewart, F., ed. 2008. Horizontal Inequalities and Conflict: Understanding Group Violence in Multiethnic Societies. Houndmills, UK: Palgrave Macmillan.

Posner, D. 2004. Measuring Ethnic Fractionalization in Africa. *American Journal of Political Science*, Vol 4:849-863.

Pearl, J. 2000. Causality. Cambridge University Press. .

Reynal-Queyrol, M. and Montalvo, J.G (2008) Ethnic Polarization, Potential

Conflict and Civil Wars. UPF Economics and Business Working Paper.

Wimmer, A. Cederman, L-E. and Min, B. 2009. Ethnic Politics and Armed Conflict: A Configurational Analysis of a New Global Data Set. *American Sociological Review*, Vol. 74: 316-337.

8 Regression tables

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
$Segregation (measure \mathrm{S})$	0.574*	0.256	0.602*	0.436	0.115	0.940*	0.421
	(0.216)	(0.205)	(0.232)	(0.262)	(0.485)	(0.292)	(0.442)
b	0.600	1.114*	0.209	0.768	1.321	−0.708	−0.645
	(0.511)	(0.479)	(0.607)	(0.660)	(1.726)	(1.046)	(1.341)
$\#Excluded groups$	−0.031*		−0.030*	−0.029*	−0.103*	−0.037*	−0.024
	(0.009)		(0.009)	(0.010)	(0.037)	(0.010)	(0.016)
$Zone$		0.331	0.188		−0.384	0.107	0.040
		(0.173)	(0.192)		(0.443)	(0.237)	(0.382)
				0.380			0.393
				(0.443)			(1.107)
$lag democracy$					1.880*		
					(0.767)		
$lag anocracy$					0.632		
					(0.595)		
$lag conflict$					1.417*		
					(0.652)		

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
Segregation	0.506*	0.391	0.400	0.558*	0.220	0.891*	1.261*
(%areasharedw/included)	(0.217)	(0.225)	(0.225)	(0.277)	(0.333)	(0.285)	(0.422)
b	0.766	0.989*	0.973	1.591*	2.724*	0.118	0.741
	(0.526)	(0.471)	(0.557)	(0.684)	(0.844)	(1.162)	(1.908)
#Excludedgroups	−0.022*		−0.020*	−0.021*	−0.022	−0.020*	−0.017
	(0.008)		(0.009)	(0.009)	(0.014)	(0.010)	(0.013)
Zone		0.291	0.263			0.147	−0.273
		(0.177)	(0.193)			(0.241)	(0.420)
lagdemocracy					0.436		
					(0.463)		
laganocracy					0.109		
					(0.343)		
lagconflict					0.503		
					(0.344)		
Prop.mountainous						0.510	−0.661
						(0.415)	(0.715)