EC3320

2016-2017

Michael Spagat

Lecture 14.

The Nunn and Qian paper studies the impact of foreign aid on conflict.

To a large extent this paper is motivated by stories suggesting that foreign aid might actually stimulate or intensify conflict.

This idea is in tension with the finding of the Berman et al. paper, covered in lecture 4, which found that small-scale aid dispersed by military commanders in Iraq (the CERP programme) created a more peaceful environment in the receiving areas of Iraq – but that was just about one country

Nunn and Qian use a cross-country approach, in contrast to Berman et al's single country approach. Still the two teams confront many of the same problems.

Probably the biggest challenge to Nunn and Qian is our old friend, reverse causation – conflict can drive aid decisions but aid decisions can also drive conflict. Think of the following two mechanisms:

1. Aid may flow into war zones as part of an attempt to pacify these zones. Thus, conflict and aid can be *positively* correlated with each other without aid actually causing conflict.

2. Aid may be withheld from war zones because of a fear that aid will be wasted (or even fuel conflict) unless stability is first established in these zones. If this is the case then aid will be *negatively* correlated with conflict without aid actually causing peace.

Nunn and Qian have a clever idea to identify a causal effect of aid on conflict.

Their key observation is that the US has an elaborate system of subsidies for wheat production according to which the government buys up large quantities of wheat during years of high production. This wheat cannot be stored indefinitely, so the US tends to donate a lot of wheat to poor countries in years immediately following these big production years.

Thus, US food aid quantities are subject to a substantial random shocks, largely determined by the weather. Of course, as we have already seen, the weather is not subject to reverse causation concerns.

Imagine a single country in two different time periods.

1. The first time period is a year after a big wheat crop in the US. The country gets a lot of food aid in this year because the US has a big surplus that it wants to unload.

2. The second year is a year after a small wheat crop in the US. The country gets little or no US food aid in this second year because the US is not sitting on a big and decaying wheat stock.

This comparison resembles a controlled experiment in which you randomly split people into two groups - one group receives a treatment (such as an experimental drug) and the other group does not receive the treatment.

The discussion so far should have reminded you of our discussion of the Miquel et al. paper which used the technique of instrumental variables.

Miguel et al. focus on just the component of (rain-driven) growth that is not subject to reverse causation.

Similarly, Nunn and Qian focus on just the component of food aid that is not subject to reverse causation.

Actually, Nunn and Qian go a step further than what I have described so far.

They also distinguish between countries that regularly receive food aid from the US and countries that usually do not receive food aid from the US.

The idea is that random weather fluctuations affecting the US wheat crop should have a bigger impact on regular recipients of US food aid than on countries that rarely receive US food aid.

This is a differences-in-differences type of idea (lecture 8)

Why might food aid fuel conflict?

1. Armed groups might steal the aid and use it to support themselves in the field. In fact, they may steal more than the aid itself, for example taking vehicles that are meant to be used to distribute aid.

2. Armed groups might generate money off of the aid effort by, for example, demanding payments in exchange for letting aid get through checkpoints or kidnapping aid workers for ransom. (Note that stolen aid (point 1 above) does not necessarily have to be consumed by the thieves – it can be sold for money instead.)

Furthermore, in many cases it is appropriate to view the government as one of the “armed groups” that could use food aid to fuel its activities.

Here are Nunn and Qian's key equations:

$$C_{irt} = \beta F_{irt} + \mathbf{X}_{irt}\Gamma + \varphi_{rt} + \psi_{ir} + \nu_{irt}, \quad (3)$$

$$F_{irt} = \alpha(P_{t-1} \times \overline{D}_{ir}) + \mathbf{X}_{irt}\Gamma + \varphi_{rt} + \psi_{ir} + \varepsilon_{irt}. \quad (4)$$

stage. The index i denotes countries, r denotes six geographic regions, and t denotes years.[11] The sample we analyze is a panel of 125 non-OECD countries between 1971 and 2006.

The dependent variable, C_{irt}, is an indicator variable that equals one if there is conflict in country i during year t. F_{irt} is the endogenous variable of interest, the quantity of wheat aid shipped from the US to recipient i in year t. $\mathbf{X}_{irt}$ is a vector of country-year covariates that we motivate and discuss when we present the results. $\delta_r Y_t$ denotes

P_{t-1} is the quantity of wheat produced in period t-1.

Let D_{irt} be an indicator variable that takes a value of one if country i receives any US food aid in year t. Then, $\overline{D}_{ir} = \frac{1}{36}\sum_{t=1971}^{2006} D_{irt}$ denotes the fraction of years between 1971 and 2006 that a country receives any US food aid. φ_{rt} denotes region-year fixed effects. All other variables are defined as before.

The notation is somewhat confusing so let me clarify the main ideas from slide 8.

The strategy is:

1. Use equation 4 to predict US food aid flowing to a particular country during a particular time period. This prediction uses US production in the previous period, the historical tendency of the US to give food aid to that country and a bunch of control variables.

2. You then plug these predictions into equation 3 to predict whether or not there will be ongoing conflict (conflict incidence), also controlling for various things.

The two figures on slide 10 show that US wheat aid does, indeed, respond to the previous year's wheat production. The two figures on slide 11 show that the link between US wheat production and conflict is weak for countries that are not regular recipients of US food aid and strong for countries that are regular recipients of US food aid.

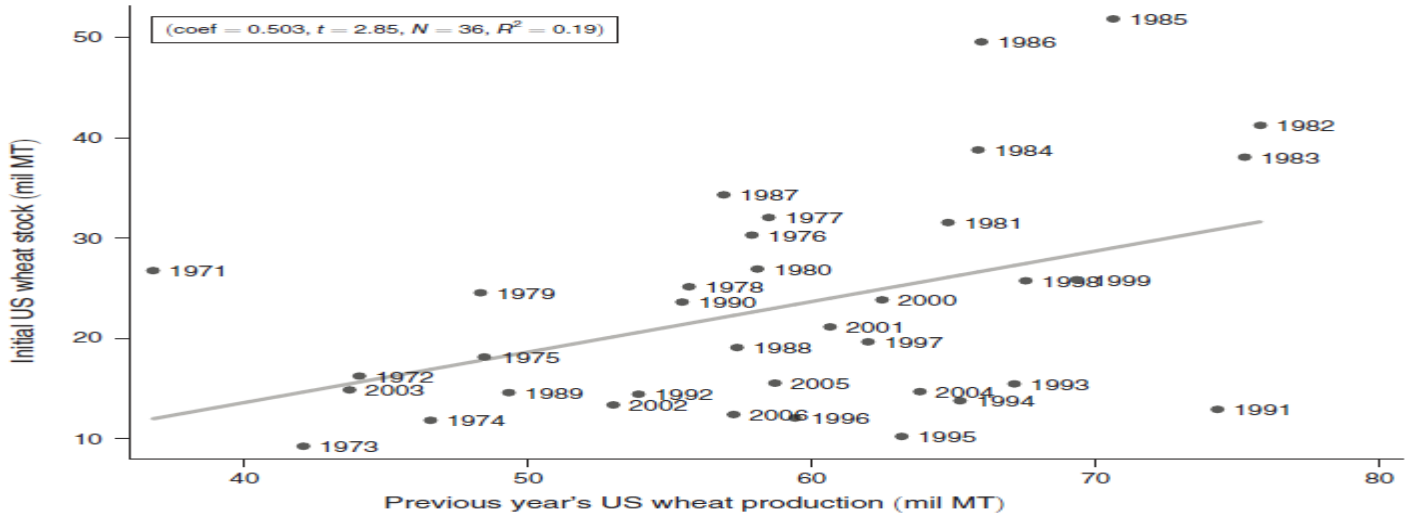

FIGURE 1. US WHEAT RESERVES AND LAGGED US WHEAT PRODUCTION

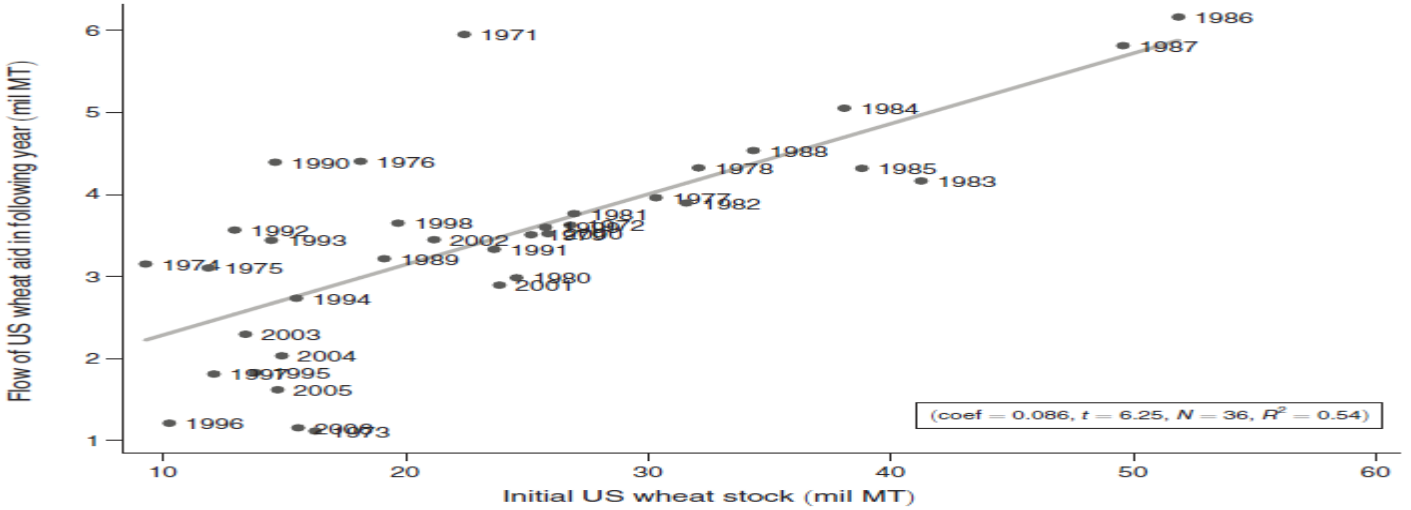

FIGURE 2. US WHEAT AID AND INITIAL US WHEAT RESERVES

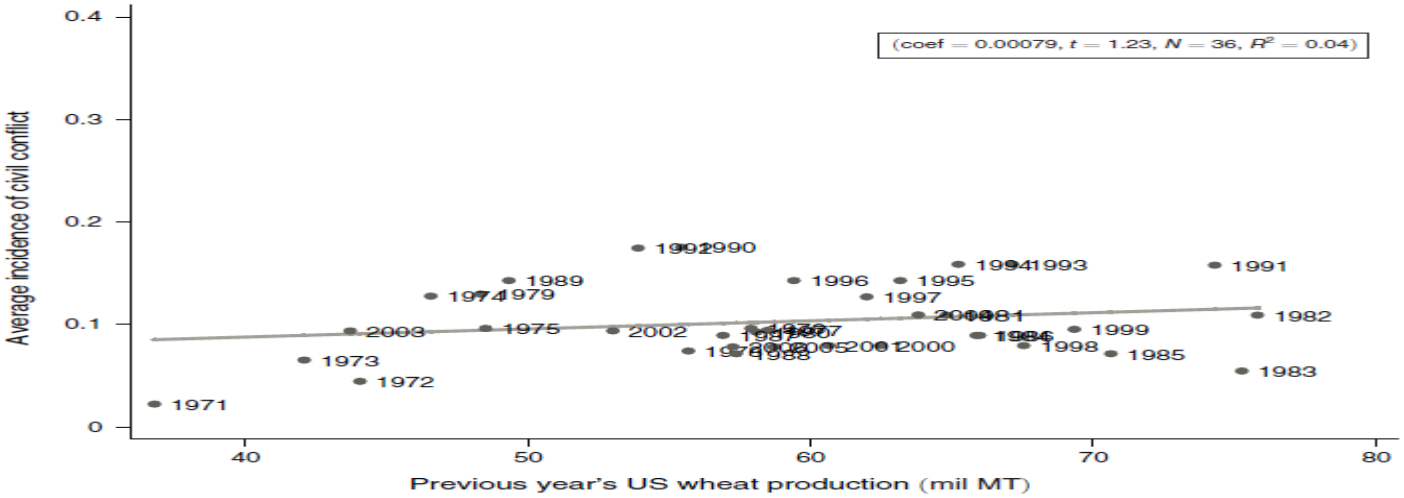

FIGURE 3. AVERAGE CIVIL CONFLICT INCIDENCE AND LAGGED US WHEAT PRODUCTION, IRREGULAR RECIPIENTS: $D_{ir} < 0.30$

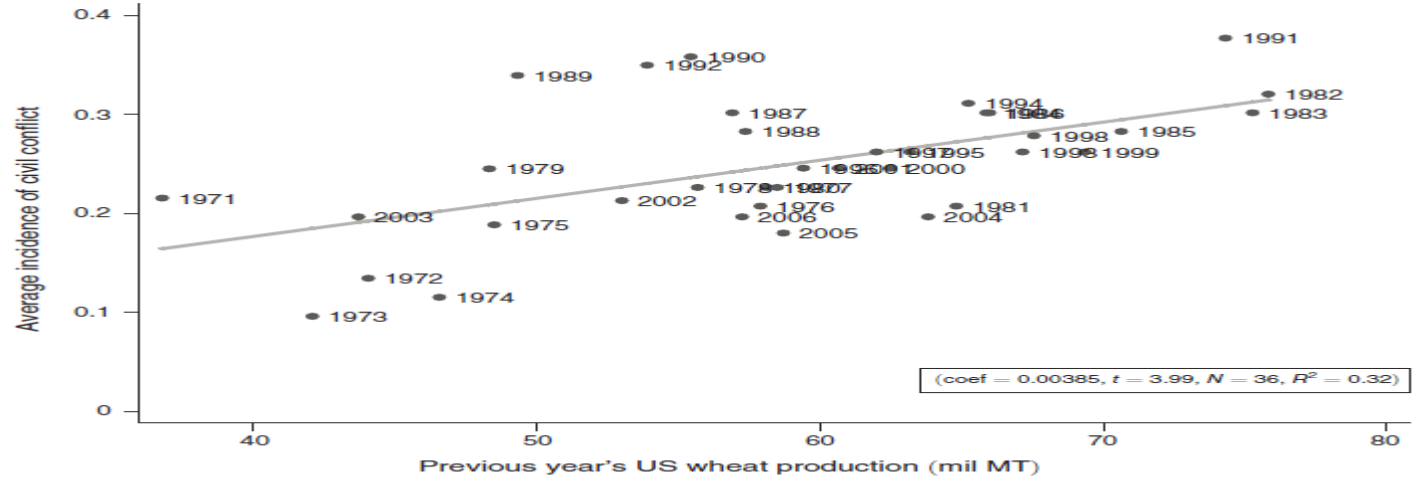

FIGURE 4. AVERAGE CIVIL CONFLICT INCIDENCE AND LAGGED US WHEAT PRODUCTION, REGULAR RECIPIENTS: $D_{ir} \geq 0.30$

The next slide gives a table with a lot of information that can be confusing so here are the key points.

1. Panel D tells us that the previous year's US wheat production together with the tendency of the US to aid a particular country does a good job of explaining the current year's wheat aid to that country (equation 4 above).

2. Panel C tells us that feeding these wheat-aid predictions into equation 3 does a good job of explaining whether or not a country is involved in armed conflict in the corresponding year.

3. Panels A and B are less important. Panel A says that wheat aid itself, as opposed to predicted wheat aid taken from equation 3, is not useful for explaining armed conflict. Panel B says that the instrument used in equation 4 does a good job on its own of explaining conflict.

TABLE 2—THE EFFECT OF FOOD AID ON CONFLICT: BASELINE SPECIFICATION WITH $P_{i-1} \times D_{ir}$ AS THE INSTRUMENT

	Parsimonious specifications				Baseline specification		
Dependent variable (panels A, B, and C):	Any conflict (1)	Any conflict (2)	Any conflict (3)	Any conflict (4)	Any conflict (5)	Intrastate (6)	Interstate (7)
Panel A. OLS estimates							
US wheat aid (1,000 MT)	−0.00006	−0.00007	−0.00005	−0.00007	−0.00011	−0.00005	−0.00011
	(0.00018)	(0.00018)	(0.00017)	(0.00017)	(0.00017)	(0.00017)	(0.00004)
R^2	0.508	0.508	0.518	0.534	0.549	0.523	0.385
Panel B. Reduced form estimates (× 1,000)**							
Lag US wheat production (1,000 MT)	0.00829	0.01039	0.01070	0.01133	0.01071	0.00909	−0.00158
× avg. prob. of any US food aid	(0.00257)	(0.00263)	(0.00262)	(0.00318)	(0.00320)	(0.00322)	(0.00121)
R^2	0.511	0.512	0.521	0.536	0.551	0.525	0.382
Panel C. 2SLS estimates							
US wheat aid (1,000 MT)	0.00364	0.00303	0.00312	0.00343	0.00299	0.00254	−0.00044
	(0.00174)	(0.00125)	(0.00117)	(0.00106)	(0.00096)	(0.00088)	(0.00033)
Dependent variable (panel D):	US wheat aid (1,000 MT)						
Panel D. First-stage estimates							
Lag US wheat production (1,000 MT)	0.00227	0.00343	0.00343	0.00330	0.00358	0.00358	0.00358
× avg. prob. of any US food aid	(0.00094)	(0.00126)	(0.00120)	(0.00092)	(0.00103)	(0.00103)	(0.00103)
Kleibergen-Paap *F*-statistic	5.84	7.37	8.24	12.76	12.10	12.10	12.10
Controls (for all panels):							
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
US real per capita GDP × avg. prob. of any US food aid	No	Yes	Yes	Yes	Yes	Yes	Yes
US democratic president × avg. prob. of any US food aid	No	Yes	Yes	Yes	Yes	Yes	Yes
Oil price × avg. prob. of any US food aid	No	Yes	Yes	Yes	Yes	Yes	Yes
Monthly recipient temperature and precipitation	No	No	Yes	Yes	Yes	Yes	Yes
Monthly weather × avg. prob. of any US food aid	No	No	Yes	Yes	Yes	Yes	Yes
Avg. US military aid × year FE	No	No	No	Yes	Yes	Yes	Yes
Avg. US economic aid (net of food aid) × year FE	No	No	No	Yes	Yes	Yes	Yes
Avg. recipient cereal imports × year FE	No	No	No	No	Yes	Yes	Yes
Avg. recipient cereal production × year FE	No	No	No	No	Yes	Yes	Yes
Observations (for all panels)	4,089	4,089	4,089	4,089	4,089	4,089	4,089

Notes: An observation is a country and a year. The sample includes 125 non-OECD countries for the years 1971–2006. The controls included are indicated in the table by Y (yes) or N (no). Coefficients are reported with standard errors clustered at the country level in parentheses. **In panel B, the point estimates and standard errors are multiplied by 1,000 for presentation purposes. In panel D, we report first-stage Kleibergen-Paap *F*-statistics.

Slides 12 & 13 (above) gives the essence of Nunn and Qian's results. The authors do a lot more than I present in this lecture in which I just want to touch on a couple more points.

1. The above results are about the *incidence* of conflict. Nunn and Qian also look at onset but do not find a significant effect.

2. They also look at offset (ending) of conflict and find that US food aid decreases the probability of offset.

To summarize, *US food aid seems to contribute to prolonging existing conflicts but not to starting new ones.*

The Crost, Felter and Johnston (CFJ) paper also addresses the aid-and-conflict issue but within the context of a single country – the Philippines.

Again, CFJ's main concern is the causality issue. Their solution is a method we have not yet encountered in the course. It is called a "*regression discontinuity approach*".

The key observation is that municipalities in the Philippines are classified according to how rich they are and those below a threshold are eligible for an aid programme that the richer ones are not eligible for.

This means that those just above the threshold are similar to those just below the threshold except that the ones above do not get aid and the ones just below do. (It is slightly more complicated than this – some above the threshold do actually get the aid and some below do not but we will leave this nuance aside for now.)

Once again, the idea is to create a situation that resembles a controlled experiment.

The “treated” municipalities are the ones just below the threshold and the “control” municipalities are the ones just above the threshold. Note that the conflict variable for this study is not a 0-1 variable denoting “no-conflict” or “conflict” but, rather, a measure of conflict casualties assembled by the Philippine’s military.

CFJ estimate the following equation:

$$Y_{ipt} = \beta_0 + \tau D_{ip} + \beta_1 X_{ip} + \beta_2 D_{ip} X_{ip} + \alpha_p + \gamma_t + \varepsilon_{ipt}.$$

Y_{ipt} denotes the number of conflict casualties suffered by municipality i in province p in month t. X_{ip} denotes the municipality’s relative poverty rank. D_{ip} is an indicator that takes the value 1 if the municipality is eligible for the program and 0 if it is not. α_p and γ_t are province and time fixed effects. The RD estimator is given by the parameter τ.

Here are some key points:

1. The relative poverty rank is on a scale from -6 to +6

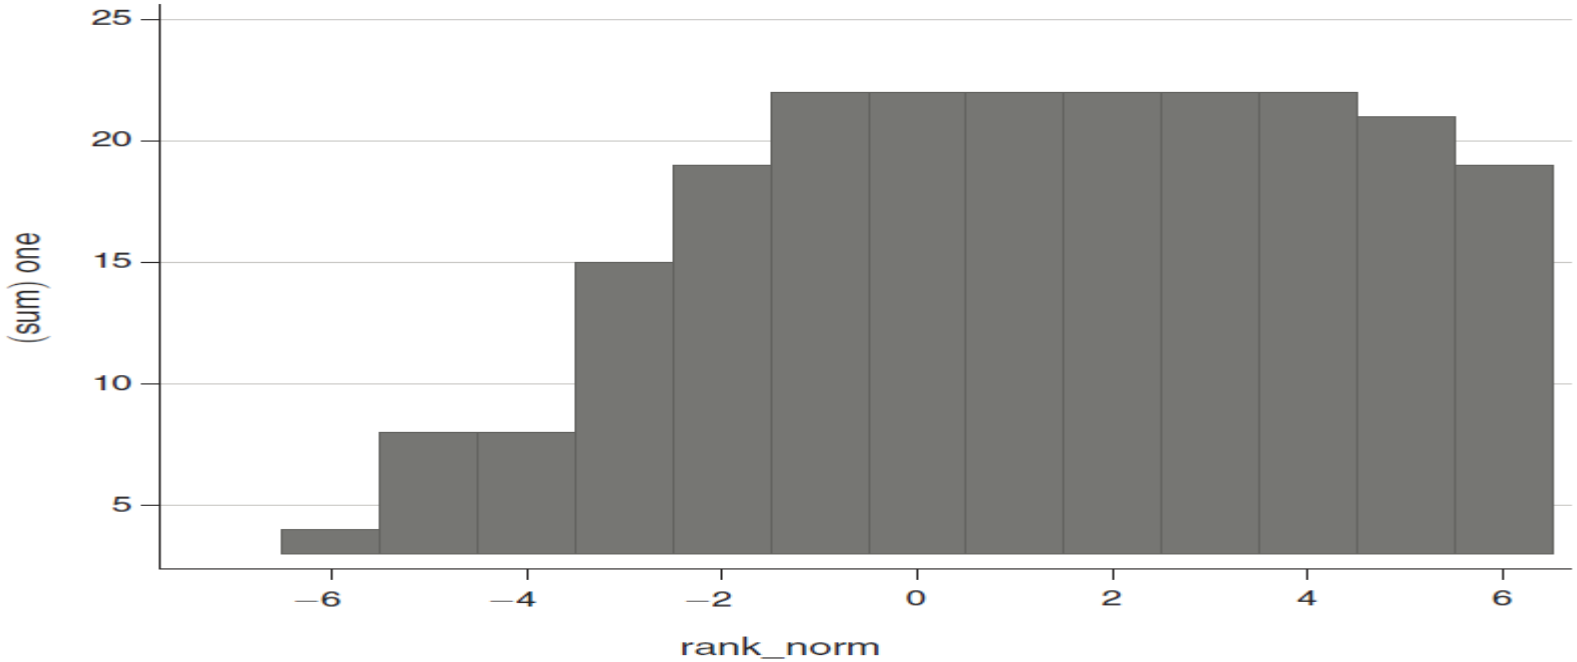

FIGURE 2. NUMBER OF MUNICIPALITIES BY RELATIVE POVERTY RANK

Notes: The figure displays the number of municipalities with each relative poverty rank. A relative poverty rank of zero indicates the richest municipality in a province that is still eligible for the KALAHI-CIDSS program. A rank of one indicates the poorest ineligible municipality.

2. The regressions are slightly non-standard in that they give greater weight to municipalities that are near 0 than they do to municipalities that are far from 0.

The usual regression minimizes the sum of squared deviations of the predicted value of the left-hand-side variable from the actual values. The regressions that CFJ use are similar to the standard ones but work harder than the standard regressions to get a close fit for the observations near 0 and work less hard to fit closely the observations far from zero.

This weighting system is consistent with the regression discontinuity idea of comparing very similar municipalities, specifically ones just above and just below the threshold, with each other.

3. CFJ run separate regressions for separate time periods.

Importantly, under the aid programme there is a “social preparation phase” at the beginning of which it is announced that certain municipalities are being considered for aid. The community then works together to develop proposals on how to use the aid and then a final decision is made on whether or not to go forward with the aid programme in that municipality.

The fact that the preparation phase is followed by a decision which may or may not be positive means that there is not a perfect match between the eligibility threshold and whether or not the aid programme goes forward. Some eligible municipalities do not get the aid in the end and some initially ineligible municipalities do get the aid in the end after some eligible municipalities drop out.

The next slide shows, however, that there is a strong link between initial eligibility and participation in the end.

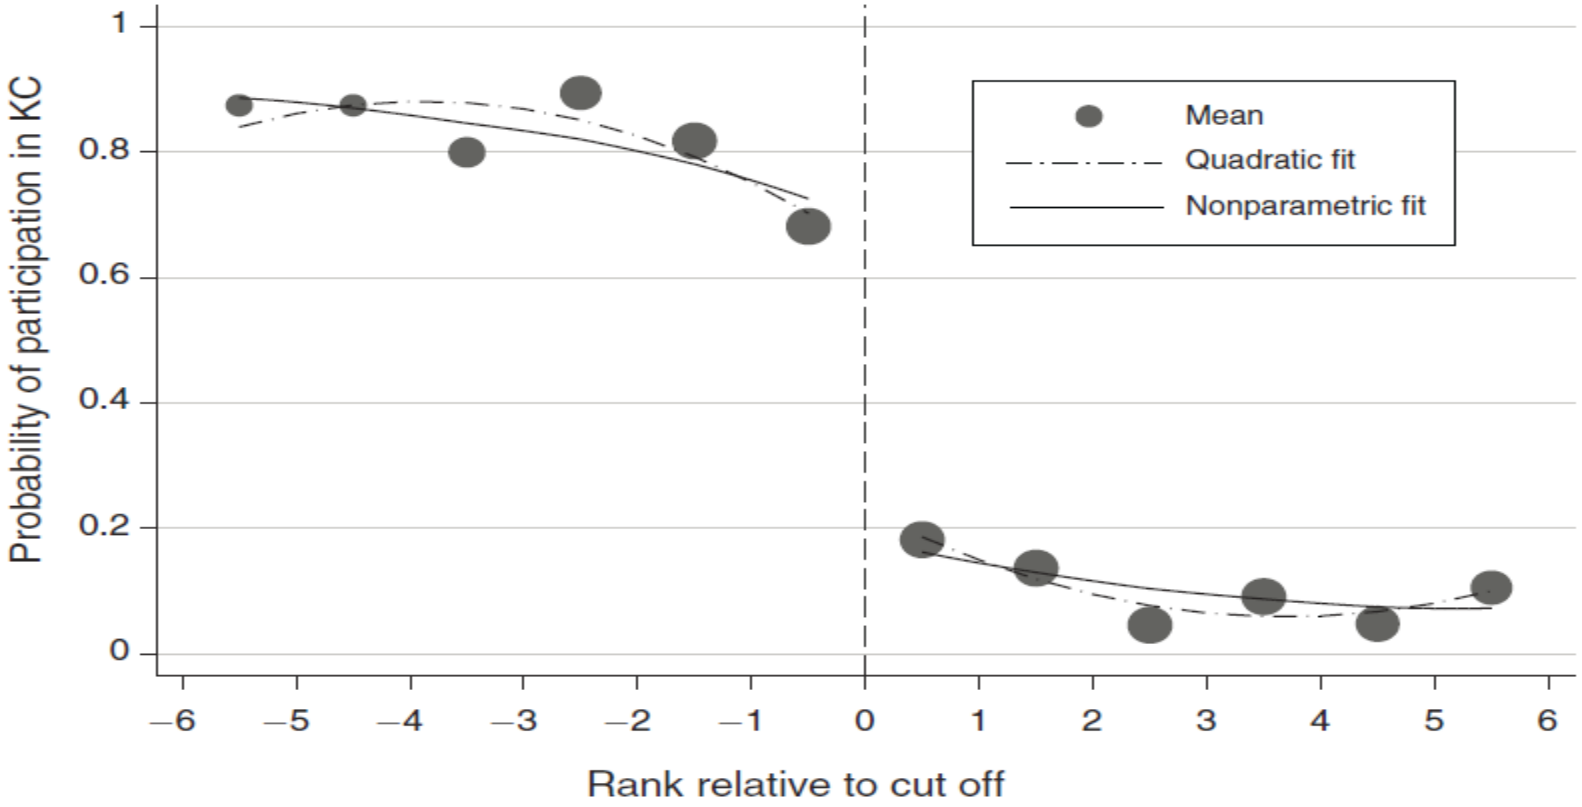

FIGURE 3. THE EFFECT OF ELIGIBILITY ON PARTICIPATION

Notes: The figure presents the relationship between the probability of participating in the KALAHI-CIDSS program and the running variable of the RD design, which is the distance between the municipality's poverty rank and the provincial eligibility threshold. Scatter dots represent means. Dashed lines are quadratic fits, separately estimated on both sides of the eligibility threshold. Solid lines are nonparametric fits from a local linear regression that uses triangular kernels with a bandwidth of 6, separately estimated on both sides of the eligibility threshold.

The picture on the next slide captures the essence of the paper.

Look at the vertical dotted lines at the 0 points and search for discontinuities, i.e., cases where municipalities just below the threshold are very different from municipalities just above the threshold.

The one picture where you see a strong discontinuity is the one for the social preparation period. Thus, *violence increases in eligible municipalities while communities make plans for how they will spend their aid if they actually receive some*. This effect is most pronounced for municipalities that are at the borderline of eligibility.

Before the social preparation phase (“preprogram”) and after the aid is disbursed (“remaining program period”) there is no real discontinuity. There is a discontinuity for the entire program period but this comes from the social preparation phase.

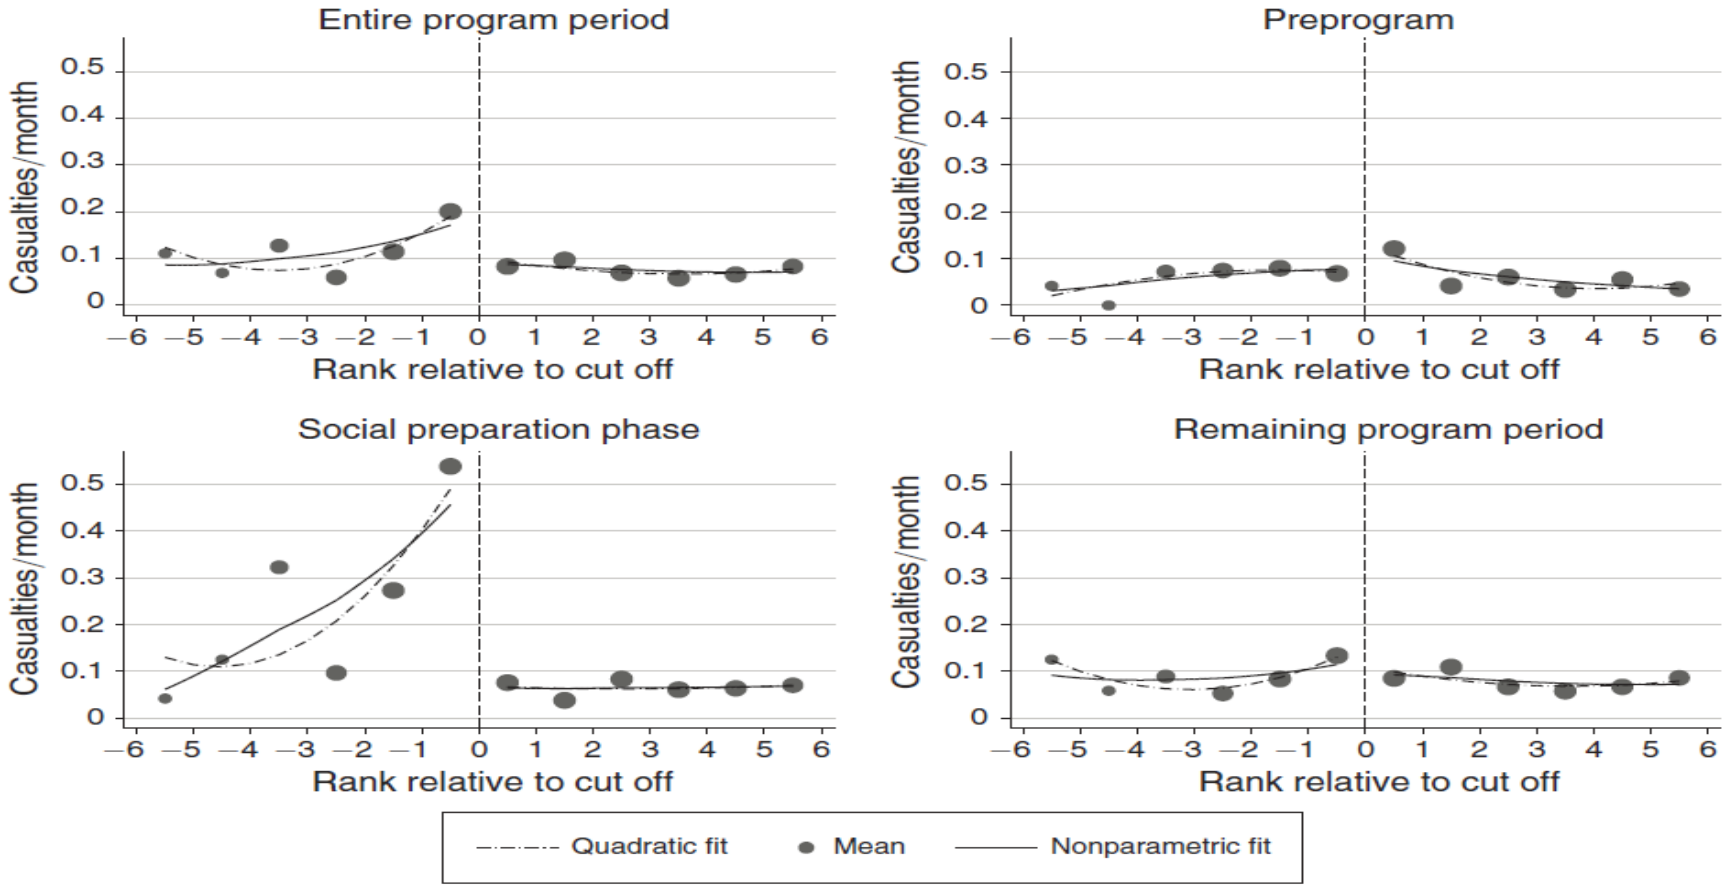

FIGURE 4. THE EFFECT OF ELIGIBILITY ON CASUALTIES

Notes: The figure presents the relationship between the number of casualties experienced during the program period and the running variable of the RD design, which is the distance between the municipality's poverty rank and the provincial eligibility threshold. Scatter dots represent means. Dashed lines are quadratic fits, separately estimated on both sides of the eligibility threshold. Solid lines are nonparametric fits from a local linear regression that uses triangular kernels with a bandwidth of 6, separately estimated on both sides of the eligibility threshold.

CFJ suggest that insurgents see aid as a threat and increase their violent attacks to prevent aid from happening. This could be true if people decrease their support for insurgents when local governments do a good job of providing services.

The table below suggest that CFJ may be right about this. Eligible municipalities that do not participate in the end experience more violence than municipalities that go forward with the aid programme.

TABLE 5—CONFLICT INTENSITY AND PROGRAM PARTICIPATION

	Participating (1)	Nonparticipating (2)	*p*-value of difference (3)
Preprogram	0.057 (0.011)	0.097 (0.031)	0.43
Social preparation phase	0.145 (0.042)	0.824 (0.292)	0.052
Remaining program period	0.073 (0.011)	0.167 (0.040)	0.136
Municipalities	76	18	94

Notes: The table reports means of monthly casualties in municipalities eligible for KALAHI-CIDSS. Standard errors in parentheses, clustered at the municipality level.

The pictures on the next slide also support this view.

They show that the extra violence in municipalities that are relatively rich but still eligible for the programme is initiated by insurgent groups.

Actually, just the fact that it is the municipalities on the borderline of eligibility that suffer the most violence is generally consistent with the CFJ view since these are more likely to walk away from aid than the very poorest municipalities are.

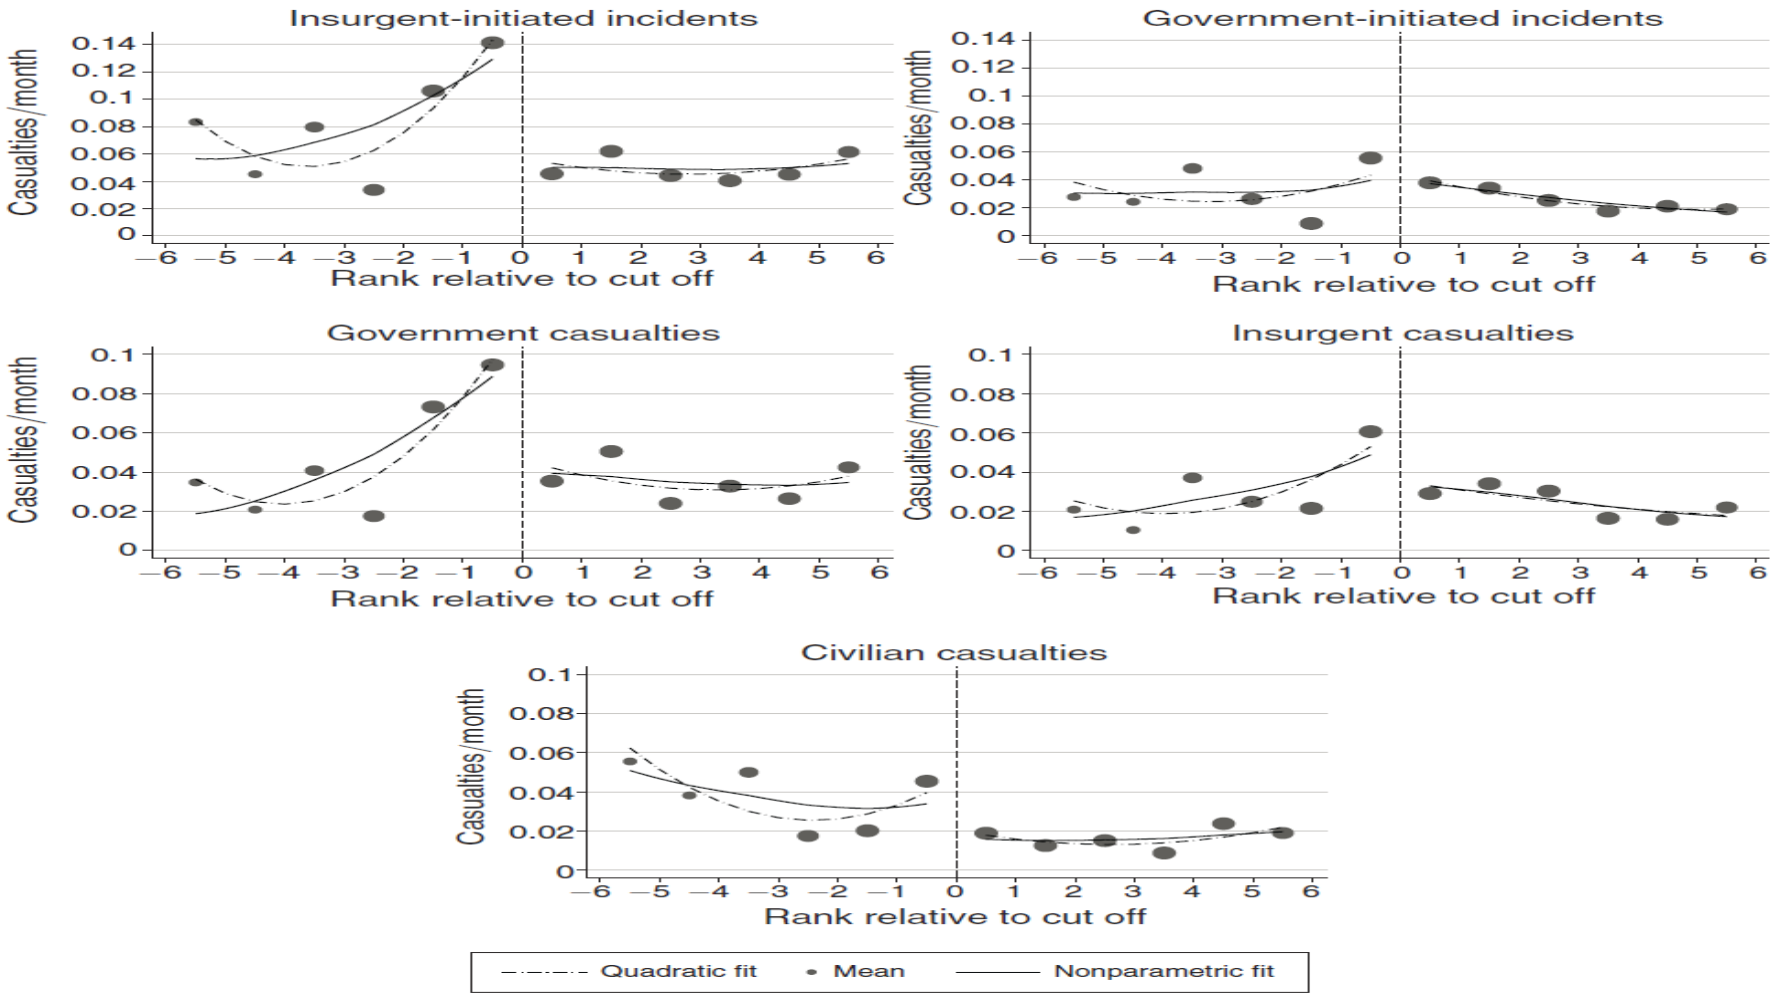

FIGURE 6. WHO INITIATES THE VIOLENCE, AND WHO SUFFERS THE CASUALTIES?

Notes: The figure presents the relationship between the number of casualties experienced by different groups during the program period and the running variable of the RD design, which is the distance between the municipality's poverty rank and the provincial eligibility threshold. Scatter dots represent means. Dashed lines are quadratic fits, separately estimated on both sides of the eligibility threshold. Solid lines are nonparametric fits from a local linear regression that uses triangular kernels with a bandwidth of 6, separately estimated on both sides of the eligibility threshold.

CFJ agrees with Nunn and Qian on the general point that aid leads to conflict.

However, the two studies are rather different in the end.

1. Nunn and Qian suggest that insurgent groups want aid to flow so that they can steal the aid.

2. CFJ suggests that insurgents work specifically to prevent aid from flowing so that they can prevent local governments from winning over its citizens by providing them with good services.

Now we change gears completely and look at the question of rape during Civil War, basing our discussion on this paper by Dara Kay Cohen.

It is worthwhile to bear in mind three things during this discussion:

1. There are various forms of sexual violence of which rape is generally the most extreme form. Cohen focuses just on rape.

2. The Cohen paper is about rape *within the context of civil war*, not about rape in general which is actually more common than rape within the context of civil war. However, gang rape is more common during war than it is outside of war and Cohen stresses gang rape in her analysis.

3. It is easy to slip into an incorrect idea that only females are raped. That said, females are certainly raped much more than males are.

Cohen is interested in testing some established ideas that are put forward in the literature for explaining rape within civil war. They are:

1. What Cohen calls “Opportunism/Greed”.

a. The opportunism idea is that war creates a chaotic, lawless environment within which there are greater opportunities to rape than are possible within a smoothly functioning society. (This idea has little plausibility in my view for single rapes since we know that these are rather common outside of war. However, the opportunism idea may have more substance for gang rape.)

b. The greed idea is that if insurgent groups are able to finance themselves easily through readily available mineral resources then they have little need to be on good terms with local populations so discipline might break down leading to rape and other atrocities committed against civilians.

2. Ethnic hatred. This idea is self-explanatory.

3. Gender inequality. Here the idea is that rape will be more prevalent to the extent that females are viewed as inferior to males. (Of course, this idea assumes that we are discussing the rape of females.)

4. Forced recruitment, cohesion, and gang rape. Cohen suggests that when fighters are forcibly recruited then it is difficult to mould them into a cohesive fighting force. Committing gang rape could create such cohesiveness. In addition, if a fighter has committed a public crime by participating in a gang rape it will become more difficult to escape the group in the future.

Number 4 is what Cohen is mainly interested in and her results support this theory. She gives some detailed background and references on this subject in her paper so please have a look if you want to go into depth on this subject, e.g., for your next essay.

Cohen puts together a new dataset on rape in civil war which she uses to support her analysis. We will discuss this dataset in detail in the seminars next week. For now I just describe its bare essentials.

1. Cohen goes through human rights reports of the US State Department for every country in conflict during the years 1980 to 2009.

2. For each country year she assigns a rating of 0, 1, 2 or 3. If a report says nothing about war rape then she assigns “0”. If the report suggests that war rape is widespread and systematic then she assigns a “3”. Cases in between get coded as 1’s or 2’s.

3. There are separate codings for state and non-state perpetrators.

Cohen collects other variables from a variety of sources. I will not go into detail on these (again, you can read the paper). But the key ones are:

1. Abduction

2. Pressganging

In reality these are just two words for violent forced recruitment of fighters. Cohen uses the word "pressganging" when *the State* essentially kidnaps people to fight for it while she uses the word "abduction" when *insurgent groups* do the same thing.

The results of Cohen's modelling are summarized in the pictures on the next two slides.

FIGURE 1. Probability of Insurgent-Perpetrated Wartime Rape with and without Abduction

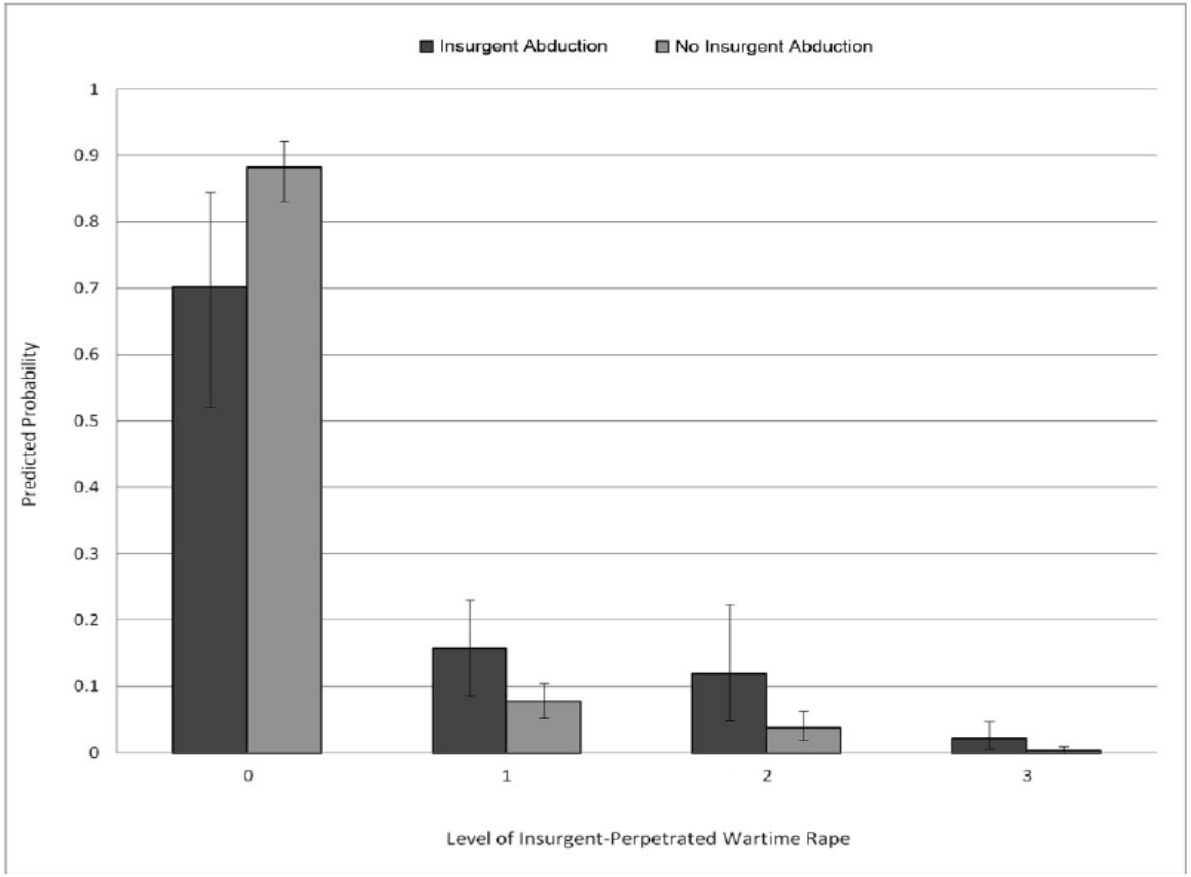

Note: Ordered probit model with standard errors clustered by conflict. Each simulation includes ethnic war, magnitude of state failure, aim, fertility rate, genocide by insurgents, contraband, Polity2, duration, year, and population (log) (all set at their mean values). The error bars represent the 95% confidence interval for each predicted probability value. Estimates calculated using CLARIFY. Two-tailed t-tests show that the differences in the mean predicted probabilities at levels 0, 1, and 2 are statistically significant at the 5% or 10% level. The difference between the mean predicted probability values for level 3 wartime rape falls just short of statistical significance ($p = 0.13$).

FIGURE 2. Probability of State-Perpetrated Wartime Rape with and without Pressganging

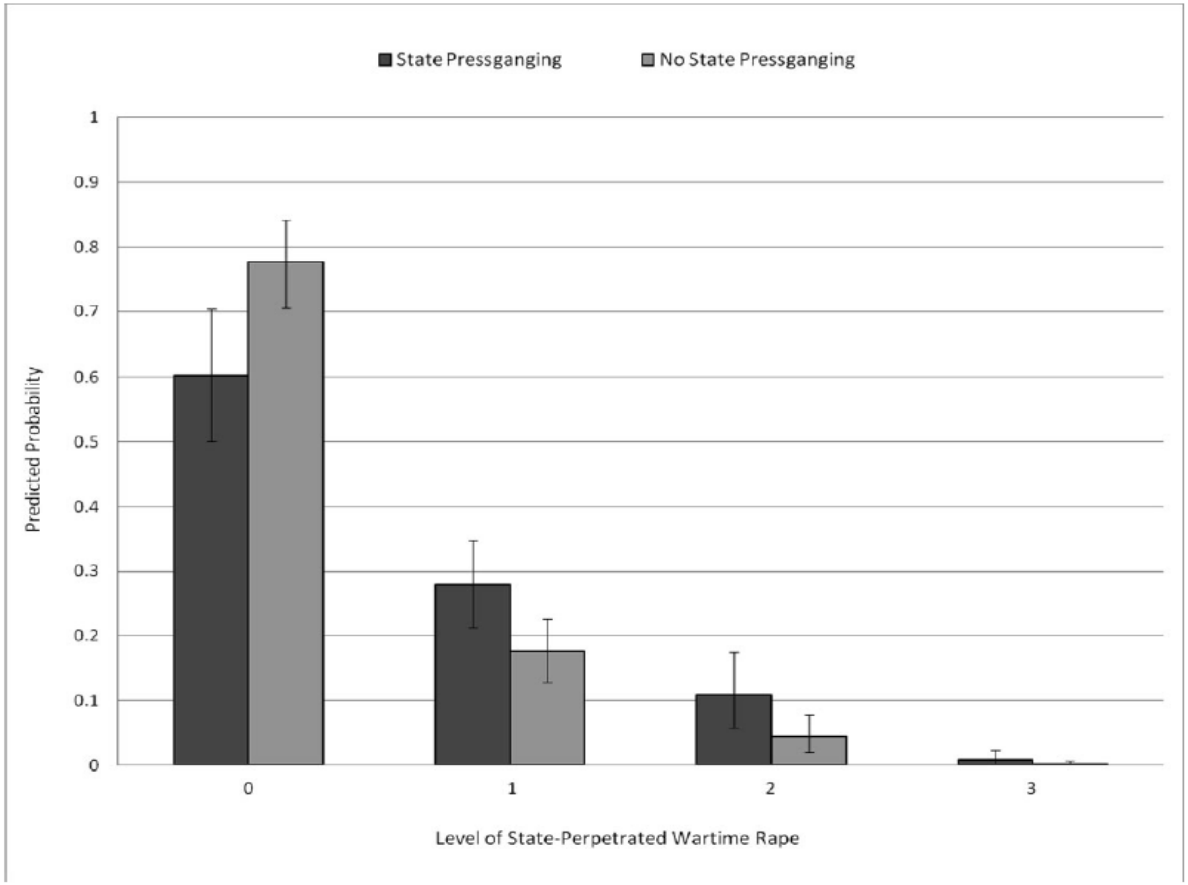

Note: Ordered probit model with standard errors clustered by conflict. Each simulation includes ethnic war, magnitude of state failure, aim, fertility rate, genocide by state actors, troop quality (log), Polity2, duration, year, and population (log) (all set at their mean values). The error bars represent the 95% confidence interval for each predicted probability value. Estimates calculated using CLARIFY. Two-tailed t-tests show that the differences in the mean predicted probabilities at levels 0, 1, and 2 are statistically significant at the 1% or 5% level. The difference between the mean predicted probability values for level 3 wartime rape is not statistically distinguishable from zero ($p = 0.26$).

The pictures on the previous two slides show that the no-rape scenario (0) is less likely when groups, either state or non-state, do not forcibly recruit their soldiers.

All higher levels of rape (1, 2 and 3) are more likely for both state and non-state groups that forcibly recruit their soldiers.

The main take home point Cohen wants to leave us with is that forced recruiting is associated with rape in war.