

# Identifying the effect of climate variability on communal conflict through randomization

Jonas Nordkvelle<sup>1</sup> · Siri Aas Rustad<sup>1</sup> ·  
Monika Salmivalli<sup>2</sup>

Received: 15 February 2016 / Accepted: 28 January 2017

© The Author(s) 2017. This article is published with open access at Springerlink.com

**Abstract** In recent years, the focus of quantitative climate-conflict research has shifted from studying civil wars to studying different types of conflicts, particularly non-state and communal conflicts, based on the argument that these local-level conflicts are a more likely consequence of climate variability than civil war. However, the findings from previous research do not paint a consistent picture of the relationship between climate and communal conflict. We posit that a research design treating the climate variable as randomized is a better and more convincing strategy for estimating the relationship between climate variability and communal conflict compared to the conventional control method to account for confounders. In this paper, we ask two questions: (1) what type of research design allows us to treat climate variability as randomized and (2) what can we say about the relationship between climate variability and communal violence using this new design? To answer these questions, we analyze six large subnational areas, at a monthly time scale, and calculate the standardized precipitation index for each area for each month. We find that both short, unusually dry intervals and long, unusually wet intervals increase the likelihood of a communal conflict event.

## 1 Introduction

In recent years, the focus of quantitative climate-conflict research has shifted from studying civil wars to studying different types of conflicts—including non-state conflicts in general and communal conflicts in particular. It has been argued that climate variability is more likely to

---

**Electronic supplementary material** The online version of this article (doi:10.1007/s10584-017-1914-3) contains supplementary material, which is available to authorized users.

✉ Siri Aas Rustad  
sirir@prio.no

<sup>1</sup> Peace Research Institute Oslo, Oslo, Norway

<sup>2</sup> Norwegian University of Life Science, Ås, Norway

lead to communal conflicts than civil war because threshold to engage in a conflict with another group is lower than the threshold for challenging the state (Fjelde and von Uexkull 2012; Theisen 2012). Further, it is clearly very costly for any group to engage in a civil war because the government is likely to have a relatively strong military force (Hendrix and Salehyan 2012). Thus, weak groups, which are most likely to suffer from resource scarcity induced from unfavorable climate, should not be expected to start civil wars (Raleigh 2010). Rather, we should expect—if anything—that such groups fight other weak groups over scarce resources. However, based on the current status of the literature on climate variability and non-state/communal conflict, it is difficult to draw any conclusion about the relationship (Fjelde and von Uexkull 2012; Hendrix and Salehyan 2012; O’Loughlin et al. 2012; Landis 2014 and Raleigh and Kniveton 2012).

The lack of consistency in findings could be attributed to the large variety of research designs, using different units of analysis, study sites, measures of climate, and definitions of conflict. However, they all share the feature that they do not treat the climate variable as a randomized treatment variable, even though climate variables have been argued to have properties conducive to such treatment (Miguel et al. 2004). Instead, they use control variables to account for confounders to establish a causal relationship.

The innovation in this paper is that we explore how climate variability can be used as a randomized variable, and let this guide our design choices. We think this leads us to better and more principled design solutions, compared to a strategy of controlling for omitted variable, although these strategies try to solve the same problem (Angrist and Pischke 2008: 60ff, 80ff). In this paper, we ask two questions: (1) what type of research design allows us to treat climate variability as randomized and (2) what can we say about the relationship between climate variability and communal violence using this new design?

To address the first question, we create a research design where we arguably treat our climate variability variable as randomized, through analyzing six large sub-national areas, at a monthly time scale, calculating a standardized precipitation index for each area-month based on monthly average precipitation in those areas, and adding 12 fixed effects for each area (one for each month). Using such a research design answering the second question, we find that unusually dry short intervals increase the likelihood of a communal conflict event within an area, while short unusually wet intervals do not have the same effect for our sample. On the other hand, long unusually wet intervals increase the likelihood of a communal conflict event. While our research design is particularly suited for studying communal conflict, we argue that our substantial conclusion could have broader implications and be relevant for other types of violence such as civil wars.

In the following, we first we review the literature on communal conflict and climate variability; second, we discuss how climate variability can be treated as a randomized variable and present our research design and data. Finally, we discuss our findings.

## 2 Communal conflict and climate variability

The literature argues that climate change will lead to conflict through scarcity of resources in interaction with social, political, and economic factors (e.g., Homer-Dixon 1994). Supply-induced scarcity may worsen people’s livelihoods by affecting their access to food and water. Thus, in countries where people are dependent on income from agriculture or herding, the potential for climate-induced conflict, such as conflict over grazing land and water sources is

considered high. While the bulk of the climate and conflict literature focuses on civil war, some studies argue that climate-induced conflicts are more likely to be in the form of local communal conflict.

First, communal conflicts and non-state conflicts are more frequent than civil wars and the threshold to engage in a conflict with another group is lower than the threshold for challenging the state, due to costs, the relative strength toward the government, and the relative distance to other communal groups (Fjelde and von Uexkull 2012; Theisen 2012; Hendrix and Salehyan 2012; Raleigh 2010). Further, communal conflicts are arguably also more likely to be motivated by climate variability than other types of non-state conflicts, such as conflicts between highly organized rebel groups. Having a highly organized rebel organization requires extensive resources. Poor and marginalized groups, which are more likely to suffer from the consequences of climate variability, often do not have such resources. Hence, they are unlikely to engage in fighting with highly organized rebel groups in order to ease their resource scarcity, but rather less organized communal groups. Several studies have tested this relationship, but currently there is not enough evidence to make any claims about the relationship.

Witsenburg and Adano (2009) found there to be more violent incidents related to cattle raiding in Northern Kenya during times of abundant rainfall. They noted that the environment in years with abundant rainfall makes raiding easier; for instance, there is better access to water during raiding trips. In contrast, in times of scarce rainfall and consequently scarce resources, the pastoralists do not have the resources to spend on raiding.

Studying sub-Saharan Africa, Fjelde and von Uexkull (2012) found, using communal conflict from the UCDP non-state conflict dataset, that exceptionally dry years are positively correlated with communal conflicts, thereby supporting the environmental scarcity thesis, which predicts that droughts may lead to resource scarcity, which in turn may lead to conflict. Fjelde and von Uexkull (2012) also found some evidence to support their hypothesis that areas characterized by political and economic marginalization see a higher risk for communal violence than other areas in exceptionally dry years.

Using a different conflict dataset (Social Conflict in Africa Database), Hendrix and Salehyan (2012) found that all types of social conflict are more likely to occur both in abnormally dry and wet years, thus partly supporting the findings by Fjelde and von Uexkull (2012). In a global sample study, Salehyan and Hendrix (2014) found that political violence, including non-state conflicts when measured separately, was more frequent at times when resources were abundant.

Studying East Africa, O'Loughlin et al. (2012) found that abnormally wet years are more peaceful than years with average rainfall, whereas abnormally dry years are not related to conflict. These results contrast with the continent-wide studies by Hendrix and Salehyan (2012) and Fjelde and von Uexkull (2012). Moreover, O'Loughlin et al. (2012) also used temperature as a climate variable and found that violence occurs more in warm years than in years with average temperature. The study by O'Loughlin et al. (2014) confirms these results for temperature, finding that high temperature extremes increased the risk of conflict in sub-Saharan Africa in 1980–2009. When studying seasonal changes in weather, Landis (2014) found that prolonged periods of warm weather are positively associated with both civil war and non-state conflicts.

Contrary to O'Loughlin et al. (2012), Raleigh and Kniveton (2012) found that violent events connected to civil wars and communal conflict in East Africa are more frequent both in times of drought and in times of excessive rain. Nevertheless, civil war events are most likely during droughts, whereas communal conflict events are most likely in times of abundant rainfall.

The current literature is characterized by inconsistent results, which may be related to a wide variety of research designs. In particular, we think that trying to solve the causal inference problem by controlling for confounders to ensure conditional independence is problematic in these studies and likely a source of the varying results. We suggest an alternative causal inference strategy, where our research design arguably can treat climate variability as a randomized variable.

### 3 Causal identification strategy

The potential outcomes definition of causality says that a causal effect is the difference between an actual outcome, and how it would have been if and only if the causal agent had been different in some specific way (Gerber and Green 2012). While this effect is fundamentally unobservable, it can be shown that the *average* causal effect between larger groups can be estimated. One identification strategy with a proven track record is randomized experiments. The randomization process ensures that the expected average outcomes between groups decided by the randomized variable are equal. Any differences in outcomes must be attributed to the difference in treatment in these groups.

Experiments are commonly designed by the researcher. However, it is also possible to find naturally occurring experiments or naturally occurring processes that can arguably be treated as experiments (see Gerber and Green 2012:15ff). Nonetheless, regardless of what kind of experiment you have, the same criteria must be fulfilled. Most important of these is a randomized treatment assignment: Assignment must be *stochastic* and with *known probabilities*. While you generally would know these two criteria to be correct when designing your own experiment, in any natural experiment, we must somehow discover these properties.

Previous literature has tried to identify a causal effect of climate variability of conflict through controlling for (un)observables through a panel structure. However, we believe that a more promising identification strategy is the randomized experiment, as it is based on a much simpler statistical procedure (in principle no need for controls), and relies more on a sound research design. Our central argument is that precipitation, treated in a particular way, can be seen as a randomized treatment. Further, we describe how we get to the naturally occurring randomization process and discuss the two additional core assumptions in experiments: non-interference and excludability.

#### 3.1 Randomization of treatment

The definition of a randomized treatment is that treatment is assigned with a known probability. The problem with experiments where nature randomizes and assigns treatment is that we do not know the probability of treatment assignment. Thus, we need to approach natural experiments through finding processes where the treatment can be argued to be assigned with equal probabilities. With the case of precipitation, it is clear that the probability distribution of precipitation varies across space and over different times of the year. It rains more some months of the year than others and more in some areas than in others. It is only for specific area-month sets (e.g., Darfur, in January between 1948 and 2013) that it is convincing to assume equal probability of precipitation.

Following this assumption, we could do an experiment, using monthly average precipitation within one area as our causal variable. However, a problem arises when we try to combine

that experiment with experiments from other area-months. To find an unbiased causal estimate in such an analysis, we have to know the probability distribution from which precipitation was assigned for each area-month and incorporate each such distribution into our analysis.

The Standardized Precipitation Index (SPI) tries to fix this exact problem (McKee et al. 1993; Guttman 1999). SPI is estimated through fitting the probability distribution of precipitation for a given area-month and then standardizes that fit (with expected mean 0 and a standard deviation of 1). The SPI should be interpreted as the normality of observing the given precipitation in that area for that month. Large negative numbers mean it was abnormally dry, and vice versa for large positive numbers.

In principle, if we select  $-1.5$  SPI or lower to be our treatment assignment variable, then the probability of treatment should be expected to be equal for all area-months, meaning that we could treat all area-months as a large pooled experiment. However, SPI normalization is not perfect; some area-months are vastly more likely to have an SPI of  $-1.5$  than others. For instance, in a very dry area such as Darfur, where the normal outcome during the driest months of the year is near zero precipitation, large negative SPI values for those months are not observed at all because they are not possible. Therefore, even after standardization, we need to take the blocked randomization into account when doing our analysis. We do it through adding area-month fixed effects.

Fjelde and von Uexkull (2012) also use SPI and employ an administrative unit-year design. However, rather than estimating SPI for each unit-year, they apply the annualized discrete coding used in PRIO-GRID (Tollefsen et al. 2012). This means that they have no means to correctly control for different treatment probabilities, which should be expected to lead to biased results.

Our research design needs 12 fixed effects for each areal unit; hence, we need to ensure that the probability of observing a conflict event within each unit is fairly large. Therefore, rather than using administrative units, or grid-cells, we identified six spatially segregated areas.

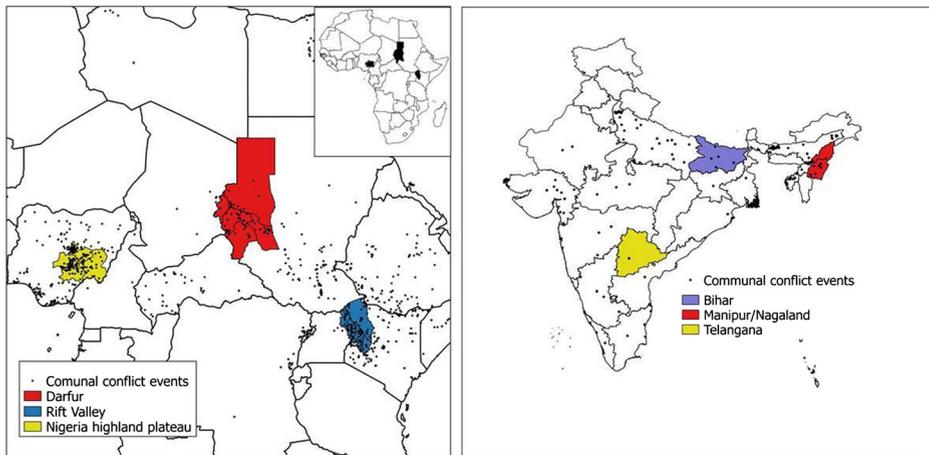
The areas we selected are by no means a random sample, rather we have selected areas where we believe we might observe the hypothesized effect. All the areas had a large number of communal conflict events between 1989 and 2013, and the livelihoods in these areas are likely to be affected by climate variability. Since we did not select randomly, we must be cautious about generalizing. Our areas are the Nigeria highland plateau, Rift Valley in Kenya including neighboring areas in Uganda, Darfur in Sudan, and three areas in India (Telangana, Bihar, and Manipur/Nagaland) (Map 1).

### 3.2 Non-interference

Non-interference is the assumption that the treatment or outcome of another unit is independent of the treatment of any other unit (Gerber and Green 2012: 43f).

A popular approach in the climate-conflict literature has been to use grid-cells or administrative units and add lags to control for spatial dependencies (e.g., Fjelde and von Uexkull 2012). However, this type of control does not fix the problem that treatment is being applied non-randomly over a larger area. By treating each grid-cell as independent, the estimate of the standard error will be lower than it should. The modelling solution to this would be to somehow cluster cells depending on the monthly weather patterns. Another solution is to add fixed effects on each area-month, as we already do to ensure a randomized treatment assignment.

Rather than modelling spatial dependence, we have selected areas that contain whole conflicts. Our assumption here is that communal conflicts are local phenomena, between



**Map 1** The areas of analysis as well as the spread of communal violence

relatively stationary groups. By increasing the area size and selecting areas containing whole groups and conflicts, we reduce the likelihood that outside interference affects the outcomes.

Another source of interference in our setup is temporal dependence. It might be that past conflicts that partly were a result of droughts increase the likelihood of subsequent conflicts. If this is the case, we will overestimate the effect of the treatment. A problem with our setup is that the treatment is applied over a longer period of time (1, 3, 6, 12, and 24 months). If we control for conflict events that happened in the period after treatment is being applied, we might inadvertently control for outcomes that happened due to the treatment. Our solution is to control for conflict events before treatment is being applied (2 months before measuring conflict events when using 1-month SPI [since we measure outcome 1 month after the last treatment month], 4 months before when using 3-month SPI, etc.). The control we use is  $\log(n + 1)$  of the number of events ( $n$ ) that happened in the 6- and 12-month periods before starting treatment.

### 3.3 Excludability

Excludability refers to the point that random assignment should only affect the outcome through the assigned treatment, and not something else in addition (Gerber and Green 2012: 39ff). To discuss excludability, we must be clear about what the treatment is. While the treatment assignment variable ( $Z$ ) is the SPI, our treatment of interest is not SPI, rather it is scarcity of resources, lost revenues, and lost livelihoods that may result from the extreme weather ( $X$ ). One problem arising from this is “non-compliance,” simply the point that some individuals who are assigned to unusually dry weather “fail” to lose their livelihoods or vice versa. “Non-compliance” should be expected to bias the estimate of the average treatment effect of  $X$  on communal conflict toward zero.

Further, we cannot exclude the possibility that the government or NGOs provide assistance whose intensity is correlated with the treatment assignment variable ( $Z$ ). We can imagine that the effect of SPI on conflict approaches zero if the government has sound insurance policies for all farmers who lose their income and investments under floods and droughts, if food prices are subsidized when prices surge due to droughts, and if markets are well connected so that the

needed resources are available to the population. This scenario would bias our estimate toward zero as assistance is increased.

Another scenario is that both droughts and floods change the terrain, which could make fighting more or less difficult (Miguel et al. 2004). If roads are impossible to travel due to floods, then the likelihood of conflicts might go down. However, floods could make it more difficult to pursue attackers with heavy weaponry such as government forces, possibly making floods a good cover to initiate a raid.

We have no way of controlling for these influences in our current setup. Thus, what we estimate is the so-called intention-to-treat effect (ITT) (Gerber and Green 2012: 138f). The ITT effect is the causal effect of experiencing unusual precipitation on the likelihood of a communal conflict. Not being able to estimate the average treatment effect of resource scarcity, however, means that we must be careful in our interpretation of the ITT. In the following section, we describe the data and the statistical interference method that we will use to test the relationship between communal conflict and climate variability.

## 4 Data

Our outcome variable is geo-coded communal conflict events between 1989 and 2013 by combining the UCDP GED and the UCDP non-state conflict datasets (Sundberg and Melander 2013; Sundberg et al. 2012). The definition of a communal conflict is “conflict [stands] along lines of communal identity,” where communal identity is defined as “[g]roups that share a common identification along ethnic, clan, religious, national or tribal lines. These are not groups that are permanently organized for combat, but who at times organize themselves along said lines to engage in fighting” (Pettersson 2014).

These events are then aggregated into the six chosen sets of administrative areas geographically defined in the GADM v.2.5 database of global administrative areas. Conflict events that span more than 1 month are coded in all those months. Our outcome variable is coded whether (1) or not (0) at least one communal conflict event in a given month intersects with the area of interest.

We get our treatment assignment variable by calculating the SPI from the average monthly precipitation in the areas we study. The precipitation measurement we use is the GPCC v.7.0 0.5° grid resolution product (Schneider et al. 2015). This data is based on station gauges and interpolated between stations. To calculate SPI, we use the SPEI R package with standard options (Beguería and Vicente-Serrano 2013). We include all months from 1948 to 2014 when estimating the SPI.

From the SPI, we define three different treatment assignment categories: “drought” (SPI  $\leq -1.5$ ), normal climate (SPI between  $-1.5$  and  $1.5$ ), and “flood” (SPI  $\geq 1.5$ ). A drought should here be understood as unusually dry climate compared to the precipitation in that area for that month or interval of months (e.g., an unusually dry January compared to normal precipitation in January).

SPI can be defined for different temporal intervals. Rather than fitting a distribution for the set of January months, we could fit a distribution for the sum of precipitation for January, December, and November, for all years. Called SPI-3, this 3-month interval is calculated for each month and measures whether the precipitation the last 3 months is unusual compared to that same interval of months in other years.

Different SPI intervals are relevant because they measure different types of “floods” and “droughts.” Below normal rainfall in the short term can lead to the soil drying up and crop stress. Below normal rainfall over a whole year might not materialize as immediate crop stress, but could reduce stream flows and reservoir levels. It is possible that any given area is both exposed to a short-term drought and a long-term flood at the same time. We test a variety of SPI time scales (1, 3, 6, 12, and 24 months) to capture these related-but-different phenomena.

To some degree, it is also possible to interpret the different SPI intervals as capturing the sustained versus the short-term effect of unusually dry or wet weather. Floods might have a more immediate effect than droughts and thus affect conflict in the short term (Raleigh and Kniveton 2012; von Uexkull 2014). However, longer SPI intervals do not necessarily capture short-term floods or droughts. Sustained large floods might better be operationalized through measuring consecutive months of high SPI-1, rather than, e.g., SPI-12. Within this research design, however, that kind of effect must be regarded as a possible interference problem, rather than something we can estimate the causal effect of. The reason is that while we can have confidence that the probability of some level of precipitation in a specific place in January is equal across time, we do not know whether the probability of having consecutive droughts in January to April is comparable across time and space.

## 5 Statistical inference

Since we cannot exclude other treatments from happening given our treatment assignment, and since the population will to a varying extent take the treatment if assigned (e.g., lose their livelihoods in a flood), we are estimating the average intention-to-treat effect, not the average treatment effect. We run OLS with fixed effects for each area-month, dummies for the two treatment conditions (holding normal weather as the base) as independent causal variables, and communal conflict as outcome.

$$cc_{it} \sim \beta_1 * \text{drought}_{it} + \beta_2 * \text{flood}_{it} + a_i * \text{area-month}_i$$

There is an ongoing discussion whether OLS give just as good estimates of marginal effects when the outcome is binary, as nonlinear models such as logit or probit (Beck 2011). The benefit of using OLS is that we get the marginal effects that we are interested in directly, rather than having to calculate these from the output of the nonlinear model (Angrist and Pischke 2008, 107). However, since Beck (2011) shows that the Angrist and Pischke’s “folk theorem” only is true under specific circumstances, we also estimate conditional logit models to check for the robustness of this assumption, and the results hold (see [Supplementary Materials](#)).

Because the core assumptions could be wrong (especially problematic is temporal dependence of past treatment), we control for the log sum of number of events in the area before treatment started in the last 6 and 12 months in separate regressions.

$$\begin{aligned} & \text{time} = 1, 3, 6, 12, 24 \\ & \text{eventcount}_{it} = \text{sum}(\text{ccevents}_{(it-\text{time}-6, it-\text{time}-1)}) \\ & cc_{it} \sim \beta_1 * \text{drought}_{it-1} + \beta_2 * \text{flood}_{it-1} + a_i * \text{area-month}_i + c_1 * \log(\text{eventcount}_{it} + 1) \\ & \text{time} = 1, 3, 6, 12, 24 \\ & \text{eventcount}_{it} = \text{sum}(\text{ccevents}_{(it-\text{time}-12, it-\text{time}-1)}) \\ & cc_{it} \sim \beta_1 * \text{drought}_{it-1} + \beta_2 * \text{flood}_{it-1} + a_i * \text{area-month}_i + c_2 * \log(\text{eventcount}_{it} + 1) \end{aligned}$$

In addition to the dummy treatment, we also consider the more parsimonious model using absolute SPI. The assumption here is that the direction of the average effect of unusual weather is the same whether it is unusually wet or unusually dry.

$$\begin{aligned}
 & \text{time} = 1, 3, 6, 12, 24 \\
 & \text{eventcount}_{it} = \text{sum}(\text{ccevents}_{(it-\text{time}-6, it-\text{time}-1)}) \\
 & cc_{it} \sim \beta_3 * | \text{SPI} |_{it-1} + a_i * \text{area-month}_i + c_1 * \log(\text{eventcount}_{it} + 1)
 \end{aligned}$$

We test the null hypothesis that the intention-to-treat effect is zero through normal approximation using the *t* test. We consider the two-tailed test appropriate because we consider both negative and positive deviations from zero theoretically interesting. We reject the null hypothesis at  $p < 0.05$  level of significance.<sup>1</sup>

Table 1 shows the results for the four equations using the five different SPI intervals. In the uppermost section, we show the model without controlling for previous conflicts. In the next two, we control for log of the number of past events during the last 6 and 12 months respectively. In the lowermost section, we have swapped the dichotomous measures of drought and flood with absolute SPI, and control for log of the past 6 month events. All regressions have fixed effects for each specific area-month. There are 20 different regressions in the table. For all drought intervals except the 24-month interval, the best estimate is above zero. Unusually dry climate the last three months is estimated to increase the likelihood of conflict with 10 percentage points in the study areas in the model without controls. Only the estimate for 3- and 6-month droughts is statistically significant. Unusually wet climate over short periods does not increase the likelihood of communal conflicts, but unusually wet years and two-year periods are estimated to increase the likelihood of communal conflicts with 6 and 7 percentage points respectively.

When we control for previous conflict events, the estimates for droughts approach zero—more so when using the last 6 months (before treatment starts) control than when using the last 12 months control. However, the best estimate is still positive for the drought intervals 1 and 12 and significantly so for 3- and 6-month intervals.

The estimates for unusually wet yearly and 2-year periods stay similar to the base model, and the 2-year interval is statistically significant. One reason for this can be that the previous conflict control is less effective, as we cannot control for conflicts happening in the last year and 2 years, respectively, due to the long exposure time. From  $R^2$ , it is evident that the model where we are controlling for communal conflicts closer in time to the outcome is a better fit to the data, than when controlling for communal conflicts longer back in time.

Since the direction of effects for both flood and drought are theorized to be the same, and we also see it in our estimates, we calculate a more parsimonious model with absolute SPI as treatment. One increase in absolute SPI-6 or SPI-12 increases the probability of a communal conflict event in a month with 4%. This latter has a *p* value of 0.0065, meaning that even after Bonferroni correction for testing five different SPI treatments, it is still significant.

Since we argue that SPI is a randomized variable, *future* SPI should be independent of past conflict events. A placebo test can be a way to test the validity of the causal model. We calculated the same models as in Table 1, only that all SPI values were shifted 24 months into the future. The full results are shown in the [Supplementary Materials](#). We find that the effects

<sup>1</sup> Since we are testing five different but related treatments, there is an argument that we must adjust our target *p* value accordingly. A conservative approach is to do a Bonferroni correction, which gives us a new target *p* value of  $0.05/5 = 0.01$ .

**Table 1** Regressions (fixed effects not shown, but included in all models)

	1 month	3 months	6 months	12 months	24 months
(Intercept)	0.54*** (0.07)	0.54*** (0.07)	0.53*** (0.07)	0.54*** (0.07)	0.54*** (0.07)
Drought	0.07 (0.04)	0.10* (0.04)	0.11** (0.04)	0.07 (0.04)	-0.02 (0.04)
Flood	-0.00 (0.04)	-0.00 (0.04)	0.03 (0.04)	0.06 (0.04)	0.07* (0.03)
$R^2$	0.19	0.19	0.19	0.19	0.19
Adj. $R^2$	0.15	0.15	0.15	0.15	0.15
Num. obs.	1728	1728	1728	1728	1728
RMSE	0.36	0.36	0.36	0.36	0.36
(Intercept)	0.23** (0.07)	0.30*** (0.07)	0.36*** (0.08)	0.44*** (0.08)	0.46*** (0.08)
Drought	0.04 (0.04)	0.07* (0.03)	0.07* (0.04)	0.05 (0.04)	-0.03 (0.04)
Flood	-0.02 (0.03)	-0.01 (0.03)	0.02 (0.03)	0.07* (0.04)	0.09** (0.03)
Log(# past 6-month events)	0.20*** (0.01)	0.17*** (0.01)	0.12*** (0.01)	0.07*** (0.01)	0.05*** (0.01)
$R^2$	0.39	0.33	0.26	0.22	0.21
Adj. $R^2$	0.37	0.30	0.23	0.18	0.17
Num. obs.	1698	1698	1698	1698	1698
RMSE	0.32	0.33	0.35	0.36	0.36
(Intercept)	0.27*** (0.07)	0.29*** (0.07)	0.32*** (0.08)	0.43*** (0.08)	0.47*** (0.08)
Drought	0.06 (0.04)	0.08* (0.04)	0.09* (0.04)	0.05 (0.04)	-0.03 (0.04)
Flood	-0.01 (0.03)	-0.00 (0.03)	0.02 (0.04)	0.07 (0.04)	0.10** (0.04)
Log(# past 12-month events)	0.15*** (0.01)	0.12*** (0.01)	0.09*** (0.01)	0.05*** (0.01)	0.04*** (0.01)
$R^2$	0.35	0.30	0.26	0.22	0.22
Adj. $R^2$	0.32	0.26	0.22	0.19	0.18
Num. obs.	1662	1662	1662	1662	1662
RMSE	0.33	0.34	0.35	0.36	0.36
(Intercept)	0.28*** (0.07)	0.29*** (0.07)	0.32*** (0.08)	0.43*** (0.08)	0.48*** (0.08)
SPI	-0.01 (0.01)	0.02 (0.01)	0.04* (0.01)	0.04** (0.01)	0.01 (0.01)
Log(# past 6-month events)	0.20*** (0.01)	0.16*** (0.01)	0.12*** (0.01)	0.07*** (0.01)	0.05*** (0.01)
$R^2$	0.39	0.33	0.26	0.22	0.21
Adj. $R^2$	0.37	0.30	0.23	0.18	0.17
Num. obs.	1698	1698	1698	1698	1698
RMSE	0.32	0.33	0.35	0.36	0.36

\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$

are estimated to be not different from zero for all operationalizations, except for the 12-month drought specifications. There are two reasonable points to make. First, barely significant results are expected to happen once in a while even while the average effect is zero. Thus, we do not think that randomization has failed, based on the placebo test. Second, while our results show that the estimated effect is significantly different from zero, the results are for the most part not significantly different from a world in which the effects of drought, flood, and |SPI| were not significantly different from zero.

## 6 Discussion: causal inference

The core problem in quantitative causal inference is to find sets of units that work as each other's counter-factual, suitable for comparison. In statistics, there are in principle many ways to get such a setup. However, the gold standard is randomized experiment. In this paper, we argue that we can compare conflict outcomes in months within area-month sets, randomized into groups through the SPI. We can combine several such experiments as long as we take into account differing probabilities of being assigned treatment in the different area-month sets.

The setup has several advantages compared to the traditional control setup. There is a simple rule for when identification is successful, compared to control setup, where there is always room for another potential confounder. We can compare our results with results from a placebo test. And perhaps most important, all the parts that go into the identification process are transparently laid out in the research design, rather than being sealed within a multivariate regression scheme.

There are some obvious short-comings of naturally occurring experiments. Most important is that not all research questions come with a convenient and relevant naturally occurring randomized treatment, and climate variability is a case in point. The causal effect of absolute rainfall or temperature cannot be determined through naturally occurring experiments because they are not randomized. We suggest that one of the few such randomized variables is SPI compared across months within area-month sets. Thus, there is a limit to the types of causal effects we can identify from this approach.

The other shortcoming is that it is still doubtful whether we manage to properly identify the stochastic process that determines average monthly precipitation. This is the core argument our results are resting on, and proof that this identification is flawed would undermine our results. However, this problem is true for all naturally occurring experiments. We think the placebo test supports the core argument and that we actually are making fair comparisons.

Results seem to be biased due to temporal dependence. Past conflicts that partly were a result of droughts increase the likelihood of subsequent conflicts. Controlling for the number of past events in the last 6 months moves the effect closer to zero. There is a possibility that an even better control variable can bring the average effect lower. We think the gain of doing so, however, would be minimal as there would still need to be *some* effect in order for treatment interference through temporal dependence to be a problem.

Another possible interference problem comes from being affected by neighboring conflicts affected by droughts. In this study, we have taken a pragmatic design choice, by selecting quite large areas which contain whole communal conflicts and groups living within these areas. We think this choice has removed the need to control for spatial interference problems.

Further, communal conflicts are more local and static than state-based conflicts, which often move across the country. We would therefore not recommend using this research design when studying state-based conflicts. An alternative for sub-country studies of state-based conflicts is to match climate variability with the living areas of conflict actors, rather than spatially matching with the conflict event itself (Nordkvelle 2016; von Uexkull et al. 2016).

Before interpreting the effect as substantially being due to scarcity of resources, lost revenues and/or lost livelihoods, we would need to exclude several other theories, such as varying levels of humanitarian assistance and tactical considerations. The long-term "floods," where we get a significant effect, are not especially likely to affect infrastructure, as this is not a situation where water is necessarily coming quickly in large amounts. It might be, however, that the effects of shorter term floods would have been higher, if it had not also been the case that such floods made it more difficult to travel. Humanitarian assistance is also expected to

reduce the effect of droughts and floods toward zero. Therefore, we believe that the estimates we find are there *despite* any humanitarian efforts in the areas we studied.

## 7 Discussion: substantial results

The results show that both droughts and floods increase the likelihood of communal conflict events in the areas we investigated. There seem to be qualitative differences between the two, however. Unusually wet 1- to 6-month periods are estimated to have zero effect on communal conflicts. Unusually wet 12- to 24-month periods, on the other hand, do seem to have an effect. The converse seems to be true for droughts, where long-term droughts have no effect, but short-term droughts do—which partly supports Fjelde and von Uexkull (2012).

This result could give support to both resource abundance and resource scarcity theories. Possibly, there needs to be some resource abundance for communal conflicts to “pay off” (long, unusually wet periods), while droughts that destroy harvests (short, unusually dry periods) pull resource-scarce groups into violence.

It is important to notice that the prevalence of droughts as we have defined them will not increase in the future. Therefore, these results cannot be used to argue that we will observe more or less communal conflicts in the future. If the future will see more extreme weather, it means that extreme weather will become more usual, making such weather having a lower SPI. However, from a theoretical perspective, we can argue that to the extent climate change puts strains on the ability to adapt to unusual weather, or creates more unusual weather before we manage to adapt, we should expect more communal conflict events in the study areas due to climate change.

There is not a clear link between the result for communal conflicts and that for state-based conflicts (civil wars). While the results indicate that changes to the resource base do affect the willingness of marginalized communal groups to use violence against other marginalized groups, it is not obvious that they would increasingly choose to target the government. The government is a much stronger adversary, the risk is higher, and the end goal of the violence is usually quite different.

At the same time, some of the most convincing empirical tests of the economic causal mechanisms of why people use violence, such as the opportunity cost mechanism, can be found in studies of non-state conflicts. Since communal conflicts are more locally situated, methodological issues such as treatment interference should be less of a problem than in state-based conflicts. Studies of communal conflict can therefore be used to motivate theoretical explanations of violence even in the civil war setting, perhaps particularly when studying local dynamics of civil war (Kalyvas 2006). However, we should be careful to directly extrapolate causal effects of climate variability from the communal conflict setting to the civil war setting.

**Acknowledgements** This work is supported by the US Army Research Laboratory and the US Army Research Office via the Minerva Initiative grant no. W911NF-13-1-0307, the Norwegian Research Council project CAVE, grant no. 240315/F10 and

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

## References

- Angrist JD, Pischke J-S (2008) Mostly harmless econometrics: an empiricist's companion. Princeton University Press
- Beck N (2011) Is OLS with a binary dependent variable really OK?: estimating (mostly) TSCS models with binary dependent variables and fixed effects. Unpublished manuscript. Available from <http://politics.as.nyu.edu/docs/IO/2576/pgm2011.pdf>
- Beguerra S, Vicente-Serrano SM (2013) Calculation of the standardised precipitation-evapotranspiration index. SPEI R package version 1.6
- Fjelde H, von Uexkull N (2012) Climate triggers: rainfall anomalies, vulnerability and communal conflict in sub-Saharan Africa. *Polit Geogr* 31:444–453
- Gerber AS, Green DP (2012) Field experiments: design, analysis, and interpretation. W.W. Norton & Company
- Guttman NB (1999) Accepting the Standardized Precipitation Index: a calculation algorithm. *J Am Water Resour Assoc* 35(2):311–322
- Hendrix CS, Salehyan I (2012) Climate change, rainfall, and social conflict in Africa. *J Peace Res* 49:35–50
- Homer-Dixon T (1994) Environmental scarcities and violent conflict: evidence from cases. *Int Secur* 19:5–40
- Kalyvas SN (2006) The logic of violence in civil war. Cambridge University Press
- Landis ST (2014) Temperature seasonality and violent conflict: the inconsistencies of a warming planet. *J Peace Res* 51:603–618
- McKee TB, Doesken NJ, Kleist J, et al. (1993) The relationship of drought frequency and duration to time scales. In Proceedings of the 8th Conference on Applied Climatology, 17:179–183. 22. American Meteorological Society Boston, MA, USA
- Miguel E, Satyanath S, Sergenti E (2004) Economic shocks and civil conflict: an instrumental variables approach. *J Polit Econ* 112(4):725–753
- Nordkvelle J (2016) Randomized rain falls on political groups: discovering an average causal effect of climate variability on armed conflict onsets. Presented at the ISA Annual Conference in Atlanta
- O'Loughlin J, Witmer FDW, Linke AM, Laing A, Gettelman A, Dudhia J (2012) Climate variability and conflict risk in East Africa, 1990–2009. *Proc Natl Acad Sci* 109:18344–18349. doi:10.1073/pnas.1205130109
- O'Loughlin J, Linke AM, Witmer FDW (2014) Modeling and data choices sway conclusions about climate-conflict links. *Proc Natl Acad Sci* 111:2054–2055. doi:10.1073/pnas.1323417111
- Pettersson T (2014) UCDP non-state conflict codebook, version 2.5-2015. [http://www.pcr.uu.se/research/ucdp/datasets/ucdp\\_non-state\\_conflict\\_dataset/](http://www.pcr.uu.se/research/ucdp/datasets/ucdp_non-state_conflict_dataset/). Accessed 14 Jan 2015
- Raleigh C (2010) Political marginalization, climate change, and conflict in African Sahel states. *Int Stud Rev* 12: 69–86
- Raleigh C, Kniveton D (2012) Come rain or shine: an analysis of conflict and climate variability in East Africa. *J Peace Res* 49:51–64. doi:10.1177/0022343311427754
- Salehyan I, Hendrix CS (2014) Climate shocks and political violence. *Glob Environ Chang* 28:239–250
- Schneider U, Becker A, Finger P, Meyer-Christoffer A, Rudolf B, Ziese M (2015) GPCC full data reanalysis version 7.0 at 0.5°: monthly land-surface precipitation from rain-gauges built on GTS-based and historic data. doi:10.5676/DWD\_GPCC/FD\_M\_V7\_050
- Sundberg R, Melander E (2013) Introducing the UCDP georeferenced event dataset. *J Peace Res* 50(4):523–532
- Sundberg R, Eck K, Kreutz J (2012) Introducing the UCDP non-state conflict dataset. *J Peace Res* 49:351–362
- Theisen OM (2012) Climate clashes? Weather variability, land pressure, and organized violence in Kenya, 1989–2004. *J Peace Res* 49:81–96. doi:10.1177/0022343311425842
- Tollefsen AF, Strand H, Buhaug, H (2012) PRIO-GRID: A unified spatial data structure. *J Peace Res* 49(2):363–374
- von Uexkull N (2014) Sustained drought, vulnerability and civil conflict in sub-Saharan Africa. *Polit Geogr* 43: 16–26. doi:10.1016/j.polgeo.2014.10.003
- von Uexkull N, Croicu M, Fjelde H, Buhaug H (2016) Civil conflict sensitivity to growing-season drought. *Proc Natl Acad Sci* 113:12391–12396. doi:10.1073/pnas.1607542113
- Witsenburg KM, Adano WR (2009) Of rain and raids: violent livestock raiding in Northern Kenya. *Civil Wars* 11:514–538