doi:10.1038/nature10311

Supplementary Tables

Supplementary Table 1: Summary statistics of primary variables

variable	mean	SD	min	max	years
ENSO index [°C] *May-Dec average*					
NINO12	0.088	1.02	-1.49	3.53	1950-2008
NINO3	0.117	0.829	-1.38	2.75	1950-2008
NINO4	0.117	0.583	-1.36	1.11	1950-2008
Annual probability of conflict onset [%/yr]					
Whole world	3.08	2.51	0	11.2	1950-2004
Teleconnected region only	4.14	2.43	0	11.2	1950-2004
Weakly-affected region only	1.90	2.02	0	9.09	1950-2004
Agricultural output in teleconnected group					
Cereal yields [kg/ha] *country-level obs.*	1500	924	0	8900	1961-2007
Value-added per capita [1990 US$] *country-level obs*	168	105	12.4	771	1970-2007
Value-added per capita [1990 US$] *group*	142	12.8	128	181	1970-2007
Agricultural output in weakly-affected group					
Cereal yields [kg/ha] *country-level obs.*	2880	1720	0	9190	1961-2007
Value-added per capita [1990 US$] *country-level obs*	425	355	40.1	2500	1970-2007
Value-added per capita [1990 US$] *group*	254	45.2	202	352	1970-2007

Supplementary Table 2: Validating the global partition with temperature, precipitation and agriculture responses to NINO3 (1950-2004)

			Independent Variable: NINO3 averaged May-Dec. (°C)	
	Dependant variable	**Model**	**Teleconnected**	**Weakly-Affected**
(1)	**Temperature**	country-level panel	0.048***	-0.017
	(°C)	country-specific trends	[0.009]	[0.011]
		country-specific constants	n = 4067	n = 3461
(2)	**Precipitation**	country-level panel	-0.12***	-0.00
	(mm/day)	country-specific trends	[0.02]	[0.01]
		country-specific constants	n = 2323	n = 1835
(3)	**Cereal yields**	country-level panel	-1.00**	0.68
	(%)	country-specific trends	[0.40]	[0.45]
		country-specific constants	n = 3381	n=2435
(4)	**Agricultural Income**	country-level panel	-0.46**	-0.25
	(%)	country-specific trends	[0.18]	[0.21]
		country-specific constants	n = 2878	n = 2195
(5)	**Agricultural Income**	group total	-0.85***	-0.08
	(%)	trends	[0.28]	[0.35]
			n = 34	n = 29

Standard error in brackets, *** p<0.01, ** p<0.05, * p<0.1. All trends have linear and quadratic terms. SE in rows 1-4 are clustered by country. SE in row 5 are robust to heteroscedasticity. In rows 3-5, coefficients are in units of %/1°C, 1.0 means yields (or income) decline 1% for each 1°C in NINO3. Models 3-5 include one lag dependent variable. Temperature and precipitation are spatially averaged. 1990-94 dropped in W.A. row 5 due to unreliability.

Supplementary Table 3: Monthly correlation coefficients for NINO3 (1950-2008)

	early season			"spring barrier"		late season (Year 0)						
	J	F	M	A	M	J	J	A	S	O	N	D
	(high coherence)											
J	1											
F	0.96	1										
M	0.85	0.92	1									
	(transition)											
A	0.62	0.73	0.85	1								
M	0.39	0.50	0.66	0.86	1							
	(little/no correlation)			*(transition)*		*(high coherence)*						
J	0.15	0.25	0.41	0.64	0.88	1						
J	-0.00	0.09	0.25	0.47	0.71	0.90	1					
A	-0.08	0.01	0.20	0.44	0.68	0.85	0.96	1				
S	-0.05	0.03	0.20	0.44	0.68	0.80	0.88	0.94	1			
O	-0.10	-0.03	0.12	0.35	0.63	0.79	0.86	0.90	0.96	1		
N	-0.10	-0.03	0.10	0.31	0.59	0.77	0.825	0.86	0.92	0.98	1	
D	-0.11	-0.06	0.07	0.29	0.59	0.77	0.81	0.84	0.90	0.96	0.98	1

Supplementary Table 4: The importance of accounting for ENSO dynamics for signal detection

Dependent Variable: Conflict Risk (%/yr)

	Independant variable	
(1)	**Jan-Dec. NINO3**	0.51
	all countries	[0.40]
		n = 54
(2)	**Jan.-Dec. NINO3**	0.76
	teleconnected only	[0.59]
		n = 54
(3)	**May-Dec. NINO3**	0.46*
	all countries	[0.24]
		n = 54
(4)	**May-Dec. NINO3**	0.85**
	teleconnected only	[0.40]
		n = 54

Heteroscedasticity robust S.E., ** p<0.05, * p<0.1.; 1989 dropped. Models all contain a linear trend.

Supplementary Table 5: Survival analysis for peaceful periods between civil conflicts (1950-2004)

	Independent Variable: NINO3 May-Dec	
	teleconnected	weakly affected
(1) Δ proportional hazard per 1°C	25%** [11%]	6% [12%]
(2) Δ average survival time per 1°C	-0.22 yr** [0.09 yr]	-0.06 yr [0.12 yr]
Observations	4929	4346

The peaceful periods between civil conflicts are modeled with survival analysis. The hazard function is assumed to be invariant in the length of time since the last failure (i.e. conflicts are a Poisson point-process) but has an exponential form in response to ENSO (it is non-homogenous). Row (1) describes the proportional change in the hazard rate associated with a 1°C increase in NINO3. A value of "1%" implies that the probability a peaceful period ends is 1.01× the baseline hazard rate. Row (2) describes the same model, but interpreted in terms of survival time. A value of "-0.1 yr" implies that the average peaceful period (across all countries in a sample) decreases by 0.1 years for a 1°C increase in NINO3. All models include country-specific constants and a common linear trend (the model does not converge with country-specific trends). 1989 is dropped. Standard errors clustered by country in brackets: ** $p<0.05$, * $p<0.1$.

Supplementary Table 6: Changes in *annual conflict risk* in response to ENSO (1950-2004)

Independent variable:	NINO12	NINO3	NINO4	NINO12	NINO3	NINO4
	(a) Global time series					
	Dependent Variable: conflict onsets / countries					
Jan-Dec NINO coeff.:	0.33	0.57	0.87*	-	-	-
	[0.24]	[0.39]	[0.49]	-	-	-
Observations:	54	54	54	-	-	-
	teleconnected group			weakly-affected group		
	(b) Regional time-series					
	Dep. Var.: conflict onsets in region / countries in region					
May-Dec NINO Coeff	0.59*	0.762*	0.842	-0.022	0.160	0.623
	[0.304]	[0.390]	[0.535]	[0.215]	[0.307]	[0.495]
Observations:	54	54	54	54	54	54
	(c) Same as (b) plus linear trend					
May-Dec NINO Coeff	0.625**	0.852**	1.07*	-0.071	0.059	0.430
	[0.310]	[0.396]	[0.555]	[0.202]	[0.300]	[0.470]
	(d) Same as (c) plus post-1989 dummy					
May-Dec NINO Coeff	0.634**	0.813**	0.879*	-0.065	0.036	0.319
	[0.252]	[0.318]	[0.482]	[0.208]	[0.305]	[0.446]
	(e) Same as (d) plus 1989 obs					
May-Dec NINO Coeff	0.589**	0.737**	0.679	-0.054	0.054	0.360
	[0.249]	[0.320]	[0.507]	[0.208]	[0.305]	[0.450]
	(f) Ignoring the "spring barrier" obscures signal					
Jan-Dec NINO coeff.:	0.542	0.762	0.901	-	-	-
	[0.374]	[0.586]	[0.622]	-	-	-
	(g) Country-level longitudinal-data linear probability model					
	Dep. Var.: conflict onset in a country (binary)					
May-Dec NINO coeff.:	0.674%	0.893	1.0	-0.108	0.038	0.414
	[0.38]*	[0.415]**	[0.522]*	[0.154]	[0.266]	[0.471]
Observations:	3978	3978	3978	3400	3400	3400

Heteroscedasticity-robust standard errors in brackets, *** p<0.01, ** p<0.05, * p<0.1. Coefficients are probability responses in units of %/1°C, 1.0 means the probability of conflict in a given year rises 0.01 for each 1°C increase in NINO averaged May-Dec. All panels except (e) drop 1989. Panel (f) includes a linear trend. Panel (g) includes county-specific constants (fixed-effects) and country-specific linear trends with estimated standard errors clustered by country.

doi:10.1038/nature10311

Supplementary Table 7: Results are not driven by patterns of serial correlation in ACR

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Teleconnected					**Weakly-Affected**				
	Dependent Variables					Dependent Variables				
	ACR_t	$Residuals_t$	ACR_t	ACR_t-ACR_{t-1}	ACR_t	ACR_t	$Residuals_t$	ACR_t	ACR_t-ACR_{t-1}	ACR_t
$NINO3_t$	0.813**		0.719**		0.737**	0.036		0.009		0.054
	[0.318]		[0.357]		[0.319]^a	[0.305]		[0.321]		[0.321]^a
					[0.302]^b					[0.336]^b
					[0.285]^c					[0.339]^c
$Residuals_{t-1}$		0.142					0.328**			
		[0.161]					[0.136]			
ACR_{t-1}			0.151					0.326**		
			[0.152]					[0.141]		
$NINO3_t$-$NINO3_{t-1}$				0.609*					-0.092	
				[0.307]					[0.249]	
year	-0.098***		-0.091***	-0.005	-0.107***	-0.014		-0.011	-0.006	-0.012
	[0.024]		[0.028]	[0.032]		[0.023]		[0.025]	[0.030]	
post 1989 const.	3.311***		2.893***	-0.400	3.954***	1.950*		1.412	0.355	1.799
	[0.885]		[1.040]	[1.295]		[1.063]		[0.993]	[1.225]	
Constant	197.125***	-0.025	181.989***	10.415	213.510***	29.840	-0.009	21.911	11.056	25.979
	[46.336]	[0.264]	[55.251]	[62.153]		[45.740]	[0.259]	[48.460]	[59.935]	
Durbin-Watson d-statistic:	1.656					1.339				
Observations	54	52	53	53	55	54	52	53	53	55

Models 2 and 7 regress residuals from Models 1 and 6 (respectively) on their lagged values. Models 5 and 10 present three estimates for the SE of the main coefficient of interest: each is a Newey-West[29] estimate with lag lengths of a: 3 yrs, b: 5 yrs and c: 10 yrs. Greene[30] (2003) recommends a lag of at least $T^{0.25}$ (where T is the length of the time series) which is 2.7 in this case. All other SEs are White estimators.[31]

doi:10.1038/nature10311

Supplementary Table 8: Results are not driven by sample selection: Teleconnected response to ENSO

	(1)	(2)	(3)	(4)	(5)	(6)
	BASELINE					FULL SAMPLE
Sample:	$t \geq 1950$	$t \geq 1950$	include 1946-49	$t \neq 1948$	$t \geq 1975^a$	include 1946-49
	$t \neq 1989$	include 1989	$t \neq 1989$	$t \neq 1989$	include 1989	include 1989
NINO3 May-Dec.	0.81**	0.74**	0.75**	0.77**	0.83**	0.68**
	[0.32]	[0.32]	[0.32]	[0.32]	[0.35]	[0.32]
Year	-0.10***	-0.11***	-0.12***	-0.10***	-0.36***	-0.13***
	[0.02]	[0.02]	[0.03]	[0.02]	[0.07]	[0.03]
Post Cold War Const.	3.31***	3.95***	3.90***	3.42***	7.51***	4.47***
	[0.88]	[1.03]	[0.97]	[0.87]	[1.41]	[1.07]
Constant	197.12***	213.51***	249.80***	206.30***	723.02***	261.14***
	[46.34]	[48.85]	[60.00]	[43.77]	[136.99]	[60.32]
Observations	54	55	58	57	30	59

Heteroscedasticity robust standard errors in brackets *** p<0.01, ** p<0.05, * p<0.1. [a]After 1974, the set of countries in the teleconnected group stabilized, reaching 87 countries (the maximum of 91 was achieved 2002-4).

Supplementary Table 9: Fixed effects or country-specific trends do not drive the result

Dependent Variable: Conflict Risk (% / yr)
Independent Variable: May-Dec NINO3 (°C)

	Panel model	**Teleconnected** (%/yr°C)	**Weakly-Affected** (%/yr°C)
(1)	**No controls**	0.85*	0.20
		[0.44]	[0.33]
(2)	**Country fixed effects**	0.89**	0.13
		[0.40]	[0.33]
(3)	**Country-trends**	0.90**	0.13
		[0.39]	[0.33]
(4)	**Country fixed effects**	0.89**	0.04
	Country-trends	[0.39]	[0.32]
	Observations	3978	3400

Conley[32] SE in brackets, *** p<0.01, ** p<0.05, * p<0.1.
1989 dropped.

Supplementary Table 10: ENSO influences ACR holding local temperature and rainfall constant

Dependent Variable: Conflict Risk (% / yr)
Independent Variable: May-Dec NINO3 (°C)

	Panel model	**Teleconnected** (%/yr°C)	**Weakly-Affected** (%/yr°C)
(1)	**No weather**	0.89**	0.04
		[0.39]	[0.32]
		n=3978	n=3400
(2)	**Include temperature**	1.02***	0.13
	(monthly, 12 vars)	[0.39]	[0.30]
		n=3978	n=3400
(3)	**Include temp & rain**	1.66***	0.36
	(monthly, 24 vars)	[0.48]	[0.33]
		n=2234	n=1774

Conley[32] SE in brackets, *** p<0.01, ** p<0.05, * p<0.1. 1989 dropped. All models include country specific constants (fixed effects) and country-specific trends.

doi:10.1038/nature10311

Supplementary Table 11: Controlling for common time-varying controls

	Dependent Variable: Conflict risk (% / yr)										
	Teleconnected Group										***W. A.***
	All Countries								***Africa***	***Not Africa***	***All***
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
NINO3$_t$	0.893**		1.212***		1.046**		0.976**	1.229***	1.399*	1.004*	0.020
May-Dec.	[0.389]		[0.453]		[0.409]		[0.407]	[0.477]	[0.807]	[0.563]	[0.313]
log GDP$_{t-1}$ **/ capita**$_{t-1}$		3.348	3.224					3.988	10.765***	-11.025***	11.563***
(lagged)		[2.170]	[2.158]					[2.635]	[3.280]	[3.335]	[3.965]
Polity IV score$_{t-1}$				0.203**	0.200**			0.291**	0.418**	0.301	0.451***
(lagged)				[0.098]	[0.097]			[0.146]	[0.205]	[0.242]	[0.133]
Polity IV score$^2_{t-1}$				-0.033	-0.033			-0.039	0.016	-0.103**	0.031
(lagged)				[0.023]	[0.023]			[0.027]	[0.038]	[0.043]	[0.035]
Log population$_{t-1}$						10.822	10.466	4.118	-1.956	17.706*	21.308***
(lagged)						[10.393]	[10.348]	[13.413]	[20.822]	[9.839]	[6.533]
Observations	3,978	2,793	2,793	3,450	3,450	3,555	3,555	2,546	1,398	1,148	1,897

Conley[32] SE in brackets, *** p<0.01, ** p<0.05, * p<0.1. 1989 dropped. All models include country-specific constants (fixed-effects) and country-specific trends. Models 1-10 are for the teleconnected group only. Model 11 is for the weakly affected group only.

Supplementary Table 12: The "kitchen sink" model

Dependent Variable: Conflict Risk (% / yr)
Independent Variable: May-Dec NINO3 (°C)

	Sample	Teleconnected (%/yr°C)	Weakly-Affected (%/yr°C)
	No Weather Controls		
(1)	**All Countries**	1.35***	0.03
		[0.47]	[0.31]
		n = 2464	n=1827
(2)	**Africa Only**	1.63**	0.27
		[0.76]	[0.53]
		n=1349	n=96
(3)	**Not Africa**	0.96	0.02
		[0.62]	[0.33]
		n=1115	n=1731
	With Weather Controls		
(4)	**All Countries**	1.83***	0.46
		[0.53]	[0.33]
		n=1973	n=1467
(5)	**Africa Only**	1.95*	0.54
		[1.01]	[1.79]
		n=1083	n=75
(6)	**Not Africa**	1.49**	0.50
		[0.66]	[0.34]
		n=890	n=1392

Conley[32] SE in brackets, *** p<0.01, ** p<0.05, * p<0.1. 1989 dropped in all models. Controls include country fixed-effects, country-specific linear trends, lagged log income per capita, lagged per capita income growth, lagged Polity IV score (with quadratic term), lagged agriculture industry share (%), lagged percent urbanized, lagged population, lagged percent female, lagged percent below 15 years old, lagged percent above 65 years old. Models 4-6 include 12 monthly temperature variables, 12 monthly precipitation variables and annual average maximum tropical cyclone windspeed.

Supplementary Table 13: Robustness analysis under different definitions of conflict risk (1950-2004)

Dependent Variable: Binary conflict indicator for teleconnected group countries

	Onset2	Onset3	Onset4	Onset5	Onset6	Onset7	Onset8	Onset9
"Small" Conflicts	0.454	0.372	0.342	0.379	0.411	0.411	0.449	0.449
"Large" Conflicts	0.453	0.291	0.210	0.121	0.093	0.064	0.102	0.102

"Small" conflicts had between 25 and 1000 battle related deaths. "Large" conflicts had more than 1000 battle related deaths. Each coefficient is a separate regression. Coefficients are percent probability responses in units of %/1°C, 1 means the probability of conflict in a given year rises 0.01 for each 1°C increase in NINO3 sea surface temperature. Onset$N$ indicates that a conflict must be dormant for N years before renewed violence marks the "onset" of a new conflict. NINO3 SSTs are measured using May-December means. Models contain country constants (fixed effect) and country-specific time trends. 1989 dropped.

doi:10.1038/nature10311

Supplementary Table 14: Conflict risk response to past and future ENSO

Dependent Variable: *annual conflict risk* (1950-2004)

	teleconnected group					weakly-affected group				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$NINO3_{t+1}$					0.472					0.101
					[0.342]					[0.223]
$NINO3_t$	0.813**	0.711**	0.707**	0.815**		0.036	0.085	0.077	0.190	
	[0.318]	[0.328]	[0.340]	[0.332]		[0.305]	[0.306]	[0.294]	[0.350]	
$NINO3_{t-1}$		-0.319	-0.305	-0.324			0.197	0.232	0.228	
		[0.317]	[0.313]	[0.309]			[0.317]	[0.336]	[0.315]	
$NINO3_t$			0.141					0.362		
$\times NINO3_{t-1}$			[0.365]					[0.513]		
$NINO3_{t-2}$				0.086					0.182	
				[0.306]					[0.289]	
Observations	54	53	53	52	53	54	53	53	52	53
R-squared	0.277	0.300	0.301	0.336	0.221	0.123	0.128	0.140	0.133	0.113
F-test p-value		0.457		0.384			0.549		0.362	

Heteroscedastic robust errors in brackets. ** p<0.05, * p<0.1. Coefficients are percent probability responses in units of %1/°C, 1 means the probability of conflict in a given year rises 0.01 for each 1°C increase in NINO3 sea surface temperature averaged May-Dec. Models contain a linear trend and a post Cold War dummy. 1989 obs. dropped. F-stat tests the null hypothesis that all ENSO associated conflicts are displaced in time: in all cases we fail to reject this null.

Supplementary Figures

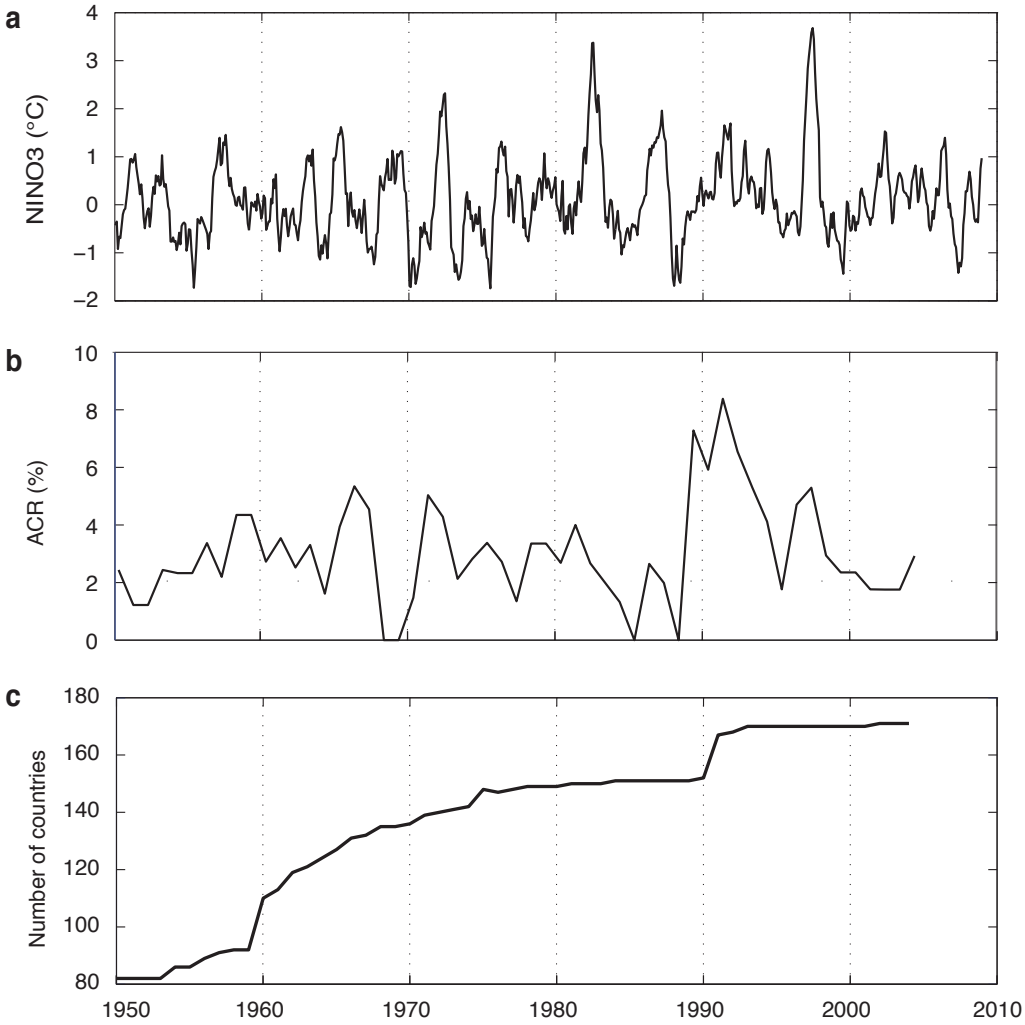

Supplementary Figure 1: (a) Monthly sea surface temperature anomaly over the NINO3 region (5°S-5°N, 150°W-90°W). (b) ACR for the whole world. (c) The number of countries in the dataset is growing over time, more than doubling over the period of observation. Because *conflict onset* is coded as a binary variable for each country-year observation, it is necessary to normalize the total number of observed *conflict onsets* by the number of distinct countries being observed. To check that this trend is not driving any of our results, we re-estimate the model while restricting the sample to the period following 1975 (inclusive) when the sample of teleconnected countries was stable; we also estimate a model with the raw panel data using country-specific constants in the regression (country fixed-effects).

doi:10.1038/nature10311

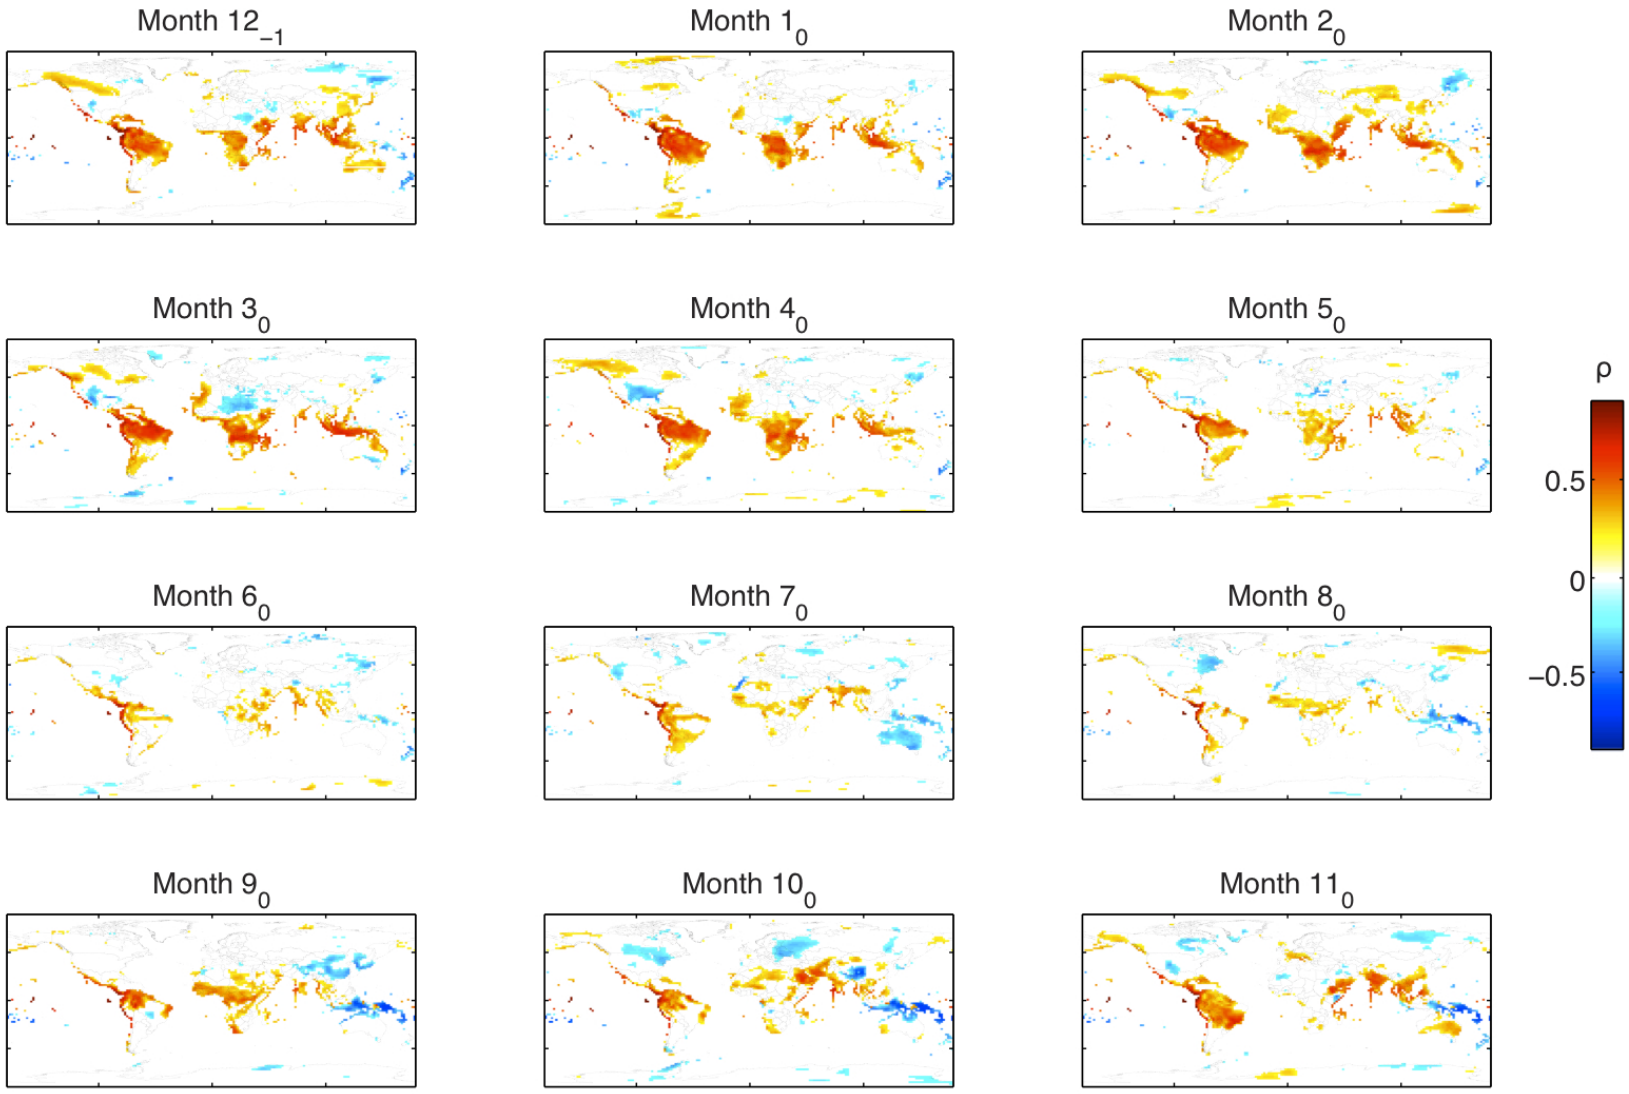

Supplementary Figure 2: The pixel-level correlation between surface temperatures and NINO3 two months earlier. Only statistically significant ($\alpha = 0.1$) correlations over land are colored.

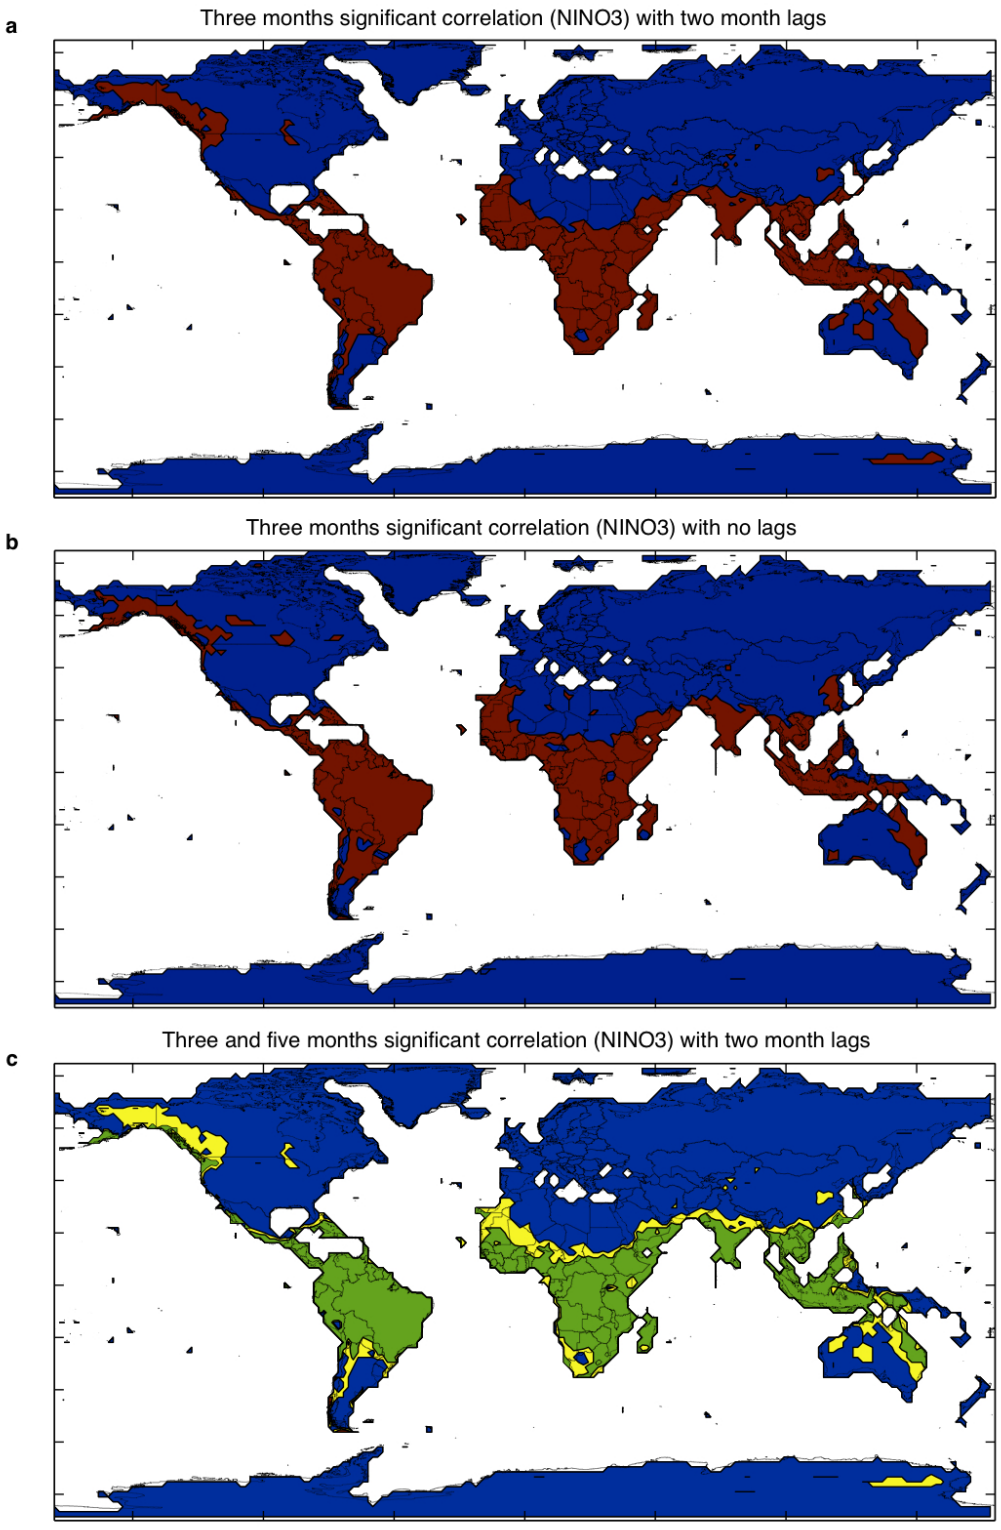

Supplementary Figure 3: (a) Global ENSO teleconnection partition used in main analysis. Red: pixels coded as teleconnected when surface temperatures are positively correlated with NINO3 two months earlier ($L = 2$) for at least three months ($R = 3$). Blue: weakly-affected pixels failing this criteria. (b) Same, but $L = 0$. (c) Green: teleconnected pixels when $R = 5$ and $R = 3$ (and $L = 2$). Yellow: teleconnected pixels when $R = 3$ but not when $R = 5$. Blue: weakly-affected pixels when $R = 3$ and $R = 5$.

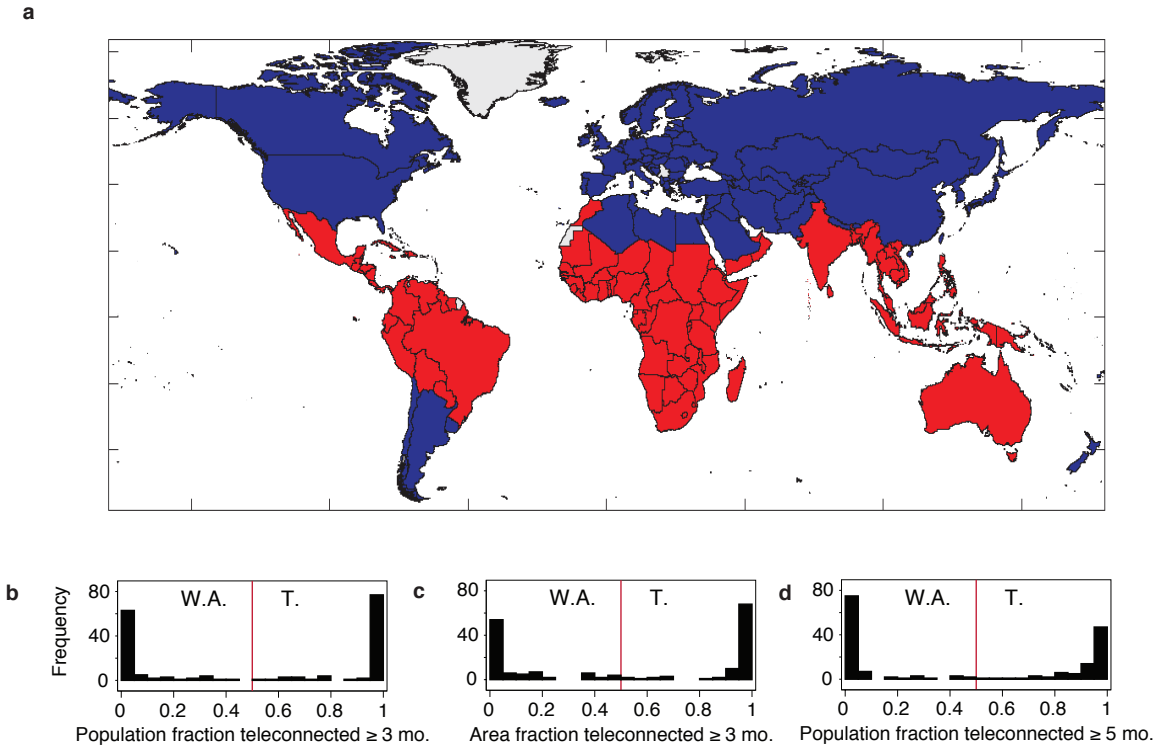

Supplementary Figure 4: (a) Countries assigned to be teleconnected when pixels are weighted by population ($w_x = population_{x,t=2000}$) are red, weakly-affected countries are blue. Grey countries have no conflict data. (b) Countries are coded as weakly-affected (W.A.) or teleconnected (T.) based on the fraction of their population that inhabited teleconnected cells in 2000. (c) The distribution of countries changes little if teleconnected area is used instead of populations. (d) The distribution of countries changes little if temperatures in teleconnected cells are required to correlate with NINO3 for five months instead of three.

doi:10.1038/nature10311

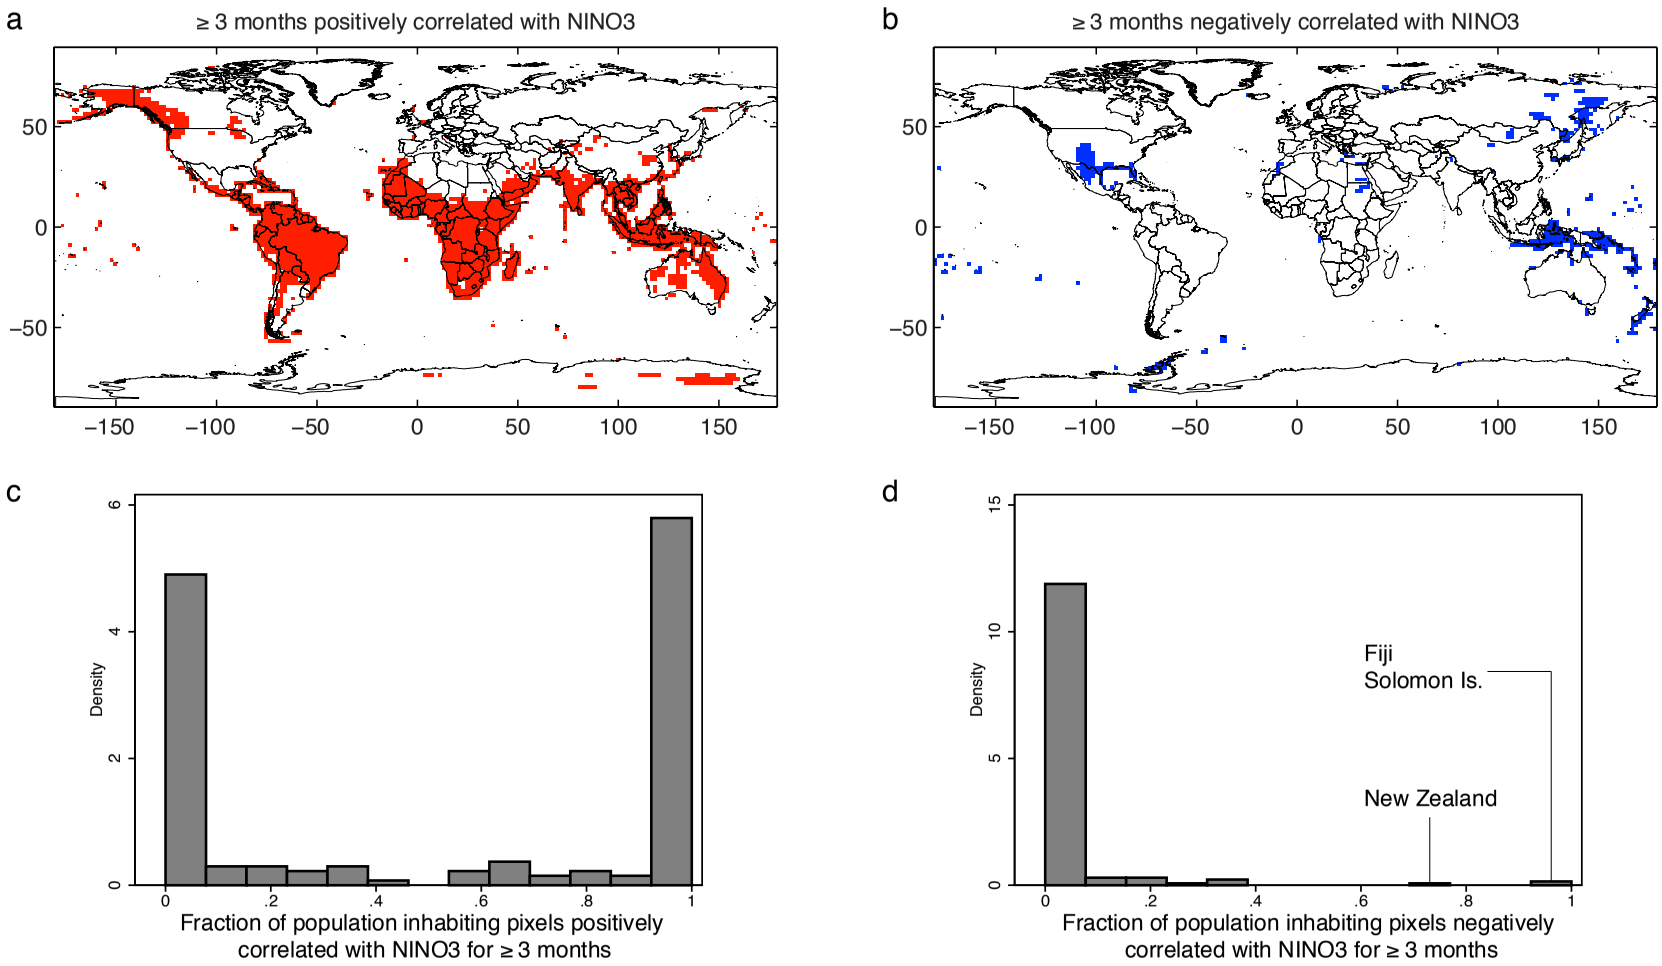

Supplementary Figure 5: Few continental locations have temperates negatively correlated with NINO3, i.e. few locations experience cooling when the NINO3 region warms. (a) Pixels with surface temperature significantly ($\alpha = 0.1$) and positively correlated with NINO3 for at least 3 months are red ($L = 2$). (b) Pixels with surface temperature significantly and negatively correlated with NINO3 for at least 3 months are blue ($L = 2$). As established in previous literature, cooling in the western United States, Mexico, Indonesia and Russia is apparent, but not enough individuals live in "negatively teleconnected" pixels to justify coding these countries differently. (c) The distribution of countries in 2000 by their population fractions that inhabit positively teleconnected pixels from Panel (a); in the main analysis, countries were coded as teleconnected if > 0.5 of the population lived in teleconnected pixels. (d) Same, but for the negatively teleconnected pixels from Panel (b). Only in Fiji, the Solomon Islands and New Zealand do a majority of individuals live in negatively teleconnected pixels, although none of these countries have experienced civil conflict during the period of observation.

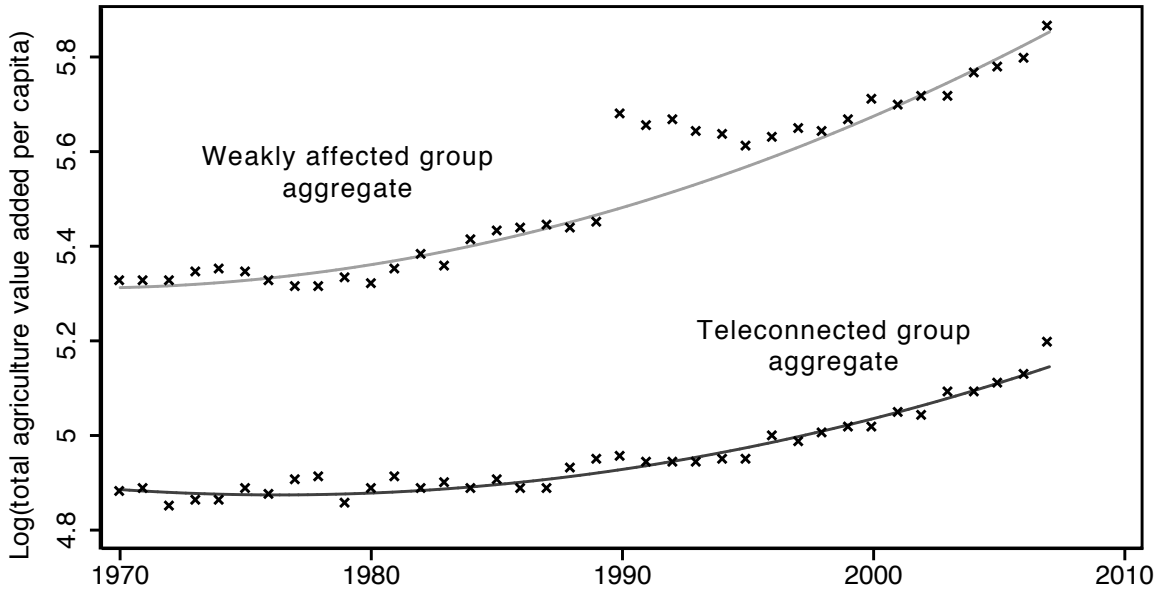

Supplementary Figure 6: Log total agricultural revenue for the weakly affected group (top) and teleconnected group (bottom), both with a quadratic fit. During 1990-1994 values for the weakly affected group are unreliable, probably due to the breakup of the Soviet Union, and are omitted from the analysis.

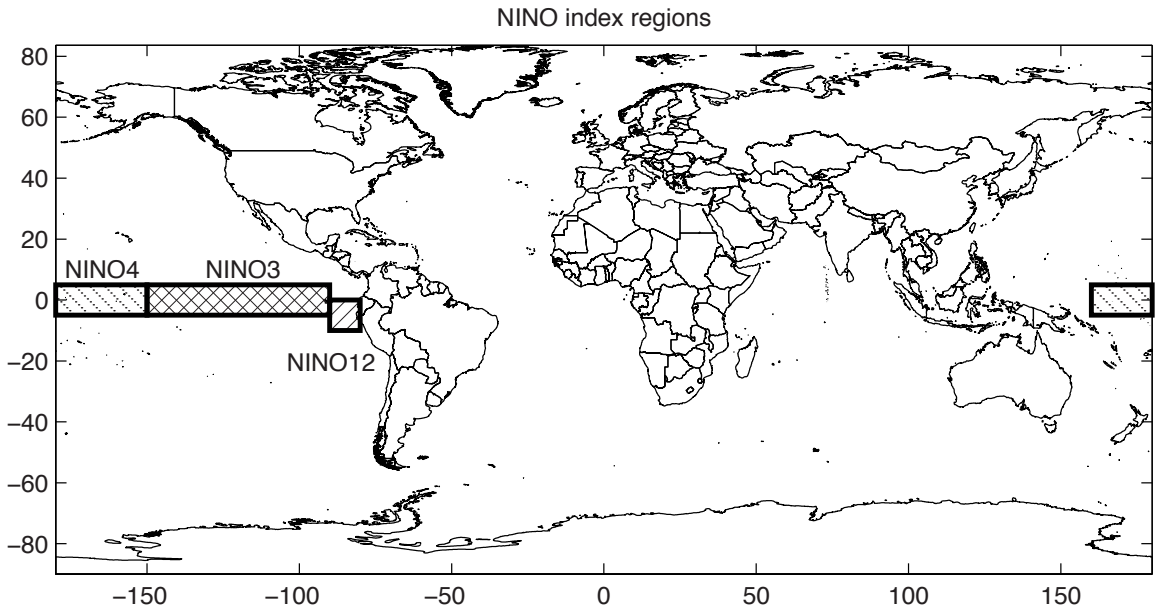

Supplementary Figure 7: The NINO12, NINO3 and NINO4 regions in the equatorial Pacific Ocean. NINO index values are defined as the average sea surface temperature over a NINO region minus the long-term mean sea surface temperature in that region.

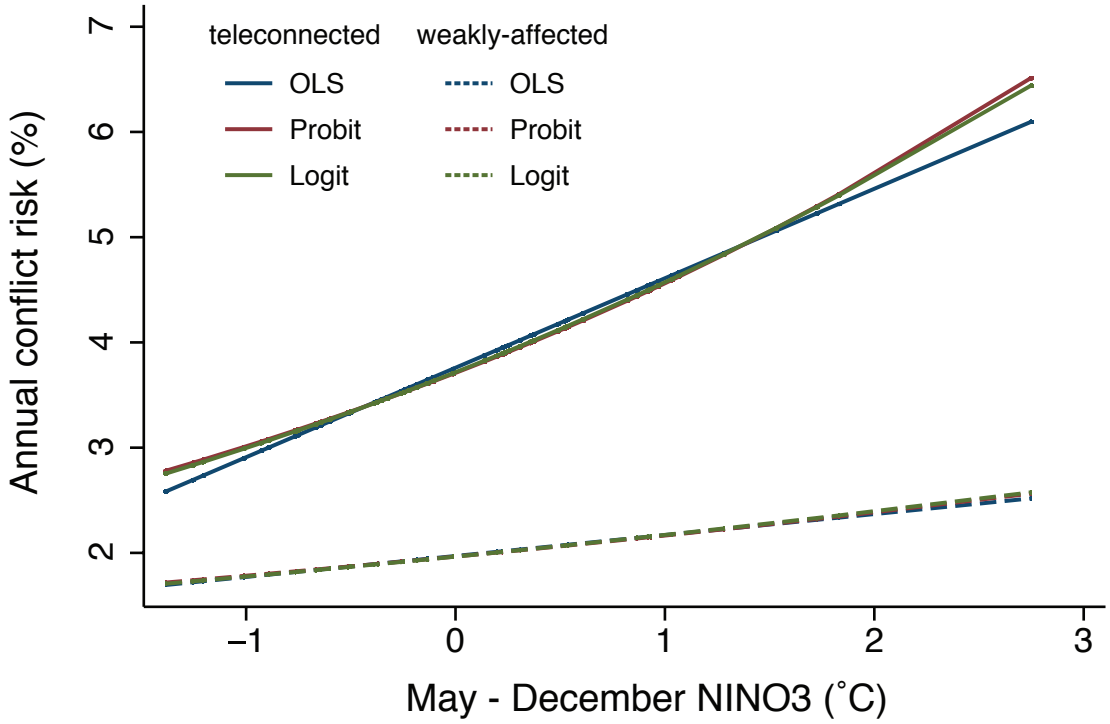

Supplementary Figure 8: The linear probability model used for our country-level analysis (see main text Table 1 row 5) is virtually indistinguishable from non-linear probability models. For both regions, logit and probit models produce almost identical results to the linear probability model used in the main analysis. We prefer (and present) the linear probability model because it is simpler.

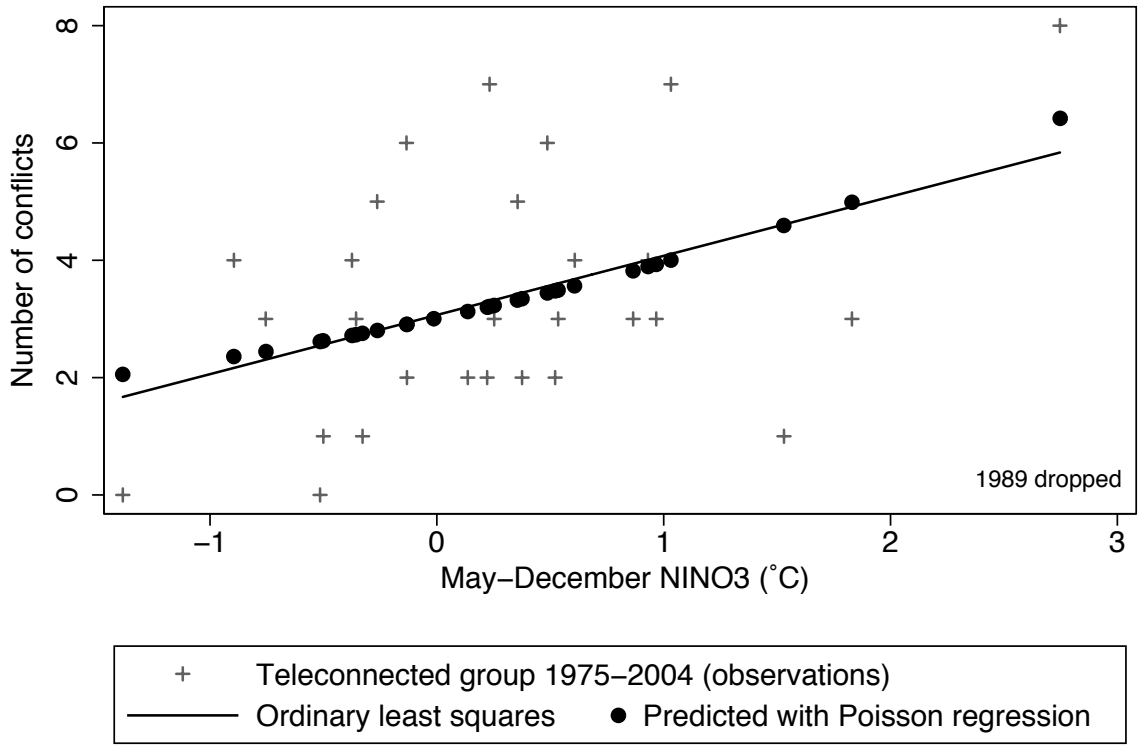

Supplementary Figure 9: Our model for ACR could instead be modeled with standard statistical techniques for "count data" such as a Poisson regression. Unfortunately, such an approach is difficult to interpret when the sample size of countries is changing over time (See Fig. 1). Nonetheless, we can use a Poisson regression to model the number of conflicts for the period 1975-2004, when the sample of countries is almost constant. However, we find there is no obvious gain in prediction over ordinary least squares.

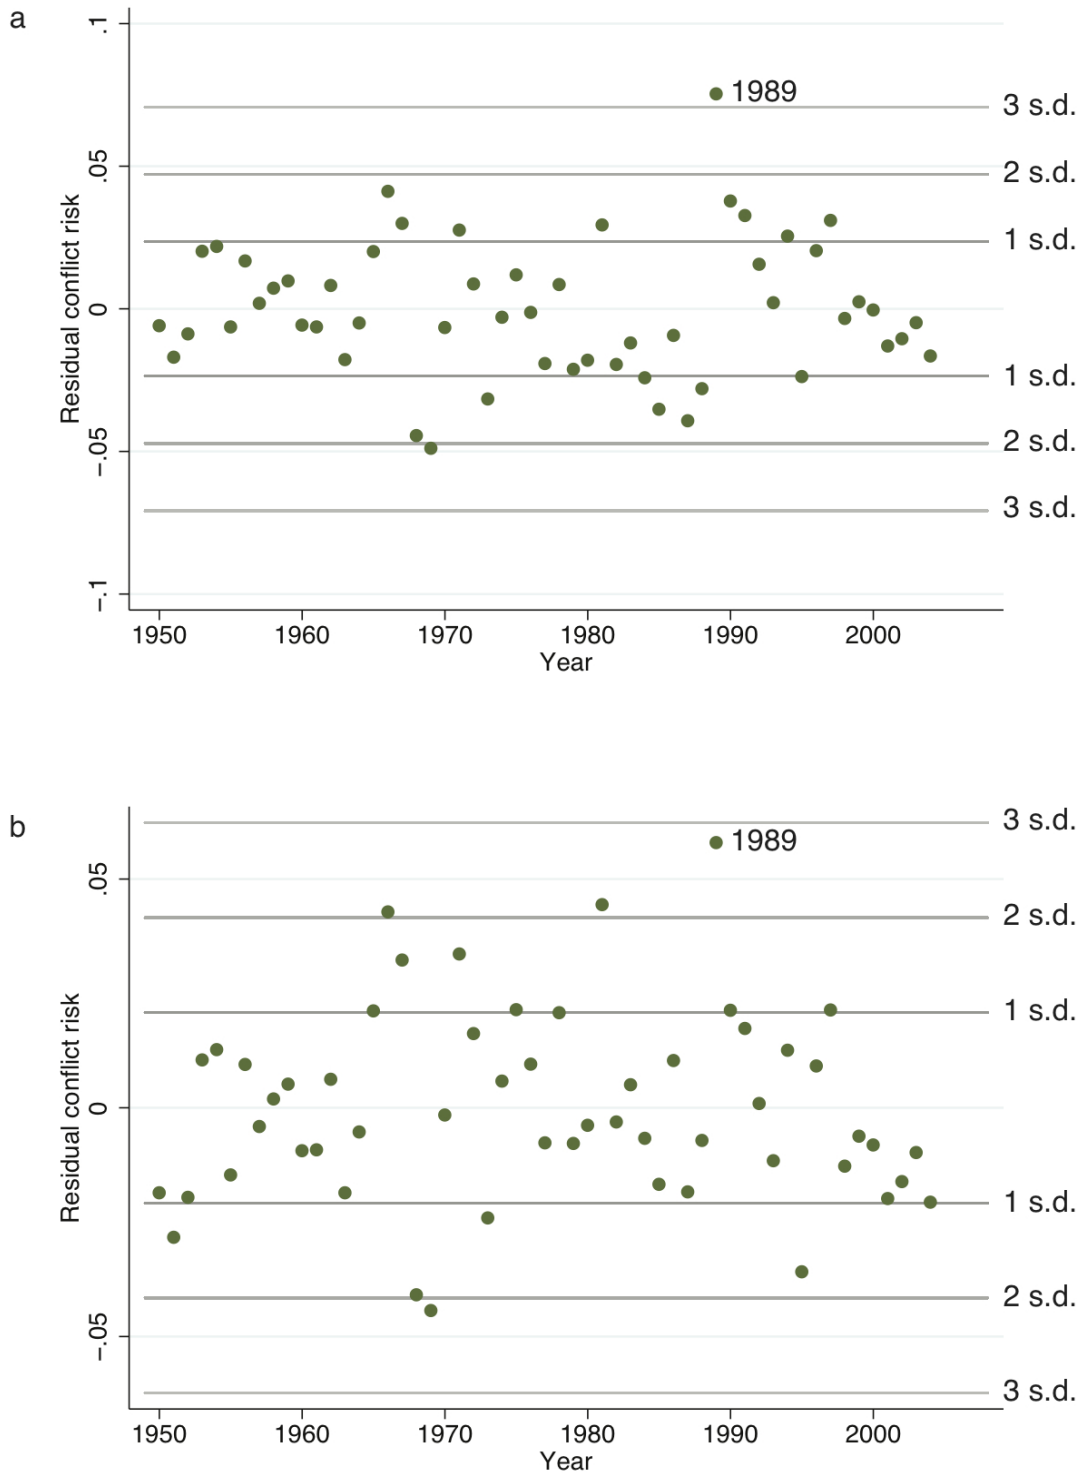

Supplementary Figure 10: (A) Residuals ($\epsilon_i(t)$) estimated with a linear trend (see main text Table 1 row 2 and Equation 6). 1989 is a 3 standard deviation outlier and therefore has been dropped from all models, for both the teleconnected and weakly-affected groups. (B) When a post-1989 (inclusive) constant term is added to the regression (see Table 1 panel d and Equation 6), 1989 is less of an outlier. However, for consistency in models, it is still dropped from the estimation. We re-introduce the 1989 observation in Table 6 Panel e using this specification and find our results are largely unchanged.

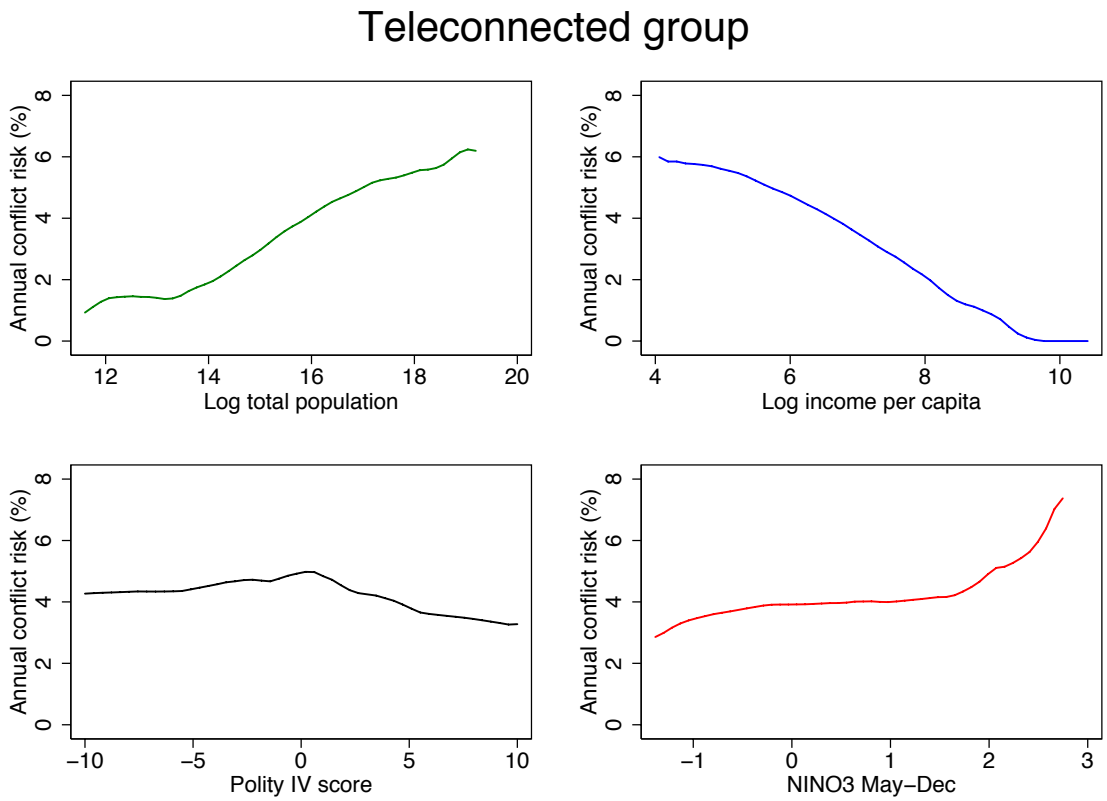

Supplementary Figure 11: Bivariate non-linear fits of ACR against three standard correlates and NINO3. While population and income are correlated with conflict risk across countries, when country means (fixed-effects) are included, the correlation is not significantly different than zero (Table 11). Only Polity IV (a measure of democratic institutions) remains statistically significant, however it explains very little variation in conflict risk over time when compared to NINO3.

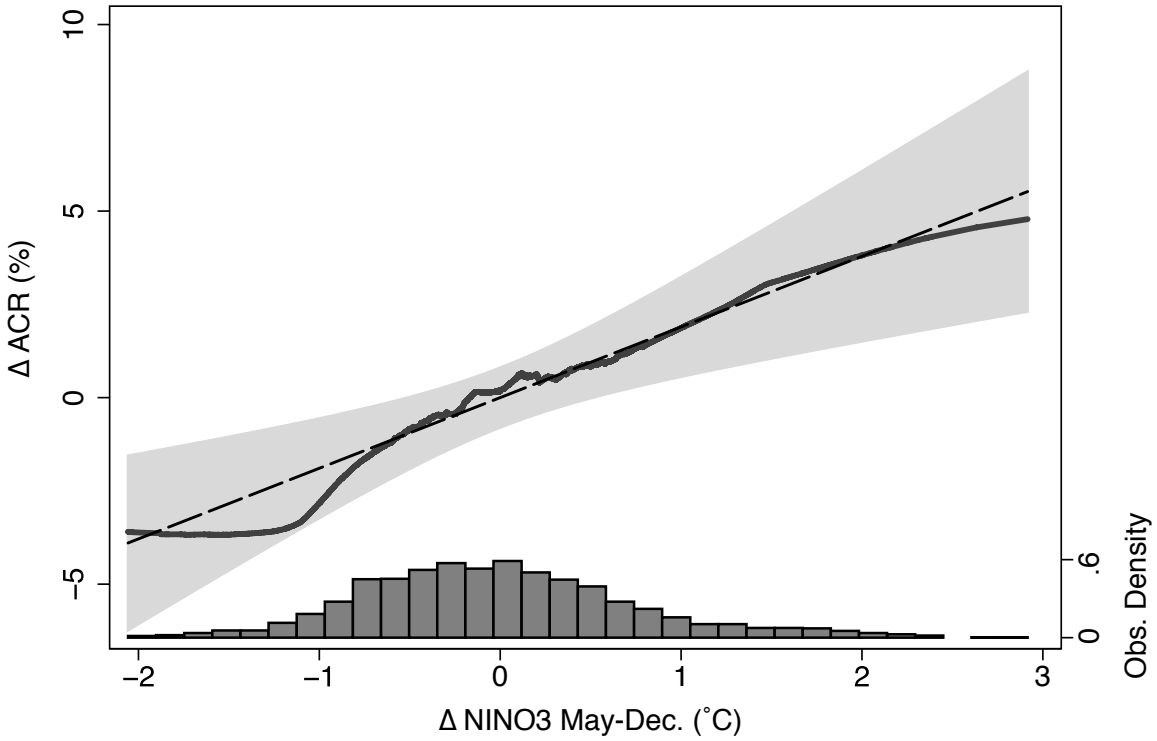

Supplementary Figure 12: Linearity of teleconnected ACR regression on NINO3 when both are partialled-out on 35 time-varying controls, country fixed-effects and country-specific time trends. Dashed line is the OLS fit (with 95% confidence interval), and solid line is a non-linear fit. The density of observations following this partialling-out transformation is the histogram at the bottom.

Supplementary Methods

This supplementary methods section contains

1. Data sources and definitions
2. The ENSO-teleconnection partition (construction, robustness and validation)
3. ENSO timing and measurement
4. Main statistical models
5. Robustness of the main result

1 Data sources and definitions

Table 1 contains summary statistics for the main variables in this study.

Climate data To determine the relative teleconnectedness of local climates to ENSO, we use the National Centers for Environmental Prediction (NCEP) Climate Data Assimilation System 1 (CDAS1) reanalysis of monthly surface temperatures during 1949-2009.[33] ENSO variations can be detected using different indices, with the most commonly used being equatorial Pacific sea surface temperature (SST) anomalies. We utilize monthly means of three such indices: NINO12 (10°S-Eq, 90°W-80°W), NINO3 (5°S-5°N, 150°W-90°W), and NINO4 (5°S-5°N, 160°E-150°W).[34] These indices measure SST in different locations in the Pacific Ocean, with NINO12 measured furthest east and NINO4 measured furthest west (see Fig. 7). These three indices differ both in the magnitude and timing of their variations, but are correlated with one another. For example, the standard deviation of our NINO12 measure (1.02°C) is substantially larger than that of our NINO4 measure (0.58°C). Our main text presents results for NINO3 because it is less influenced by coastal perturbations (compared to NINO12) and captures "medium-scale" events reliably (compared to NINO4), although we find that the choice of NINO index is inconsequential.

Conflict data We use the Onset and Duration of Intrastate Conflict dataset compiled by the Uppsala Conflict Data Program (UCDP) at Uppsala University and the International Peace Research Institute in Oslo (PRIO).[35,36] This dataset contains information on the magnitudes and recurrence periods of conflicts. A *conflict* is defined as "a contested incompatibility that concerns government and/or territory where the use of armed force between two parties, of which at least one is the government of

a state, results in at least 25 battle-related death".[35,36] A *conflict onset* is the date on which a certain fatality threshold has been crossed in specific conflict between a unique government-opposition dyad (note that a country can experience *conflict onset* in sequential years if its government fights with different opposition groups in sequential years). The dataset uses fatality thresholds of 25 and 1000 battle-related deaths to distinguish events of different intensities. Should a conflict that has subsided be reinitiated, it is counted as a new conflict only if the last time it surpassed the 25 battle-deaths cutoff was more than X years ago, where $2 \leq X \leq 9$. For example, "Onset2" records a new *conflict onset* only if that conflict is new or if that specific conflict has been quiescent for at least two years. Our main results utilize this 25 fatality and 2 year intermittency definition of *conflict onset*. We check the robustness of this choice in a later section.

A subsample of this data set (approximately half of the full data set) has monthly dates associated with the timing of conflict onset.[35,36] This subsample is used to construct Fig. 2c in the main text.

Agricultural data To validate our partition of the world into ENSO teleconnected and weakly-affected groups, we estimate the differential impact of ENSO on agricultural yields and agricultural revenue in each. Cereal yields are obtained from the Food and Agricultural Organization's (FAO) FAOStat database[37] and are available for 1961-2007. Agriculture revenue is acquired from the United Nations National Accounts database[38] and is available for 1970-2008. The total annual agricultural revenue for each group is simply the sum of agricultural revenue across all countries in the given year (see Fig. 6).

Income data Income data is obtained from the United Nations National Accounts database.[38] Incomes are available for 1970-2008 in 1990 US dollars per capita.

Demographic data Population data is obtained from United Nations World Development Indicators[39] and are available for 1950-2008.

Political institutions data Polity IV data describes the level of democratization embodied by the political institutions of a country and is obtained from the Polity IV Project sponsored by the Political Instability Task Force[40] and are available for 1950-2008. Polity IV scores take integer values ranging from -10 (hereditary monarchy) to +10 (consolidated democracy).

Rainfall data Data merged from gauges, satellite observations and numerical simulations is obtained from the Climate Predicition Center (CPC) Merged Analysis of Precipitation[41] (CMAP) and are

available for 1979-2008.

Tropical Cyclone data Hurricane, typhoon, cyclone and tropical storm data are obtained from the Limited Information Cyclone Reconstruction and Integration for Climate and Economics (LICRICE) model[42] and are available for 1950-2008. The data describe the annual maximum windspeed spatially averaged over each country.

2 The ENSO-teleconnection partition

2.1 ENSO teleconnections

A central contribution of this work is to develop a simple and robust partition of the continents based on their teleconnectedness to ENSO. Previous analyses have identified different types of ENSO teleconnections using a number of statistical techniques on different data sets.[43–51] Because of the heterogeneity in these approaches, there are variations in the patterns of teleconnections that they characterize. Nonetheless, there are many common patterns and we attempt to summarize this agreement with a simple approach that partitions the world into locations that are either strongly teleconnected or weakly affected. While this a crude simplification of the high dimensional and continuous structure of ENSO teleconnections, it provides a surprising amount of power to global analyses of ENSO impacts, despite its simplistic nature.

El Niño events are associated with abnormally warm sea surface temperatures in the central and eastern equatorial Pacific, releasing large fluxes of thermal energy into the atmosphere. This warming of the tropical Pacific free troposphere induces warming throughout the tropical free troposphere, generally stabilizing the air column to vertical motions, inhibiting rainfall and warming the surface. For a fully developed discussion of these dynamics, see Sarachik and Cane[51] (2010). While there are many other local impacts of ENSO variations with complex structure in space and time, the simplest and most general fingerprint throughout impacted regions, in comparison to rainfall[43] or cloud cover,[47] is near-surface warming induced by the increased static stability.[48] Furthermore, using surface temperatures to identify teleconnections benefits from the fact that temperature data is the most reliably collected atmospheric statistic and it is well modeled in reanalyses.[33]

2.2 Construction method

To construct our global partition of ENSO teleconnections, we examine whether reanalysis grid cells exhibit surface temperatures that are positively correlated with NINO3 on a monthly basis (with a two

month lag) for at least three months out of the year[i]. The details of this construction are as follows:

Let $NINO(m, y)$ = the value of the ENSO index in calendar month $m \in \{1, ..., 12\}$ and year $y \in Y \equiv \{1950, ..., 2004\}$. Similarly, let $T(x, m, y)$ be the surface temperature at location x, month m and year y. Let $\rho(x, m, L)$ be the correlation coefficient (over all years $y \in Y$) of $NINO(m, y)$ and $T(x, m + L, y)$ (ρ is plotted for each month in Fig. 2). Note that T lags NINO3 by $L > 0$ months to account for the fact that signals from the tropical Pacific where NINO3 is measured need some time to propagate and influence the rest of the world.[48] In the main analysis $L = 2$. Now let $\tilde{\rho}(x, m, L) = 1$ if $\rho(x, m, L)$ is positive and statistically different from zero at $\alpha = 0.1$; $\tilde{\rho} = 0$ otherwise. Then

$$M_{xL} = \sum_{m=1}^{12} \tilde{\rho}(x, m, L)$$

is the number of months that grid point x exhibits interannual surface temperature variations that are significantly correlated at a level ≤ 0.1 with NINO3 L months earlier.

Even with no mechanistic connection between surface temperatures at x and NINO3, most values of M_{xL} should be ≥ 1 since 1 in 10 random draws will be correlated at the $\alpha = 0.1$ level. If the monthly draws of temperature were all completely independent, $M_{xL} \geq 3$ would occur randomly 11% of the time[ii] ($M_{xL} \geq 5$ would occur 0.4% of the time). Independence is not a valid assumption in this case, but it is a useful benchmark. For a cutoff value R, a point x is denoted ENSO "teleconnected" if $M_{xL} \geq R$. We define a binary measure of "teleconnectedness" $V_x(L, R) = 1$ if $M_{xL} \geq R$ and $= 0$ otherwise. In the main analysis, $R = 3$. Fig. 3a displays V_x (for $L = 2$ and $R = 3$, the values used in the main analysis).

To estimate country-level teleconnection index $\bar{V}_i$ for country i, we take a weighted average of V_x over all points in country i, denoted by $x \in X_i$:

$$\bar{V}_i = \frac{1}{(\sum_{x \in X_i} w_x)} \sum_{x \in X_i} (V_x \times w_x).$$

Values for $\bar{V}_i$ will range from 0 in the least teleconnected countries to 1 for the most teleconnected countries. The histogram in Fig. 4b displays the distribution of $\bar{V}_i$ for the values used in the main analysis ($L = 2$, $R = 3$ and w_x = the population of pixel x in the year 2000[53]). Because it is so strongly bimodal, a partition into just two groups appears to be an excellent approximation. It is surely an attractive simplification. We assign those countries with $\bar{V}_i > 0.5$ to the teleconnected group

[i] In generating our partition we avoid using socioeconomic variables since these might themselves influence ACR, confounding our analysis.[52]

[ii] $\sum_{m=3}^{12} \left[0.1^m 0.9^{12-m} \binom{12}{m}\right] = 0.1109$

and those with $\bar{V}_i \leq 0.5$ to the weakly-affected group.

2.3 Robustness of the ENSO-teleconnection partition

In this subsection, we examine the sensitivity of our partition under different choices for w, L, and R. We find that the gross features of the global partition are unchanged for a range of reasonable parameter choices and we document that it agrees well with previous analyses of regional ENSO teleconnections.

Pixel Assignment Fig. 3 illustrates changes to the partition that occur when L and R are modified. Panel a depicts teleconnected pixels ($V_x(L,R) = 1$) as red for $L = 2$ and $R = 3$, the values used in the main analysis. Panel b is the same, except $L = 0$. Almost no features change, suggesting that the selection of L is inessential. Panel c displays the changes that occur when L is held fixed at 2 and R is changed. As described before, setting $R < 3$ would admit too many chance assignments, but perhaps $R = 3$ is also too low. When $R = 5$ is used, those pixels that are teleconnected are marked as green. Pixels that are teleconnected when a cutoff $R = 3$ is used but are no longer teleconnected when $R = 5$ are yellow. Pixels that are weakly-affected when either $R = 3$ or $R = 5$ are blue. Note that the red region in Panel a is the union of the green and yellow regions in Panel c. When R is increased from 3 to 5, only the group assignment of the yellow region changes. This region is small and represents the boundary of the moist tropics and the arid subtropics; it is the region where the annual monsoonal rains stop in their meridional translation and reverse direction. The primary effect of increasing R from 3 to 5 is to simply omit relatively dry pixels that do not have five rainy months in a normal year.

Country Assignment Fig. 4a displays countries assigned to the teleconnected group in red and the weakly-affected group in blue, using the values from the main analysis, $L = 2, R = 3$. The histogram in Panel b displays the distribution of the country-level ENSO teleconnection index $\bar{V}_i$ when pixels are aggregated using weights w_x that reflect the population of each pixel. The histogram in Panel c displays how this distribution changes if the weights are changed to reflect the area of each pixel; the overall distribution is virtually unchanged. The histogram in Panel d displays how the distribution changes if population weights are used but $R = 5$ instead of 3. Again, the structure hardly changes.

Negative temperature correlations When designating pixels and countries as teleconnected, we only examined whether surface temperatures were positively correlated with NINO3 or not. One might be concerned that we ignore significant negative correlations with NINO3 because such correlations are well known,[51] however most locations with negatively correlated temperatures are oceanic and not

continental.[48] If large continental regions were negatively correlated with NINO3, we might expect that ACR in those countries would have a response to ENSO that was opposite the response in the teleconnected group. To check that such a pattern does not influence our results, we document that only three countries in our sample ($< 2\%$) have most of their population in locations where surface temperatures are negatively correlated with NINO3 (see Fig. 5). Fiji, the Solomon Island and New Zealand are the three countries that could plausibly be coded as "negatively teleconnected," yet none experience any civil conflict during the period of observation. In addition, Northern Mexico, the western United States, eastern Russia and eastern Indonesia also contain locations where surfaces temperatures anticorrelate with NINO3, however these locations contain only a small fraction of these countries' total populations.

Agreement with previous analyses Several previous studies have illustrated different types of ENSO teleconnections using different environmental variables and statistical techniques.[43–46,48–50] The pixels we designate as teleconnected to ENSO (Fig. 1a in main text) is approximately the union of regions previously found to be teleconnected to ENSO. The seminal work by Ropelowski and Halpert[43,44] demonstrated that several regions in the tropics and subtropics exhibited rainfall anomalies (of both signs) in association with ENSO (see Fig. 21 p. 1625 in [43]). Nicholls[45] showed the dependence of rainfall throughout eastern Australia (see Fig. 4b p. 969 in [45]). Nicholson and Kim[46] illustrated impacts throughout sub-Saharan Africa by utilizing more complete data sets than earlier work could access (see Fig. 7 p. 125 in [46]). Chiang and Sobel[48] explained the propagation of teleconnections with Kelvin wave dynamics and demonstrated their results using 1000-200mb temperature anomalies, which were correlated with ENSO throughout the tropics and some of the subtropics (see Fig. 2 p. 2618 in [48]). Giannini, Saravanan and Chang[49] demonstrated the dependance on ENSO of rainfall and surface temperatures in the Sahel (see Fig. 4E-F p. 1029 in [49]). A large number of other studies have illustrated flooding or ecological responses in more limited regions, many of which are summarized in Rosenzweig and Hillel.[50]

Our binary teleconnection partition is a dramatic simplification of the rich relationships studied in earlier work. For many purposes it would be an oversimplification, but it serves our purposes here by providing a structurally simple inclusive description of worldwide ENSO impacts. The binary assignment is justified *a posteriori* by the sharpness of the division. However, this strong generalization does not allow us to identify the specific mechanisms that are driving our main findings.

2.4 Partition validation using weather and agricultural outcomes

As a validation exercise of our global partition of ENSO teleconnection, we explore the effect of ENSO on surface temperatures, rainfall and agricultural output (Table 2) because these relationships are well established.[43, 48, 50, 51, 54]

Temperature and rainfall First we check that temperature and rainfall at the country-level are correlated with interannual variations in ENSO in the teleconnected group and not in the weakly affected group. The partitioning technique was designed to isolate those countries that are strongly influenced by ENSO, however it is possible that countries that have temperatures positively correlated with NINO3 for three months also exhibit negative correlations in other months or do not exhibit annually averaged signals for some other reason. Moreover, it is important to verify that rainfall patterns are negatively correlated with NINO3 in the teleconnected countries because rainfall was not explicitly used in the construction of the partition. (Rainfall was not used in the partition construction because rainfall signals are more variable and rainfall data is both less complete and noisier.)

To verify our partition with weather data, we estimate

$$W_i(t) = \beta NINO3(t) + \gamma_i + \theta_{i1}t + \theta_{i2}t^2 + \epsilon_i(t) \tag{1}$$

where $W_i(t)$ is either temperature or rainfall for country i in year t, $NINO3$ is averaged May-December, γ_i is a country-specific constant (fixed effect) and θ_{i1} and θ_{i2} are country-specific linear and quadratic time trends. This equation is estimated once for each country group and the resulting values for β are in rows 1 and 2 of Table 2. In the teleconnected group, annual average temperature (precipitation) is significantly and positively (negatively) correlated with the dominant NINO3 signal in each year. In the weakly affected group, the estimated values of β are near zero and statistically insignificant.

Cereal yields Next we compare inter-temporal variations in cereal yields for individual countries in the two groups with inter-temporal variations in ENSO. Using a longitudinal dataset of all countries between 1961-2007 we estimate the following dynamic model:

$$\log(Y_i(t)) = \beta_1 NINO3(t) + \beta_2 \log(Y_i(t-1)) + \gamma_i + \theta_{i1}t + \theta_{i2}t^2 + \epsilon_i(t) \tag{2}$$

where $Y_i(t)$ is the *cereal yield* for country i in year t. The trend terms are intended to account for technological innovation such as the "green revolution". Equation 2 is estimated separately for the teleconnected and weakly-affected groups. Standard errors are clustered by country[55–57] to account

for unknown patterns of within-country autocorrelation. Row 3 of Table 2 presents the regression coefficient β and associated standard errors for the two groups. For the teleconnected group, the coefficient is significantly negative while for the weakly-affected group the coefficients are positive but insignificant, a result that is consistent with observed increases in rainfall for some locations in the mid-latitudes.[50]

Agricultural income Finally, we check our teleconnection partition against agricultural outcomes by using agricultural income data from another source.[38] In row 4 of Table 2, we estimate equation 2 using country-level agricultural income and find that relative to the impact of ENSO on the teleconnected group, the drop in agricultural income for weakly affected countries is smaller and not significant. We also aggregate a single time-series of total agricultural income per capita A for the entire teleconnected and weakly-affected groups for 1970-2007. For each group, this number represents the total value of all agricultural output for roughly half of the world population (see Fig 6). Specifically, for each group, we estimate the following auto-regressive model

$$A(t) = \beta_0 + \beta_1 NINO3(t) + \beta_2 A(t-1) + \theta_1 t + \theta_2 t^2 + \epsilon(t). \quad (3)$$

Row 5 of Table 2 shows the regression coefficient β_1. Consistent with the other panels in Table 2, we find that in the teleconnected group, an increase in NINO3 leads to large and significant negative impacts in agricultural income while the coefficient for the weakly-affected group is smaller and not statistically significant. (We have verified that these different agricultural responses are also robust across NINO12 and NINO4 indices, results that are available on request.)

These results, which are consistent with previous regional and local-scale analyses,[43,48,50,51,54] broadly validate our global partition of ENSO teleconnections.

3 ENSO timing and measurement error

Annual conflict onset data is organized by years that begin in January and end in December.[35,36] NINO indices are collected for individual months and therefore must be aggregated into years in order to match the conflict data. The simplest approach would be to average NINO indices over years that begin in January and end in December, producing measures of ENSO that are exactly contemporaneous with the conflict onset observations. However, ENSO "events" do not begin in January nor end in December. ENSO events generally begin in May/June and persist until they break down in March/April of the

following year. For this reason, the period April-May is often termed the "spring barrier" and separates "tropical years" from one another. Fig. 1b in the main text illustrates the spring barrier by plotting monthly correlations between NINO3 values in December and other months. Importantly, the values of the NINO3 index in January-April of *Calendar Year t* are unrelated to the dominant event observed in *Tropical Year t*.

To demonstrate that ENSO events are organized this way, Table 3 tabulates the correlation coefficients for monthly NINO3 values over the period 1950-2008. Entries describe the correlation between the NINO3 values that are observed in two months of the same calendar year. The first three columns, marked *early season*, have values near one for the first three rows, indicating that months prior to the spring barrier have NINO3 values that are highly correlated with one another. The bottom seven rows have values close to zero, indicating that NINO3 values before and after the spring barrier are not correlated. In any given year, the annual average NINO value is dominated by measurements following the spring barrier. Thus, if an El Niño or La Niña "event" occurs, the calendar year in which it begins (referred to as "Year 0" in the climate literature) is the calendar year that exhibits the largest annually averaged signal. These events are generally coherent across the months June-December, as indicated by the high values in the last seven columns of Table 3. The correlation coefficients associated with spring barrier months (April and May) have intermediate values as this is the transition period between "tropical years".

Because El Niño or La Niña events are dominated by signals following the spring barrier, including early season measurements that are uncorrelated with these late season signals is equivalent to introducing noise into estimates of a given year's dominant climatology. Therefore, we omit early season months and estimate the ENSO state of the global climate by averaging NINO index values over May-December only. The NINO index values for January-April of the following year are omitted, despite their occurrence in the same "tropical year," because they follow all conflicts that are recorded in the matching calendar year.

Table 6 demonstrates the result when the spring barrier is ignored in the "naive" approach described above. Panel f displays the coefficients and standard errors when ACR for the teleconnected group is regressed on annual average NINO indices using all months, ignoring the existence of the spring barrier. Contrast this with Panel c displaying the results when index values for January-April are omitted from the annual averages of NINO. Including January-April NINO values reduces the magnitude of the coefficients by about 11-16% and increases the size of the standard errors by 13-49%. Thus, a "naive" estimate ignoring the timing of ENSO events would suggest that there is no statistically significant relation between ENSO and conflict onset. Tables 3, 4 and 6 suggest that the inclusion

of January-April NINO values lead to "attenuation bias," a well known statistical problem associated with classical (additive) measure error in a regressor variable.[30,52]

In addition, recall that Fig. 2c of the main text displays the within-year distribution of additional conflicts associated with El Niño-like conditions. It demonstrates that conflict onsets preceding the spring barrier are not driving our main result.

4 Main statistical models

Table 1 in the main text presents the main results of this study: a shift in the global climate from a strong La Niña to a strong El Niño increases the probability of conflict onset in the teleconnected group from 3% to 6% whereas the probability of conflict onset in the weakly-affected group remains unchanged at 2%. Estimating the magnitude of these changes requires the use of statistical models, which are detailed here.

Framework: Let $C_i(t)$ be a random variable that takes a value of one if country i experiences *conflict onset* in year t and zero otherwise. We define the *annual conflict risk* of country i (ACR_i) to be the probability that i experiences a conflict in year t conditional on the state of the world in year t:

$$ACR_i(t) = \mathrm{E}[C_i(t)|t].$$

For a set of years, we can estimate the conditionally expected value of $ACR_i(t)$, given an observable statistic $X(t)$ describing the state of the world in year t,

$$\mathrm{E}_t[ACR_i(t)|X] = f(X(t))$$

by applying data to the regression

$$ACR_i(t) = f(X(t)) + \epsilon_i(t) \tag{4}$$

where $\epsilon_i(t)$ is the component of $ACR_i(t)$ that is not predicted by $X(t)$. By substitution and iterated expectations we also have

$$\mathrm{E}_t[C_i(t)|X] = f(X(t))$$

which we can estimate with data in the regression

$$C_i(t) = f(X(t)) + \epsilon_i(t). \tag{5}$$

Equation 4 is the basis for our time series models and Equation 5 is the basis for our longitudinal (panel) models. In theory, both approaches estimate the same function $f(X(t))$.

Time series model of *conflict risk*: In our primary results, presented in main text Table 1 row 3, we estimate a group's average risk as

$$\begin{aligned} ACR(t) &= \frac{conflict\ onsets\ in\ group(t)}{countries\ in\ group(t)} \\ &= \frac{\sum_i C_i(t)}{N(t)} \end{aligned}$$

and then estimate the time-series model

$$ACR(t) = \alpha + \beta NINO(t) + \theta t + post\_Cold\_War(t) + \epsilon(t) \tag{6}$$

where α is a constant, θ is a linear trend and $post\_Cold\_War(t)$ is a constant term for all years following 1989 (inclusive) and zero otherwise. The coefficient of interest is β, which describes how many more conflict onsets per country are associated with a 1°C increase in NINO. This model is estimated once for the teleconnected group and once for the weakly-affected group. In each case, the observational unit is a "group-year" and the comparison is between the time series of NINO and the time series of a group's *annual conflict risk*, each of which has 54 observations. The linear trend captures slow increases in overall conflict risk and NINO. It also ensures that the estimated coefficients only represent high-frequency variations in ENSO and not correlated trends. The introduction of a post Cold War constant is a common practice in statistical analyses of conflict[58] because overall conflict levels changed qualitatively following the collapse of the Soviet Union. The standard errors presented for these time-series models are White standard errors and are robust to heteroskedasticity of arbitrary form.[31]

One drawback of Equation 6 is that country-level trends in conflict cannot be accounted for (eg. some countries may have become more conflict-prone over the observation period), although the stationarity of NINO indices suggest that this is not a major threat to unbiased estimation of β. However, the entry of new countries into our sample (see Fig. 1) may bias our estimates if new countries are systematically different from older countries (eg. they may be more violent) and they enter the sample

differentially during different ENSO states (eg. during more El Niño-like conditions). It is for these reasons that we check our estimates with the following second estimation procedure.

Longitudinal linear probability model of *conflict risk*: Rows 5 and 6 of Table 1 in the main text presents the coefficients from a country-level panel data linear probability model. These results come from the following specification:

$$C_i(t) = \beta NINO(t) + \mu_i + \theta_i t + \epsilon_i(t) \tag{7}$$

where $C_i(t)$ is unity if country i begins a conflict in year t and zero otherwise, μ_i is a country-specific constant and θ_i is a country-specific trend. Unlike the time-series model in Equation 6, the unit of observation for this model is a "country-year". For each country and year combination, there is one observation for a total of 3978 observations in the teleconnected group and 3400 observations in the weakly-affected group. The advantage of this technique is that it allows us to remove country-level trends in violence with the term θ_i because each country is observed multiple times. It also allows us to control for new countries that enter the sample during different ENSO states: the country-specific term μ_i removes any time-invariant characteristics of a country that might make it more or less conflict prone.

One disadvantage of the longitudinal data approach is that it is more difficult to calculate appropriate standard errors for our estimates of β. First, there may be serial correlation in conflict onset that are observed for an individual country. For example, countries may go through violent periods when they experience several conflict onsets over a short period. Second, all countries in a region are exposed to the same NINO index in a given year. If there are other factors that give rise to spatial correlation in conflict onsets[iii], then our observations will not be independent and we will underestimate the size of our standard errors.[52,57] We address these concerns by computing standard errors

[iii] Previous analyses[59] have examined "contagion" as a process that may generate spatial correlations in conflicts. If such contagion occurs in response to ENSO, it will be captured by our parameter β in Equations 6-7. To see this, let the ACR of country i be a function of both ENSO and the conflict state of i's neighbors:

$$ACR_i(t, \mathbf{C}_{\neq i}(t)) = \psi_i NINO(t) + \sum_{j \neq i} f(C_j(t), D_{ij}) + \epsilon_i(t)$$

where $\mathbf{C}_{\neq i}$ is a vector describing the conflict state of all countries other than i, $f(.)$ is some function describing how conflicts spread across countries and D_{ij} is some "distance" metric that describes how easily conflicts in j can influence ACR in i. Because $\mathbf{C}_{\neq i}$ is stochastic, we take expectations over its possible states conditional on t

$$\begin{aligned} \mathrm{E}_{\mathbf{C}_{\neq i}}[ACR_i(t, \mathbf{C}_{\neq i}(t))|t] &= \mathrm{E}_{\mathbf{C}_{\neq i}}\left[\psi_i NINO(t) + \sum_{j \neq i} f(C_j(t), D_{ij}) + \epsilon_i(t)|t\right] \\ &= \psi_i NINO(t) + \sum_{j \neq i} \mathrm{E}_{\mathbf{C}_{\neq i}}[f(C_j(t), D_{ij})|t] \end{aligned}$$

since ϵ is mean zero by construction. Let us denote $ACR_i(t) = \mathrm{E}_{\mathbf{C}_{\neq i}}[ACR_i(t, \mathbf{C}_{\neq i}(t))|t]$. Taking the total derivative of

using the generalized method of moments proposed by Conley[32] (1999) to account for unknown forms of spatial correlation. Because ENSO is regional to global in scale, we allow for spatial correlations in errors over distances up to 5000 km. In addition, we allow for serial correlation over periods less than six years and heteroskedasticity of unknown form.[31,42]

A second issue associated with this longitudinal data model is that the variable $C_i(t)$ can only take on the values of zero or one (i.e. it is a "binary response" model).[30] Several methods have been developed to estimate probabilities when only a binary outcome is observable. The linear model in Equation 7 represents the simplest of these models and is the easiest to interpret. However, when using this type of model, two issues must be examined to ensure that a linear model is appropriate. First, the predicted probabilities of an event should not be lower than zero nor higher than one. Figure 8 plots the predicted ACR for both regions (labeled "OLS") over the values of NINO3 that are observed. All predicted ACR values are well within the unit interval. Second, the probability response function should be well approximated by a linear function. Probit and logit models are two commonly used models that have been developed to deal with the type of non-linearity commonly observed in probability response functions.[30] Fig. 8 also plots the predicted ACR using these two other methods (labeled "Logit" and "Probit"). These response functions are indistinguishable from the linear model, suggesting that linearity is a good approximation.

The similarity of the results between rows 3 and 5 of Table 1 in the main text provides strong support for our main results. Two estimation procedures with different strengths and weaknesses

$ACR_i(t)$ and linearizing we obtain

$$\begin{aligned}\frac{\mathrm{d}ACR_i(t)}{\mathrm{d}NINO(t)} &= \frac{\partial ACR_i(t)}{\partial NINO(t)} + \sum_{j\neq i}\frac{\partial f(C_j,D_{ij})}{\partial C_j}\frac{\partial \mathrm{E}_{\mathbf{C}_{\neq i}}[C_j(t)|t]}{\partial NINO(t)} \\ &= \frac{\partial ACR_i(t)}{\partial NINO(t)} + \sum_{j\neq i}\frac{\partial f(C_j,D_{ij})}{\partial C_j}\frac{\partial ACR_j(t)}{\partial NINO(t)} \\ &= \psi_i + \sum_{j\neq i}\frac{\partial f(C_j,D_{ij})}{\partial C_j}\psi_j \\ &= \beta_i\end{aligned}$$

Thus, the estimated effect of ENSO on the ACR of i (β_i) is the direct effect of ENSO on i (ψ_i) plus the indirect effects of ENSO that are transmitted to i from i's neighbors j via contagion. Previous studies[59] that were concerned with isolating [an analogue to] the direct effect ψ_i had to make additional assumptions regarding D and $f(.)$, however our approach (Equations 6-7) does not require these structural assumptions and is suitable when the total average effect of ENSO ($\beta = \mathrm{E}_i[\beta_i]$) is the parameter of interest. In addition, it is worth noting that once β_i is estimated in our regression model

$$\begin{aligned}ACR_i(t) &= \beta_i NINO(t) + \epsilon_i(t) \\ &= \left[\psi_i + \sum_{j\neq i}\frac{\partial f(C_j,D_{ij})}{\partial C_j}\psi_j\right]\cdot NINO(t) + \epsilon_i(t)\end{aligned}$$

the residuals $\epsilon_i(t)$ do not contain any "contagion effect" that is driven by ENSO, since β captures these effects. Thus, while $\epsilon_i(t)$ may exhibit spatial correlation due to other factors, such as regional geopolitics, this spatial correlation will not be driven by ENSO. Nonetheless, this non-ENSO-related spatial correlation reduces the effective number of degrees of freedom in our data, motivating our use of standard errors that are robust to unknown forms of spatial autocorrelation.[32]

provide almost identical results.

Note on 1989 In the regressions presented in the main text, observations from 1989 are dropped. Fig. 10a illustrates why we have done this. It plots the residuals for a regression in row 2 of Table 1 from the main text. The number of conflicts in 1989 for the teleconnected group was three standard deviations from the conditionally expected value. It seems likely that the large number of conflicts in that year were related to events associated with the end of the Cold War. Yet, we note that once a constant for the post-Cold War period is included in Equation 6 (a common technique in the literature[58]) our main result appears robust to the reintroduction of 1989 to the sample (see Table 6 Panel e). Fig. 10b plots the estimated residuals using this second model.

Non-parametric estimates of conflict risk: Fig. 2b in the main text provides non-parametric estimates of our time-series model. Non-parametric techniques are employed to (1) validate that linearity between NINO3 and conflict risk is a reasonable assumption, and (2) obtain a better sense of local behavior around different parts of the NINO3 distribution. For Fig. 2b, a Nadaraya-Watson estimator[60,61] was fit to the data using a Epanechnikov kernel with a 1°C bandwidth. The 90% confidence intervals are bootstrapped.

Hierarchical regression model of ENSO teleconnection and income: Fig. 3 in the main text uses a hierarchical regression model to decompose conflict risk for each country into ENSO unaffected (α) and ENSO-affected (β) components. Specifically, for each country i, we run the regression:

$$C_i(t) = \alpha_i + \beta_i NINO3(t) + \epsilon_i(t) \tag{8}$$

where the variables are defined as they were in Equation 7. We next conduct a local non-parametric estimate for the relationship between income per capita and the two estimated terms α_i and β_i. Fig. 3a in the main text plots α_i against log income per capita in 2007 using a Epanechnikov kernel with a bandwidth of 0.6. The 90% confidence intervals are bootstrapped. Fig. 3b is the same, but for β_i.

5 Robustness of the main result

Earlier, we showed that the bimodal distribution of the country-level teleconnection index is largely insensitive to parameters in the construction of the global teleconnection partition. Here we further test the robustness of our results by using alternative NINO indices, explicitly modeling potential serial

correlations, altering the sample of years used in the model, including a large number of potentially confounding control variables, estimating a distributed lag model, altering the definition of ACR, examining whether our results hold outside of Africa and modeling conflict onset as a Poisson point process. Our main findings survive all of these checks.

NINO index Table 6 provides results when NINO12, NINO3 and NINO4 are used to check that our results are robust to different measures of the global ENSO state. While the coefficients in thise table may look as if they change substantially as the NINO index changes, this is partly driven by differences in the scale of variation observed for each index. These coefficients are in units of $\frac{\%\ year^{-1}}{1^{\circ}\mathrm{C}}$, however the range of degrees Celsius over which each index varies is not fixed. Table 1 presents the standard deviations for the three NINO indices. When the coefficients from our preferred specification (panel d in Table 6) are converted to units of $\frac{\%\ year^{-1}}{1\ standard\ deviation}$, they look similar, ranging from 0.51-0.47.

Serial correlation In our time series analysis of ACR and NINO3, one might be concerned that patterns of serial correlation in one or both variables could be affecting our central findings, however we do not find that this is the case. Table 7 presents our baseline results in models 1 and 6. The Durbin-Watson d-statistic for the teleconnected group rejects the null hypothesis of serial correlation in errors, however the same test for the weakly affected group fails to reject the null. When the residuals from the baseline model are regressed on their lagged values (models 2 and 7) we again find no serial correlation in the teleconnected group but some serial correlation in the weakly affected group. When we explicitly add lagged values of ACR to the baseline model (models 3 and 8) we find no substantive change in our coefficient of interest. We also model changes in ACR as a response to changes in NINO3 (models 4 and 9) and find that our results are not statistically different from the baseline model. Finally we re-estimate our standard errors for our baseline model using a Newey-West estimator[29] that is robust to both serial correlation and heteroscedasticity (rather than the White estimator[31] that we used previously and is only robust to heteroscedasticity). Regardless of whether we use a lag cutoff of three, five or ten years[iv], our estimated standard errors and significance levels do not change in a meaningful way.

Sample selection In our main analysis, we focus on the period 1950-2004 and omit the observation in 1989. We begin our analysis in 1950 because the post-war years were extraordinary, containing two extreme outliers (1946 and 1948). In addition, no other data sets except NINO reconstructions are

[iv] Greene[30] (2003) recommends using a cutoff length that is at least as large as the fourth-root of the number of observations, which in this case is 2.7 years.

available prior to 1950. We omit 1989 because it is also an influential outlier, probably related to the end of the Cold War. Retaining these obvious outliers in our analysis would likely bias our estimates, however we examine whether their inclusion in our analysis substantially changes our finding. In model 1 of Table 8 we present our baseline sample. In model 2 we add 1989 only and in model 3 we add the years 1946-1949 only. In model 4 we exclude only 1948 and 1989 (the largest outliers) and in model 5 we restrict the sample to years following 1975 (inclusive), when the sample of teleconnected countries stabilized. In model 6 we include all years in the Conflict Onset and Duration Dataset. In all of these samples, we find that the coefficient on NINO3 is statistically significant and statistically indistinguishable from our baseline model. However, these outliers exert a large and probably unreasonable influence on our model, so they remain omitted in our primary analysis.

Semi-parametric control variables Buhaug[58] (2010) suggested that the compelling longitudinal-based results in studies such as this should not depend critically on using country-specific constants (fixed-effects) or country-specific trends. Burke et al.[62] (2010) respond that such controls are essential to reliable statistical inference. It is our view that such controls are generally appropriate, however we check that our results are robust to their omission. Table 9 presents results from our longitudinal linear-probability model (Equation 7) with and without country fixed-effects (μ_i) and country-specific trends (θ_i). None of these alterations affects our estimated coefficient of interest, however the omission of country fixed-effects increases our estimated standard errors slightly.

Other control variables Because previous studies have identified many parameters that are correlated with conflict, it may be worth it to try and include these variables in our longitudinal analyses. We do this to check the robustness of our results in light of previous work, but do not present heavily-controlled regression results as our main findings because we feel that such control is not methodologically sound,[30,52] especially when our independent variable of interest (ENSO) is unquestionably exogenous. Moreover, our preferred model generates results that appear weaker than the following heavily-controlled models, suggesting that our preferred approach is the more conservative one.

We begin by introducing idiosyncratic country-level temperature and precipitation shocks that others[63–66] have suggested influence civil conflicts. ENSO may influence conflict through temperature and precipitation, however it may also affect conflict through a variety of other mechanisms including (but not limited to) the timing of rainfall, altered wind patterns, humidity, cloud cover, disasters, ecological events, or other environmental changes. It may also be the case that ENSO induced changes in temperature and precipitation have a fundamentally different impact on conflict than idiosyncratic

weather shocks because ENSO induced changes are experienced by a large number of countries. Thus, it is not surprising that when monthly temperature or rainfall are included in our longitudinal model, the coefficient on NINO3 remains large and statistically significant (Table 10). When rainfall is introduced to the model our coefficient of interest becomes larger, however this is partially an artifact of the subsample of years for which reliable global rainfall data is available. We do not believe that this is a well-specified model because ENSO is known to affect temperature and rainfall,[52] however we present it here for completeness.

We now turn to three of the most commonly suggested correlates of civil conflict: income,[64,67,68] population[66,69] and political institutions.[58] When all of our data is pooled together (Fig. 11) it appears that ACR increases with population, decreases with income and is greatest for countries that are "anocracies" (Polity IV scores near zero). We include lagged values for these controls one by one and jointly into our fixed-effects model and present the results in Table 11. In all of these models, the effect of ENSO on ACR is large and statistically significant. This contrasts with the weakly affected region, where correlations between ACR and NINO3 continue to be absent even in the presence of these controls (model 11). Further, when we stratify the teleconnected sample according to whether countries are in Africa or not (models 9-10), both regions exhibit similar results. This contrasts with the effects of income, population and Polity IV, all of which change in magnitude and/or sign between the African and non-African subsamples. Polity IV is the only control that exhibits a reasonably consistent and significant correlation with ACR, however the amount of variation it explains is small (Fig. 11) and it is known to be extremely endogenous.[52,62,64]

Finally, we estimate a "kitchen sink" model that contains the following controls. The paper that motivates each variable's inclusion is listed in as a citation.

1. country fixed-effects[70]
2. country-specific time trends[62]
3. log income per capita[67] (lagged)
4. income growth[64] (lagged)
5. Polity IV score[58] [linear & quadratic] (lagged)
6. agriculture industry share (%)[65] (lagged)
7. percent urbanized[70] (lagged)
8. log population[66] (lagged)
9. percent female[69] (lagged)
10. percent below 15 yrs old[69] (lagged)
11. percent above 65 yrs old[69] (lagged)

12. cyclone maximum windspeed[71] (area average)
13. monthly temperature[65] (12 variables)
14. monthly rainfall[63,64] (12 variables)

Results from this heavily controlled model are presented in Table 12. The full-sample estimates with or without the weather controls (rows 1 and 4) exhibit large and statistically significant correlations between ACR and NINO3 in the teleconnected group only. We check that a linear model for the ACR response to NINO3 is a good approximation for the data in this "kitchen-sink" model (Fig. 12) and observe that these results are not driven by outliers. When the sample is split into African and non-African countries, we again see large coefficients for all groups. Only one teleconnected coefficient (model 3) is not significant despite being larger in magnitude that the analogous coefficient in our baseline model (Table 1 in the main text, row 3). This is hardly surprising and occurs because our standard errors grow substantially with such dramatic "over-fitting" in our statistical specification.[30] Again, it is our view that including so many endogenously determined and/or irrelevant control variables is not the correct approach to causal inference since ENSO is known to vary over time exogenously. We only present these results as a robustness exercise.

Lag and lead NINO terms As discussed in the main text, we conduct tests to determine whether or not (1) the conflicts induced by ENSO in the teleconnected group would have occurred in its absence and (2) our main results might be spurious.

To check if ENSO simply advances inevitable conflicts, we add a one-year lagged NINO term into Eq. 6. Model 1 of Table 14 replicates our main result for the teleconnected group. Model 2 includes a lagged NINO3 term while model 3 also includes an interaction term between current NINO3 and lagged NINO3 because sequential ENSO events may have compounding impacts. Model 4 includes two lagged NINO3 terms. Observe that the only significant coefficients in the four models are for current NINO3. Columns 2 and 4 show that when lagged NINO3 terms are included, the coefficient on the lagged terms are of lesser magnitude than the current NINO3 coefficient. The point estimate suggests that about 40% of the observed conflicts might be displaced, however we cannot reject the null that no displacement occurs.

As an additional check against potentially spurious results, model 5 of Table 14 includes a future NINO3 term which, as expected, is not a statistically significant predictor of current ACR.

Columns 6-10 replicate the analysis for the weakly-affected group and do not yield any statistically significant coefficients.

Conflict size Next, we check whether our main result is driven by large or small conflicts and find that neither is dominating our result. The UCDP/PRIO Onset and Duration of Intrastate Conflict database separates conflict onsets into three categories: *large conflicts* that exhibit more than one-thousand battle-related deaths in a single year, *small conflicts* that exhibit more than twenty-five battle-related deaths in the year of onset but never exceed one-thousand cumulative deaths throughout the conflict, and *intermediate conflicts* that exhibit more than one-thousand cumulative battle related-deaths throughout the conflict, but never exceed one-thousand deaths in a given year. Because there are a very small number of *intermediate conflicts*, they are dropped from the following analysis (although they are included in the main analysis). Also note that the scale of conflicts is determined by the absolute number of battle related-deaths, not by the fraction of individuals in a country that are killed. Thus, a large conflict in a large country may not be of the same "intensity" as a large conflict in a small country. Nonetheless, we follow the existing literature and use these cutoffs[36,58,65] despite their imperfections.

To check that our main results are not driven by only small or large conflicts, we re-estimate Equation 7 for *small conflicts* only and then again for *large conflicts* only (Table 13 column 1, titled "Onset2"). The sum of the coefficients for small and large conflicts should be approximately equal to the coefficient of NINO3 in row 5 of Table 1 in the main text. If the coefficient for small conflicts were much larger (smaller) than the corresponding coefficient for large conflicts, than that would indicate that small (large) conflicts were driving our result in Table 1 of the main text. We find no evidence this is the case. The increase in conflict onset risk associated with a 1°C increase is identical for large (0.45%) and small (0.45%) conflicts.

Intermittency threshold Previous statistical analyses of conflict have demonstrated that altering the definition of "conflict onset" may substantially influence their results.[35,36,72] One parameter that has been shown to affect results is the "intermittency threshold" for conflicts.[72] Throughout our main analysis, if a specific conflict has been quiescent for at least two years and then becomes active again, it is coded as a new conflict onset. Borrowing the terminology of the UCDP/PRIO database, we call the binary variable that results from this coding rule $Onset2$.[35,36] We allow the number of quiescent years X to change when estimating Equation 7. Thus, $OnsetX$ denotes a new conflict onset that occurs only after X years of inactivity. Note that by construction all conflict onsets recorded in $OnsetX$ will also be recorded in $Onset(X-1)$, but the reverse is not true.

Table 13 shows how the results for large and small conflicts change when the intermittency threshold

is increased[v]. As the intermittency threshold is increased, there is almost no effect on the coefficients for small conflicts. However, increasing the intermittency threshold for large conflicts causes the NINO coefficients to fall in magnitude. The coefficients are about half as large when *Onset*4 is used and one-fourth as large when *Onset*9 is used, compared to *Onset*2. These results suggest that all types of small conflicts are affected by the global ENSO state, while there may be two types of large conflicts: one of which occurs infrequently and is associated with ENSO and one of which is more frequent and exhibits cycles of "flaring up" and "cooling down" in association with the global ENSO state.

Non-African countries Most previous studies examining the relationship between weather and conflict have focused on Africa.[58,64,65,73,74] We find that while the effect of ENSO on conflict is strong in Africa, it appears similarly strong and statistically significant on other continents as well. Row 6 of Table 1 in the main text displays the baseline longitudinal regression model using non-African teleconnected countries only and the estimated coefficient is almost identical to the average value (row 5). Models 9-10 in Table 11 display results for African and non-African countries in the longitudinal model with commonly used control variables and the coefficients for NINO3 in both subsamples are similar to the average value (model 8). Contrast this with the coefficients for log(GDP/capita) which are opposite in sign and significant for the two subsamples. Finally, Table 12 displays the "kitchen sink" models for both subsamples and finds that the coefficients on both subsamples remain larger than the baseline model (Table 1 row 4 in the main text). In one of the specifications (row 3) the coefficient for non-African countries is not statistically significant, but this appears to be due to an inflation in standard errors. This is probably a result of over-fitting the model with too many irrelevant variables.[30]

[v]One might be concerned that the existence of any intermittancy threshold may generate cyclical behavior in the ACR which could generate artificially large responses to the quasi-cyclical timing of ENSO events. To see why this type of bias is incapable of dramatically affecting our results, let

C_t = true ACR at time t

N_t = number of countries at time t

Assume conflict risk and the sample are stationary until the following year:

$$C_t = C_{t+1}, \quad N_t = N_{t+1}$$

We observe $C_t N_t$ conflicts in the data at time t. However, we only observe $C_{t+1}(N_{t+1} - C_t N_t)$ conflicts the following year because the $C_t N_t$ countries that had conflicts the preceding year were artificially excluded from observation (this is not true, as explained in the text, but suppose it were).

In year t we would accurately estimate conflict risk to be

$$\hat{C}_t = \frac{number\ of\ conflicts}{number\ of\ countries} = \frac{C_t N_t}{N_t} = C_t$$

however in the following year we would estimate

$$\hat{C}_{t+1} = \frac{number\ of\ observable\ conflicts}{number\ of\ countries} = \frac{C_{t+1}(N_{t+1} - C_t N_t)}{N_{t+1}} = C_t - C_t^2.$$

Therefore, the potential bias that would be introduced is of the magnitude C^2. Given that we estimate an average ACR of 4% in the teleconnected region, the introduced bias would be 0.16%. This bias is only $\frac{1}{25}$th of the signal's magnitude and is well below the estimated uncertainty of our preferred models ($\pm$0.84%).

Conflict as a Poisson point process An alternative approach to modeling conflict onsets is to assume that conflict onsets can be represented by a non-homogenous Poisson point-processes. Two techniques can be applied in this setting: count analysis and survival analysis. We implement both and verify that our central findings remain unchanged.

In our count analysis, we model the number of conflicts observed in the teleconnected group as if these values were Poisson distributed, with a mean (variance) parameter that changed in response to NINO3. We estimate the model:

$$\mathrm{E}\left[\sum_i C_i(t)\right] = e^{(\alpha+\beta NINO3(t))} \tag{9}$$

via maximum likelihood (notation is the same as in Equation 7). Here, we assume that the structure of disturbances is such that $\sum_i C_i(t)$ can be represented by a Poisson distribution. Figure 9 shows $\sum_i C_i(t)$ in the teleconnected group against $NINO3(t)$ for the period 1975-2005 when the sample of countries is roughly constant. The regression line is an ordinary-least-squares fit to the data while the circles display predicted values using the Poisson regression described in Equation 9. Observe that movements in the predicted mean are the same, regardless of which statistical model is used.

In our survival analysis, we model how long peaceful periods "survive" before a "failure," i.e. a conflict, occurs. Assuming that conflict onsets can be represented by a non-homogeneous Poisson point-process, we allow the hazard rate h (instantaneous ACR) to change over time in response to ENSO. We estimate the model:

$$h(t) = e^{(\beta NINO3(t)+\mu_i+\theta t)} \tag{10}$$

again by maximum likelihood (notation is the same as in Equation 7). Table 5 presents our estimates for both the teleconnected and weakly affected groups. In row 1 we present our estimates of the hazard rate increase associated with a 1°C increase in NINO3. We find that the hazard rate response in the teleconnected group is significant and matches our main result: a 1°C increase in NINO3 increases the risk of transition from a peaceful state into conflict by 25%. If one uses the temperature swing from La Nina to El Nino (3°C), the total risk of conflict doubles ($e^{0.25\times 3} = 2.1$). This magnitude matches our main results that we estimated with ordinary least squares (recall Fig 2 in the main text).

Finally, for a sense of scale, in row 2 of Table 5 we estimate how NINO3 influences the length of uninterrupted peaceful periods. This is an alternative but equivalent interpretation of Equation 10. We estimate that a 1°C increase in NINO3 reduces the length of an average peaceful period in the teleconnected group by 0.22 years. Thus, a shift from El Niña to La Niño (3°C) reduces the average

peaceful period by 0.66 years. Summed over all eighty countries in the teleconnected group, this same shift in NINO3 reduces the global number of peaceful country-years by 53 ($0.22 \frac{years}{1°C} \times 3°C \times 80\ countries = 52.8\ country\text{-}years$).

References

[29] Newey, W. K. & West, K. D. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* **55**, 703–708 (1987).

[30] Greene, W. H. *Econometric Analysis, Fifth Edition* (Prentice Hall, 2003).

[31] White, H. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* **48**, 817–838 (1980).

[32] Conley, T. GMM estimation with cross sectional dependence. *J. Econometrics* **92**, 1–45 (1999).

[33] Kalnay, E. *et al.* The NCEP/NCAR 40-year reanalysis project. *Bulletin of the American Meteorological Soc.* **77**, 437–471 (1996).

[34] Kaplan, A. *et al.* Analyses of global sea surface temperature 1856-1991. *J. Geophysical Research* **103**, 18,567–18,589 (1998).

[35] Gleditsch, N. P., Wallensteen, P., Eriksson, M., Sollenberg, M. & Strand, H. Armed conflict 1946-2001: a new dataset. *J. Peace Research* **39**, 615–637 (2002).

[36] Strand, H. Onset of armed conflict: A new list for the period 1946-2004, with applications. Tech. Rep., Center for the Study of Civil War (2006). URL `http://www.prio.no/CSCW/Datasets/Armed-Conflict`.

[37] Food and Agriculture Organization. FAOStat database (2009). `http://faostat.fao.org`.

[38] *National Accounts Statistics: Main Aggregates and Detailed Tables, 2007* (United Nations, 2007).

[39] *World Development Indicators 2008* (World Bank, 2008).

[40] Marshall, M. G., Gurr, T. R. & Jaggers, K. Polity IV project. Tech. Rep., Center for Systemic Peace (2009). URL `www.systemicpeace.org/polity/polity4.htm`.

[41] Xie, P. & Arkin, P. A. Analyses of global monthly precipitation using gauge observations, satellite estimates, and numerical model predictions. *J. Climate* **9**, 840–858 (1996).

[42] Hsiang, S. M. Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *PNAS* **107**, 15367–15372 (2010).

[43] Ropelewski, C. F. & Halpert, M. S. Global and regional precipitation patterns associated with the El Niño/Southern Oscillation. *Monthly Weather Rev.* **115**, 1606–1626 (1987).

[44] Ropelewski, C. F. & Halpert, M. S. Precipitation patterns associated with the high index phase of the Southern Oscillation. *J. Climate* **2**, 268–284 (1989).

[45] Nicholls, N. Sea surface temperatures and Australian winter rainfall. *J. Climate* **2**, 965–973 (1989).

[46] Nicholson, S. E. & Kim, J. The relationship of the El Niño-Southern Oscillation to African rainfall. *International J. Climatology* **17**, 117–135 (1997).

[47] Klein, S. A., Soden, B. J. & Lau, N.-C. Remote sea surface temperature variations during ENSO: Evidence for a tropical atmospheric bridge. *J. Climate* (1999).

[48] Chiang, J. C. H. & Sobel, A. H. Tropical tropospheric temperature variations caused by ENSO and their influence on the remote tropical climate. *J. of Climate* **15**, 2616–2631 (2002).

[49] Giannini, A., Saravanan, R. & Chang, P. Oceanic forcing of Sahel rainfall on interannual to interdecadal time scales. *Science* **302**, 1027–1030 (2003).

[50] Rosenzweig, C. & Hillel, D. *Climate Variability and the Global Harvest: Impacts of El Niño and Other Oscillations on Agro-Ecosystems* (Oxford University Press, 2008).

[51] Sarachik, E. S. & Cane, M. A. *The El Niño-Southern Oscillation Phenomenon* (Cambridge University Press, 2010).

[52] Angrist, J. D. & Pischke, J.-S. *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press, 2008).

[53] CIESIN, Columbia University. Gridded population of the world V3 (2009). URL `http://sedac.ciesin.columbia.edu/gpw`.

[54] Cane, M. A., Eshel, G. & Buckland, R. W. Forecasting Zimbabwean maize yield using eastern equatorial pacific sea surface temperature. *Nature* **370**, 204–205 (1994).

[55] Liang, K.-Y. & Zeger, S. L. Longitudinal data analysis using generalized linear models. *Biometrika* **73**, 13–22 (1986).

[56] Arellano, M. Computing robust standard errors for within-groups estimators. *Oxford Bulletin of Economics and Statistics* **49**, 431–434 (1987).

[57] Bertrand, M., Duflo, E. & Mullainathan, S. How much should we trust differences-in-differences estimates? *The Quarterly J. Economics* **119**, 249–275 (2004).

[58] Buhaug, H. Climate not to blame for African civil wars. *PNAS* **107**, 1647716482 (2010).

[59] Sandholt, J. P. & Gleditsch, K. S. Rain, growth, and civil war: The importance of location. *Defense and Peace Economics* **20**, 359–372 (2009).

[60] Nadaraya, E. A. On estimating regression. *Theory of Probability and Its Applications* **9**, 141–142 (1964).

[61] Watson, G. S. Smooth regression analysis. *Sankhya* **26**, 359–372 (1964).

[62] Burke, M., Dykema, J., Lobell, D., Miguel, E. & Satyanath, S. Climate and civil war: is the relationship robust? *NBER working paper 16440* (2010). URL `http://www.nber.org/papers/w16440`.

[63] Levy, M. A., Thorkelson, C., Vorosmarty, C., Douglas, E. & Humphreys, M. Freshwater availability anomalies and outbreak of internal war: Results from a global spatial time series analysis. *International Workshop for Human Security and Climate Change* (2005).

[64] Miguel, E., Satyanath, S. & Sergenti, E. Economic shocks and civil conflict: An instrumental variables approach. *J. Political Economy* **112**, 725–753 (2004).

[65] Burke, M., Miguel, E., Satyanath, S., Dykema, J. & Lobell, D. Warming increases risk of civil war in Africa. *PNAS* **106**, 20670–20674 (2009).

[66] Bruckner, M. Population size and civil conflict risk: is there a causal link? *The Economic Journal* **120**, 535–550 (2010).

[67] Collier, P. & Hoeffler, A. Greed and grievance in civil war. *Oxford Economic Papers* **4** (2004).

[68] Blattman, C. & Miguel, E. Civil war. *J. Economic Literature* **48**, 3–57 (2010).

[69] Urdal, H. Population, resources and violent conflict: A sub-national study of India 1956-2002. *J. Conflict Resolution* **52**, 590–617 (2008).

[70] Fearon, J. D. & Laitin, D. D. Ethnicity, insurgency and civil war. *American Political Science Rev.* **97**, 75–90 (2003).

[71] Barron, P., Kaiser, K. & Pradhan, M. Local conflict in Indonesia: Measuring incidence and identifying patterns. Tech. Rep. 3384, World Bank (2004).

[72] Hegre, H. & Sambanis, N. Sensitivity analysis of empirical results on civil war onset. *J. Conflict Resolution* **50**, 508–535 (2006).

[73] Hendrix, C. S. & Glaser, S. M. Trends and triggers: Climate, climate change and civil conflict in Sub-Saharan Africa. *Political Geography* **26**, 695–715 (2007).

[74] Meier, P., Bond, D. & Bond, J. Environmental influences on pastoral conflict in the Horn of Africa. *Political Geography* **26**, 716–735 (2007).