EC3320

2016-2017

Michael Spagat

Lecture 12

The paper by Ward, Greenhill and Bakke (WGB) is one of my favourite papers in the course. It is very readable and useful for you.

So far we have focused on statistical and practical significance ("practical" significance is about whether effect sizes are big enough so that we should care about them).

Many people assume that coefficients that are both statically and practically significant are also important for making predictions - ***but this is not true***. The WGB paper develops this important point.

WGB analyses the Fearon and Laitin (FL) and Collier and Hoeffler (CH) papers covered in lecture 11 to make their points. The second section of the WGB paper has a nice summary of the “story time” discussion that surrounds FL and CH so I will briefly summarize the WGB summary before moving on to prediction.

First we look at areas where FL and CH agree:

1. They both argue that the longstanding emphasis of conflict scholars and journalists on the importance of *ethnic conflict* and *income inequality* for explaining armed conflict is misplaced. These variables do not come out significant in the FL-CH (logistic) regressions. CH use the term “grievance” and conclude that grievances are not important for explaining why conflicts break out. FL and CH point out that grievances are ubiquitous yet armed conflicts are not ubiquitous. This is an awkward fact for theories that try to use grievance to explain conflict.

2. They both try to resolve this tension by arguing that the *opportunity* to rebel, rather than a desire to rebel, is the determining factor on whether or not a conflict breaks out.

Then FL and CH split:

1. CH think that natural resources are the key. They argue that having an accessible supply of valuable resources that can be easily extracted and sold gives rebel groups a feasible way to finance themselves.

2. FL think that the opportunity for rebellion arises when a state is weak. Various factors can make a state weak including:

a. Low GDP which makes it hard for a government to raise revenue to support, e.g., an army.

b. Lots of territory or a large population which makes it hard for a government to control a whole large country.

c. Mountainous terrain which, again, makes territory hard to control.

We now return to the prediction question.

The table on slide 6 shows that both models are bad at predicting civil wars. You need to understand a few things before you can understand this table as well as possible so we now pause for a moment to acquire this knowledge.

First, when people make predictions using a logistic model they normally plug numbers in (just like we did for the Rwanda genocide paper). When the predicted probability is greater than 0.5 they predict that the event will happen and when the predicted probability is less than 0.5 they predict that the event will not happen.

However, this 0.5 *threshold* is arbitrary. We could predict war when the model spits out a number bigger than 0.3, or any other number between 0 and 1 if we like.

Next, after we make a prediction we can think of there being four possible outcomes:

1. We predict “war” and “war” happens. We can call this a *true positive*.

2. We predict “no war” and “no war” happens. We can call this a *true negative*.

3. We predict “war” but war does not happen – “false positive”.

4. We predict “no war” but war happens – “false negative”.

This is the main table in the WGB paper.

Table III. Number of correctly predicted onsets and false positives at varying cut-points

	Fearon & Laitin model		*Collier & Hoeffler model*	
Threshold	*Correctly predicted*	*False positives*	*Correctly predicted*	*False positives*
0.5	0/107	0	3/46	5
0.3	1/107	3	10/46	20
0.1	15/107	66	34/46	110

When the threshold is 0.5 then FL never predict war, i.e., each case they feed into their model gives a predicted probability of war that is less than 0.5. This means that they fail to predict all 107 civil war onsets that occur in their dataset. On the plus side they have no false positives, i.e., they never predict a war that does not happen.

CH do better on correctly predicted civil wars, 3 predicted wars out of 46 occurring in their data, but at the cost of 5 false positives (predictions of wars that do not happen).

Note on terminology – WGB’s column called “Correctly predicted” corresponds to what I call “true positives”.

The table also shows that both true and false positives increase as you lower the threshold. This movement is intuitive and a mathematical necessity – when you predict war more frequently you will make more good calls but at the expense of raising more false alarms.

The table does not display false and true negatives but you can calculate these once you know that FL have a total of 6295 cases of “no war” and CF have 642 cases of “no war”. (Recall that CF use 5-year periods whereas FL use 1-year periods. The resulting difference accounts for much of the difference in sample size.

The picture on the next slide summarizes the trade-off between the true and false positive rates for FL-CH.

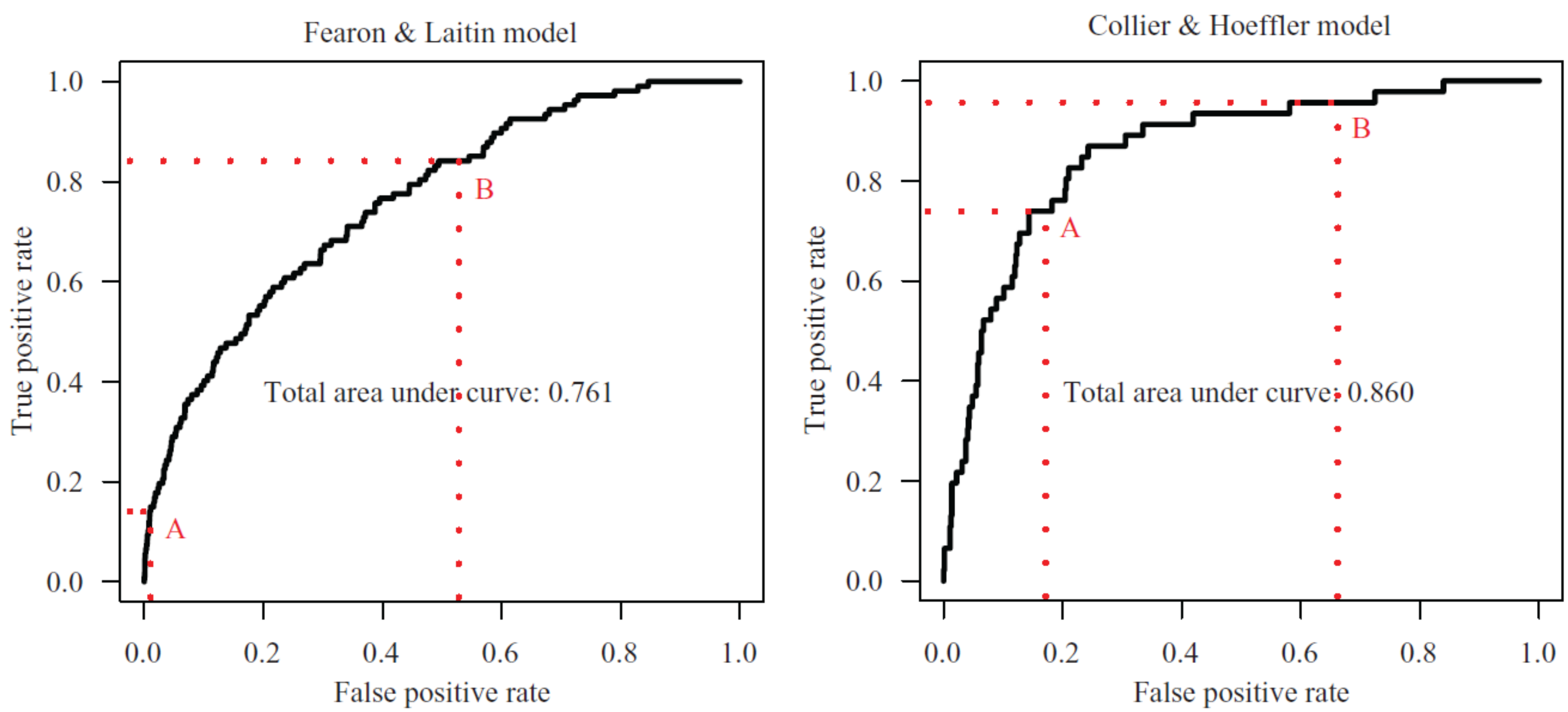

Figure 1. ROC plots

The X axes give the fraction of cases of peace for which the models predict war. The Y axis gives the fraction of cases of war for which the model predicts war.

CH do a little better on prediction than FL do. For example, when CH have a false positive rate of 0.2 their true positive rate is around 0.8 while FL only make it up to around 0.5.

One way to summarize the performance of a model at prediction is *total area under the curve*.

You can see this intuitively by looking at the picture on the previous slide. You want the true positive rate to be as close to 1.0 as possible for every value of the false positive rate. If it is always near 1.0 then the area under the curve will be close to 1.0.

The area under the curve decreases as the true positive rates decrease.

The following picture gives the relationship between statistical significance and performance at prediction for variables in the FL and CH models where the latter is measured by how much the area under the curve decreases if that variable is removed from the logistic regression. You can see that the relationship between the two criteria is not very tight. They are really two different things.

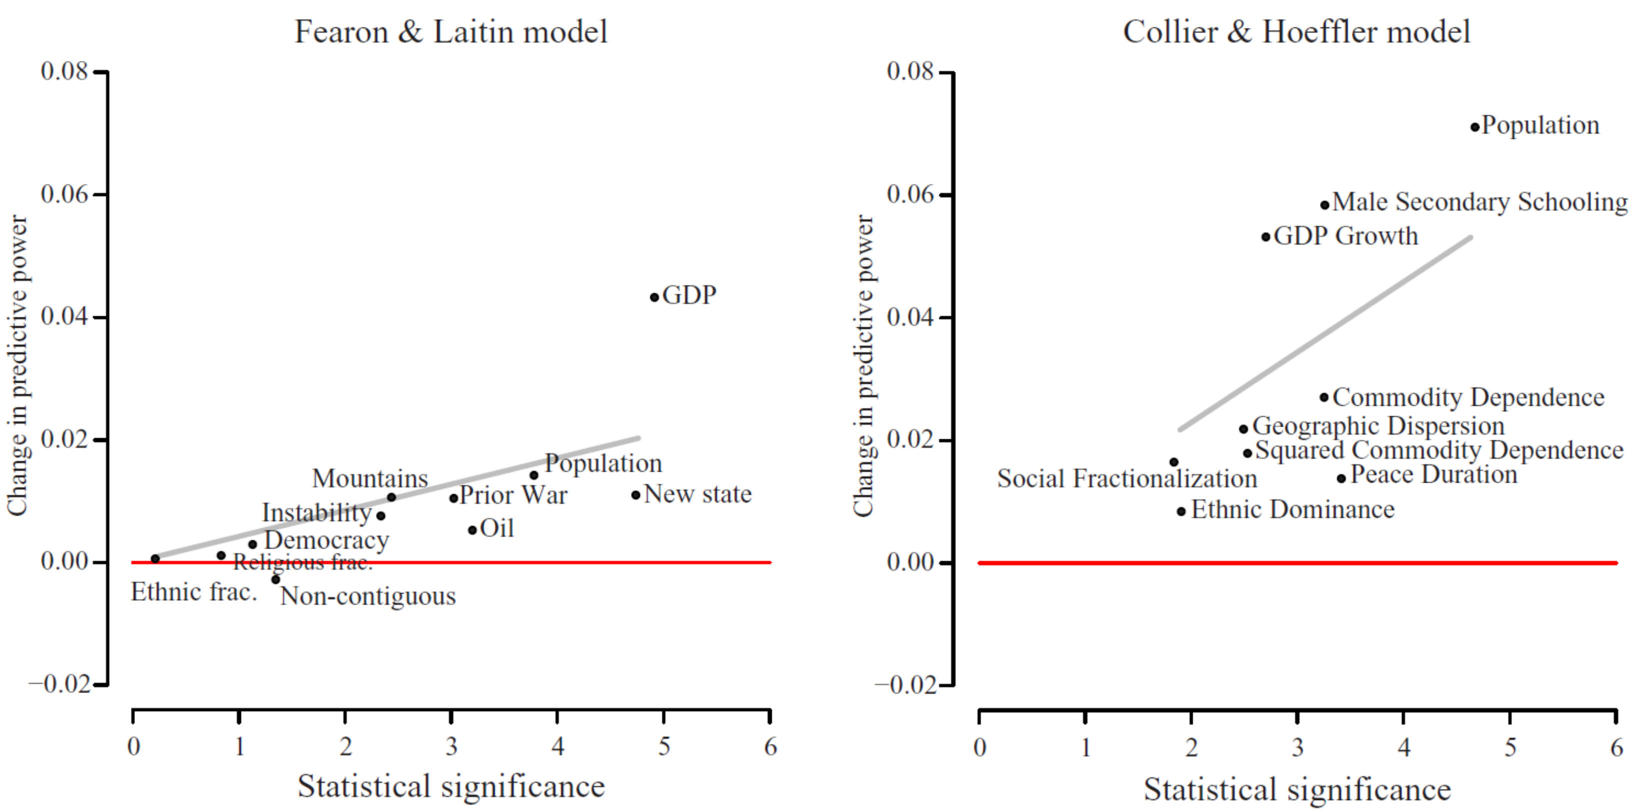

We now move from the topic of the prediction to the topic of causality.

Recall that CH and FL can, at best, establish correlations between various factors and conflict onset. But we are interested in causality as well as correlation.

The paper by Miguel, Satyanath and Sergenti attempts to make a causal connection between economic growth and conflict.

Here is the basic issue.

It is problematic to put economic growth on the right-hand-side of a conflict regression of the sort discussed in lecture 11. The main difficulty is that there is a very plausible channel of reverse causation.

That is, while it is true that low or negative growth might cause conflict it is also true that conflict might cause low or negative growth.

Let's take a concrete example that makes it pretty clear that reverse causation is a problem for estimating the relationship between growth and conflict.

The UN sends peacekeepers into an unstable and economically stagnant country. The peacekeepers prevent a conflict from starting and, subsequently, there is strong economic growth, in part because of the peace.

In this story peace is among the causes for high economic growth but high economic growth does not cause peace.

Suppose such cases are common. Then if you run a regression with conflict as the left-hand-side variable you might get a negative and significant coefficient on growth without there necessarily being any causation running from economic growth to peace (i.e., no conflict).

The idea of Miguel et al. is to isolate causation by focusing on only a part of economic growth that is not plausibly caused by whether there is peace or conflict.

Specifically, the trick is to zero in on the impact of rainfall in Sub-Saharan Africa. Since agriculture is important in this region it is plausible that rainfall would cause some economic growth. At the same time conflict (or peace) cannot plausibly have any effect on rainfall.

Thus, although only part of all economic growth can be linked to rainfall this part of economic growth is not plagued by reverse causation.

In technical terms we say that rainfall is an “instrument” for economic growth.

The table on slide 16 shows that the growth in rainfall is, indeed, correlated with economic growth in Sub-Saharan Africa.

This means that rainfall growth is a plausible instrument for economic growth. Of course, if rainfall and economic growth were uncorrelated with each other then there would be no point in exploring the Miguel et. al idea any further.

Notice that Miquel et al. use the *growth* of rainfall rather than rainfall itself. We will return to this fact later in the lecture.

TABLE 2
RAINFALL AND ECONOMIC GROWTH (First-Stage)
Dependent Variable: Economic Growth Rate, t

Explanatory Variable	Ordinary Least Squares				
	(1)	(2)	(3)	(4)	(5)
Growth in rainfall, t	.055***	.053***	.049***	.049***	.053***
	(.016)	(.017)	(.017)	(.018)	(.018)
Growth in rainfall, $t-1$	.034**	.032**	.028**	.028*	.037**
	(.013)	(.014)	(.014)	(.014)	(.015)
Growth in rainfall, $t+1$				.001	
				(.019)	
Growth in terms of trade, t					−.002
					(.023)
Log(GDP per capita), 1979		−.011			
		(.007)			
Democracy (Polity IV), $t-1$		.0000			
		(.0007)			
Ethnolinguistic fractionalization		.006			
		(.044)			
Religious fractionalization		.045			
		(.044)			
Oil-exporting country		.007			
		(.019)			
Log(mountainous)		.001			
		(.005)			
Log(national population), $t-1$		−.009			
		(.009)			
Country fixed effects	no	no	yes	yes	yes
Country-specific time trends	no	yes	yes	yes	yes
R^2	.02	.08	.13	.13	.16
Root mean square error	.07	.07	.07	.07	.06
Observations	743	743	743	743	661

NOTE.—Huber robust standard errors are in parentheses. Regression disturbance terms are clustered at the country level. A country-specific year time trend is included in all specifications (coefficient estimates not reported).
* Significantly different from zero at 90 percent confidence.
** Significantly different from zero at 95 percent confidence.
*** Significantly different from zero at 99 percent confidence.

Table three provides the essence of Miquel et al's results.

TABLE 3
RAINFALL AND CIVIL CONFLICT (Reduced-Form)

	DEPENDENT VARIABLE	
EXPLANATORY VARIABLE	Civil Conflict ≥25 Deaths (OLS) (1)	Civil Conflict ≥1,000 Deaths (OLS) (2)
Growth in rainfall, t	−.024 (.043)	−.062** (.030)
Growth in rainfall, $t-1$	−.122** (.052)	−.069** (.032)
Country fixed effects	yes	yes
Country-specific time trends	yes	yes
R^2	.71	.70
Root mean square error	.25	.22
Observations	743	743

NOTE.—Huber robust standard errors are in parentheses. Regression disturbance terms are clustered at the country level. A country-specific year time trend is included in all specifications (coefficient estimates not reported).
* Significantly different from zero at 90 percent confidence.
** Significantly different from zero at 95 percent confidence.
*** Significantly different from zero at 99 percent confidence.

The results on this slide and the next one are for conflict *incidence*, not onset (lecture 11). So these results are not directly comparable with those discussed in Lecture 11 since those were about conflict onset but the results for onset are similar to the results for incidence.

The first column of table 3 is for conflict (at least 25 battle deaths) and the second column is for war (at least 1,000 battle deaths).

There are negative and significant correlations between rainfall growth and the probability of both conflict and war.

For conflict the significant effect shows up only at period t-1 which gives the growth of rainfall between period t-2 and period t-1.

Table 4 shows results from a more elaborate specification that gives basically the same results as the less elaborate one on slide 17. The differences are:

1. Now there are control variables

2. Rainfall is now used as an ***instrument*** for economic growth in columns 5, 6 and 7. This means that rainfall is not used directly to predict conflict/war. Instead, rainfall predicts economic growth and then the part of economic growth that is predicted by rainfall is used to predict conflict.

The interesting results are the significant ones for economic growth at time t-1 in columns 5, 6 and 7.

TABLE 4
ECONOMIC GROWTH AND CIVIL CONFLICT

Explanatory Variable	Dependent Variable: Civil Conflict ≥25 Deaths						Dependent Variable: Civil Conflict ≥1,000 Deaths
	Probit (1)	OLS (2)	OLS (3)	OLS (4)	IV-2SLS (5)	IV-2SLS (6)	IV-2SLS (7)
Economic growth rate, t	−.37	−.33	−.21	−.21	−.41	−1.13	−1.48*
	(.26)	(.26)	(.20)	(.16)	(1.48)	(1.40)	(.82)
Economic growth rate, $t-1$	−.14	−.08	.01	.07	−2.25**	−2.55**	−.77
	(.23)	(.24)	(.20)	(.16)	(1.07)	(1.10)	(.70)
Log(GDP per capita), 1979	−.067	−.041	.085		.053		
	(.061)	(.050)	(.084)		(.098)		
Democracy (Polity IV), $t-1$	.001	.001	.003		.004		
	(.005)	(.005)	(.006)		(.006)		
Ethnolinguistic fractionalization	.24	.23	.51		.51		
	(.26)	(.27)	(.40)		(.39)		
Religious fractionalization	−.29	−.24	.10		.22		
	(.26)	(.24)	(.42)		(.44)		
Oil-exporting country	.02	.05	−.16		−.10		
	(.21)	(.21)	(.20)		(.22)		
Log(mountainous)	.077**	.076*	.057		.060		
	(.041)	(.039)	(.060)		(.058)		
Log(national population), $t-1$	.080	.068	.182*		.159*		
	(.051)	(.051)	(.086)		(.093)		
Country fixed effects	no	no	no	yes	no	yes	yes
Country-specific time trends	no	no	yes	yes	yes	yes	yes
R^2	...	.13	.53	.71	...	...	...
Root mean square error	...	.42	.31	.25	.36	.32	.24
Observations	743	743	743	743	743	743	743

NOTE.—Huber robust standard errors are in parentheses. Regression disturbance terms are clustered at the country level. Regression 1 presents marginal probit effects, evaluated at explanatory variable mean values. The instrumental variables for economic growth in regressions 5–7 are growth in rainfall, t and growth in rainfall, $t-1$. A country-specific year time trend is included in all specifications (coefficient estimates not reported), except for regressions 1 and 2, where a single linear time trend is included.

* Significantly different from zero at 90 percent confidence.
** Significantly different from zero at 95 percent confidence.
*** Significantly different from zero at 99 percent confidence.

Here is a table of results in simple specifications using onset, rather than incidence, of conflict as the left-hand-side variable.

TABLE 6
ECONOMIC GROWTH AND CONFLICT ONSET

Explanatory Variable	Dependent Variable: Onset, Civil Conflict ≥25 Deaths (IV-2SLS) (1)	Dependent Variable: Onset, Civil Conflict ≥1,000 Deaths (IV-2SLS) (2)
Economic growth rate, t	−3.15*	−2.85*
	(1.87)	(1.45)
Economic growth rate, $t-1$	−1.84	−.80
	(1.48)	(1.25)
Country fixed effects	yes	yes
Country-specific time trends	yes	yes
Root mean square error	.28	.24
Observations	555	625

NOTE.—Huber robust standard errors are in parentheses. Regression disturbance terms are clustered at the country level. The instrumental variables for economic growth are growth in rainfall, t and growth in rainfall, $t-1$. A country-specific year time trend is included in all specifications (coefficient estimates not reported).
* Significantly different from zero at 90 percent confidence.
** Significantly different from zero at 95 percent confidence.
*** Significantly different from zero at 99 percent confidence.

Again, there is a negative relationship between conflict and growth, although this time it shows up at time t rather than at time t-1.

The paper by Antonio Ciccone is a critique of the Miguel et al. paper. Ciccone:

1. re-examines the Miguel et al. data

2. analyses an extended dataset, covering ten additional years not covered in the Miquel et al. paper (which was written much earlier than Ciconne's paper was written).

The main difference in approach is that Ciccone looks at rainfall *levels* at each time period rather than the *growth rates* of rainfall from period to period as Miguel et al. do. This may seem like an inconsequential technical detail but, in fact, it is potentially quite important.

To get the general idea of the critique, let’s look at some hypothetical numbers.

Consider a country for which 100 cm of rainfall in a year is normal and fine for agriculture.

1. Suppose that the rainfall is 140 cm, 120 cm and 100 cm in three successive years. Thus, there is negative growth in rainfall for two successive periods. Still, the country never experiences below-normal rainfall.

2. Now let’s reverse the situation. Rainfall levels of 60, 80 and 100 will be recorded as two successive periods of positive growth although rainfall would be below average in this example until the final period.

The point is that rainfall levels, and how these compare to average levels, probably convey more useful information than do growth rates of rainfall.

We can look at this also from a slightly more technical perspective (Ciccone, page 6). Suppose that the growth rate of rainfall is what really matters for the probability of conflict. Nevertheless, we can still fully capture these growth effects using a model specified entirely in terms of rainfall levels.

Consider the following simple model (following Ciccone's notation):

$$Pconflict_t = \beta RGr_{t-1}$$

According to this model, the probability of conflict at period t is determined by a constant, β, times the growth rate of rainfall between periods t-2 and t-1.

We can rewrite this equation as:

$$Pconflict_t = \beta(\log R_{t-1} - \log R_{t-2}) = \beta \log R_{t-1} - \beta \log R_{t-2}$$

(Note that the growth of rainfall between period's t-2 and t-1 is equal to (approximately) the logarithm of the rainfall level at period t-1 minus the logarithm of the rainfall level at period t-2. I will put this mathematical fact on the worksheet for next week.)

We can actually run a regression and estimate the coefficients on rainfall at time t - 1 and t – 2.

$$Pconflict_t = \beta_1 \log R_{t-1} - \beta_2 \log R_{t-2}$$

If the conflict probability is truly determined by the growth rate of rainfall then the estimated coefficients on R_{t-1} and R_{t-2}, i.e., β_1 and β_2 should turn out to be equal to each another.

Slide 25 shows that anything that can be done empirically with *growth rates* of rainfall can also be done in terms of *rainfall levels*.

In other words, the theory that growth rates of rainfall determine conflict probabilities is really a special case of the theory that rainfall levels determine conflict probabilities.

An important implication of the analysis on slide 25 is that, if it is rainfall growth rates that determine conflict then an increase in rainfall at time t-1 should have the same effect as a decrease in rainfall at time t-2.

The table on slide 28 will show that the estimated coefficients on rainfall levels at times t – 1 and t – 2 do not satisfy the condition of the above paragraph.

The next slide is the key table in the Ciconne paper. Notice that the dependent variable is conflict onset.

The first column replicates the result from the Miguel et al. paper that rainfall growth at period t-1 is negatively and significantly associated with conflict onset.

The second column switches from using rainfall growth to using rainfall levels. There is a significant and, surprisingly, *positive* relationship between conflict and the level of rainfall in period t-2.

This strange relationship seems to imply that more rainfall is associated with more conflict as long as the extra rain occurred long enough in the past, i.e., two years ago.

Tables

Table 1. Rainfall and civil conflict onset

	MSS (2004) data, which are for 1979-1999		latest data, which are for 1979-2009	
	(1)	(2)	(3)	(4)
Rainfall Growth, t	-0.063 (0.044) [0.048]		-0.037 (0.029) [0.031]	
Rainfall Growth, t-1	-0.120* (0.062) [0.068]		-0.052 (0.033) [0.036]	
Log Rainfall, t		-0.073 (0.078) [0.086]		0.005 (0.041) [0.044]
Log Rainfall, t-1		-0.026 (0.069) [0.075]		0.023 (0.042) [0.044]
Log Rainfall, t-2		0.156** (0.068) [0.074]		0.074 (0.052) [0.056]
Country FE and Trend	Yes	Yes	Yes	Yes
Observations	555	555	873	873

Note: The left-hand-side variable is an indicator variable capturing civil conflict onset (see p.7 in the main text). The method of estimation is least squares. Standard errors in parentheses are robust for arbitrary heteroskedasticity and clustered at the country level. Standard errors in square brackets also apply the STATA small-sample adjustment (see p.8 in the main text). *Significantly different from zero at 90 percent confidence, ** 95 percent confidence, *** 99 percent confidence. When the asterisks are next to the least-squares point estimate, the confidence level applies no matter which of the two standard errors is employed. When the asterisks are next to the standard error, the confidence level applies to that standard error only.

This positive relationship, shown in column 2, between the rainfall level at time t-2 and conflict at time t explains why Miguel et al. get a negative relationship between the growth rate of rainfall at time t-1 and conflict (shown in Column 1).

In fact, the column-2 results show that the column 1 relationship is not really about rainfall *growth*.

If the link between rainfall and peace is really due to rainfall growth then there should be a negative and significant relationship on the rainfall level at time t-1 in addition to the positive and significant coefficient on the rainfall level at time t-2. Moreover, these coefficients should be roughly equal in absolute value and opposite in sign.

That is, if it is the growth rate of rainfall that matters then increases in rainfall at time t-1 should be equivalent to decreases in rainfall at time t-2 and vice versa. Yet the estimated coefficients do not behave this way.

The second important point of the table comes from columns 3 and 4.

These columns report results for the extended dataset covering 1979-2009 rather than 1979-1999 (the years covered by Miguel et al.).

None of the rainfall coefficients are significant in these regressions.

So, with the benefit of more data, it now seems that the significant coefficients in the first two columns are just statistical anomalies dues to not having enough data. It appears that rainfall does not actually matter for conflict onset.

In other words, the Miguel et al. findings look like spurious correlations.