Climate, Conflict and Growth in Africa

Reverse Causality and Predictive Power

Sebastian Krantz

August 30, 2022

Abstract

In this paper I investigate two aspects of the climate-economy conflict nexus that have thus far been neglected in the literature: Reverse causality and predictive power. With regards to predictive power, I employ a machine learning approach and show that climate moments can successfully predict a good part of African growth rates and conflict incidence. In the second part of this paper, I attempt to model the reverse causality between conflict and growth using two different instrumentation approaches. The results confirm the existence of a strong channel from conflict outbreak to negative growth, but could not provide convincing evidence of a strong channel from negative growth to conflict outbreak.

With the rise of climate change related social issues such as the effects of droughts, floods, extreme temperatures, storms etc. becoming subject at the highest levels of policy discourse, recent years have seen a surge of interest in economics and other social sciences to link socioeconomic outcomes to climate change and extreme weather events. A prominent place in this exponentially growing body of research is taken by the task of relating climate change and extreme weather events to economic growth and social conflict. Besides these new strands of research involving climate, the links between conflict and economic conditions have been of interest for a long time to a broad range of researchers, and research in the area has been very productive since the rise in popularity of peace and conflict studies in the late 90's.

It shall in the following not be upon me to review this literature. Excellent reviews of the climate-economy-conflict literature are given by Burke et al. (2015) and Koubi (2017) and most of the articles cited in this paper. More general reviews of research on the socioeconomic impacts of climate change are provided by Dell et al. (2014) and Carleton & Hsiang (2016). Good reviews of the conflict-economy literature are included in Berman & Couttenier (2015) and Miguel et al. (2004) and in many of the latest papers in the Journal of Peace Research owned by the The Peace Research Institute Oslo (PRIO).

The purpose of this paper is to address two issues that the literature has not given much attention to: modelling the reverse causality between conflict and economic outcomes (growth), and assessing the predictive power of climatic factors regarding these outcomes. Both issues are important: Refraining from explicitly modelling and estimating specifications that account for reverse causality between conflict and economic outcomes, and thus not allowing for general-equilibrium effects in the form of vicious cycles, is likely to downward bias the empirical results obtained from partial specifications. In particular, a lot of research to date has been devoted to understanding the one-way partial equilibrium effect of economic conditions on conflict. Predictive power is a crucial test to evaluate the added value of new theories and studies relating conflict and economic outcomes to climate change/climatic conditions. Research in this area has exploded in recent years. Significant effects of climate have been found in many studies, but the mechanisms through which climate affects conflict and the economy are not well understood, and the empirical practice in terms of moments of climate variables being plugged into the right hand side of regressions is very diverse. An assessment of these new developments in terms of predictive power is thus urgently needed.

Predictive power is chosen as the preferred metric because significance levels and effect sizes are too easily manipulatable and models are prone to overfitting the data. Furthermore the theoretical uncertainty regarding climate impacts invites the use of nonparametric methods. Machine learning tools have been developed to produce outstanding predictions from highly nonlinear and nonparametric models while effectively addressing the overfitting problem through model averaging and out-of-sample prediction (cross-validation). In this paper I will employ these tools to assess the predictive power of climate.

The remainder of the paper is structured as follows: Section 1 assesses the predictive power of climate on two datasets: A country-year dataset covering Sub-Saharan Africa from 1971-2006, and a highly disaggregated (0.5×0.5) gridded dataset covering the whole of Africa from 1980-2016. Section 2 addresses reverse causality using the same two datasets and an instrumental variables approach. A conclusion follows.

1 Predictive Power

To assesses the predictive power of climate and weather events on growth and conflict in Africa, I pool data from 6 cross-country studies jointly covering the period 1971-2007, and predict conflict and growth while assessing the importance of different climate and weather related factors. I employ a machine learning approach for this task. The main reason for machine learning tools instead of the often more interpretable regression based methods such as binary and multinomial choice models or linear regression, is the large uncertainty in the literature regarding the functional form of climate based models of conflict and economic growth.

Whereas there has emerged a near consensus on the conflict measures employed[1], the literature still exhibits great ambiguity with regards to the right climate measures, interaction effects, and the way climate measures enter conflict or growth regressions[2]. The study of interactions between climatic factors affecting conflict and growth has thus far also been neglected in the literature. Machine learning algorithms come in handy in settings like this because they don't require an explicit parametric form and only make some method specific assumptions about the data-generating process. They are also geared towards predictive power and learn from associations in the data in order to attained maximal prediction accuracy, resulting in predictive performances often far surpassing those of regression based methods. Lastly, many machine learning algorithms are highly non-linear (such as tree-based methods or neural networks), and consider complex patterns and interactions among variables. These algorithms can often handle an arbitrary amount of variables and accurately assess the contribution of each variable to the prediction without running into multicollinearity problems. Overfitting is addressed through cross-validation and out of bag predictions[3].

The country dataset is constructed by combining climate, economic and conflict data from 6 cross-country panel studies examining the links between climatic events, the economy and intergroup conflict. The studies are: Miguel & Satyanath (2011) (M&S), Hendrix & Salehyan (2012) (H&S), Buhaug et al. (2015) (BBST), Koubi et al. (2012) (KOB), Bergholt & Lujala (2012) (B&L), Landis (2014) (LAN) and Couttenier & Soubeyran (2014) (C&S). From the dataset thus constructed I take 22 climate/weather indicators and 9 outcome measures that well reflect the current practice in the literature, and create a new Sub-Saharan African dataset with no missing data on the climate variables. The result is a quite balanced panel with 1497 observations covering 42 Sub-Saharan African countries from 1971-2007. The data coverage for each country is shown in Table 1.

[1] Conflict incidence and onset measures from the UCDP/PRIO, ACLED or SCAD databases.

[2] The most frequent variables are temperature and precipitations measures, mostly taken from the UDel or CRU climatic databases, but, depending on the study, these measures enter in levels, lags, polynomials, extreme values, standard deviations, growth rates, moving averages, other more complex codings and convex combinations of all of the aforementioned.

[3] See large footnote on the Random Forest algorithm a few pages down

Table 1: DATA COVERAGE: 42 SSA COUNTRIES, 1971-2007

Country	*Years*	*N*	*Country*	*Years*	*N*	*Country*	*Years*	*N*
Angola	1975 - 2007	33	Gabon	1971 - 2007	37	Niger	1971 - 2007	37
Benin	1971 - 2007	37	Gambia	1971 - 2007	37	Nigeria	1971 - 2007	37
Botswana	1971 - 2007	37	Ghana	1971 - 2007	37	Rwanda	1971 - 2007	37
Burkina Faso	1971 - 2007	37	Guinea	1971 - 2007	37	Senegal	1971 - 2007	37
Burundi	1971 - 2007	37	Guinea-Bissau	1974 - 2007	34	Sierra Leone	1971 - 2007	37
Cameroon	1971 - 2007	37	Kenya	1971 - 2007	37	Somalia	1971 - 2007	37
Central African R	1971 - 2007	37	Lesotho	1971 - 2007	37	South Africa	1971 - 2007	37
Chad	1971 - 2007	37	Liberia	1971 - 2007	37	Sudan	1971 - 2007	37
Congo	1971 - 2007	37	Madagascar	1971 - 2007	37	Swaziland	1971 - 2007	37
Cote d'Ivoire	1971 - 2007	37	Malawi	1971 - 2007	37	Tanzania	1971 - 2007	37
Democratic Republ	1971 - 2007	37	Mali	1971 - 2007	37	Togo	1971 - 2007	37
Djibouti	1977 - 2007	31	Mauritania	1971 - 2007	37	Uganda	1971 - 2007	37
Equatorial Guinea	1971 - 2007	37	Mozambique	1975 - 2007	33	Zambia	1971 - 2007	37
Ethiopia	1991 - 2007	17	Namibia	1991 - 2007	17	Zimbabwe	1971 - 2007	37

The 9 dependent variables that will be subject to prediction are summarized in Table 2. I take two measures of GDP per capita growth, one from the Penn World Table and the other one from the World Bank. There has been some debate in the literature about which of these is the better measure. Both certainly have their shortcomings as GDP accounting in Sub-Saharan Africa is known to be error-prone, especially in the past. The correlation between the two growth rates is .77, as Table 3 shows.

Table 2: PREDICTED DEPENDENT VARIABLES

Variables	*N*	*Mean*	*SD*	*Min*	*Max*	*Skew*	*Kurt*
PWT version 6.3 - GDP per capita PPP	1,497	2,290	2,361	154	23,444	3.38	20.6
PWT version 6.3 - GDP per capita growth PPP	1,497	0.91	9.29	-62.4	123	2.93	38.7
WDI - GDP per capita growth	1,393	0.96	8.40	-50.2	142	4.52	74.8
B&L - Conflict year (25+ battle deaths in current year)	1,497	0.20	0.40	0	1	1.51	3.29
LAN - UCDP Non-State Conflict	834	0.32	0.96	0	8	4.06	22.3
B&L - Onset (>25 death) after >2 years of peace	1,497	0.043	0.20	0	1	4.52	21.4
KOB - Civil Conflict Onset (>25 deaths) after >9 years of p.	1,317	0.029	0.17	0	1	5.63	32.7
KOB - Conflict onset (>1000 deaths) after >9 years of peace	1,317	0.012	0.11	0	1	8.91	80.3
BBST - PRIO battle deaths best estimate	1,497	577	2,834	0	36,250	8.63	91.3

From Table 2 it is clear that not all of the outcome measures take on the full coverage of observations. The correlation matrix presented in Table 3 shows that the 3 GDP variables correlate negatively with all 6 conflict indicators (as expected), but the loadings are quite meager. GDP and conflict variables correlate positively among one-another, with the exception of non-state conflict (9) correlating negatively with onset of a big conflict (8). The latter is nut surprising as it is presumably less likely for a non-state conflict to occur at the same time as a big conflict involving a government party begins.

Table 3: CORRELATION MATRIX OF DEPENDENT VARIABLES
(*Pairwise Correlations*)

# *Variable*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) GDP per capita PPP, PWT version 6.3	1								
(2) GDP per capita growth PPP, PWT version 6.3	.11	1							
(3) WDI - GDP per capita growth	.15	.77	1						
(4) B&L - Conflict year (25+ battle deaths in current year)	-.09	-.10	-.10	1					
(5) BBST - PRIO battle deaths best estimate	-.08	-.08	-.10	.40	1				
(6) B&L - Onset (>25 death) after >2 years of peace	-.04	-.03	-.02	.42	.01	1			
(7) KOB - Civil Conflict Onset (>25 deaths) after >9 years of p.	-.04	-.04	-.03	.29	.02	.65	1		
(8) KOB - Conflict onset (>1000 deaths) after >9 years of peace	-.04	-.02	-.05	.18	.02	.43	.64	1	
(9) LAN - UCDP Non-State Conflict	-.13	-.06	-.03	.21	.10	.02	.01	-.03	1

For the sake of completeness, Tables 4 and 5 show the summary statistics and correlations of the 22 climate variables selected. The correlation matrix shows that the climate variables, though often seemingly similar in definition, are less correlated than expected. In particular levels, growth rates and standard deviations of temperature and precipitation measures are nearly unrelated to one another and thus able to potentially relate to growth and conflict variables in very different ways. A principal components analysis of the climate variables (not reported) yields 7 components with an eigenvalue greater that 1, together explaining 67% of the variance of the 22 variables, where the first component takes 18% of the total variance. The Kaiser-Meyer-Olkin measure of sampling adequacy is 0.63, which is not magnificent and indicates that the high-dimensional predictor space is not readily reducable to a smaller number of underlying dimensions.

Table 4: Climate Variables, $N = 1497$

#	*Variables*	*Mean*	*SD*	*Min*	*Max*	*Skew*	*Kurt*
(1)	KOB - Precipitation dev from 30-year moving avg, GPCC	-47.1	152	-944	994	-0.33	7.15
(2)	KOB - Temperature dev from 30-year moving avg, GPCC	0.22	0.50	-1.43	3.15	0.42	5.38
(3)	C&S - Palmer Drought Severity Index	-0.49	1.11	-4.65	3.40	0.13	3.34
(4)	B&L - Dummy = 1 if drought year	0.050	0.22	0	1	4.12	18.0
(5)	LAN - El Nino Year dummy	0.21	0.41	0	1	1.39	2.94
(6)	B&L - Large climate related disaster 1[>10,000 affected]	0.12	0.32	0	1	2.34	6.50
(7)	B&L - Affected by floods (% of pop), time weights	0.0017	0.016	0	0.33	14.5	239
(8)	B&L - Affected by storms (% of pop), time weights	0.00052	0.0091	0	0.31	29.1	948
(9)	B&L - Population share affected by earthquakes	4.7e-06	0.00013	0	0.0043	31.2	1,024
(10)	B&L - Population share affected by volcanoes	3.1e-06	0.000078	0	0.0023	27.0	745
(11)	BBST - Mean temperature (°C), UDel data	24.2	3.37	10.2	29.4	-1.46	6.26
(12)	BBST - Mean temperature growth (%), UDel data	0.077	1.94	-9.73	8.82	-0.021	4.74
(13)	BBST - Mean precipitation (1000mm), UDel data	0.93	0.49	0.11	2.45	0.56	3.06
(14)	BBST - Mean precipitation growth (%), UDel data	1.68	19.1	-52.6	167	1.49	10.3
(15)	BBST - Mean temperature in 1961-90	23.9	3.30	11.4	28.2	-1.43	6.07
(16)	BBST - Mean precipitation in 1961-90	0.96	0.51	0.15	2.24	0.55	3.02
(17)	BBST - Temperature dev as std dev from 1961-90 mean	0.50	1.04	-2.71	3.53	0.049	2.76
(18)	BBST - Precipitation dev as std dev from 1961-90 mean	-0.23	0.95	-4.76	4.16	0.27	4.14
(19)	LAN - NCAR/NCEP Yearly Temperature Mean	23.6	3.10	14.4	28.3	-0.96	3.31
(20)	LAN - Temperature Shock	0.089	0.29	-0.84	1.79	0.88	5.69
(21)	LAN - Standard Deviation of Temperature Mean	1.88	1.08	0.61	4.37	0.75	2.33
(22)	LAN - Standard Deviation Precipitation Mean	2.18	0.99	0.62	5.00	0.73	3.49

Indeed the considerable amount of 0's (for correlations <.005) in the correlation matrix suggests that there are a number of orthogonal climatic dimensions captured by these variables. A final point worth mentioning is that there are two cases of near-collinearity in the matrix, which are the mean temperature and precipitation indicators from Buhaug et al. (2015) that are almost perfectly correlated with their 1961-90 mean, although the latter is time-invariant. I note that due to randomization of predictors, near-collinearities are not problematic for the Random Forests (RF) algorithm that will be employed.

Table 5: Correlation Matrix of Climate Variables

#	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)
(2)	-.19	1																			
(3)	.44	-.29	1																		
(4)	-.03	.08	-.03	1																	
(5)	-.04	-.02	.01	-.03	1																
(6)	.17	0	.12	.05	-.06	1															
(7)	.07	.02	.09	-.02	-.04	.29	1														
(8)	.04	-.02	.03	0	.05	.15	-.01	1													
(9)	.07	-.02	.02	0	-.02	.06	.01	0	1												
(10)	.05	0	.06	-.01	.03	-.01	0	0	0	1											
(11)	-.10	.08	-.25	.01	-.01	-.01	.01	-.02	-.02	0	1										
(12)	-.25	.42	-.18	.05	.03	-.05	-.07	0	-.03	.01	.06	1									
(13)	-.07	-.07	.05	-.11	.01	-.03	-.06	.04	.02	.04	-.01	-.06	1								
(14)	.44	-.14	.17	-.03	.03	.17	.16	-.01	.04	0	-.02	-.39	.07	1							
(15)	-.07	-.01	-.22	0	0	-.02	.01	-.02	-.02	0	.99	0	0	0	1						
(16)	-.21	-.01	-.05	-.10	0	-.06	-.07	.03	.01	.04	.01	0	.97	-.05	.01	1					
(17)	-.20	.65	-.22	.09	-.05	.09	-.03	-.01	-.01	0	.25	.47	-.01	-.12	.12	.05	1				
(18)	.61	-.26	.52	-.04	.01	.18	.09	.03	.04	.01	-.05	-.30	.04	.60	-.02	-.16	-.25	1			
(19)	-.09	.05	-.26	-.01	0	-.03	-.04	-.01	-.03	-.01	.90	.05	.13	-.03	.89	.15	.19	-.05	1		
(20)	-.08	.42	-.06	.05	-.02	.11	-.04	-.01	-.03	.03	.12	.28	.11	-.06	.05	.14	.64	-.13	.15	1	
(21)	.13	.07	.02	.04	0	0	.04	-.01	0	-.04	-.20	0	-.62	.03	-.20	-.64	-.08	.13	-.21	-.22	1
(22)	-.17	.06	-.14	-.06	0	-.04	-.04	.05	.02	-.03	.10	0	.69	-.03	.10	.72	.06	-.13	.20	.04	-.40

1.1 Identification and Prediction

Although machine learning methods are nonparametric and differently purposed that parametric statistical models, in the context of this applications they are not immune to omitted variable bias or country level individual heterogeneity in the panel setup considered, because an inconsiderate setup runs danger to overpredict outcomes in the sense that a proportion of the predictive power of the machine learning model stems not from the set of predictor variables causing (directly or indirectly) the outcome, but from factors uncaused by and correlated with predictors that are causing the outcome. This possibility must be excluded for the results to be valid, and therefore this issue warrants some careful consideration.

In the literature the question of whether one may use cross-sectional variation to identify the effects of climate on socioeconomic outcomes, formally known as the unconfoundedness assumption, remains controversial. One side to this controversy, represented by Burke et al. (2015) among others, holds that there are countless country-specific characteristics shaping socioeconomic outcomes, and that causal identification of climatic effects would require the researcher to account for all of these differences (that is to create ceteris paribus homogenous populations except for climate). Since this task is impossible, these researchers have abandoned cross-sectional designs in favor of panel-data designs where the frequency-identification trade-off represents a sizable, but, in the view of these authors, more manageable challenge.

The alternative to this view is authors that claim that climate and weather are fundamentally exogenous constructs[4]. Another way to put it is that this view assumes that climate is external and given, and that socioeconomic characteristics are either ultimately a consequence of climate (thus controlling for them would be bad control), or random relative to climate (i.e. produced by processes fundamentally unrelated to climate and therefore uncorrelated with it). Which of the two views one espouses depends, in part, on the time horizon one is considering and the type of effects one wishes to detect. In a study opting to detect the effects of climate fluctuations on conflict over very short time horizons, it appears a reasonable choice to include cross-sectional unit fixed effects, while in a study aiming to uncover the effects of climate on long run economic development, including fixed effects would be considered "bad control". In this study, the effects of different "levels" of climate on different base levels of economic growth and conflict across Sub-Saharan African countries is of interest in assessing the predictive power of climate, therefore the stance taken will be closer to the second view, but more discussion on this point will follow later.

The second main empirical challenge in time series or panel data studies, as already hinted at, is the frequency identification trade-off implied by the unit-homogeneity assumption for time series. This assumption holds that discrete populations can serve as their own baseline, and respond in the same way to climatic changes over time. Formally, this is stated as:

$$E[Y_{i\tau+\Delta\tau}|C,\mathbf{x}_{i\tau+\Delta\tau}] = E[Y_{i\tau}|C,\mathbf{x}_{i\tau}] \tag{1}$$

which reads: Controlling for observable characteristics $\mathbf{x}_{i\tau}$ in period τ and $\tau+\Delta\tau$ and conditioning on a fixed climate C, the outcome of interest Y will be equal in periods τ and $\tau+\Delta\tau$, such that the effect of climate change on Y may be identified as:

$$\hat{\beta}_{\Delta C} = E[Y_{i\tau+\Delta\tau}|C_{\tau+\Delta\tau},\mathbf{x}_{i\tau+\Delta\tau}] - E[Y_{i\tau}|C,\mathbf{x}_{i\tau}]. \tag{2}$$

In the setup at hand this assumption would imply that African societies in 1971 and 2007 respond similarly to climate shocks in terms of conflict and economic growth, thus the data within a given country can be pooled and treated as a kind of randomized or staged experiment: treatments in the form of climate/weather shocks are delivered to the same country at certain moments in time, but society does not adapt to them or change its response in other fundamental ways over time. Thus under this assumption, the difference in means represents the ATE of climate. The frequency identification trade-off (termed so by Hsiang (2016)) then holds that exploiting high-frequency variation in climate/weather is more likely to comply with the unit homogeneity assumption given in Eq. 1, but comes at the cost of being unable to identify the total effects of climate/weather

[4]Exempting global climate change, which is associated with the global production level and often abstracted from, and exempting other geographic factors correlated with climate and influencing the outcome of interest (such as mountainous terrain, coastlines etc.)

change which involve belief effects and adaptations. The direction of the bias resulting in 2 from a failure of Eq. 1 depends on how climate effects and belief effects/adaptations/socioeconomic change relate to one another. If belief effects are large relative to direct effects of climate change and are in the same direction (e.g. migrations, economic down turn etc.), the total effect of climate/weather change is underestimated with high frequency data and frequency-identification trade-off may represents a major challenge. However if belief effects are mainly adaptation designed to mitigate direct effects (e.g. of floods, droughts, or storms), high frequency data might be close to capturing total effect (Hsiang, 2016).

The extent to which African societies were able to adapt to climatic events over the past 40 years is uncertain. There have certainly been changes in the composition of income from strongly agricultural economies to more industry and service oriented economies in most African societies, which is supported by anecdotal empirical evidence, such as the near collapse of the strong inverse relationship between rainfall and growth studied by Miguel et al. (2004) and others in the early 2000's. It remains reasonable to assume that adaptation to climate change will further improve in Africa in the coming decades, but it equivalently appears plausible that many effects and trends, such as the waves of migration, sustained population growth and scarcities, and resource quarrels such as the effects of Ethiopia's Blue Nile dam project, have increased Africa's vulnerability.

In the context of this study there are thus fundamentally two decisions to make concerning the cross-sectional and time series unit homogeneity assumptions. As regarding the cross-sectional dimension, denying unconfoundedness and demeaning the data (or any equivalent transformation) would throw away all of the cross-sectional variation. But this would also render a good proportion of questions of interest, for example the effects of different climatic conditions on base levels of growth and conflict across different Sub-Saharan African countries unanswerable. Furthermore, the interactions of levels and changes in conflict, growth, and climatic factors, which machine-learning tools can easily detect and use, is also unrecoverable in a demeaned design. The estimate of the predictive power of climate provided in such a design would necessarily be a lower bound of the total climatic effect.

For the purposes of this exercise more value is placed on obtaining an upper bound estimate by viewing climate as given and social characteristics as either an outcome of climate or random with respect to it. Tentative threats to this assumption are especially the possibility of other geographic characteristics (such as mountainous terrain and distance to coastlines) affecting outcomes of interest and correlated with climate. In a linear setting one would just include these variables into the regression to partial them out, but partialling out is remote to the machine learning context, e.g. including variables generally increases predictive performance but does not let the researcher extract precisely the effect of some variables ceteris paribus. One can only compare the predictive power models with and without geographic predictors added, which is done when using the gridded data set. This is however also of limited value since the correlation of climate with geographic characteristics such as distance to coast is not informative about which one (climate or distance to coast) is actually causing the social outcome.

I thus refrain from demeaning the data in the knowledge that attributing all predictive power to climate yields an upper bound. Empirical support for this decision is provided by a series of Hausman tests in which I regressed each of the dependent variables in Table 2 on all time-varying climate variables (18 of 22) in levels and lagged once (thus 36 predictors in total), and compared the fixed to the random effects estimator. In all 9 cases, I fail to reject the null that the differences in coefficients is non-systematic[5]. In the same accord I also run 9 Breusch and Pagan Lagrangian multiplier tests for random effects, and in all cases fail to reject the null that the variance of the individual heterogeneity is 0. The linear panel model of choice would thus be pooled OLS.

Regarding the time series dimension, the problem is confronted in the same way. In a linear setup, a natural step to strengthen the time series unit homogeneity assumption would be to include time fixed effects or country specific time trends, possibly interacted with climate measures. In partialling out one again faces a bias-variance tradeoff, since including time fixed effects or time

[5]For most of the models, the R-squared is between 5 and 15%, and there are a lot of variables, thus these test results should be taken with extreme caution.

trends will remove the effects of global climate change and trends in growth/conflict. Partialling out time effects linearly would also compromise on some of the nonlinear abilities of machine-learning methods. It is clear that Sub-Saharan Africa has changed in some very fundamental ways over the last 40 years, but, as reasoned before, it is unclear how these fundamental changes moderate the relationship between climate, economic growth and conflict. Therefore not partialling time effects and obtaining an upper bound estimate on the predictive power of climate change is again the preferred strategy.

An imperfect empirical test of the time series unit homogeneity assumption in this context can be conducted by examining correlations between climate and outcome variables over time. I implement the test by computing the correlation coefficient between each climate and outcome variable over 5 year intervals. In order for changing correlation coefficients to truly capture trends in the relationship between climate and outcome variables, and not changes in the sample of countries from each period to the next, the panel is rebalanced for each pair of climate and outcome variables[6]. The sample size in each pooled 5-year interval is between 100 and 200.

Figure 1 shows the result, together with a quadratic fit to visualize trends in each relationship. The first panel involving temperature variables in red, and precipitation variables in blue, indicates that many relationships are empirically quite stable, with greater variability for the effects of temperature moments than for precipitation moments. The conflict incidence chart shows that temperature shocks (light orange) have over the course of 35 years begun to significantly relate to conflict incidence, whereas the correlation of temperature deviations (yellow) with conflict incidence has changed sign. The relationship of temperature and precipitation with growth has been declining over time in most cases, which is roughly consistent with empirical findings, e.g. of (Miguel & Satyanath, 2011). Other trends can be interpreted into the data, but should be treated with caution.

The second panel shows that natural disasters and recurring events such as drought and El-Nino display much smaller and less volatile correlations with growth and conflict. The most interesting observation here is the vary large correlation between volcanic eruptions and conflict onset in the late 1970's. This is related to the events in eastern Congo at that time. In 1977 Mount Nyiragongo erupted and partly destroyed the city of Goma, and this coincided with (and potentially caused) the onset of conflicts between rebel groups in the area in 1975-1979. The correlations are quite robust to the aggregation choice, taking three-year intervals yields a very similar result. The sample of countries (approximately 30-40 in each case) might however not be representative since countries with less than optimal coverage are dropped in each case. Furthermore, country-year studies have been criticized for failing to address local heterogeneity in adaptation and responses. Overall however this test, despite being of value in its own right, suggests that there are not many cases of systematic adaptation so climate change aside from diminishing correlations of rainfall variables with growth.

Thus, to conclude the discussion on identification and prediction, the untransformed climate variables are fed into the learning algorithm. The resulting estimates of the predictive power of climate on growth and conflict thus constitute an upper bound, with bias stirring especially from potential geographic confounders.

[6]The Rebalancing and estimation is done in the following way: (1) For each pair of dependent and climate variable, drop all missing observations (2) Find the countries with the highest yearly coverage and drop countries with a yearly coverage less than that (3) Starting from the lowest available year, take pooled samples in 5-year intervals and obtain the correlation coefficient in each sample (4) plot the coefficients against the median year of the sample.

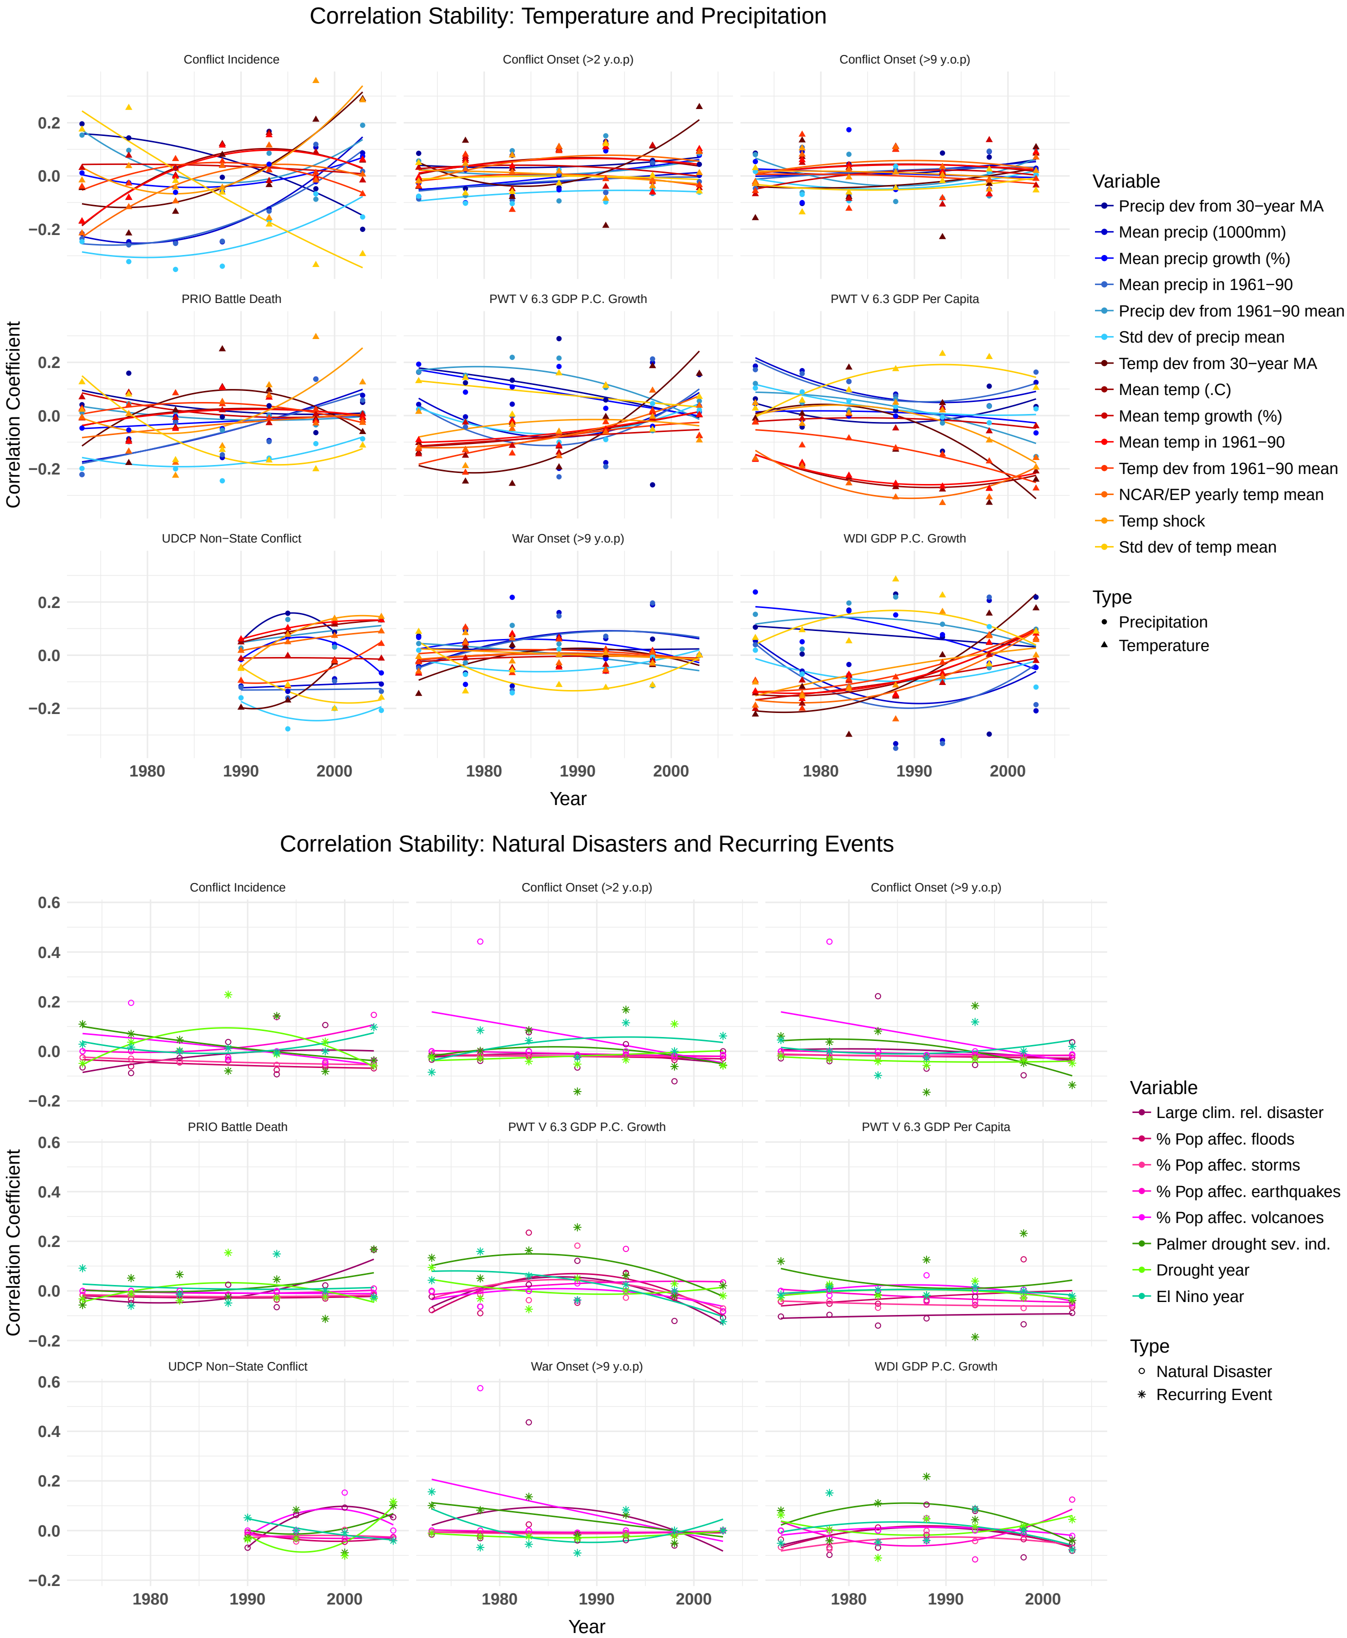

Figure 1: Correlation stability test of TS unit-homogeneity assumption

Variables: Temperature (Red), Precipitation (Blue), Natural Disasters (Pink) and Recurring Events (Green). 5-year intervals, quadratic fit. Panel re-balanced for each pair of dependent and climate variables

1.2 Model and Method

The most important choice in any machine learning task is the choice of algorithm and the way it is tuned to reach optimal performance. The preferred algorithm for this problem is the Random Forest (RF) algorithm (Breiman, 2001), which has proven itself capable of delivering extraordinary predictive performance in various settings, and which is also valued for its built-in variable importance selection measures. Additional benefits of the algorithm are that it works well for both continuous and categorical outcomes, and is considered very robust to overfitting through the built in cross-validation procedures. For the reader unacquainted with tree based supervised learning algorithms, a brief description addressing decision trees, Random Forests and the variable importance metrics of Random Forests is provided in the footnotes[7]. Furthermore, a short introduction to the algorithm is provided by Liaw & Wiener (2002), and a more general reference to methods of statistical learning is Friedman et al. (2001). A good reference regarding the application machine learning methods to economic problems is Varian (2014), and Bang et al. (2015) provide an example using a random forest algorithm to assess variable importance of predictors of growth. Random Forrest was however not blindly chosen, but tested in it's prediction of growth and conflict incidence against linear/logistic regression, classical decision/regression trees, support vector machine with linear or radial kernel, and against bayesian regularized and averaged neural networks. Random forest outperforms all of these other algorithms. To provide a feel for this, a simple logistic regression of conflict incidence on all 22 climate variables in levels and lags achieves a sensitivity of 24.1% (that is it rightly predicts 47/195 or 24.1% of conflict incidences based on climatic factors). The random forest on the other hand achieves a sensitivity of 59.5%

[7]The random forest is a tree-based ensemble machine learning algorithm based on Breiman, L. (2001). A decision tree is a type of supervised learning algorithm having a predefined target variable, that is mostly used in classification problems (categorical outcomes). It however also works well for continuous variables. The algorithm splits the sample into the most homogenous subsets (in terms of the values of the outcome variable) based on the most significant splitter among the input variables. For continuous outcomes this split is determined such that the average variance of the outcome in the two resulting subsets is minimized. The algorithm then predicts values by the mean in the resulting subsets, thus it is effectively minimizing the sum of squared residuals in an ANOVA type model. Procedurally this is executed as: 1. for each predictor variable x, determine the splitting point on the range of x that minimizes the outcome variance (calculated as the weighted average of the outcome in the two subsets resulting from the split). 2. Choose as the root node of the tree the variable whose split minimizes the outcome variance. 3. repeat this process in each of the resulting subsets, until either the whole outcome space is divided into homogenous subsets (e.g. no residual variance left), or until some stopping criterion (usually a restriction on the sample size of the terminal subsets or on the complexity of the tree) is reached. The advantage of tree based models is that thex perform better than linear models (e.g. regression), when there is high non-linearity and complex relationships between dependent and independent variables. The method is also fully non-parametric. A key challenge faced with trees is overfitting, that is a tree unrestricted in complexity will give a 100% accurate prediction on the dataset it was trained on but horrible predictions on a different dataset (it's like including a dummy for each observation in a linear model). Another problem with trees is that the algorithm is essentially forward looking, meaning that at every node it performs the optimal split on the optimal x variable to yield the most homogenous subsets in the resulting branches, but it does not consider that splitting on a less that optimal variable first, and then splitting on some other variable might result in a much cleaner outcome than optimizing throughout. Random Forrest is a method that effectively addresses both of these issues, and has produced superior predictions for both categorical and continuous outcomes. The secret of random forest is to average multiple trees grown with randomization on two accounts: I. random observations to grow each tree, II. A random subset of variables selected for splitting at each node. The algorithm is summarized as follows: 1. Draw a Bootstrap sample from the dataset (sampling with replacement, typically 1/3 of observations are left out, known as out of bag (OOB) samples), 2. Grow a tree on this sample while at each node selecting a subset of variables to split upon (typically $1/3 \times K$ (K = number of variables) for continuous outcomes and $\sqrt{K}$ for discrete outcomes). The tree is grown to full complexity 3. repeat this multiple times (typically 500 or more random trees are grown). 4. Predict new data by aggregating the predictions of all trees (i.e., majority votes for classification, average for regression). An unbiased estimate of the models accuracy (the error rate) is obtained as follows: 1. For each random tree, predict the data not in the bootstrap sample (the OOB data). 2. Aggregate the OOB predictions, and calculate the error rate. This double randomization leads to to key advantages of the random forest over many other predictive algorithms: 1. the methods of selecting random subsets and subsequent testing effectively guards the algorithm against overfitting the data. 2. Splitting on random subsets of variables decorrelates the resulting trees and effectively addresses the multicollinearity problem and the problems with the "greedy" optimizing nature of classical tree learners. A further feature of the algorithm is that is calculates measures of variable importance by looking at how much prediction error increases when (OOB) data for that variable is permuted while all others are left unchanged. The theory here is that if a predictor is important, then permuting that predictors values should significantly worsen the out-of bag predictions. The necessary calculations are carried out tree by tree as the random forest is constructed. For continuous outcomes this overall measure is then reported as a % increase in MSE when that variable is omitted from the model. A second metric is the mean increase in node impurity, where for each split involving a given variable, the decrease in residual sum of squares resulting from that split is calculated, and these decreases are then aggregated for that variable across all trees. As we are mainly concerned with MSE and also because not all variables produce an equal amount if splits, the first metric shall be of primary focus in this paper, although both are reported.

(while keeping a low specificity or false negative rate) and thus outperforms logistic regression more than twofold. Similarly, a linear regression of climate moments predicting GDP per capita growth achieves an R-squared of 4.2%, whereas the cross-validated random forest explains 6.9% of the variation in African growth rates. The results are presented in the figures below. The performance of the algorithm is indicated at the top of each figure (% variance explained for continuous outcomes and prediction accuracy together with false positive rate (sensitivity) and false negative rate (specificity) for binary outcomes), together with the number of trees used to build the random forest and the number of variables considered at each split. The latter was chosen to optimize the predictive performance of the algorithm. The Figure then shows the two metrics for variable importance described in the description of the algorithm in the footnotes.

Figure 2: WDI GDP per capita Growth

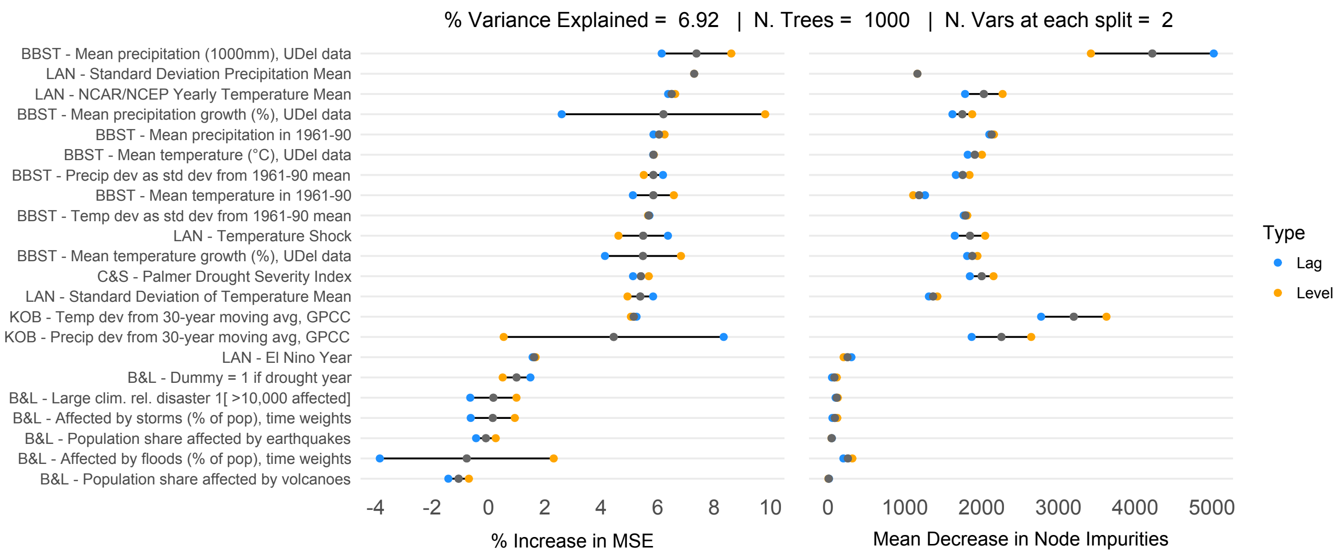

Figure 3: PWT Version 6.3 GDP per capita Growth

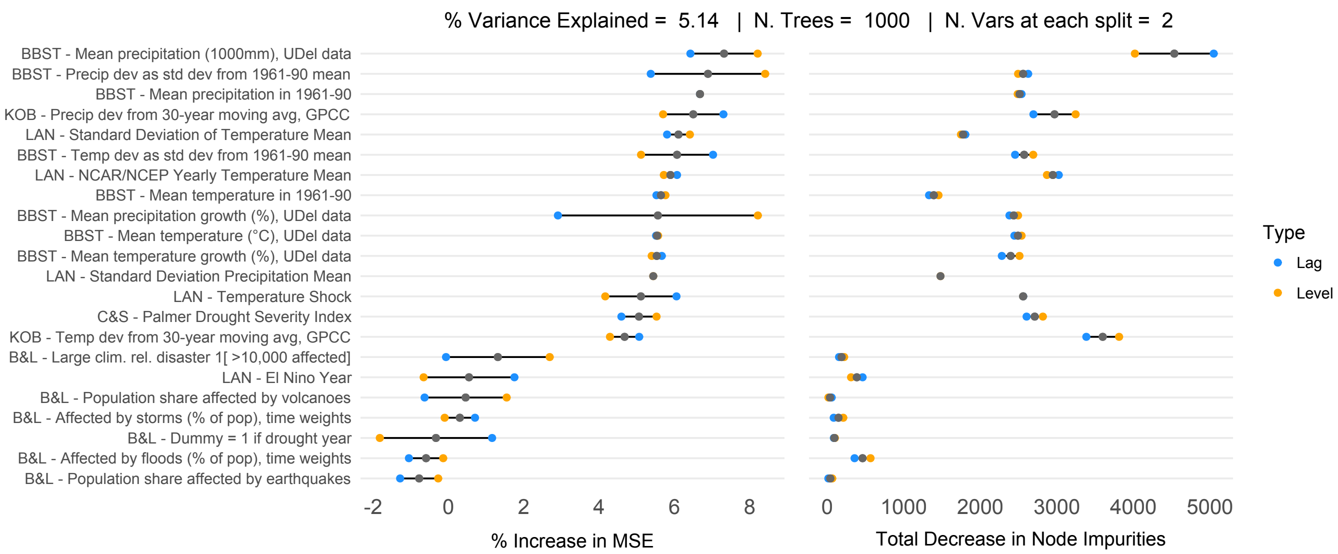

Figures 2 and 3 show that climate variables predict about 5-7% of the cross-country variation in GDP growth rates. The most important variables are mean precipitation and deviations from mean precipitation, closely followed by temperature moments. The ordering of predictive performance of climate variables is very similar in both cases, although the GDP growth measures are slightly different (corr .77). The effects in lags and levels are generally (with some exceptions) very close together. When double demeaned variables are used[8], the variance explained drops to -3.6%, suggesting that all of the important varyation is contained in the cross-sectional dimension.

Figure 4: PWT Version 6.3 GDP per capita

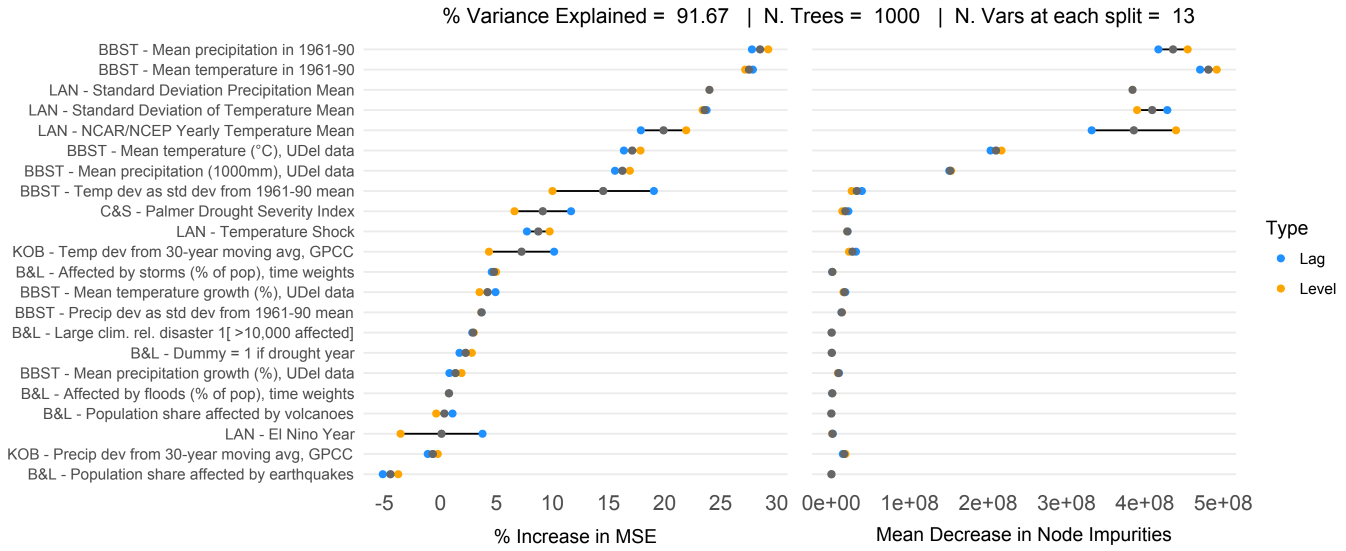

Figure 4 shows that the Random Forest is able to predict 91.7% of the variance in GDP per capita levels just based on climatic conditions. This is outstanding, but must be treated with extreme caution as this prediction is most likely to violate one of the unconfoundedness assumptions. Using double demeaned variables, the variance explained drops to 2.7%. This is, of course, a gross understatement of the impact of climatic conditions on levels of GDP.

Figure 5: Conflict Incidence

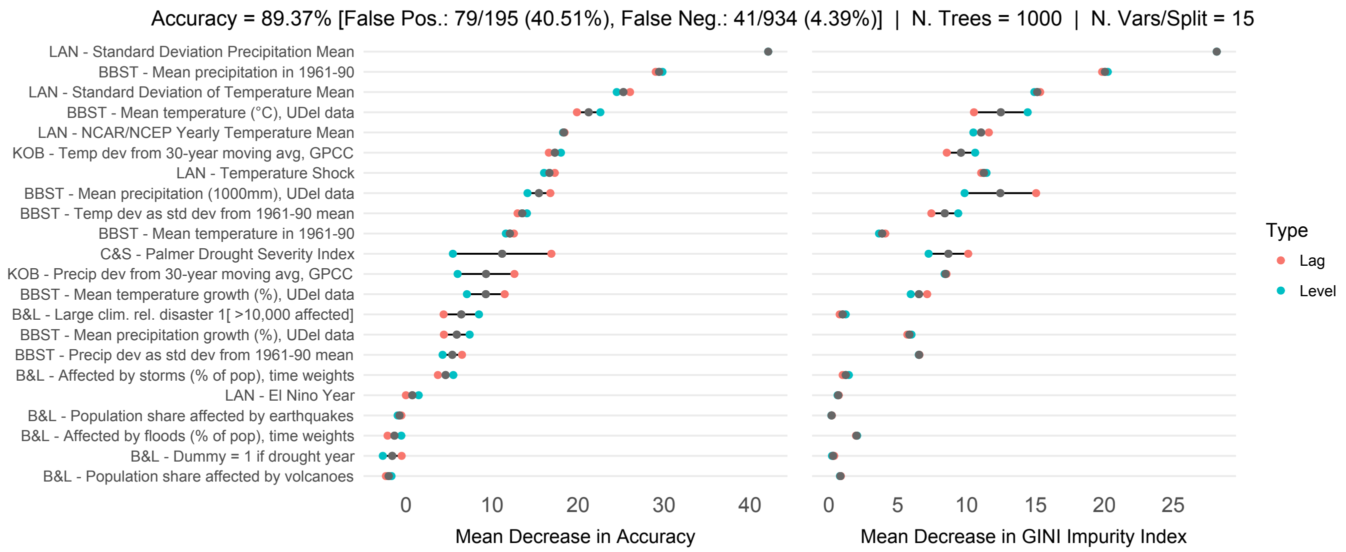

Figures 5 and 6 report the predictive power of climate on conflict incidence and non-state conflict onset at the country level. In both cases the algorithm rightly predicts about 60% of conflicts with a small false negative rate of around 4% in both cases. The standard deviation of year to year levels of precipitation is clearly the most important variable. Non-state conflict onset was a count variable and has been transformed into a dummy here. If the count variable is used and treated as continuous, the algorithm predicts 34.6% of the variance, with a very similar importance ranking. Using double demeaned climate predictors, the algorithm still rightly predicts

[8]Obtained by regressing a set of country and time dummies on each variable (dependent and independent) and taking the residuals. For binary outcomes (e.g. conflict incidence) however the dummy is maintained as it is.

27% of conflict incidences and 8% of non-state conflict onsets (sensitivity). The false negative rate (specificity) drops to 1% in both cases yielding still an overall predictive accuracy of around 85% in both cases. Not reported are the predictions for conflict onset (both with 2 year and 9 year of preceding peace thresholds) and war onset. The reason is that the algorithm was not able to predict any of these. There were 44 conflict onsets at the 2-year threshold, 32 and the 9-year threshold, and 13 war onsets over the whole period represented by >1000 observations in each case. These are very few for an algorithm to learn potential climatic features that make conflict onset more likely.

Figure 6: Non-State Conflict Onset

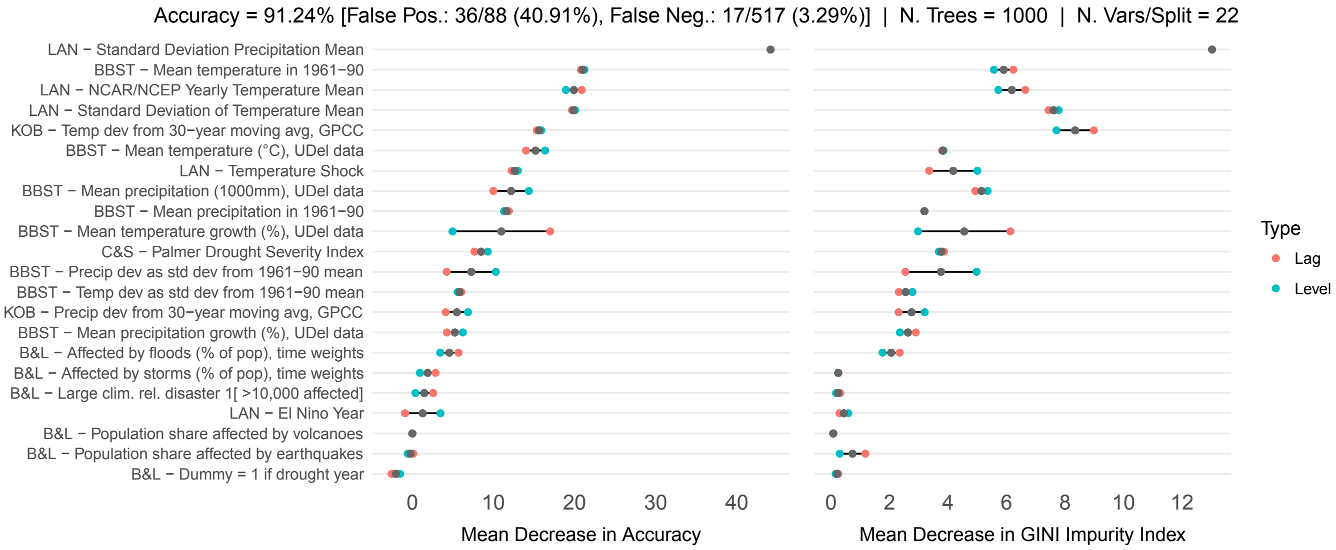

Figure 7 finally shows the results of predicting a continuous battle death estimate from the ACLED database. When double demeaned climate variables are used, the variance explained falls to 0.4%.

Figure 7: ACLED Battle Death Best Estimate

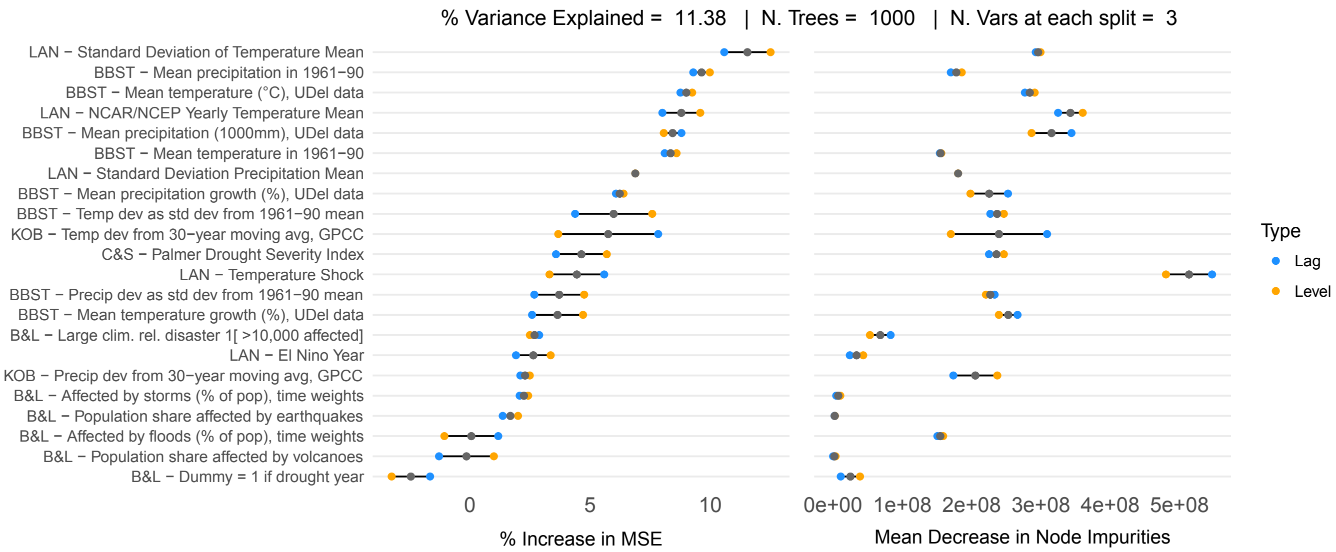

1.3 Gridded Data

The analysis with gridded data is conducted in analogous manner. This data is obtained from the PRIO grid v2.0 (Tollefsen et al., 2012), and from Berman & Couttenier (2015), and has a spatial resolution of 0.5° × 0.5°. Some variables in the grid, such as the temperature and precipitation measures, are available from 1946-2016, but since detailed conflict databases such as ACLED and UCDP/PRIO only record events from the 1980's onwards, only the period 1989-2013 is of relevance[9].

The availability of climate indicators is limited to a few sources, I take 4 drought variables that were obtained from the Standardized Precipitation and Evapotranspiration Index SPEI-1 from the SPEIbase v.2.3, based on precipitation and potential evapotranspiration data from the Climatic Research Unit of the University of East Anglia CRU v.3.22. *Droughtyr* gives the proportion of consecutive months out of 12 months where the cell experienced drought (defined as SPEI-1 values below -1.5). *Droughtstart* gives the SPEI-1 value during the first month of the cells rainy season. *Droughtend* gives the SPEI-3 drought severity value the for the last month of the cells rainy season. Further, I take as precipitation measure the yearly total amount of precipitation (in millimeter) in the cell, based on monthly meteorological statistics from the Global Precipitation Climatology Centre, and for temperature the yearly mean temperature (in degrees celsius) in the cell, based on monthly meteorological statistics from GHCN/CAMS, developed at the Climate Prediction Center, NOAA/National Weather Service. For these two variables I generate additional moments in the spirit of the cross-country variables, namely: (1) annual growth rates, (2) a 30 year moving average, (3) a 30 year moving standard deviation (4) the deviation in each year from the 30 year moving average and (5) the deviation from the 30 year moving standard deviation, defined as the absolute value when the absolute value of (4) is subtracted (3) (in short: (5) = abs[abs[(4)] - (3)]). These measures are jointly available for 10,260 grid-cells and 65 years, yielding an unbalanced panel of 544,716 observations. Table 6 shows some summary statistics.

Table 6: Spatial Climate Variables, $N = 544,716$, cells $= 10,260$, years $= 64$

#	*Variables*	*Mean*	*SD*	*Min*	*Max*	*Skew*	*Kurt*
	droughtyr_speibase	0.054	0.065	0	0.833	1.654	8.259
	droughtstart_speibase	-0.041	1.040	-5.877	7.314	0.185	2.617
	droughtend_speibase	-0.083	1.032	-5.564	4.369	0.088	2.642
	temp	24.108	4.011	7.509	40.224	-0.405	3.088
(1)	g_temp	0.124	3.090	-50.844	131.561	1.014	40.515
(2)	temp_30MA	23.879	3.905	10.223	38.893	-0.427	3.003
(3)	temp_30SD	0.611	0.388	0	7.478	5.085	44.200
(4)	temp_30MA_dev	0.229	0.774	-9.254	15.627	0.205	13.286
(5)	temp_30SD_dev	0.331	0.344	0	12.618	4.207	42.048
	prec_gpcc	158.703	156.722	0	1,478.142	1.018	3.736
(1)	g_prec_gpcc	67.314	4,447	-100	1,937,000	293.348	103,272
(2)	prec_gpcc_30MA	161.360	156.061	0.056	1,067.414	0.938	3.365
(3)	prec_gpcc_30SD	29.411	22.164	0.004	225.304	0.974	4.815
(4)	prec_gpcc_30MA_dev	-2.657	36.657	-486.116	634.246	0.117	10.293
(5)	prec_gpcc_30SD_dev	15.531	17.427	0	454.228	2.539	17.745

As dependent variables I take conflict incidence, the number of events and conflict onset from the UCDP database (from Berman & Couttenier (2015)), and conflict incidence from the ACLED database (also from Berman & Couttenier (2015)). For growth I take the 5-year growth rate of the gross-cell product per capita PPP (GCP), based on the G-Econ dataset v4.0, aggregated at the 1° level and only available for the years 1995, 2000 and 2005. I also consider the growth rate of a calibrated nightlights variable, available from 1992-2012, which measures average measured nighttime light emission from the DMSP-OLS Nighttime Lights Time Series Version 4[10]. These

[9] The grid is explained, and can further be visualized and downloaded at http://grid.prio.org/#/.

[10] Average Visible, Stable Lights, & Cloud Free Coverages.

data are calibrated for use in time series analyses using calibration values from Elvidge et al. (2014), and standardized to be between 0 and 1. These variables are summarized in Table 7.

Table 7: Spatial Dependent Variables

Variables	*N*	*Mean*	*SD*	*Min*	*Max*	*Skew*	*Kurt*
Conflict incidence, UCDP-GED	147,858	0.029	0.167	0	1	5.629	32.689
# Conflict events, UCDP-GED	141,864	0.111	1.486	0	126	36.699	1,880
Conflict onset, UCDP-GED	146,169	0.016	0.125	0	1	7.763	61.265
Conflict incidence, ACLED 1997-2010	82,125	0.042	0.201	0	1	4.563	21.821
Gross Cell Product per Capita, PPP$ USD	38,539	0.144	0.671	0	21.260	15.241	329.502
Gross Cell Product per Capita, PPP$ USD growth	27,671	20.530	25.799	-68.692	482.863	3.025	38.414
Nightlights calibrated, mean	214,748	0.040	0.034	0	0.957	8.508	129.119
Nightlights calibrated, mean, growth	196,747	45.975	770.171	-69.562	75,803	49.104	3,243
Nightlights, standard deviation	222,206	0.721	2.173	0	26.118	5.164	35.942

Pairwise correlations reported in Table 8 show that all conflict variables are positively correlated, and slightly negatively correlated with economic activity growth, measured either through GCP or nightlights. It is curious to observe that the growth rates of GCP and nightlights are negatively correlated, with $\rho = -0.392$.

Table 8: Correlation Matrix of Spatial Dependent Variables
(*Pairwise Correlations*)

# *Variable*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) Conflict incidence, UCDP-GED	1								
(2) # Conflict events, UCDP-GED	.425	1							
(3) Conflict onset, UCDP-GED	.946	.472	1						
(4) Conflict incidence, ACLED 1997-2010	.405	.270	.291	1					
(5) Gross Cell Product per Capita, PPP$ USD	.061	.127	.038	.058	1				
(6) Gross Cell Product per Capita, PPP$ USD growth	-.019	-.018	-.002	-.016	.010	1			
(7) Nightlights calibrated, mean	.078	.101	.045	.107	.669	-.005	1		
(8) Nightlights calibrated, mean, growth	-.002	-.001	-.001	.001	-.055	-.392	-.033	1	
(9) Nightlights, standard deviation	.090	.097	.056	.136	.531	-.015	.749	-.011	1

Figure 8 shows again the correlational trends test for the time series homogeneity assumption. Due to the high sample size (about 10.000 for each year in most cases), the correlations are calculated for each year in the panel, and a lowess smoother with span 1.5 is fit to the data. The correlations are smoother and trends more clearly visible than in Figure 1. Overall, Figure 8 suggest quite stable empirical relationships between climate moments and economic outcomes, and also shows an aggregate decline in importance of precipitatio, as well as a simultaneous gradual increase in the importance of temperature moments, that was already observed in the cross-country panel data. Particularity remarkable is the slow increase in the 80's/90's, and subsequent decrease in the 2000's, of the correlation of precipitation moments and conflict incidence. A decline in the importance of precipitation is also visible with GCP growth, but not with nightlights growth.

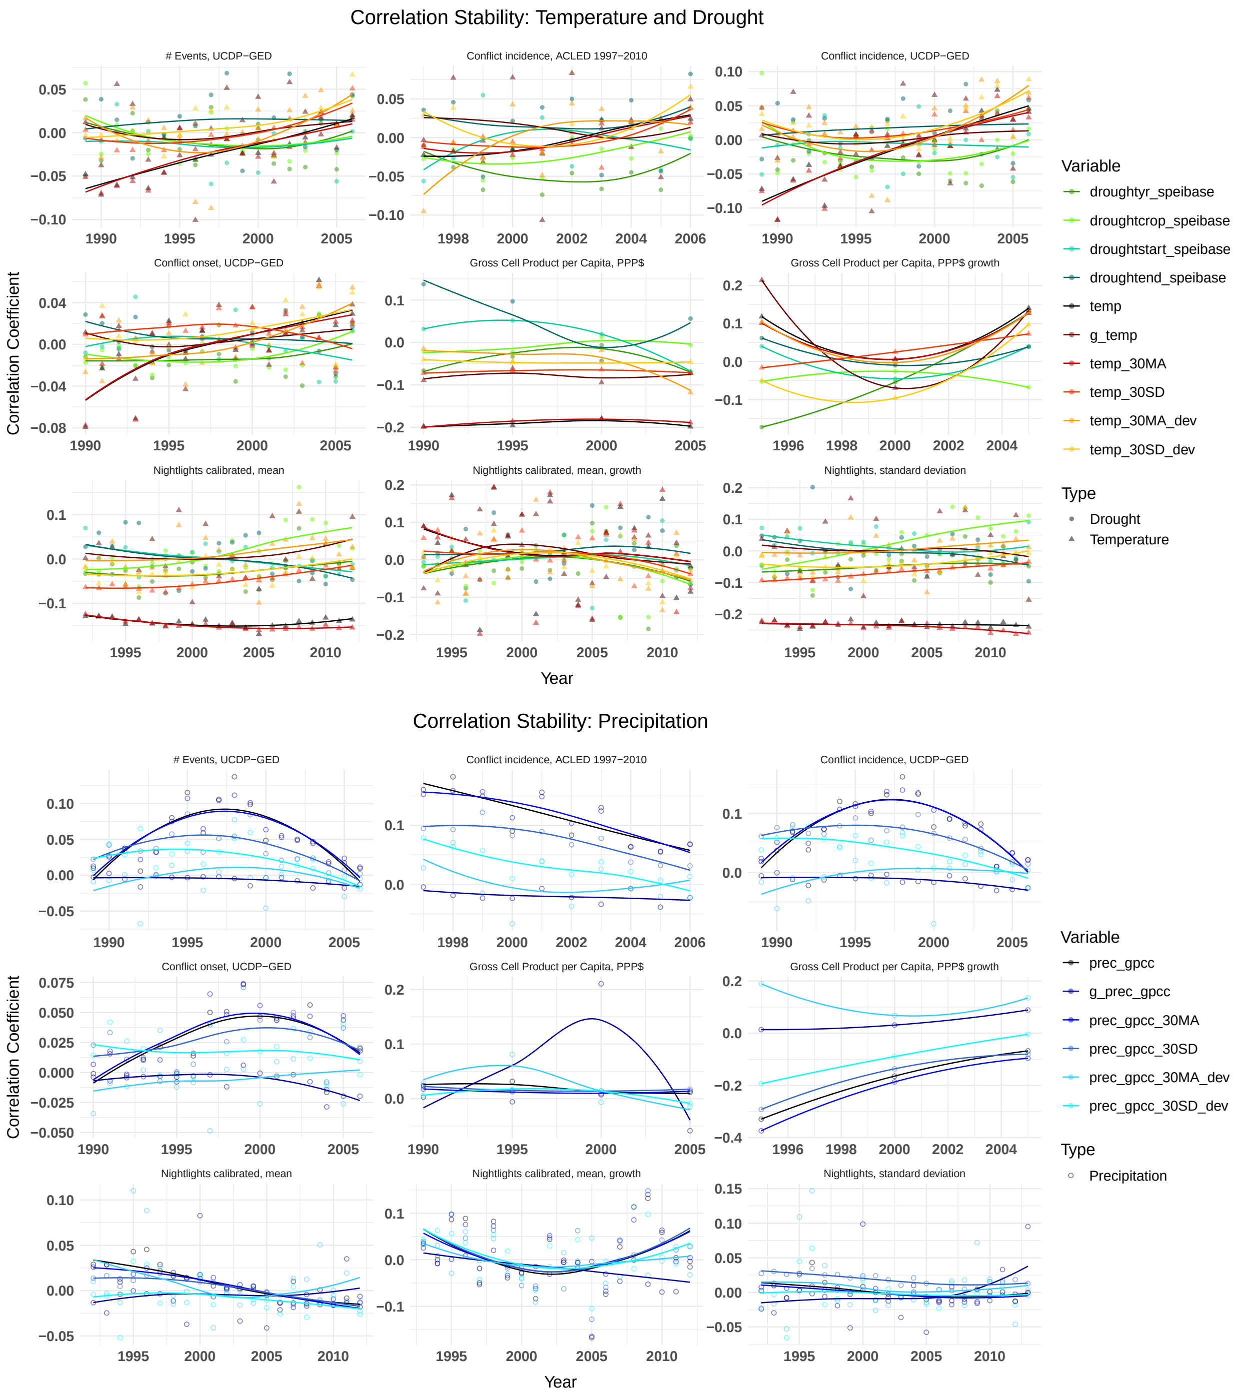

Figure 8: Correlation stability test of TS unit-homogeneity assumption, PRIO

Variables: Temperature (Red), Drought (Green) and Precipitation (Blue). 1-year intervals, lowess fit (span 1.5). Panel re-balanced for each pair of dependent and climate variables

Figure 9 shows the results for UCDP conflict incidence. The algorithm correctly predicted 224 out of 4040 incident cell-years, and wrongly predicted 66 out of 143,509 non-incident cells, yielding a sensitivity of 5.25% and an overall accuracy of 97.22%. the most important variables are the 30-year MA's of temperature and precipitation, followed by their levels and 30-year moving SD's. The same holds for the number of conflict events in each cell, where the algorithm predicts 13.3% of the variance. Nearly identical results are also achieved with the alternative incidence measure from ACLED reported in Figure 10.

Figure 9: Conflict Incidence and Number of Conflict Events (UCDP)

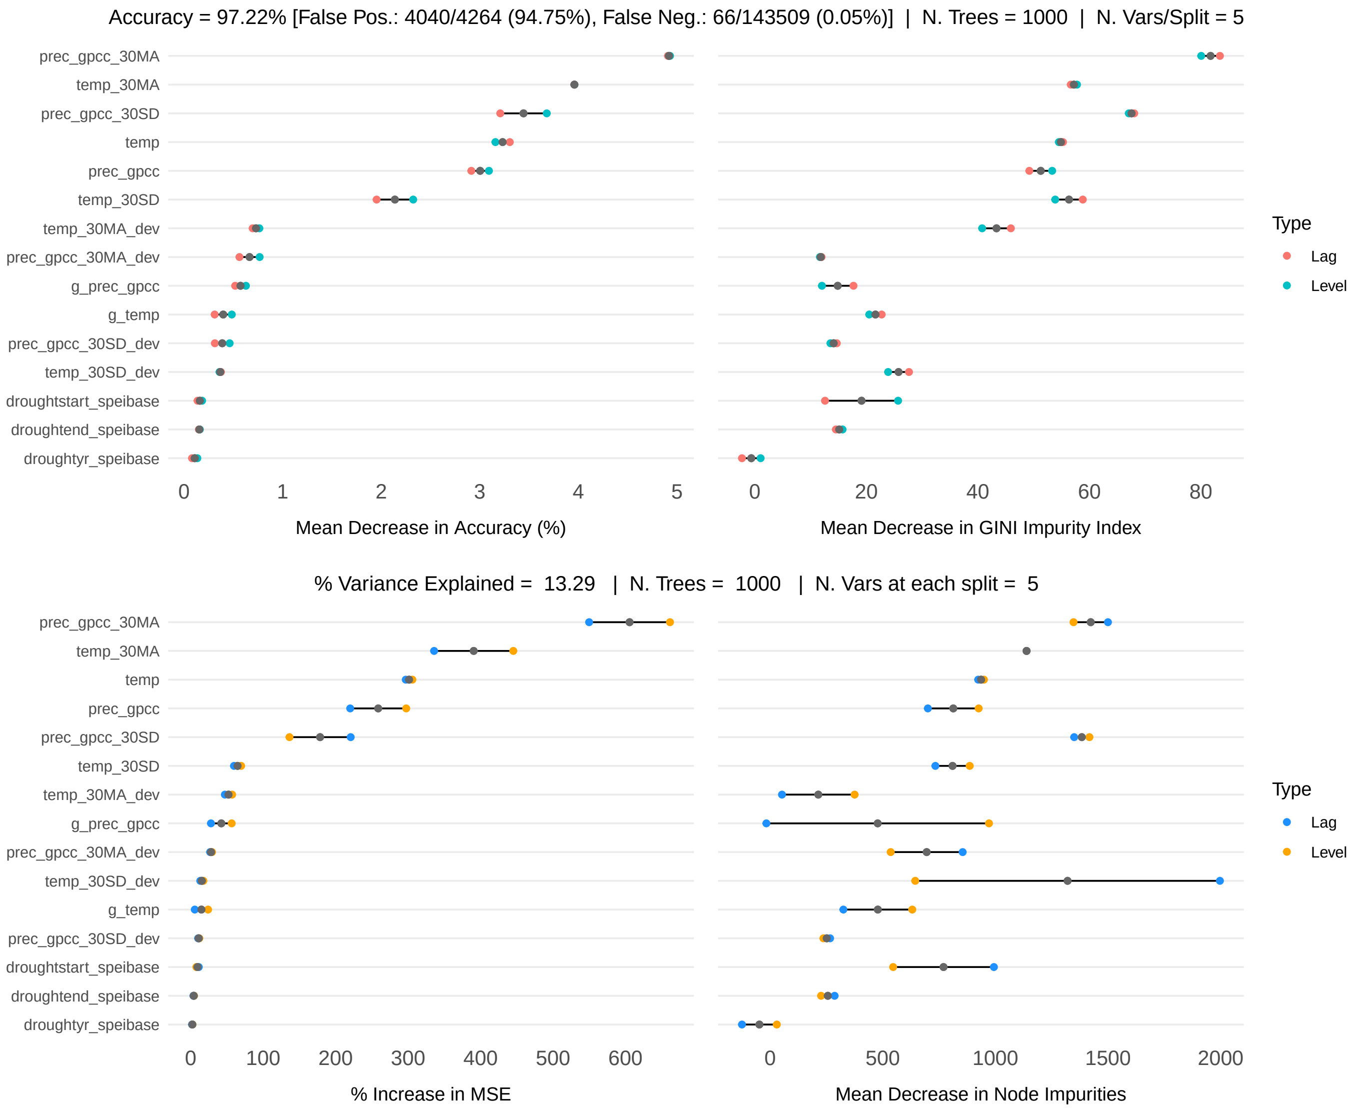

Figure 10: Conflict Incidence (ACLED)

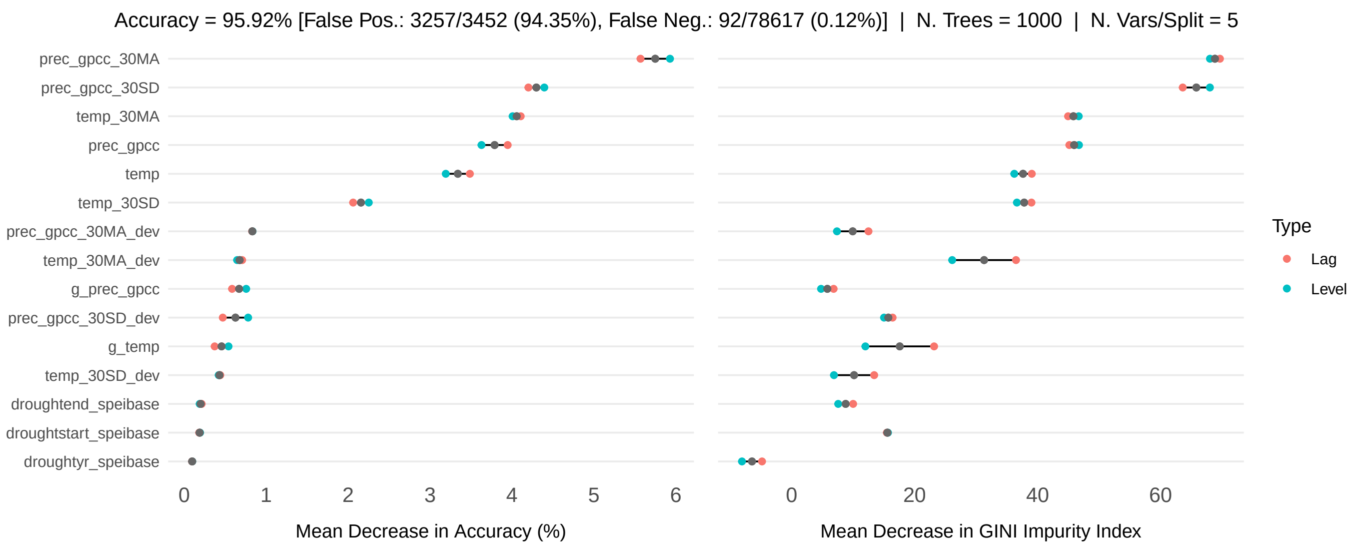

When instead considering UCDP conflict onset, reported in Figure 11, the predictive performance of the algorithm drops considerably, with only 9 cell-level conflict onsets predicted correctly, yielding a sensitivity of 0.4%, but at a specificity of 99.99%, thus the model is still of value. It is noteworthy that also here the moving averages and standard deviations of precipitation and temperature are the most important variables, indicating that weather shocks (as captured by growth rates or deviations from moving quantities, or drought events) are not helpful towards predicting either conflict incidence or conflict onset at the cell-level.

Figure 11: Conflict Onset (UCDP)

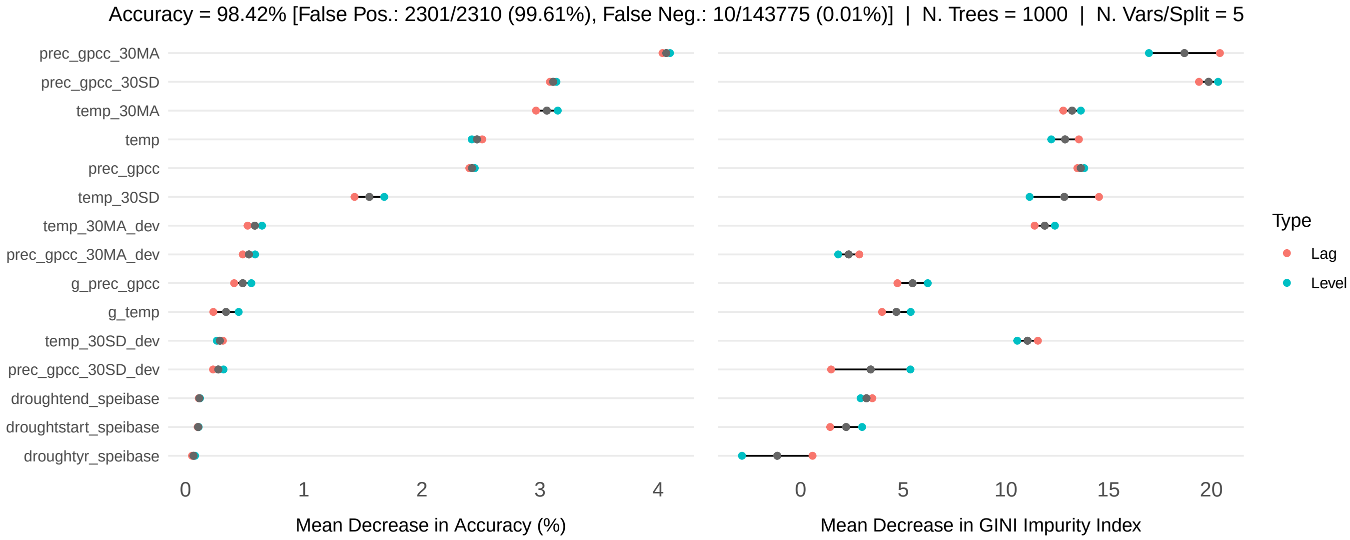

Turning to economic growth, Figure 12 shows that the RF predicted almost 70% of the variation of GCP growth in Africa, with long-term measures of precipitation the most important variables, followed by temperature shocks. Figure 13 shows the results with growth of nightlights. Here the predictive power dropped significantly to 4.5%, which can be attributed to nightlights being a coarser measure of spatial activity, and also available at higher spatial and temporal resolutions than GCP. When considering the level of nightlights, reported in Figure 14, the predictive power increases considerably to 52.3%, with temperature and precipitation 30-year MA's being again the most important variables.

Figure 12: Gross Cell Product per Capita, PPP$ growth (5-year, 1995, 2000, 2005)

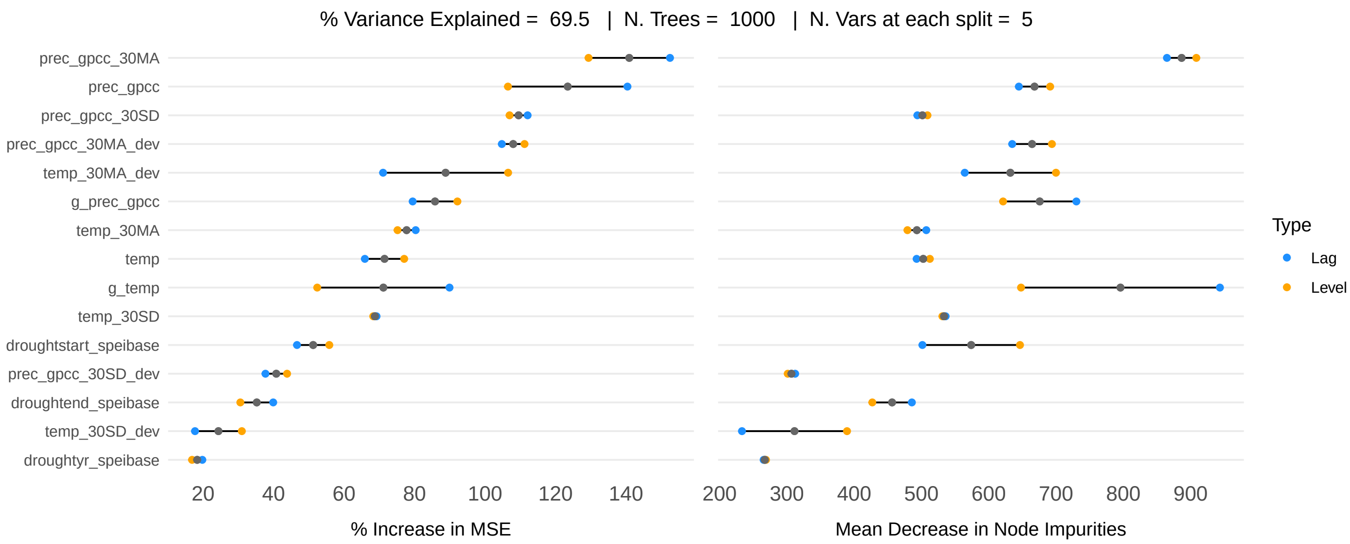

Figure 13: Nightlights Calibrated Mean, Growth

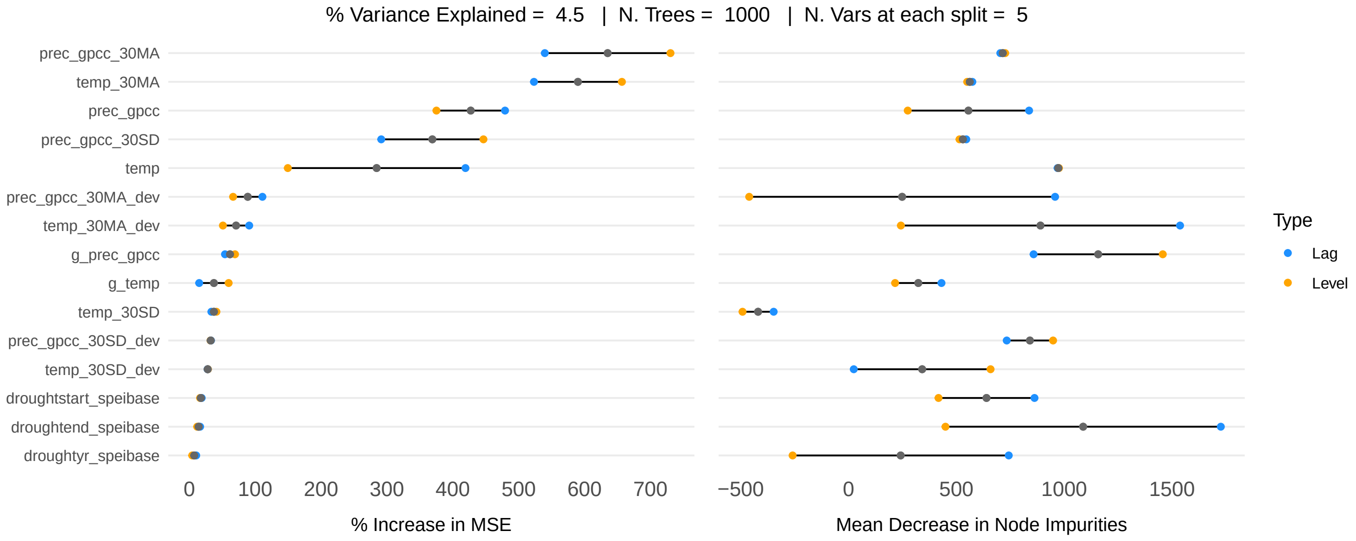

Figure 14: Nightlights Calibrated Mean

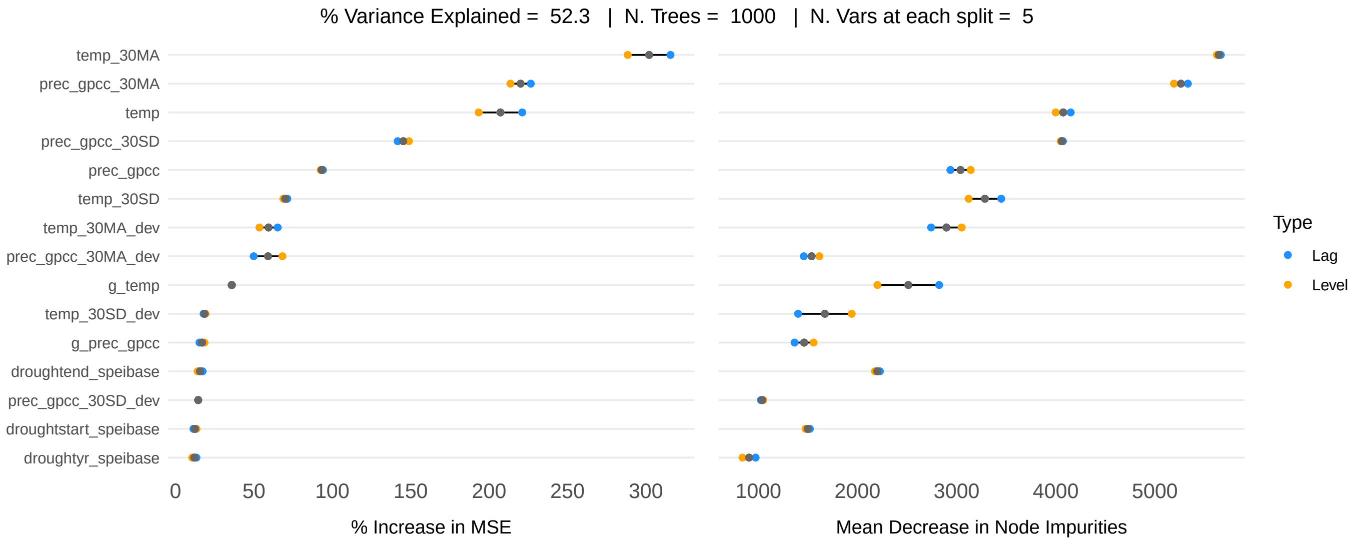

Compared to the cross-country results, the changes in predictive performance appear intuitively explicable: whereas disaggregation to the cell level allows for within-country heterogeneity in climate and weather condition, and thus enables the algorithm to better distinguish urban growth centers from lone deserts or jungles based on climatic characteristics, the disaggregation to the cell level makes pinpointing conflict incidences or onsets more difficult.

The PRIO grid and the data of Berman & Couttenier (2015) include some valuable additional natural and social characteristics of grid-cells. To gauge the strength of the unconfoundedness assumption for climate, I choose approximately 40 of these other variables, including cell distance to capital, border and nearest urban center, mountainous terrain, etc. and various political and social characteristics (none of which too closely resemble either conflict or growth, but many are correlated with climate) and then fit the RF model using only these variables, and using all variables, including the climate predictors. The results for UCDP conflict onset and GCP growth are provided in the Appendix in Figures 15, 16, 17 and 18.

Comparing the predictions from climate variables with those from other variables, it is apparent that predictions based on other variables than just climate are better. The best predictors are variables measuring the position and remoteness of the cell (namely: Distance to capital, distance to major port, distance to border, travel time from cell centroid to next urban center and distance to closest natural resource). When all other variables are present, adding the climate variables only increase predictive accuracy by a few percentage points. For conflict incidence, the random forest involving all variables actually performs slightly worse than the forest using only other predictors. This is due to the fact that the forest is built using only 250 trees, which is very low (normally for 60 predictors one would build a forest with 2000 trees, but that would take my computer a few hours to process). Thus in any case these results should be treated with caution.

2 Reverse Causality

A second issue that is of high theoretical relevance but hardly addressed by the climate-economy-conflict literature is reverse causality. Most of the literature has focused on identifying the channel from economic conditions to conflict (e.g. Berman & Couttenier (2015)), or the extended channel from climate through economic conditions to conflict (e.g. Miguel et al. (2004); Miguel & Satyanath (2011)). The main theoretical underpinnings for the effect of economic conditions on conflict are the "opportunity cost" and the "state capacity" mechanism. The former, which traces back to Collier and Hoeffler (1998, 2001, 2002), holds that when gains from conventional economic activity are low, the opportunity cost for young men to engage in illicit economic activities / join rebel groups is lower (Miguel et al., 2004). The latter describes the idea that low national income or adverse economic shocks curtail a governments ability to invest in the military, infrastructure, public administration, schools etc., and thus its capacity to maintain law and order and effective control over its territory, and to provide its citizens with adequate economic opportunities. Next to these two main mechanisms, other sometimes competing mechanisms have been thought, for example the "state as prize" mechanism discussed in Berman & Couttenier (2015), which stipulates that larger income increases the likelihood of conflict by increasing the value of the state that can be captured through rebellion. These mechanisms are marginalized though and Berman & Couttenier (2015) reject the state as prize mechanism in favor of the opportunity cost mechanism.

Most studies acknowledge that economic conditions are endogenous to conflict, which warrants the use of instrumental variables such as rainfall in the literature. Likewise most studies mention that conflict affects economic conditions in important ways, but these channels are often neither theoretically explored or estimated. This systematic omission can entail problems of its own, namely (1) if the magnitude of reverse causality between conflict and economic conditions and the channels through which conflict impacts economic conditions are not well understood, then the exclusion restriction in many IV papers estimating the effects of economic conditions on conflict may be violated, and (2) abstracting from reverse causality does not allow for virtuous or vicious cycles, and empirical results are thus likely to understate the true general-equilibrium effect of economic shocks on conflict. The absence of general equilibrium models in this literature can be explained by the inherent difficulties in causal identification. I will in the following attempt to contribute in this field by estimating a simple 2-way causal system between economic growth and conflict using the two datasets constructed in the previous section, after briefly discussing some of the causal mechanisms by which conflict affects economic conditions, which for the most part are unambiguous.

A primary impact of conflict on economic activity is that conflict displaces peoples, preventing them to engage in economic activities. At the same time conflict often destroys productive capital, such as production sites, supporting infrastructure such as electric grids, as well as state/administrative structures and social services needed to support complex economic activity. While these effects immediately hamper economic activity in the short- and medium-term, in the long-term, conflict also displaces/relocates industry and trade networks, mindering the overall economic structure and stability of regions. There is ample evidence that conflict in Africa negatively affected economic activity in several African regions in all 3 ways just outlined. I refer the reader to Cooper (2019) for discussions of post-colonial conflicts in Africa.

2.1 Estimation using Country-Year Data

I start my analysis with a simple set of variables taken from the cross-country dataset, summarised in Table 9. From Miguel & Satyanath (2011) I take conflict incidence and GDP growth as the jointly endogenous dependent variables, and from Buhaug et al. (2015) I take a revolutions dummy to instrument conflict incidence, and data on food and crop production to instrument growth. The idea behind using revolutions is that often revolutions are a consequence of political events and thus in their occurrence, timing and intensity not linked to specific economic conditions. Food and crop production are justified by being important determinants of the level and composition of income in African economies, and themselves unlikely to be much affected by conflict or to affect conflict via channels other than income. These instruments are however necessarily imperfect, revolutions could be triggered by economic conditions, and likewise it could be that food and crop production impact conflict other than through GDP or are themselves impacted by conflict. It has been a characteristic of this literature that, in part because of the theoretical ambiguity regarding causal mechanisms, strong instruments with justifiable exclusion restrictions are extremely hard to come by. I do therefore not argue that these instruments are exogenous and valid in all settings and at all time horizons, but rather, that they are useful in limiting short-term simultaneity between conflict and economic downturn, and thus provide a better identification of the growth $\leftrightarrow$ conflict nexus than a simple OLS approach could. To strengthen this short-term exclusion restriction, some specifications are also estimated with additional lags of the instruments.

Table 9: Summary Statistics

Variables	*N*	*Mean*	*SD*	*Min*	*Max*	*Skew*	*Kurt*
MIG'15 - UCDP/PRIO v4-2010, >25 deaths	1,179	0.26	0.44	0	1	1.08	2.17
MIG'15 - Per cap GDP growth from PWT 7.0	1,189	0.73	9.21	-62.1	88.7	1.78	25.2
BBST - Revolutions	2,221	0.20	0.40	0	1	1.47	3.15
BBST - % growth in food prod pc since t-1	2,101	0.30	9.72	-50.5	68.3	0.84	9.89
BBST - Crops production per capita t-1	2,101	0.095	0.049	0.0029	0.31	0.78	4.34

Table 10 shows the corresponding correlation matrix. Revolutions load strongly onto conflict incidence but only weakly (and negatively) onto growth. Food production shows a similar pattern, and crop production appears to be weak at first sight, but will turn out to be stronger in combination with food production. I estimate this system using 2SLS and 3SLS, all specifications include country fixed effects and country specific time trends (following Miguel et al. (2004); Miguel & Satyanath (2011)). Table 11 shows the results, which include the first stages (columns (1) and (2)), the reduced forms (columns (3) and (4)), the system estimated by 2SLS (columns (5) and (6)), the system estimated by 2SLS with the instruments included in levels and lagged once (columns (7) and (8)), the system estimated by 3SLS (columns (9) and (10)), and the system estimated by 3SLS with instruments included in levels and lags (columns (11) and (12))[11]. A lag was added to the instruments in (7), (8), (11) and (12), in the hope that this would yield stronger first stages / enhance identification of the system.

Table 10: Correlation Matrix
(*Pairwise Correlations*)

# *Variable*	(1)	(2)	(3)	(4)	(5)
(1) MIG'15 - UCDP/PRIO v4-2010, >25 deaths	1				
(2) MIG'15 - Per cap GDP growth from PWT 7.0	-.05	1			
(3) BBST - Revolutions	.46	-.08	1		
(4) BBST - % growth in food prod pc since t-1	-.02	.22	-.04	1	
(5) BBST - Crops production per capita t-1	-.06	.01	-.02	-.09	1

Sadly it only enhances the identification statistics for conflict. Table 11 shows that conflict incidence instrumented by revolutions has a significant negative effect on growth. The effect is stable across specifications and very large (a conflict decreases the annual GDP growth rate by 8% in this specification, which is mainly due to the fact that revolutions are used as instrument).

[11] In columns (11) and (12), the instruments in levels and 1 lag are not reported for lack of space.

Unfortunately, the effect of growth on conflict is not significant in any of the specifications, which, comparing the first-stage R-squared's, is due to weak identification. Comparing the first stage to the reduced form, it is evident that food and crop production explain more of the variation in conflict incidence than in GDP growth, which is peculiar. A failure to properly identify the system here is evident. When one believes that revolutions are exogenous to economic conditions, Table 11 at least makes a strong case for a sizeable negative impact of large conflicts on economic conditions, which is what one would expect to find. Regarding the channel from growth to conflict, this is the one the literature has been researching extensively. I refer the reader to Miguel et al. (2004) and others for some significant though contested findings.

Table 11: Cross-Country Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	First Stages		Reduced Forms		2SLS		2SLS + Lag. IV		3SLS		3SLS + Lag. IV	
Variables	CINC	%GDP	CINC	%GDP	CINC	%GDP	CINC	%GDP	CINC	%GDP	CINC	%GDP
BBST - Revolutions	0.228***	-1.670**		-2.136**					0.225***			
	(0.041)	(0.785)		(0.854)					(0.027)			
BBST - % growth in food prod pc since t-1	0.000	0.239***	-0.000							0.241***		
	(0.002)	(0.074)	(0.001)							(0.032)		
BBST - Crops production per capita t-1	1.426	60.789**	0.278							65.675**		
	(1.611)	(28.559)	(1.658)							(28.175)		
MIG'15 - Per cap GDP growth from PWT 7.0					-0.001		-0.002		0.001		-0.001	
					(0.005)		(0.005)		(0.004)		(0.004)	
MIG'15 - UCDP/PRIO v4-2010, >25 deaths						-9.277**		-8.905***		-7.398**		-7.078**
						(3.845)		(3.273)		(3.446)		(3.162)
Observations	1,096	1,096	1,179	1,106	1,179	1,096	1,178	1,095	1,096	1,096	1,094	1094
R^2	.681	.146	.655	.101	.005	-.030	.007	-.025	.655	.129	.572	.143
Kleibergen-Paap LM stat. (H0: Underid.)					6.779	13.50	9.741	13.48				
Kleibergen-Paap P-Value					0.033	0.0002	0.045	0.001				
Hansen J Overid. test (H0: Eq. Identified)					0.046		2.348	0.047				
Hansen J P-Value					0.830		0.503	0.828				
Endogeneity of Regressor (H0: Exogenous)					0.230	2.495	0.083	3.077				
Endog. P-Value					0.631	0.114	0.773	0.079				

Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1

NOTE: Country FE and Country specific time-trends included in all specifications. "2SLS + Lag. IV" indicates that the instruments were included in levels and lags, to increase the first stage r-squared. In all specifications except for 3SLS, the error matrix is cluster-robust at the country level. 3SLS estimates are obtained with iterative convergence (the 'ireg3' option in STATA).

2.2 Estimation using Gridded Data

The second approach I follow in order to identify the climate-conflict causal system builds on the paper of Berman & Couttenier (2015), who use gridded data (from the PRIO grid and other sources) to identify the effects of external shocks (agricultural commodity price shocks and financial crisis in neighbouring countries) on conflict outbreak in Sub-Saharan Africa. Their main dependent variable is UCDP conflict incidence, which was already used in the predictions in section 1.2. They construct an agricultural commodity shock by weighting the world import value of the pool of agricultural commodities produced in a specific cell at a specific time, by the share in production these commodities have in the cells agricultural output. They construct the financial crisis shock by weighting a banking crisis dummy by the average share of the countries(s) with crisis in the home countries total exports. Berman & Couttenier (2015) regress conflict on these two shocks (sometimes interacting the shock variables with distance to port and other variables), and find a significant effects of both shocks on conflict outbreak. In this they are estimating an equation similar to the reduced form for conflict incidence in Table 11, and arguably more successful in doing so. I will, in this part of the paper, extend their analysis a bit by using their external shock measures as instruments for GDP per capita growth on conflict, and construct my own two instruments based on external political shocks, to instrument for conflict in the regression on growth. My instruments are (1) based on irregular political activity that could lead to conflict, as captured by revolutions, government crisis, major constitutional changes and coups d'etat, and (2) electoral activity, which is based on national elections of legislative and executive and presidential elections. The basic idea is that political tensions producing conflict in neighbouring countries can spill over to cells in the border area of a country, especially if the population in both countries are ethnically or culturally related to one another.

I construct two sets of two instruments. The first set is just the mean of revolutions and executive elections dummies taken over all neighbouring countries (with a common land border with the country at issue), weighted by the distance of the cell to the border. The second set of interest is constructed analogously, but using (1) the sum of the neighbouring-countries mean of revolutions, government crisis, major constitutional changes and the number of coups d'etat, weighted by cell distance to border (I call this instrument "weighted irregular political activity"), and (2) the sum of the neighbouring-countries mean of national elections of legislative and executive and presidential elections weighted by cell distance to border (I call this instrument "major elections"). At times the instruments are also interacted with the mean share of neighbouring countries where at least 9% of the population speak the same language, in order to further specify under which conditions such 'spillovers' of political conflict are likely.

Table 12 shows a summary of all variables used in these estimations, and Table 13 shows the corresponding correlation matrix.

Table 12: Summary Statistics

Variables	*N*	*Mean*	*SD*	*Min*	*Max*	*Skew*	*Kurt*
Conflict incidence, UCDP-GED	150,804	0.029	0.17	0	1	5.61	32.4
Binary, 1 for event, ACLED 1997-2010	83,670	0.042	0.20	0	1	4.54	21.6
Gross Cell Product per Capita, PPP$ growth	28,785	20.6	26.1	-68.7	483	3.29	41.4
Nightlights calibrated, mean, growth	205,102	46.3	775	-84.6	75,803	48.7	3,187
ln agr. com. shock	130,500	10.0	0.93	-6.05	12.0	-3.00	30.4
ln agr. shock × ln dist. to closest port	130,500	64.0	9.75	-42.9	85.2	-1.28	7.26
Exposure to crises	223,263	0.14	0.17	0	0.88	1.86	6.07
Exp. to crises × ln dist. to closest port	223,263	0.89	1.15	0	6.13	1.92	6.26
Weighted Revolutions	279,038	0.070	0.092	0	3	6.56	90.9
Weighted Nat. Elec. for Executive	317,116	0.17	0.16	0	2	6.46	62.2
Irregular Political Activity, weighted sum	259,794	0.13	0.17	0	8	7.83	136
Major Elections, weighted sum	259,785	0.23	0.23	0	6	6.94	78.9
Common language spoken by 9% of pop	335,950	0.50	0.32	0	1	-0.20	2.02
Log distance from cell centroid to int. border	362,655	4.68	1.15	0.50	7.59	-1.07	4.20

Table 13: Correlation Matrix
(*Pairwise Correlations*)

# *Variable*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
(1) Conflict incidence, UCDP-GED	1												
(2) Binary, 1 for event, ACLED 1997-2010	.41	1											
(3) Gross Cell Product per Capita, PPP$ growth	-.02	-.02	1										
(4) Nightlights calibrated, mean, growth	0	0	-.38	1									
(5) ln agr. com. shock	-.01	-.01	.03	-.02	1								
(6) ln agr. shock × ln dist. to closest port	-.07	-.06	0	-.02	.54	1							
(7) Exposure to crises	0	0	-.20	.03	-.16	-.04	1						
(8) Exp. to crises × ln dist. to closest port	-.01	0	-.19	.02	-.17	.03	.99	1					
(9) Weighted Revolutions	.04	.06	.04	0	-.03	.02	.03	.04	1				
(10) Weighted Nat. Elec. for Executive	.04	.05	.02	0	.01	0	-.05	-.06	.44	1			
(11) Irregular Political Activity, weighted sum	.04	.07	.01	.01	-.09	-.03	.08	.08	.82	.46	1		
(12) Major Elections, weighted sum	.03	.05	.02	-.01	.01	0	-.04	-.04	.44	.92	.44	1	
(13) Common language spoken by 9% of pop	-.04	-.03	-.02	0	-.14	-.13	-.08	-.08	0	-.11	0	-.10	1
(14) Log distance from cell centroid to int. border	-.03	-.05	-.01	0	.06	.04	.01	.01	-.39	-.62	-.39	-.58	.03

The estimations are conducted analogous to the previous subsection, but I refrain from using 3SLS (because of problems with fixed effects in the data structure), and report next to first stages, reduced forms and simple 2SLS, specifications where the instruments are interacted as previously described, where I use ACLED and nightlights data instead of UCDP and G-econ, and a specification where I add all of the climate moments shown in Figure 8 to both equations, to partial out the effect of climate. Table 14 shows the results using the weighted revolutions and weighted executive election instruments, and Table 14 shows the results using the weighted irregular political activity and weighted major elections instruments. The results are broadly similar to the ones obtained in Table 11 (one must consider the 5-year growth rates with the G-Econ data, thus the coefficients on conflict incidence are considerably larger).

Table 14: PRIO Gridded Data Regressions Using Weighted Revolutions and Weighted Executive Election Instruments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	First Stages		Reduced Forms		2SLS		2SLS + Int. IV		ACLED & Nightl.		2SLS + Clim.	
Variables	%GDP	CINC	%GDP	CINC	%GDP	CINC	%GDP	CINC	%NL	CINC	%GDP	CINC
ln agr. com. shock	-1.481	-0.072***		-0.068***								
	(9.374)	(0.020)		(0.018)								
Exposure to crises	27.986***	-0.008		-0.011								
	(6.708)	(0.013)		(0.013)								
Weighted Revolutions	9.839	0.002	13.726									
	(8.025)	(0.019)	(8.624)									
Weighted Nat. Elec. for Executive	-16.906	0.036	3.767									
	(16.431)	(0.030)	(20.441)									
Conflict incidence, UCDP-GED					-317.512		-141.439				-451.427	
					(334.653)		(100.866)				(338.519)	
GCP per Capita, PPP$ growth						-0.002*		-0.000				-0.003**
						(0.001)		(0.001)				(0.002)
Conflict incidence, ACLED 1997-2010									-44.578*			
									(23.433)			
Nightlights calibrated, mean, growth										-0.012		
										(0.010)		
Observations	17,870	112,500	20,444	130,500	20,444	20,870	20,444	20,870	72,680	72,450	15,339	16,093
R^2	0.586	0.273	0.576	0.270	-3.024	-0.004	-0.328	0.008	0.935	-0.247	-8.558	-0.025
Number of gid					6,826	6,968	6,826	6,968	7,268	7,245	5,119	5,370
Kleibergen-Paap LM stat. (H0: Underid.)					1.347	15.368	4.976	18.420	6.162	19.172	1.919	11.734
Kleibergen-Paap P-Value					0.510	0.000	0.290	0.001	0.046	0.000	0.383	0.003
Hansen J Overid. test (H0: Eq. Identified)					0.027	7.153	18.679	10.178	1.676	3.786	0.164	8.283
Hansen J P-Value					0.868	0.007	0.000	0.017	0.196	0.052	0.686	0.004
Endogeneity of Regressor (H0: Exogenous)					2.934	0.330	0.227	0.151	6.137	7.075	15.012	0.244
Endog. P-Value					0.087	0.566	0.634	0.698	0.013	0.008	0.000	0.621

Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1

NOTE: Cell FE and time FE included in all specifications.

This time GCP growth is also significant, but the identification statistics suggest that nearly all equations are not well identified. The coefficients on growth suggest that a 1% increase in the

5-year growth rate decreases the likelihood of conflict by 0.002-0.003. These results seems too small and should be treated with extreme caution.

Table 15: PRIO Gridded Data Regressions using Weighted Irregular Political Activity and Weighted Major Elections Instruments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	First Stages		Reduced Forms		2SLS		2SLS + Int. IV		ACLED & Nightl.		2SLS + Clim.	
Variables	%GDP	CINC	%GDP	CINC	%GDP	CINC	%GDP	CINC	%NL	CINC	%GDP	CINC
ln agr. com. shock	-0.483	-0.071***		-0.068***								
	(9.497)	(0.020)		(0.018)								
Exposure to crises	27.644***	-0.009		-0.011								
	(6.486)	(0.013)		(0.013)								
Irregular Political Activity, weigh. s.	7.346	-0.007	2.811									
	(7.724)	(0.008)	(8.703)									
Major Elections, weighted sum	4.214	0.012*	7.348									
	(6.200)	(0.007)	(7.302)									
Conflict incidence, UCDP-GED					85.408		-172.693**				56.052	
					(156.820)		(77.116)				(86.493)	
GCP per Capita, PPP$ growth						-0.002*		-0.000				-0.003**
						(0.001)		(0.001)				(0.002)
Conflict incidence, ACLED 1997-2010									-107.024			
									(83.688)			
Nightlights calibrated, mean, growth										-0.012		
										(0.010)		
Observations	17,870	112,492	20,444	130,500	20,444	20,870	20,444	20,870	72,680	72,450	15,339	16,093
R^2	0.585	0.273	0.575	0.270	-0.040	-0.004	-0.648	0.008	0.783	-0.247	0.173	-0.025
Number of gid					6,826	6,968	6,826	6,968	7,268	7,245	5,119	5,370
Kleibergen-Paap LM stat. (H0: Underid.)					3.027	15.368	7.604	18.420	3.066	19.172	3.652	11.734
Kleibergen-Paap P-Value					0.220	0.000	0.107	0.001	0.216	0.000	0.161	0.003
Hansen J Overid. test (H0: Eq. Identified)					0.333	7.153	25.500	10.178	0.103	3.786	0.281	8.283
Hansen J P-Value					0.564	0.007	0.000	0.017	0.748	0.052	0.596	0.004
Endogeneity of Regressor (H0: Exogenous)					1.011	0.330	2.611	0.151	9.792	7.075	0.423	0.244
Endog. P-Value					0.315	0.566	0.106	0.698	0.002	0.008	0.516	0.621

Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1

NOTE: Cell FE and time FE included in all specifications. "2SLS + Lag. IV" indicates that the instruments were included in levels and lags, to increase the first stage r-squared. In all specifications except for 3SLS, the error matrix is cluster-robust at the country level. 3SLS estimates are obtained with iterative convergence (the *'ireg3'* option in STATA).

3 Conclusion

In this paper I have investigated two aspects of the climate-economy conflict nexus that have thus far been neglected in the literature: Reverse causality and predictive power. With regards to predictive power I have shown that climate moments can successfully predict a good part of African growth rates and conflict incidence. Non of the Random Forest models has however successfully predicted conflict onset. The thus obtained measures of predictive power constitute an upper bound, as other variables (geographic features in particular) are correlated with climate and also impact outcomes. Experiments involving other predictors have shown that at the cell-level, measures of remoteness are better predictors of conflict and growth than climate, although in the joint regressions some climate moments feature among the most important predictors (as shown in the Appendix). The most important climate moments for prediction at the cell-level were 30-year temperature and precipitation MA's and SD's, followed by growth rates and deviations around these values (volatility in short). At the country-level also deviations are among the top predictors. Drought events were not important for any predictive exercise conducted in this paper. In the second part of this paper, I have attempted to address the reverse causality issue between conflict outbreak and economic growth using two datasets and two different instrumentation approaches. The results confirmed the existence of a strong channel from conflict outbreak to negative growth, but could not provide convincing evidence of the channel from growth to conflict (which has been the main focus of the literature). The results through and the difficulties in identifying these channels suggest that reverse causality is an issue that deserves more attention in the new literature on climate, economy and conflict.

References

Bang, J. T., Basuchoudhary, A., & Sen, T. (2015). New tools for predicting economic growth using machine learning: A guide for theory and policy.

Bergholt, D., & Lujala, P. (2012). Climate-related natural disasters, economic growth, and armed civil conflict. *Journal of Peace Research*, *49*(1), 147–162.

Berman, N., & Couttenier, M. (2015). External shocks, internal shots: the geography of civil conflicts. *Review of Economics and Statistics*, *97*(4), 758–776.

Breiman, L. (2001). Random forests. *Machine learning*, *45*(1), 5–32.

Buhaug, H., Benjaminsen, T. A., Sjaastad, E., & Theisen, O. M. (2015). Climate variability, food production shocks, and violent conflict in sub-saharan africa. *Environmental Research Letters*, *10*(12), 125015.

Burke, M., Hsiang, S. M., & Miguel, E. (2015). Climate and conflict. *Annu. Rev. Econ.*, *7*(1), 577–617.

Carleton, T. A., & Hsiang, S. M. (2016). Social and economic impacts of climate. *Science*, *353*(6304), aad9837.

Cooper, F. (2019). *Africa since 1940: the past of the present* (Vol. 13). Cambridge University Press.

Couttenier, M., & Soubeyran, R. (2014). Drought and civil war in sub-saharan africa. *The Economic Journal*, *124*(575), 201–244.

Dell, M., Jones, B. F., & Olken, B. A. (2014). What do we learn from the weather? the new climate–economy literature. *Journal of Economic Literature*, *52*(3), 740–798.

Elvidge, C. D., Hsu, F.-C., Baugh, K. E., & Ghosh, T. (2014). National trends in satellite-observed lighting. *Global urban monitoring and assessment through earth observation*, *23*, 97–118.

Friedman, J., Hastie, T., & Tibshirani, R. (2001). *The elements of statistical learning* (Vol. 1). Springer series in statistics Springer, Berlin.

Hendrix, C. S., & Salehyan, I. (2012). Climate change, rainfall, and social conflict in africa. *Journal of peace research*, *49*(1), 35–50.

Hsiang, S. (2016). Climate econometrics. *Annual Review of Resource Economics*, *8*, 43–75.

Koubi, V. (2017). Climate change, the economy, and conflict. *Current Climate Change Reports*, 1–10.

Koubi, V., Bernauer, T., Kalbhenn, A., & Spilker, G. (2012). Climate variability, economic growth, and civil conflict. *Journal of peace research*, *49*(1), 113–127.

Landis, S. T. (2014). Temperature seasonality and violent conflict: The inconsistencies of a warming planet. *Journal of Peace Research*, *51*(5), 603–618.

Liaw, A., & Wiener, M. (2002). Classification and regression by randomforest. *R news*, *2*(3), 18–22.

Miguel, E., & Satyanath, S. (2011). Re-examining economic shocks and civil conflict. *American Economic Journal: Applied Economics*, *3*(4), 228–232.

Miguel, E., Satyanath, S., & Sergenti, E. (2004). Economic shocks and civil conflict: An instrumental variables approach. *Journal of political Economy*, *112*(4), 725–753.

Tollefsen, A. F., Strand, H., & Buhaug, H. (2012). Prio-grid: A unified spatial data structure. *Journal of Peace Research*, *49*(2), 363–374.

Varian, H. (2014). Big data: New tricks for econometrics. *The Journal of Economic Perspectives*(June 2013), 1–36. doi: 10.1257/jep.28.2.3

APPENDIX

Gridded Predictions, Additional Results

Figure 15: Conflict Incidence UCDP: Other Variables

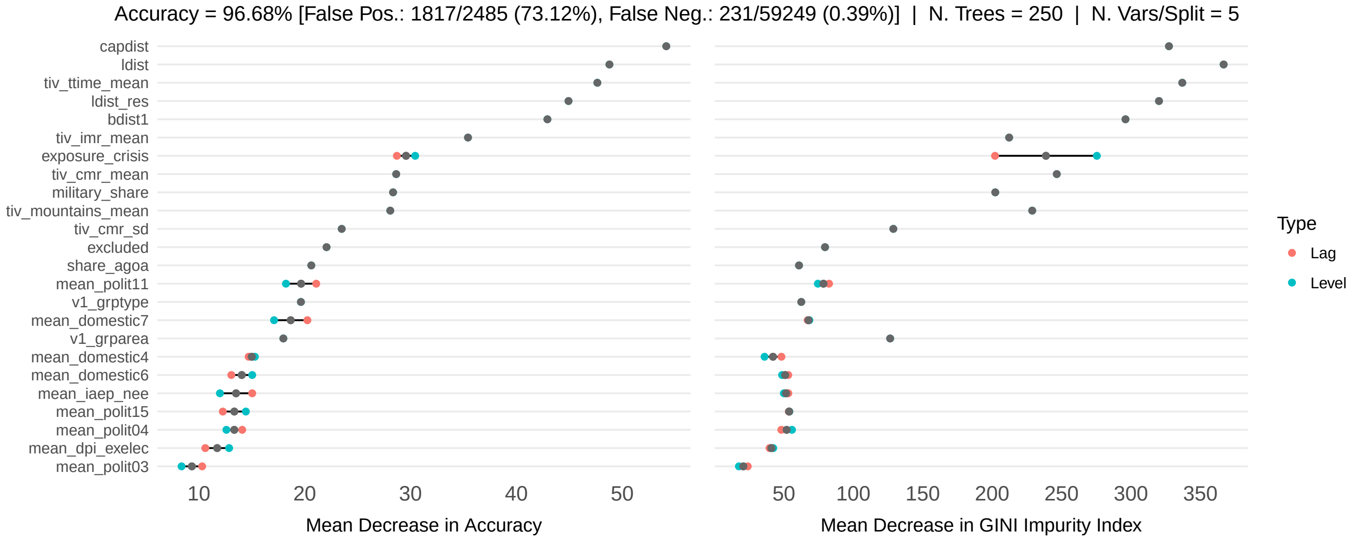

Figure 16: Conflict Incidence UCDP: Climate + Other Variables

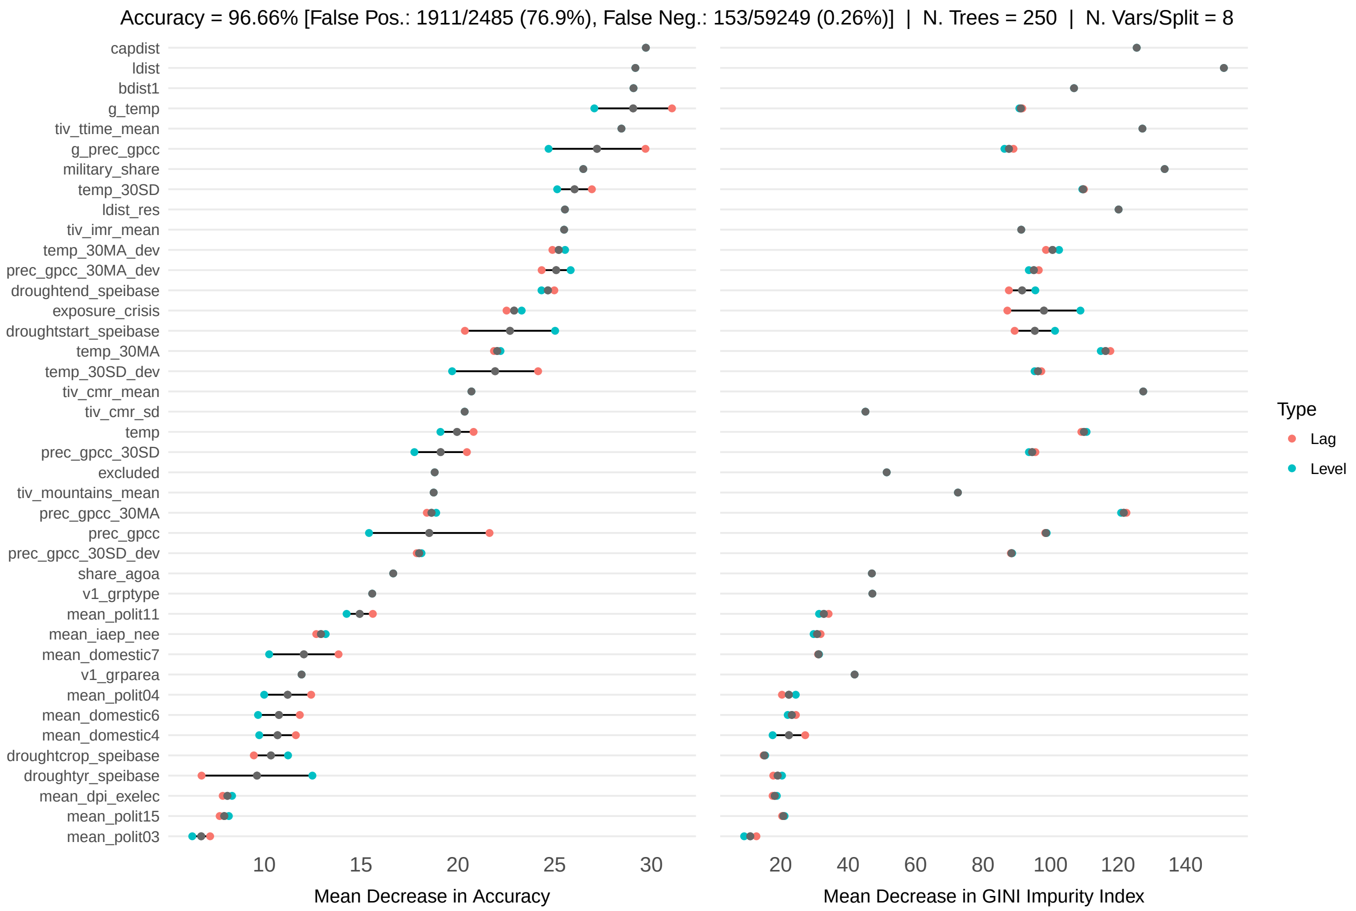

Figure 17: Gross Cell Product per Capita, PPP$ growth (5-year, 1995, 2000, 2005): Other Variables

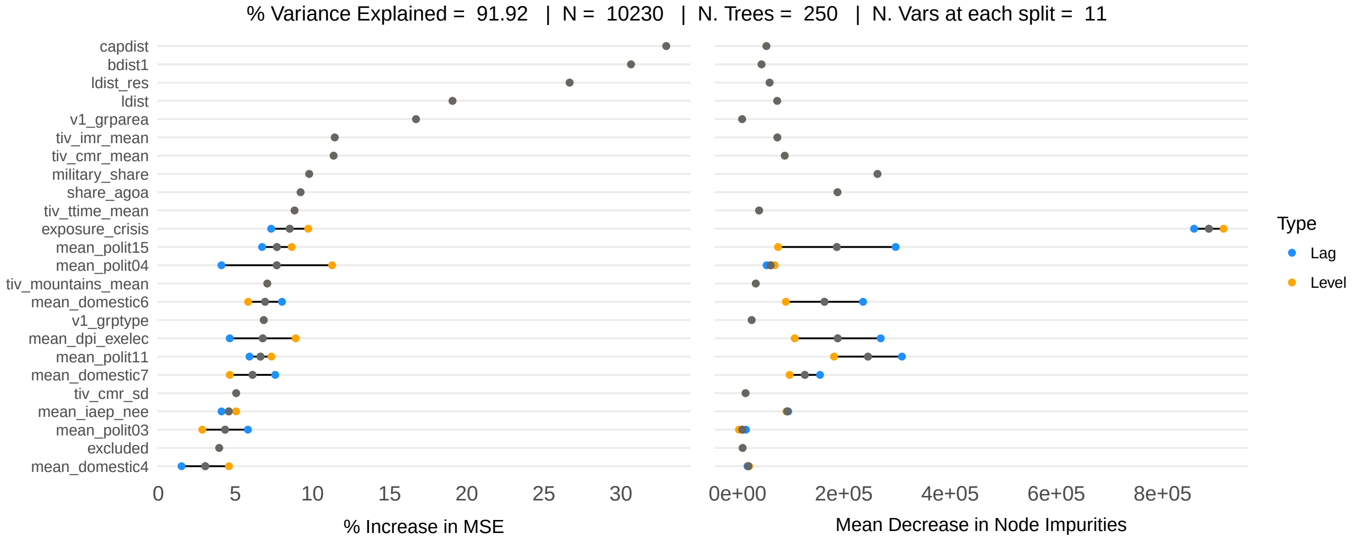

Figure 18: Gross Cell Product per Capita, PPP$ growth (5-year, 1995, 2000, 2005): Climate + Other Variables