

Incentives and opportunities: A complexity-oriented explanation of violent ethnic conflict

Online Appendix

Revised: 11 February 2014

Corinne Bara, ETH Zurich

This Online Appendix supplements the JPR article ‘Incentives and opportunities: A complexity-oriented explanation of violent ethnic conflict.’ It includes the following:

| | |
|--|-----------|
| <i>A: Coding rules for the explanatory conditions</i> | <i>2</i> |
| <i>B: Robustness test with a modified coding of political exclusion</i> | <i>3</i> |
| <i>C: Information on the quasi-sufficiency of paths after dropping the ‘absent’ conditions</i> | <i>4</i> |
| <i>D: Parameter estimates for the logistic regression model (comparison of predictions)</i> | <i>5</i> |
| <i>F: The QCA solution for the estimation data (1990-2004)</i> | <i>7</i> |
| <i>G: QCA of non-onset (explaining ‘peace’).....</i> | <i>8</i> |
| <i>E: The predictive capacity of the logit model with interaction terms.....</i> | <i>10</i> |
| <i>H: Cases not covered by the QCA solution.....</i> | <i>13</i> |
| <i>References</i> | <i>14</i> |

A: Coding rules for the explanatory conditions

Table V. Measurement of the explanatory conditions

| Condition | Label | Operationalization |
|---------------------------|--------------|--|
| Politically excluded | polx | Condition is coded 1 if in the StatusID variable of the EPR-ETH dataset, status is 5 (regional autonomy), 6 (separatist autonomy), 7 (powerless) or 8 (discriminated) in the majority of period-years (excluding the onset year). |
| Ousted from rule | oust | Condition is coded 1 if in the StatusID variable of the EPR-ETH dataset, a group was downgraded from central government rule (EPR status 1-4) to excluded (EPR status 5-8) in the course of the period. |
| Ruling group | ruler | Condition is coded 1 if in the StatusID variable of the EPR-ETH dataset, a group has status 3 (senior partner) in all years of the period. |
| Oil and gas | petrol | Condition is coded 1 if there is at least one oil or gas field with more than 500 million recoverable barrels of oil equivalent in a group's territory. This information was obtained by combining GeoEPR-ETH data on groups' settlement areas with a georeferenced petroleum dataset by Lujala, Rod & Thieme (2007). The condition is coded zero for groups with no set area (i.e., dispersed, migrant, or predominantly urban groups, see condition 'territorially concentrated'). Note: The Bakongo and Cabindan Mayombe in Angola were recoded from 0 to 1. Cabinda is clearly oil-rich, although Cabinda's oil fields are offshore fields and as such not directly in the territory of the group. |
| Previous conflict | precon | Condition is coded 1 if there are less than ten years between the start of a group-period and the end of the last ethnic conflict of that group as coded in the <i>onset_do_flag</i> variable downloaded from the GROWup data portal. Note: The Armenians in Azerbaijan (1991) were coded to have had a previous conflict, even if this conflict was in the country of which they were previously part (USSR), because it was about the same territory. |
| Tiny group | tiny | Condition is coded 1 if the relative size of an ethnic group is less than 1% of the total country population as reported in the EPR-ETH GroupSize variable, AND the absolute group size is less than 1 Mio., with information on total country population taken from the CIA World Factbook. |
| Territorial concentration | conc | Condition is coded 1 if a group has a defined settlement pattern in the GeoEPR-ETH dataset (as opposed to dispersed, migrant, or predominantly urban groups). The EPR-ETH dataset reports this in the variable <i>hasSetArea</i> . |
| Regime change | instab | Condition is coded 1 if a group lives in a country in which there was at least a 3-point change in a country's POLITY score in a period, or if there was foreign occupation (POLITY score -66), anarchy (POLITY score -77), or regime transition (POLITY score -88) at any time in the period. |
| Extreme state poverty | xpoor | Condition is coded 1 if a group lives in a country that in the first year of the group-period belonged to the poorest 10% of all countries. GDP data is from the Penn World Tables' <i>rgdpch</i> variable (PWT 71 and 56), extrapolated for missing years using World Bank growth rates. |
| Neighboring ethnic kin | havtek | Condition is coded 1 if a group has ethnic kin as defined in the TEK (transnational ethnic kin) data (Cederman et al., 2013) in a country that is connected to its host country by land borders (information on land borders from CIA World Factbook). |

| | | |
|-----------------|--------|---|
| Kin in conflict | tekcon | Condition is coded 1 if a group's ethnic kin as defined above experiences ethnic conflict in a period (for shorter periods: in the five years before the end of a period). An additional condition is that the two groups' territories are adjoining, which was assessed visually on the GROWup map. If one group was dispersed and the other concentrated, this condition was fulfilled if the concentrated group's territory borders the dispersed group's country. |
|-----------------|--------|---|

B: Robustness test with a modified coding of political exclusion

Table VI reports the results of a robustness test I conducted to assess whether the QCA results change if groups with regional autonomy (EPR-ETH status 5) are considered included rather than politically excluded. This recoding takes into account that some groups (smaller groups in particular) may be satisfied with regional autonomy status, especially in decentralized political systems, and may not consider regional autonomy a political grievance as implied in the original coding. For the robustness test, the condition *polx* was replaced with *polx2*. *Polx2* takes on the value 1 if an ethnic group is neither represented in central government nor has regional autonomy, i.e., EPR-ETH status 6-9.

Table VI. QCA solution with *polx2* instead of *polx* (intermediate solution term)

| | Solution/ configuration consistency | Solution coverage | Configuration raw coverage | Configuration unique coverage |
|--|--|----------------------|-------------------------------|-------------------------------------|
| Conditions: | ruler, polx2, oust, petrol, precon, tiny, instab, tekcon | | | |
| | (frequency cut-off: 1.00 / consistency cut-off: 0.70) | | | |
| Model parameters: | 0.87 | 0.60 | | |
| 1: precon*~tiny*~ruler*polx2 | 0.94 | | 0.32 | 0.11 |
| 2: ~instab*~petrol*precon*~tiny*~ruler | 0.84 | | 0.26 | 0.10 |
| 3: instab*~tiny*~ruler*oust*~polx2 | 0.86 | | 0.06 | 0.06 |
| 4: tekcon*instab*precon*~tiny*~ruler | 1.00 | | 0.03 | 0.01 |
| 5: tekcon*instab*~petrol*~tiny*~ruler*polx2 | 0.80 | | 0.08 | 0.08 |
| 6: ~tekcon*instab*petrol*~tiny*~ruler*polx2 | 0.86 | | 0.06 | 0.03 |

As the results indicate, the QCA solution reported in the paper is more or less robust to this change of coding. As discussed in Schneider & Wagemann (2012: 284-286), QCA solutions can be considered robust if the consistency and coverage scores of the solution remain roughly the same, and if the solution paths do not suggest a substantively different theoretical interpretation. As we see in Table VI, the consistency and coverage scores are almost the same as in the original QCA solution, but six instead of only four paths now explain conflict onset. Three paths (in bold in the table) remain exactly the same. These are the 'conflict trap,' 'ousted rulers,' and 'resource curse' configurations in rows one, three, and six. The

‘bad neighborhood’ configuration in row five, also in bold, now additionally contains the condition *polx2*, which means that the new path is in a subset relation to the original path, making this result robust as well. While all four original paths are still there, some of them have a lower coverage due to the two new paths in rows two and four. Row four explains only one conflict uniquely and may be ignored here given the large-N nature of this study. Remains row two, which is basically a ‘split-off’ from configuration one ‘conflict trap,’ as the reduced unique coverage of row one indicates. According to this path, among non-tiny and non-ruling groups, previous conflict is also dangerous if a group has no political instability and no oil and gas. This does not make much theoretical sense, but the path has a quite high coverage, so a look at the cases may be helpful. Out of the 27 cases covered by this path, eight were affected by the *polx* recoding. Of those, six are groups in India’s troubled Northeast (the remaining two are the Basques in Spain, and the Moro in the Philippines). The groups in India’s Northeast in particular demonstrate that the recoding of regional autonomy as politically included may not have been appropriate to start with. ‘The region seems distant from the hearts and minds of many Indians...,’ Baruah (2005: vii) writes, and regional autonomy may not suffice to change that and make this region and its people part of the ‘pan-Indian project’ (Baruah, 2005: 25). The same can be said of other autonomous regions: If regional autonomy does not go hand in hand with some sort of representation in decision-making at the central level, it may not suffice for alleviating political grievances. To sum up: Although the QCA solution is quite robust to the recoding, it is not fully robust, and neither should it be, given that whether we code groups with regional autonomy as included or excluded is a substantive change of our definition of political exclusion.

C: Information on the quasi-sufficiency of paths after dropping the ‘absent’ conditions

This section offers some information on the interpretation of those three solution paths that contain conditions that contribute to conflict onset in their absence, and which were ignored in the theoretical discussion of the individual paths.

The second path contains the condition $\sim petrol$, i.e., the effect of this configuration is limited to groups without oil and gas resources. However, a look at the truth table reveals that this restriction is produced because of merely one inconsistent case (the Berbers in Algeria from 1992-1996), while the other three ethnic groups with the same conditions and oil and gas all experience conflict. With an only slightly lower consistency (0.76 instead of 0.77), the configuration is thus sufficient for onset even without the $\sim petrol$ condition.

The addition of *~polx* in the third path ('ousted rulers') is partially an artifact of the coding rules: In order to be coded as politically excluded, an ethnic group had to be excluded in the majority of group years, i.e., three out of five years. For most ousted groups who rebelled, this was not the case because they rebelled within a year or maximum two of their exclusion. This itself is an important finding, especially for conflict prevention purposes, but less important for the sake of the theoretical argument made in the paper.

Similar to the second configuration, the fourth path ('resource curse') contains the qualifier *~tekcon*, meaning that this path is only sufficient for groups with no ethnic kin in conflict — a restriction that is once again caused by the same single inconsistent case (Berbers in Algeria from 1992-1996). With a slightly lower consistency of 0.85 (instead of 0.88), the combination of exclusion, oil, and instability is sufficient for conflict no matter whether ethnic kin are also rebelling or not.

D: Parameter estimates for the logistic regression model (comparison of predictions)

Table VII shows the parameter estimates for the logistic regressions that were run in order to compare the predictive power of the QCA and logit models (both for the full dataset and the estimation data). Except political exclusion (*polx*) and ruling group (*ruler*), all conditions are significant, and the signs are in the direction expected when making directional assumptions for the QCA solution. The four conditions with the strongest effects on the likelihood of conflict (significant at the 0.01 level with the exception of *oust* in the estimation dataset) are previous conflict, ethnic kin in conflict, ousted from rule, and oil and gas. Interesting enough, these are also the key conditions in each of the four paths identified in the QCA solution. Had we used logistic regression in this study, we would have thus largely found the same conditions to be important, but would have failed to identify the combinations in which the likelihood of conflict is highest. A case in point is the fact that political exclusion is not significant in either of the two datasets. This is exactly what led scholars like Collier & Hoeffler (2004) to reject grievance-based explanations of conflict: Political exclusion is (unfortunately) too common among both onset and non-onset groups to exhibit a strong *independent* effect on the conflict likelihood. In the QCA model, however, it was demonstrated that political exclusion is still an important 'ingredient' in some risk patterns that often lead to conflict — notably in the 'conflict trap' and 'resource curse' patterns.

Table VII. Logistic regression results, outcome: Onset

| <i>Variables</i> | <i>1990-2009 (full data)</i> | <i>1990-2004 (estimation data)</i> |
|---------------------------------|------------------------------|------------------------------------|
| Political exclusion (polx) | 0.417 (0.360) | 0.426 (0.380) |
| Ousted from rule (oust) | 1.618** (0.621) | 1.749* (0.678) |
| Tiny group (tiny) | -1.307* (0.570) | -1.229* (0.606) |
| Ruling group (ruler) | -0.904 (0.765) | -0.945 (0.749) |
| Previous conflict (precon) | 3.829** (0.417) | 3.387** (0.438) |
| Oil or gas fields (petrol) | 1.422** (0.399) | 1.551** (0.420) |
| Political instability (instab) | 0.808* (0.327) | 0.600† (0.347) |
| Ethnic kin in conflict (tekcon) | 1.751** (0.468) | 1.620** (0.512) |
| Constant | -2.901** (0.357) | -2.648** (0.373) |
| Observations | 500 | 394 |

Standard errors in parentheses. † p<0.1; * p<0.05; ** p<0.01

E: QCA solution for the estimation data (1990-2004)

Table VIII reports the QCA solution for the period 1990-2004 (estimation data). This analysis was conducted in order to obtain a model for the out-of-sample prediction 2005-2009. All four paths remain the same as in the analysis of the full period (1990-2009), with comparable consistency and coverage scores, making the QCA result reported in the paper robust to a slightly different specification of the time period studied.

Table VIII. QCA solution for 1990-2004 (intermediate solution term)

| | Solution/ configuration consistency | Solution coverage | Configuration raw coverage* | Configuration unique coverage* |
|---|---|----------------------|-----------------------------------|--------------------------------------|
| Conditions: | ruler, polx, oust, petrol, precon, tiny, instab, tekcon | | | |
| | (frequency cutoff: 1.00 / consistency cutoff: 0.70) | | | |
| Model parameters: | 0.88 | 0.58 | | |
| ~tiny*precon*polx*~ruler | 0.92 | | 0.38 | 0.34 |
| tekcon*instab*~tiny*~petrol*~ruler | 0.80 | | 0.09 | 0.09 |
| instab*~tiny*oust*~polx*~ruler | 0.83 | | 0.06 | 0.06 |
| ~tekcon*instab*~tiny*petrol*polx*~ruler | 0.89 | | 0.09 | 0.06 |

* Raw coverage includes cases explained by more than one configuration, while unique coverage includes only cases exclusively covered by that configuration.

F: QCA of non-onset (explaining ‘peace’)

QCA assumes that causal relationships are asymmetric, defined as a situation in which $X \rightarrow Y$ does not imply $\sim X \rightarrow \sim Y$, or as Schneider & Wagemann (2012: 81) write: ‘Asymmetry, thus, describes the fact that insights on the causal role of a condition are of only limited use for the causal role of its absence, and the explanation of the occurrence of an outcome does not necessarily help us much in explaining its non-occurrence.’ As a consequence of this, the outcome (onset) and its absence (i.e., peace or at least non-onset)¹ require separate QCA analyses. The focus in the paper was on the outcome, and space constraints prohibited a discussion of non-onset. That discussion is thus provided here.

Again, as in the analysis of onset, the first step in the analysis of non-onset is a test of necessary conditions for peace, the results of which are provided in Table IX. As in the onset analysis, I have set a consistency cut-off of 0.95 for this test. Only three single explanatory factors fulfill the condition of a quasi-necessary condition for peace: the absence of kin in conflict ($\sim tekcon$), the absence of previous conflict in the past ten years ($\sim precon$), and not being ousted from power recently ($\sim oust$).

Table IX. Test of necessary conditions for non-onset (consistency cut-off: 0.95)

| Condition | Consistency | Coverage | Number of deviant cases | Cases having this condition |
|---------------|-------------|----------|-------------------------|-----------------------------|
| $\sim tekcon$ | 0.97 | 0.82 | 14/398 | 467/500 |
| $\sim precon$ | 0.97 | 0.88 | 11/398 | 440/500 |
| $\sim oust$ | 0.98 | 0.81 | 7/398 | 484/500 |

Note that when these three quasi-necessary conditions are present rather than absent, they are the ‘key ingredients’ in three of the four sufficient paths to conflict. This is not so surprising: A condition that is sufficient for the outcome when it is present is necessary for the absence of the outcome when it is absent. Of course, the three conditions above were not individually sufficient for the outcome, but only in combination with other factors. Nevertheless, given their status as quasi-necessary conditions, I will include them in the model for the sufficiency test. For the same reason, *tiny* and *ruler* are included. In their absence, they were quasi-necessary for onset and should thus be quasi-sufficient for peace when present.

¹ It is wrong for several reasons (discussed in the paper) to assume that all groups without an onset of ethnic conflict were at peace. They may have had conflict that was either not coded as a civil conflict at all, or was coded as a civil conflict but not an ethnic one. For the sake of readability, and because peace is more intuitive than ‘non-onset,’ I will nevertheless use the term peace in the discussion of this analysis.

The model I start out with for the sufficiency test is thus *tekcon*precon*oust*tiny*ruler* → *peace*, and the results of this test are reported in Table X below. Note that because of the high number of non-onsets, a frequency cut-off of five cases is set, meaning that only combinations with five cases or more are considered for the analysis. Also, with 0.9 a higher consistency cut-off is set than in the onset-analysis, because with 398 non-onsets, allowing a 10% mistake rate already means 40 ‘false negatives,’ i.e., rebellions started by groups who were predicted to remain peaceful according to the model.

Table X. QCA solution for non-onset

| | Solution/ configuration consistency | Solution coverage | Configuration raw coverage* | Configuration unique coverage* |
|---|---|----------------------|-----------------------------------|--------------------------------------|
| <i>Consistency cut-off: 0.9; Frequency cut-off: 5</i> | | | | |
| Model 1: ~tekcon*~precon*~oust | oust, precon, tekcon, tiny, ruler 0.91 | | 0.92 | 0.92 |
| Model 2: ~tekcon*~petrol*~precon*~oust | oust, precon, tekcon, petrol, tiny, ruler 0.93 | | 0.86 | 0.86 |

* Raw coverage includes cases explained by more than one configuration, while unique coverage includes only cases exclusively covered by that configuration.

As is seen in the results for Model 1, it takes only one configuration to explain 92% of all cases of peace in the sample. While the three conditions ~tekcon, ~precon, and ~oust were individually quasi-necessary for peace, together they are quasi-sufficient. The conditions tiny and ruler prove to be redundant in this solution. However, with a consistency of 0.91, we make 35 mistakes, i.e., we predict peace for 35 groups who then had conflict. Model 2 adds the ‘key ingredient’ of the fourth path to conflict, petrol, to the analysis. This improves consistency to 0.93, so that we now make only 26 mistakes. Not having kin in conflict, not having oil and gas, not having had previous conflict, and not being ousted is quasi-sufficient for peace, explaining 86% of all cases of peace in the sample. Adding any of the other conditions, such as polx, conc, havtek, or instab, proves to be redundant, as they do not change the solution in Model 2. Recoding ‘extremely poor country’ to ‘very rich country’ (a country being among the richest 30% and 20% of all countries, respectively) in order to test whether a very high per capita income could be sufficient for peace, also does not improve this result. In this case we can conclude that the causes of peace are really somewhat the mirror image of the causes of conflict: Having none of the four key conditions for onset is sufficient for peace. Nevertheless, peace is much more easily explained than conflict, needing just one causal configuration to explain 86% of all cases.

G: The predictive capacity of the logit model with interaction terms

In the prediction section of the article, the predictive capacity of QCA was compared to the predictions generated by a simple logistic regression model that did not include any interactions. In order to assess whether the predictions generated by the logit model improve considerably when causal complexity is introduced, this section reports and compares the predictive capacities of several logistic regression models that include interaction terms.

As we are now comparing different logistic regression models, we can make use of Receiver Operator Characteristic (ROC) plots, which plot the true positive rate against the false positive rate for *all* possible thresholds and are thus suitable to compare predicted probabilities across different statistical models (Ward, Greenhill & Bakke, 2010: 366). In a single and simple statistic, the size of the area under the ROC curve reports the overall predictive performance of a model. It ranges from 0.5 (a model that is no better than chance) to 1.0 (predicts the outcome perfectly) (Ward, Greenhill & Bakke, 2010: 367). The baselines for comparison are the models reported in the paper, i.e., Logit1 and Logit2. In order to still be able to compare those predictions to the baseline QCA models, the true positive rate, number of false positives, and precision at two chosen thresholds are reported as well: Table XI compares the predicted probabilities of the logit models with interaction terms to the models QCA1 and Logit1 (predicted probability threshold for onset prediction determined by the QCA result/sensitivity, see paper). Table XII compares the predicted probabilities of the same logit models to the models QCA2 and Logit2 (at the threshold of predicted probabilities at which the logit model performs best in terms of an optimum true to false positive ratio, see paper).

Each of the four paths to conflict identified with QCA was modeled in one regression each, leaving out the trivial necessary conditions and the above-mentioned ‘absent conditions.’ The model ‘conflict trap’ in Tables XI and XII thus included the interaction of *polx*precon*, ‘bad neighborhood’ included the interaction of *tekcon*instab*, ‘ousted rulers’ included the interaction of *oust*instab*, and ‘resource curse’ included the interaction of *petrol*polx*instab*.² The ‘conflict trap’ and ‘bad neighborhood’ models were also estimated once with the necessary condition ‘not tiny AND not ruler.’ In the last model estimated, all four paths were included as interaction terms in order to partially mimic the QCA solution reported in the paper.

With the exception of two models (‘ousted rulers’ and ‘resource curse’), the inclusion of interaction terms does improve the predictive performance of the simple model reported in the paper, but none of the models achieve or surpass the predictive capacity of the QCA model. The model with the best

² All constitutive terms were also included in the models. The R code to replicate Tables XI and XII is provided in the supplementary materials to the paper.

predictive capacity is — maybe not surprisingly — the one in which all four QCA paths were included as interaction terms. Note, however, that after the inclusion of these interaction terms (and the constitutive terms) only three coefficients remain significant. At a rate of 61 true onset predictions (of total 102 onsets), the inclusion of interaction terms makes the number of false positives drop from previously 16 to only 10 (see Table XI). At the logit's optimum rate of 47/102 true positives, the number of false positives drops from 6 to 4 (see Table XII).

That none of the models tested here is able to exactly match the predictive capacity of QCA may have several reasons: Even if the inclusion of some interaction terms was an improvement, the full complexity offered by QCA could not be matched (see also Grendstad, 2007: 127). In the presence of multiple paths and a situation of necessary and sufficient conditions, the resulting equation would have been far too complex. Necessary conditions in particular need to be interacted with all other predictors, because the assumption is that if a necessary condition takes on the value zero, all other effects should be zero as well. Finally, QCA also differs in that it looks for the best fit in terms of sufficiency, not the best overall fit (Grofman & Schneider, 2009: 668). What is worth remembering, however, is that logistic regression was somewhat 'disadvantaged' in this comparison because only binary variables were used. QCA with a binary outcome requires that all other conditions are dichotomized as well, whereas logistic regression allows the analyst to use both binary and continuous variables as explanations.

Table XI. Predictive capacities of logit models using interaction terms (comparison with QCA1 and Logit1)

| Model: | Area under ROC | Sensitivity (true positive rate) | False positives: | Precision: |
|-----------------------------|----------------|----------------------------------|------------------|------------|
| QCA1 | --- | 61/102 (0.60) | 8 | 0.88 |
| Logit1 | 0.8589 | 62/102 (0.61) | 16 | 0.79 |
| Conflict trap | 0.8628 | 61/102 (0.60) | 13 | 0.82 |
| Conflict trap w/nec.con. | 0.8625 | 61/102 (0.60) | 13 | 0.82 |
| Bad neighborhood | 0.8611 | 62/102 (0.61) | 15 | 0.81 |
| Bad neighborhood w/nec.con. | 0.8603 | 62/102 (0.61) | 14 | 0.82 |
| Ousted rulers | 0.8588 | 62/102 (0.61) | 20 | 0.76 |
| Resource curse | 0.8578 | 66/102 (0.65) | 17 | 0.80 |
| All four paths | 0.8629 | 62/102 (0.61) | 10 | 0.86 |

Table XII. Predictive capacities of logit models using interaction terms (comparison with QCA2 and Logit2)

| Model: | Area under ROC | Sensitivity (true positive rate) | False positives: | Precision: |
|----------------------------|---------------------------|---|-------------------------|-------------------|
| QCA2 | --- | 47/102 (0.46) | 3 | 0.94 |
| Logit2 | 0.8589 | 47/102 (0.46) | 6 | 0.89 |
| Conflict trap | 0.8628 | 47/102 (0.46) | 4 | 0.92 |
| Conflict trap w/nec.con | 0.8625 | 47/102 (0.46) | 4 | 0.92 |
| Bad neighborhood | 0.8611 | 47/102 (0.46) | 6 | 0.89 |
| Bad neighborhood w/nec.con | 0.8603 | 51/102 (0.50) | 8 | 0.86 |
| Ousted rulers | 0.8588 | 49/102 (0.48) | 6 | 0.89 |
| Resource curse | 0.8578 | 50/102 (0.49) | 7 | 0.88 |
| All four paths | 0.8629 | 47/102 (0.46) | 4 | 0.92 |

H: Cases not covered by the QCA solution

Table XIII lists all the onset cases not covered by the QCA solution reported in the paper, i.e., about which no causal statements are made, as discussed in the conclusion.

Table XIII. Onsets not covered

| |
|--|
| Indigenous peoples in Mexico, 1994 |
| Peoples of the Caucasus in Russia, 2007 |
| Albanians in Macedonia, 1997 |
| Croats in Yugoslavia, 1991 |
| Albanians in Yugoslavia, 1998 |
| Slovenes in Yugoslavia, 1991 |
| Armenians in Russia, 1990 |
| South Ossetians in Georgia, 1992 |
| Abkhazians in Georgia, 1992 |
| Tuareg in Mali, 1990 |
| Tuareg in Mali, 2007 |
| Arabs/Moors in Mali, 1994 |
| Diola in Senegal, 1990 |
| Toubou in Niger, 1995 |
| Northerners (Mande and Voltaic/Gur) in the Ivory Coast, 2002 |
| Ijaw in Nigeria, 2004 |
| Yakoma in the Central African Republic, 2001 |
| Sara in Chad, 1992 |
| Tutsi-Banyamulenge in the DR Congo, 2006 |
| Hutu in Burundi, 1991 |
| Tutsi in Rwanda, 1990 |
| Afar in Djibouti, 1991 |
| Oroma in Ethiopia, 1998 |
| Other Muslims in Eritrea, 1997 and 2003 |
| Fur in Sudan, 2003 |
| Southern Shafi'i in Yemen, 1994 |
| Uzbeks in Tajikistan, 1998 |
| Assamese in India, 1990 and 1994 |
| Naga in India, 1992 |
| Baluchis in Pakistan, 2004 |
| Mons in Myanmar, 1990 |
| Wa in Myanmar, 1997 |
| Karenni (Red Karens) in Myanmar, 1992, 1996, and 2005 |
| Dalits (both Hill & Tarai) in Nepal, 1996 |
| Adivasi/Janajati in Nepal, 1996 |
| Malay Muslims in Thailand, 2003 |
| Achinese in Indonesia, 1990 |

References

Baruah, Sanjib (2005) *Durable Disorder: Understanding the Politics of Northeast India*. Oxford: Oxford University Press.

Cederman, Lars-Erik; Kristian Skrede Gleditsch, Idean Salehyan & Julian Wucherpfennig (2013) Transborder ethnic kin and civil war. *International Organization* 67(02): 389-410.

Collier, Paul & Anke Hoeffler (2004) Greed and grievance in civil war. *Oxford Economic Papers* 56(4): 563-595.

Grendstad, Gunnar (2007) Causal complexity and party preference. *European Journal of Political Research* 46(1): 121-149.

Grofman, Bernard & Carsten Q Schneider (2009) An introduction to crisp set QCA, with a comparison to binary logistic regression. *Political Research Quarterly* 62(4): 662-672.

Lujala, Päivi; Jan Ketil Rod & Nadja Thieme (2007) Fighting over oil: Introducing a new dataset. *Conflict Management and Peace Science* 24(3): 239-256.

Schneider, Carsten Q & Claudius Wagemann (2012) *Set-Theoretic Methods for the Social Sciences: A Guide to Qualitative Comparative Analysis*. Cambridge: Cambridge University Press.

Ward, Michael D; Brian D Greenhill & Kristin M Bakke (2010) The perils of policy by p-value: Predicting civil conflicts. *Journal of Peace Research* 47(4): 363-375.