Angry or Weary?
The effect of personal violence on attitudes towards peace in Darfur

Chad Hazlett – Massachusetts Institute of Technology

November 2012

Abstract

Does an individual's exposure to personal violence during civil war tend to make them more "angry" and anti-peace, or more "weary" and pro-peace? Despite the fundamental nature of this question and its potential consequences for conflict duration, termination, and recurrence, it remains largely unanswered. This paper examines the effect of exposure to direct violence faced by Darfurian civilians on their attitudes toward peace four years later, using a novel dataset from a survey of Darfurian refugees in eastern Chad in 2008. Among individuals of a given gender and village, I argue that violence was committed indiscriminately, allowing its causal effect to be estimated. The results support the "weary" hypothesis: more violence leads to stronger reported beliefs that peace is possible and decreased desire to see perpetrators of violence executed. While generalizations to other conflicts would be premature, the implication in this case is that during conflict resolution and reconciliation efforts, it could be particularly fruitful to include those who have been directly harmed, as they may be advocates for peace rather than vengeance.

Chad Hazlett, Department of Political Science, 77 Massachusetts Avenue, Cambridge, MA 02139. E-mail: hazlett@mit.edu. I thank Jonathan Loeb, Ethan Siller, Benjamin Naimark-Rowse, and the entire *24 Hours for Darfur* research team for their tireless work in generating the dataset used here. Thanks to Jens Hainmueller, Adam Berinsky, Fotini Christia, Daniel Posner, Dan de Kadt, Teppei Yamamoto, and Daniel Hidalgo for invaluable comments and discussion.

I. Introduction

Mass violence against civilian groups is a common feature of civil wars, especially when relatively weak states attempt to defeat insurgencies embedded in or supported by civilian communities (Valentino et al.; 2004; Colaresi and Carey; 2008). Beyond the immediate and horrific human consequences of targeting civilians, these atrocities may also shape how long conflicts go on, how they end, and how stable the resolution can be. The effects of these atrocities on conflict trajectories could flow through direct reasons of military capacity if, for example, the attacks are successful in "draining the sea" and forcefully removing support for rebel groups. However, attacks on civilians can also have wide-ranging effects on the course of conflict through mechanisms involving the perceptions and attitudes of civilians. For example, attacks could generate increased support for the non-perpetrating side (Lyall et al.; 2011) in terms of recruitment, material support, or safe haven, through strengthening grievances, reducing opportunities for ordinary economic life, or generating the need for protection from the perpetrating force. Mass violence against civilians could also generate fears of future attacks, which can have long-lasting effects. Such concerns can make otherwise reasonable political arrangements impossible because they may not be able to identify credible arrangaements that ensure the safety of civilian groups once each side lays down arms. The demonstrated lack of security can also provide fertile grounds for dangerous security dilemmas to develop between groups (Posen; 1993), which can generate pressures to escalate violence in response to minor incidents between groups that would otherwise have been resolved peacefully.

Despite the many pathways by which mass violence against civilians can influence civilian attitudes and in turn shape conflict, little is known about the effects of violence on individual attitudes. In this paper I focus on arbitrating among two reasonable but directly opposing hypotheses regarding how civilians react to violence. First, the "anger hypothesis" states what most experts and non-experts alike might expect: that direct exposure to violence during conflict ought to make individuals more likely to less likely to seek peace or believe it is possible, and more likely to be angry or vengeful. Many mechanisms could produce this result. Most obviously, violence may harden divisive ethnic identifications, and generate new, stronger grievances. Alternatively, mass violence against civilians may act as a type-separating event, revealing that the perpetrating group is of a type that is willing to use such violence. It is reasonable for the victims of these attacks to expect that these types do not change quickly. As a result, counter-violence or preventive strikes become justifiable security measures, which can lead to further violence. Moreover, peace becomes apparently less possible, as it is difficult to imagine living in safety with former perpetrators of mass violence nearby. Violence may also increase recruitment into insurgent groups by generating demands for justice or reprisal (though see Lyall; 2009), by reducing opportunities for ordinary economic activity and thus lowering the opportunity cost of taking up arms, by showing that neutrality is not a viable strategy for avoiding violence, and by triggering a need for protection against the perpetrator of that violence. While these many paths differ in their details, they all imply that exposure to violence lead to a reduced desire for or belief in the possibility of peace.

Alternatively, the opposite outcome may be reasonably expected. The "weary" hypothesis states that exposure to violence makes one wish for peace more strongly or believe it more achievable.[1] Individual exposure to violence may make the costs of fighting a war more apparent, may alter calculations of whether it is worthwhile to pursue political aims rather than protect the pre-war status quo, and may produce psychological effects such as depression or post-traumatic stress (Pham et al.; 2004; Vinck et al.; 2007; Pham et al.; 2009), making individuals less interested in fighting.

Using a novel dataset from a random sample of Darfurian refugees in eastern Chad, and exploiting a conditionally exogenous component of the violence in Darfur in 2003-2004, this paper aims to make two contributions. First, by focusing on Darfur, it examines a conflict of great interest and significance – the first (and so far, only) conflict this century to be deemed a genocide by the U.S. government and International Criminal Court – but which has remained under-studied in empirical work due to severe logistical constraints. Second and most importantly, it seeks to answer whether exposure to direct violence, in this case, predominantly makes individuals appear more "angry" or "weary". It does so using a causal identification strategy based on conditionally exogenous variation in exposure to violence. This study thus

[1]Note that the terms "angry" and "weary" are merely shorthand for these hypotheses, each of which is the observable implication of multiple possible mechanisms. The terms are not intended to suggest that emotions as such, whether rational or not in their origin, are to blame.

has implications for theories considering the duration, termination, and recurrence of civil wars, as well as the social legacy of mass violence. While narrowly focused on the particular case of Darfurian refugees in eastern Chad, it also adds to a small but growing literature seeking to causally identify the effects of violence in a range of conflict zones.

Background

A small but growing literature examines the effects of violence on civilians during wars. While no other study has specifically examined the causal effect of exposure to violence on attitudes toward peace, studies have examined a wide-range of other possible effects of violence.

Due to the strategic nature in which violence is usually deployed, it is unsurprising that most past work on the effects of violence belong to a category of observational studies that provide strictly associational information. One productive area of this literature has focused principally on the psychological condition of those exposed to violence Pham et al.; 2004; Vinck et al.; 2007; Pham et al.; 2009, finding that exposure to violence is associated with higher levels of post-traumatic stress disorder and depression in post-conflict populations. Other studies of this type have examined how support for different warring parties is influenced by exposure to violence. Lyall et al. (2011) and Bullock et al. (2011) directly examined how violence committed in Afghanistan by either the the Taliban or the International Security Assistance Force (ISAF) may affect levels of support for these groups. They find an asymmetric effect: violence by ISAF forces appears to shift civilians into supporting the Taliban more often, while violence by the Taliban does little to increase the popularity of ISAF. They do not report, however, how exposure to violence by either group might influence perceptions towards the desirability or likelihood of achieving peace.

A second category of papers has placed greater emphasis on causal identification of the effects of violence during conflict, relying primarily on model-based inference to adjust for (observable) confounders. Several of these focus on social costs of violence in terms of education and earnings. Foltz and Opoku-Agyemang (2010) is unique in that it examines ongoing, low-intensity violence. Using data from multiple districts in Uganda, it compares changes in education rates over time in districts that were exposed to ongoing low-level (but systematic) violence to changes in education over time in districts that did not experience this violence. They find that exposure to ongoing, low-level violence after major conflict results in lower education rates. Similarly, Akresh and De Walque (2008) examines how differences in the intensity of violence during the Rwandan genocide relate to schooling outcomes. Using a variety of controls including pre-war education levels, in each district they compare cohorts that were of school-age during the violence to those who were not. They find a strongly negative effect of exposure to violence on schooling, with the more exposed children completing half a year less of education, or 18.3% less on average. Bellows and Miguel (2009) also attempts to make a causal claim about the effects of violence by using statistical controls in a regression framework. Postwar economic outcomes and social engagement are measured at the household level using a nationally representative survey after the conclusion of the 1991-2002 civil war. Using a series of statistical controls to attempt to eliminate selection bias, they find that individuals from households experiencing more intense war violence were more likely to attend community meetings, join local political and community groups, and to vote.

Finally, the third group includes "quasi-experimental" studies, in which violence is argued to have been exogenous, at least within certain sub-groups (i.e. conditional on certain covariates). Lyall (2009) examined the effects of mortar fire on villages by Russian soldiers in Chechnya, arguing that among villages with similar characteristics (such as distance from Russian bases), this mortar fire was effectively random. Variation in violence is thus measured at the village level rather than individually, as is the outcome: the number of reprisal attacks launched from a given location. Shelled villages experience a 24% reduction in attacks, relative to matches control villages that were not shelled.

Of particular relevance is the work of Blattman and Annan (2010) and Blattman (2009). In Blattman and Annan (2010), the authors argue that since recruitment into the Lord's Resistance Army (LRA) was forced, recruits did no self-select into joining. Among those individuals with the same village, age, and gender, there should be no important differences between those who get abducted and those who do not. Blattman finds that abduction is associated with approximately one less year of education, though this corresponds closely to how long they were forced to serve the LRA. Skilled employment is also much less common among those who were abducted, and incomes are lower by a third. Using the same dataset, Blattman (2009) also examines the legacy of LRA abductions on political engagement. Again, they find that strong and convincing evidence of the opposite effect: abduction led to large increased in both voting and community leadership.

Altogether, the previous literature cited here thus sheds little light on the basic question of whether experiencing violence predominantly makes individuals more “weary”, or more “angry”. Nevertheless, a striking consensus is beginning to emerge. Some (mainly observational) studies have found that exposure to violence is frequently associated with worse psychological outcomes in observational studies (Pham et al.; 2004; Vinck et al.; 2007; Pham et al.; 2009), and direct costs in terms of years of schooling and future earnings (Blattman and Annan; 2010) in studies of all types. However, the common theme finding from studies using model-based causal inference or quasi-experimental approaches is the more counter-intuitive finding that those exposed to greater violence tend to be *more* socially and politically engaged. While not directly translatable into “weary” versus ”angry” attitudes, this paper mirrors that growing consensus in so far as the increase in pro-peace attitudes seen here corresponds to increasingly pro-social attitudes and behaviors identified elsewhere.

II. Methods

Data

The primary data source is a survey conducted in 2009 by the “Darfurian Voices” project, funded in part by the U.S. Department of State. It sought to systematically document the views held by Darfurian refugees in Chad on issues of peace, justice, and reconciliation, and to accurately transmit these views to policymakers, mediators, negotiating parties, and other key stakeholders. The project consisted of two components: (1) a random-sample survey representative of the entire Darfurian refugee population living in the refugee camps in eastern Chad, and (2) in-depth interviews with tribal, civil society, and rebel leaders also living in Chad. The full report and other materials from this project can be downloaded at `http://www.darfurianvoices.org`. In this paper, I rely mainly on results from the random-sample survey. Briefly, the sample includes 1,872 individuals from the target population of adult (18 years or older) refugees from Darfur living in all 12 Darfurian refugee camps in eastern Chad. We used a stratified random sampling method, with geographic location (camp and block) and gender as strata. Further details of the sampling methodology are given in the Appendix.

Sample Inclusion for Analysis

To ensure the comparability of those exposed to different degrees of violence, this analysis considers only those who report leaving Darfur due to direct violence, ruling out those who left long before violence occurs. This eliminates approximately 20% of the sample. Second, for purposes of restricting attention to non-leadership civilians specifically, only those who report being non-leaders both in Darfur and while in the camps are considered. Those who were members of the tribal administration in Darfur are considered to be leaders, and thus excluded. The remaining sample size is 1345. Practically speaking however, the requirement of conditioning on village effectively drops individuals who are the only representatives from their village. This reduces the effective sample size to approximately 1100 depending on the exact model.

Measurement

Regarding the "treatment" – exposure to violence – I focus exclusively on whether or not the respondent was the victim of direct violence during this conflict, which is coded as a binary variable $DirectHarm$ indicating injury or maiming during an attack.[2] The benefit of this measure is that it speaks directly to an individual's exposure to violence above and beyond the experiences of those around them. All violence-related questions come at the end of the survey to avoid possible priming effects. People are not asked to describe the violence, and in particular, women are not asked whether the violence was of a sexual nature. Thus while there may be under-reporting, particularly among women who are reluctant to report surviving sexual violence, this is mitigate somewhat.

The survey offers many possible questions regarding attitudes that could be influenced by exposure to violence. To narrow the focus of the analysis and avoid an explosion of the

[2] The exact question was: "Have you suffered violence, or have you been physically maimed in an attack related to the current conflict? (a) yes; (b). no; (c/d/e) uncertain/refused/not understood"

family-wise error rates for hypothesis tests, only several outcomes are examined here. The first is *EnemiesPeaces*, which asked if it is possible for former enemies to live peacefully together after war.[3] The second is *PeaceWithJJIndividuals*, which asks if it is possible to live in peace with individual members of the *Janjaweed*[4], and the third is *PeaceWithJJTribes*, which asks whether it is possible to live in peace with the tribes from which the *Janjaweed* came.[5] These questions differ in important ways – one suggests living together with the very people who have committed violence against the respondent's group; the other suggests living together with members of the same group, but less likely with individuals culpable for violence. All three response items are transformed into binary responses, with the split point chosen to maximize variance (i.e. the point nearest the median). Note that the direction of coding is such that the more "weary" or peaceable answer is always positive, while the less peaceable answer is always negative.

Fourth, *ExecuteGoS* is a variable that asks what punishment is appropriate for Government of Sudan soldiers involved in the conflict. This is coded as a "1" when the answer was "execution", and "0" for any other (lesser) punishment, and so points in the opposite direction to the previous three (the more "weary" answer now being the lower value). These four measure are highly inter-related, with pairwise correlations coefficients ranging from 0.22 to 0.61 in absolute value. A factor analysis on these four variables retains only a single factor, with the expected signs on the factor loadings.[6] Using these loadings, and then rescaling by the sum of the weights, I create the variable *PeaceFactor*, for use when a single measure of the outcome is desired. Note that the re-scaling allows $PeaceFactor$ to be interpreted as a weighted average of the four indpenednet variables, with the same overall scale, such that a 10% shift in $PeaceFactor$ is comparable to a 10% change in one of the component variables.

[3]Some people say that it is possible for former enemies to live peacefully together after a war. Some people say that it is not possible for former enemies to live peacefully together after a war. Do you believe (a) strongly that it is possible; (b) somewhat that it is possible; (c) somewhat that it is impossible; (d) or strongly that it is impossible?

[4]In the future, I can see myself living peacefully with actual members of the *Janjaweed*: (Strongly agree/ somewhat agree/ somewhat disagree/ strongly disagree).

[5]In the future, I can see myself living peacefully with the tribes from which the *Janjaweed* came: (Strongly agree/ somewhat agree/ somewhat disagree/ strongly disagree).

[6]Principal factor analysis, no rotation. Factor loadings were 0.68,0.63, and 0.79 for *EnemiesPeaces*, *PeaceWithJJIndiv*, and *PeaceWithJJTribes* respectively, and -0.35 for *ExecuteGoS*, which would be expected to have the opposite sign as the higher value is the more "angry" answer in this case

While the loading of these four variables onto a single factor suggest that they measure a single underlying construct, what this construct actually represents is not known. These questions – particularly the three relating to living in peace – are difficult ones. They require the participant to evaluate counterfactual circumstance and estimate the chances of a complex process leading to a particular outcome. Moreover, they are emotionally charged, and come towards the end of a gruelling two hour interview. For these reasons, I expect that a "substitution" effect (Kahneman; 2011) occurs. Specifically, in order to answer a difficult and emotional questions about the *possibility* of living in peace, respondents are likely to instead answer an easier and more intuitive question such as "would I like to live with these groups?" or "how would I feel about living with these groups?". The fact that *ExecuteGoS* differs widely in meaning from the other three variables and yet they load onto a single factor is suggestive that an underlying construct is involved that relates more to individuals' intuitive or emotional reaction to the idea of peace rather than to a cool-headed assessment of what each question actually asks.

Identification Assumption: Conditional Exogeneity

The central aim of this paper is to identify micro-level effects of exposure to violence during atrocities in Darfur. The most critical assumption is that of "selection on observables": conditional on some set of observed covariates, whether an individual experiences violence must not depend on the outcome an individual would have if exposed to violence, or the outcome they would have if not exposed to violence. This assumption is more precisely stated in terms of the potential outcomes framework or Rubin Causal Model (Rubin; 1974). Let y_i^1 designate individual $i's$ (possibly counterfactual) outcome had they been exposed to violence; let y_i^0 be the same individual's (possibly counterfactual) outcome under non-exposure to violence. The causal effect for unit i is then defined as $y_i^1 - y_i^0$, and the average treatment effect (ATE) over the population is $E[y_i^1 - y_i^0]$. Let D_i be an indicator of exposure to violence for individual i, while $gender_i$ and $village_i$ designate the gender and village of respondent i. The identification assumption here is then expressed as $(y_i^1, y_i^0) \perp D_i | gender_i, village_i$. That is, among individuals of a given gender and village, those who experience violence and those who do not look "the same", in that their non-exposure outcomes (y_i^0) or exposure-outcomes (y_i^1) do not depend on the actual treatment assignment. For convenience, I refer to this as the "conditional exogeneity" assumption, since it implies that violence is exogenous conditional on village and gender.

The Distribution of Violence

Justifying the conditional exogeneity assumption requires first some background on the nature and purpose of the violence involved. In February 2003, rebel groups began attacking government military outposts in Sudan's eastern Darfur region. The attacks were met with brutal retaliation against civilian populations presumed to be supportive of the insurgency, including but not limited to the largest groups – the Fur, Zaghawa, and the Masalit. During the height of violence in Darfur in 2003/4, wide-spread violence against civilians was employed throughout Darfur, including the state of West Darfur, from which almost all the survey respondents in this study originate. This violence has been described as genocide by some, including the U.S. State Department and Congress and in indictments by the International Criminal Court. Despite the use of ethnic or tribal affiliations to mobilize violence on both side, many of these one-time ethnic alliances have broken down, and it is widely accepted that the use of violence of this type was largely instrumental, orchestrated by the Government of Sudan as "counter-insurgency on the cheap", a bloody tactic habitually used by the Government in Darfur as well as elsewhere in the country (De Waal; 2004).

Critically, the aims of these attacks were not to selectively route out rebel or political leaders as is the premise in much work on violence against civilians during civil war Kalyvas; 2006. Instead, it was to punish or destroy the *communities* behind rebel groups, through both direct violence against the populace and forced displacement. Displacement served a second purpose of incentivizing members of the *Janjaweed* militia, whose tribes have long sought more reliable and secure access to grazing lands. Violence was thus not generally targeted at specific individuals, but at communities.

According to satellite imagery collected by the U.S. State Department, 2964 villages have been confirmed as destroyed between February 2003 and December 2009, with 1700-1800 of these being destroyed in 2003-2004 (U.S. State Department, 2010). When a village was under attack, it would typically involve one or both of the following: first, government of Sudan planes would often begin crude and indiscriminate aerial bombardment. Second, militias, generally referred to by respondents in our survey as *Janjaweed* would charge into the village, during which time many would be killed and many women were raped. Often times, both strategies were employed together, in this order. Other times only one strategy would be used.

In the case of government bombing of villages, within a given village it is relatively straightforward to claim that one's chances of dying or losing family members is largely random (excepting family size). These villages are relatively small, allowing for little variation in targeting. The bombing techniques involved methods as crude as pushing bombs, scrap metal, and barrels full of shrapnel out of aircraft. This simply does not allow for any kind of strategic targeting within the village level.

The *Janjaweed* attacks, too, produced effectively exogenous exposure to violence within a given village, conditional on gender. Women and girls were targeted for rape on a massive scale. Beyond the use of different types of violence for males versus females, the *Janjaweed* not only appeared to be indiscriminate in their use of violence, but also were unlikely to have any knowledge of what individuals in the village were potentially more or less politically or militarily active. Villages are ethnically very homogenous, and thus while certain villages may be targeted, within village, there was no basis for ethnic targeting. Men and women, the old and young, were all apparently subject to injury and killing.

Those in the village, whether sleeping or attempting to flee, were subject to attack. Even those fleeing to nearby hiding places were frequently pursued. Livestock and belongings were often stolen (97% of respondents in our sample reported losing all or most of their livestock, crops, and belongings), and villages were almost always burned to the ground. One immediate concern is that some individuals would have been more likely to have resisted or counter-attacked, and also more likely to experience violence. This is relatively unproblematic for two reasons. First, during the phase of violence experienced by those in the survey, resistance within the village had become extremely rare. One reason is that once the government had clearly joined the effort using its aircraft, this was no longer a war among tribes, and the would-be resisters among the Fur, Massalit, and Zaghawa tribes realized that protecting the village was not an option.[7] Relatedly, those who did wish to resist in this area had already left to join rebel groups operating outside the villages (and do not enter our sample). Second, it is important to note that those who hid or attempted to flee were not evidently shown mercy. Testimony describes how those who fled or hid in the moment of the attack were often chased and were also likely to be killed, raped, or otherwise attacked.

In short, some villages may be targeted more or less than other villages, but within a given village, bombs cannot discriminate among those with different potential attitudes, while the *Janjaweed* militia lack both the desire and the knowledge to target civilians selectively. Within a given village and gender sub-group, whether or not an individual experiences violence appears to be unrelated to what their attitudes would have been under non-exposure. Note that age is another characteristic that would be evident to the attackers, so it is possible that selection on this basis occurred as well. For example, the *Janjaweed* could have been given instructions to make efforts to kill those of potential fighting age. However (a) examination of testimony given by those present at attacks does support an age-specific targeting approach; (b) the data themselves show no such age-related effect in exposure to violence (at least among those that made it to the refugee camps where our survey was conducted).

Interviews collected during our survey process provide an enormous amount of information about the nature of attacks and largely support this view of the nature of the attacks. I provide only one such example here, though much can be learned by examining such testimony at great length. One respondent recounted, "The government came with antonovs (aircraft), and targeted everything that moved. They made no distinction between the civilians and rebel groups. If it moved, it was bombed. It is the same thing, whether there are rebel groups (present) or not...They shoot everyone when they see them from a distance, and they have any doubt about him, they shoot him. The government Antonovs survey the area from time to time to see if there is anything moving and if it is a human or an animal...The government

[7] Anecdotal evidence suggests that by the time the government of Sudan had become involved the fighting so clearly, villagers realized that fighting in the villages was not an option. "Once the government appeared with it's aircraft," one respondent reported, "we knew this was not another war against the Arabs", meaning that this was not a war among well-matched powers which they were honor-bound to stay and fight, but rather a slaughter by a superior government force, requiring them to flee.

bombs from the sky and the *Janjaweed* sweeps through and burns everything and loots the animals and spoils everything that they cannot take." Such statements are corroborated by very similar reports collected by other organizations at other times, such as those collected in Human Rights Watch (2006).

It remains possible to challenge this claim. One possible exception is that while rare, testimonials indicate the on occasion individuals were rounded up and interrogated in an effort to get information on rebel leaders. None of the cases, however, suggest that any information of this kind was able to be employed during village attacks. Those who were interrogated seem to have been chosen at random or by convenience, so even if they experienced more (or less) violence, it would not bias the apparent effect of violence on attitudes. Another possible exception is that not all violence occurred in the villages. Markets were also the cite of violence against traders and merchants in particle. Other potential violations of the identification assumption are examined in the discussion section.

This pattern of violence is broadly similar across localities, though there are some variations. Virtually all attacks involved both militia and Sudanese Army personnel (see de Waal et al., 2010 for similar findings from later in the conflict). Aerial bombardment sometimes preceded attacks, and sometimes did not. The militia sometimes occupied the village, but generally moved on immediately. Villagers sometimes had warning information of immediately impending attacks, but generally did not. While men were selectively killed and women selectively raped, there seems to have been little selection on age - neither infants nor elders were spared. In some cases women and children benefited from hiding out in their hut; in other cases this strategy led to tragedy as soldiers or militias entered these huts or burned them down with people still inside.

III. Results

Balance

The first step in the analysis is to establish the plausibility of the identification assumption that individual exposure to violence is exogenous conditional on village and gender. A traditional balance test does this by testing whether the treated and untreated at least look similar on covariates measured prior to treatment. However in this case, we only expect treated and untreated units to look similar within each village and gender sub-group. I therefore test "conditional" balance, first splitting the sample by gender, and then within each, regressing the treatment indicator on each covariate and a series of dummies representing each village. This is done in multivariate analyses, simultaneously using all the pre-treatment covariates to model the probability of exposure to violence. If a covariate is conditionally unrelated to the probability of receiving treatment, then its coefficient in this regression will be statistically indistinguishable from zero. Moreover, if *all* the covariates are have the same distribution among the treated and untreated after controlling for village and gender, then a test for the joint significance of these covariates will show they do not have predictive power.

Covariates are included in this analysis if they are certain to be pre-treatment (measured prior to treatment or clearly not altered by the treatment), and if there is sufficient variance that the less frequent outcome occurs at least 10% of the time (for binary variables). These include age, wether they were a farmer, herder, merchant, or trader in Darfur, their household size in Darfur, and whether or not they had voted in the past. All the results here are linear probability models, with heteroscedasticity-robust standard errors.

The results of these balance tests are promising (tables 1 and 2). The only covariate with a $p-value$ of less then 0.10 is $herder\_dar$: herders appear to be more likely to experience direct harm. This is concerning, and will require conditioning on herder status to ensure this is not acting as a confounder, though herders make up only 15% of the sample, and dropping them does not effect the results reported below. The primary concern is whether the covariates other than village are *jointly* predictive of who experienced violence or not. Comparing each full model to a restricted one with only the village fixed effects, I find that these covariates are not jointly predictive of treatment status for either men ($F(8,338) = 1.10, p = 0.37$) or women ($F(6,321), p = 0.43$).

Distributions of Treatment Probabilities

In experiments with (unconditional) randomization of the treatment, balance can be examined visually by estimating propensity scores and comparing the distribution of these propensity scores (the probability of exposure to direct harm given an individual's values on pre-treatment covariates) for the treated and the untreated. Here, I seek to perform an analogous test, however the treated and untreated are only expected to have similar distributions of propensity scores within each stratum of gender and village. Conditioning on gender is easily achieved by separately plotting male and female propensity scores. Conditioning on village can be achieved by a re-weighting procedure. For each participant exposed to violence, a weight of "1" is assigned (though units are dropped if there is no corresponding control unit from the same village). For each participant who reports not being directly harmed, I weight the data such that the post-weighting number of untreated participants from each village is the same as the number of treated units. This is accomplished by assigning weights $\omega_i = \frac{P(Village=village_i|D=1)}{P(Village=village_i|D=0}$, where $village_i$ is the village from which participant i originates. This ensures that the post-weighting distribution of villages is the same for the treated as the untreated, thus any differences in the distribution of propensities to treatment are not due to differences in villages of origin.

Figure 1 shows the gender-specific distributions of propensity scores *prior* to this re-weighting by village. Clearly, the balance is not good, reflecting that some villages experienced much more complete violence than others. However once the untreated observations are re-weighted to adjust for differences in village of origin, the balance is extremely good, with very similar distributions of propensity scores for the treated and untreated (figure 2). This boosts our confidence that those units within a single village and gender group are exposed to violence in ways unrelated to any of the observed pre-treatment covariates.

OLS Estimates

Treatment effects were estimated using linear regression models (OLS), Mahalonobis matching (matching), and regression with weights determined by entropy balancing (ebal). For OLS, the first question is how to achieve the necessary conditioning on village and gender. Conditioning on village is handled by village fixed effects. Conditioning on gender would be achieved with the fewest assumptions by either splitting the sample on gender or, equivalently, including a gender dummy while also interacting it with every other covariate included in the model. However, this requires effectively cutting the sample size in half when estimating each gender-specific effect, which is too damaging to statistical power given that the number of village fixed effects to estimate is already very large. Therefore conditioning on gender is implemented simply by including a gender dummy in the model.

Given the identification assumptions, it should be necessary only to regress the outcome on the treatment and village and gender dummies. I refer to this model as the "short" specification. Adding further covariates to the model is not required for unbiased estimation, but allows these (pre-treatment) covariates to explain additional variation, possibly improving the precision of all estimates. I refer to this as the "long" specification. In either case, estimated coefficients are interpreted as the best linear approximation to the sample average treatment effects (SATE).

Table 3 shows the results from short and long OLS models. Both show the same pattern of treatment effects: those who report being directly harmed are approximately 10% more likely to say it is possible to live in peace with "former enemies", to live in peace with individual members of the *Janjaweed*, or to live in peace with the tribes from which the *Janjaweed* were drawn. All results on these three outcomes, under either model, fall in the 8-11% range. Those directly harmed are also 9-11% less likely to penalize Government of Sudan soldiers to death.

Unsurprisingly, the factor created by a weighted sum of these four, $PeaceFactor$, is also significantly affected by $DirectHarm$, rising by 13% among the treated. Together, these results consistently point towards the hypothesis that exposure to violence stimulates a *greater* desire for or belief in the possibility of peace, and lesser desire to punish enemies to death. Evidence thus appears to lie in favor of the "weary" argument, rather than the "angry".

Finally, note that the effects found here are not apparently due to one or two heavily represented villages that happen to show large effects. This is best demonstrated by examining the distribution of village-level effects. Given the identifying assumption, the casual effect of violence is identified within each sub-group of individuals of the same gender from the same village. Here I examine the distribution of village level effects, by regressing the outcome (either $PeaceFactor$ or $EnemiesPeace$) on $DirectlyHarmed$ and a gender dummy. This estimates the effect of $DirectHarm$ per village. In figure 4, I first show the effects of $DirectHarm$ on two outcome in panels (a) and (b), giving each village equal weight. As expected, there are some extreme values, largely coming from villages that are poorly represented in the analysis, and thus producing imprecisely estimated village-level treatment effects. I then re-weight these villages by how often they appear in the analysis (panels (c) and (d)), to indicate how the village-level effects drive the veral average. These weighted histograms indicate that the average effect is driven by village-level effects on the order of 0-0.10 and 0.10-0.20 (close to the average of roughly 0.10). The average effect over the whole sample is thus not due simply to a few extreme villages, nor is it the result of highly polarized village-level effects.

Entropy Balancing

I also employ entropy balancing (Hainmueller; 2012) as an additional means of estimating effects with reduced model dependency. This approach chooses weights for the control units such that after weighting, the marginal distributions of covariates is the same for the treated and untreated up to a specified number of moments, while keeping the weights as close as possible to equality. These weights can then be used in later analyses such as selective averaging or regression to estimate a sample average treatment effect on the treated. Here I require that post-weighting distribution for covariates among the controls matches the distributions of those covariates among the treated units. Entropy balancing is successful in equating the means and variances of the covariate distributions. I then employ these weights in regressions with village fixed effects to complete the required conditioning.[8] Again this is done with (a) the shortest possible specification to achieve identification – using only village fixed effects – and (b) a longer specification where covariates are included in the regression stage for additional robustness.

The results are very similar to those produced by earlier analyses: respondents directly harmed by violence are 8-12% more likely to give the pro-peace or "weary" response to all questions, all highly significant. The common factor, $PeaceFactor$, rises by 14-16% among those exposed to direct violence, commensurate with the increases on the individual items after re-scaling.

[8]Conditioning on village through the fixed effects regression is sufficient here to estimate within-village effects and aggregate across these, as desired. While it is possible in principal to include village dummies in the entropy balancing stage, this would effectively add over 400 additional variables to balance on.

Matching

Finally, matching offers another alternative to estimate the effect of the treatment while conditioning on village, gender, and other covariates. Here the aim is less to improve balance on observables (other than village and gender), since the conditional balance already proves to be largely satisfactory. Rather it's benefit is to allow for conditioning on covariates in a way that is less model dependent than linear modeling. Mahalanobis matching was used, with 1-to-1 matching without replacement. The estimand is the sample average treatment effect on the treated (SATT). The variables matched on were the same as those in the multivariate models above: all available pre-treatment variables with enough variation such that at least 10% of the participants fall in the smaller group. Matching is exact on all variables except age and household size in Darfur. Post-matching balance tests showed no statistically significant imbalances on any covariates.

Table 5 shows estimates from the matching analyses. The findings are consistent with the regression estimates, though larger and more significant in some cases. While the number of observations is substantially lower due to the strict matching requirements, the more precise estimates also allow gender-specific effects to be estimated more precisely than under regression. The effects all lie in the same direction for men and for women, however the effects for men tend to be larger. The only outcome on which the effects dramatically differ by gender is $PeaceWithJJTribes$, on which men see a large 20% increase in positive responses after exposure to violence, while women see an insignificant change of only 3%. $GoSExecute$ is also significantly negative for men (as it is in the overall sample), and negative but non-significant among women. Otherwise, all the effects that were significant for men or for the overall sample are significant among women as well.

Results from direct OLS, OLS on the entropy-balanced data, and matching are summarized in figure 3.

IV. Discussion

The findings of these analyses are highly consistent: across the different modeling approaches shown, those exposed to direct violence are approximately 10% more likely to report the more "weary" or pro-peace answer. Given that these variables load onto a single factor despite being about topics ranging from the possibility of peace to the appropriate punishment of Government soldiers, and given the difficulty of answering these questions and their emotional content, I interpret these effects as revealing principally an increased desire for peace among this group, relative to those refugees who are otherwise comparable but who were not directly harmed. Here I discuss several factors that either threaten the validity of these findings, or suggest alternate interpretations.

Robustness of Identifying Assumption

Possible violations of the identification assumption are a primary concern. In addition to qualitative evidence for the indiscriminate use of violence, the balance tests, and the distribtion of propensity scores, I also consider the likely direction of bias caused by a failure of this assumption, and conduct placebo tests and sensitivity analyses.

First, confidence in the reliability of the results is boosted by the direction of the findings, which is opposite to what might be expected if the result was due to most imaginable confounders. Specifically, if some unobserved characteristic makes some people more likely to experience violence, we would expect those people to also be those who are *less* interested in peace and less likely to believe it is possible. This would cause a bias in the direction of the "angry" hypothesis, opposite to the observed effect. Nevertheless it remains possible that some confounder makes more conciliatory individuals more likely to expose themselves to violence, thought this seems less plausible, and sensitivity analyses below show that such a confounder would have to be substantially more confounding than any of the included covariates would have been had they been unobserved.

Second, the variable *pastvoted* is a useful "placebo" outcome because it occurred prior to violence and so should not be affected by it, but if there is some selection into violence or targeting of politically active individuals in the villages, it would likely be correlated with *pastvoted*, and this would create an apparent effect of $DirectHarm$ on *pastvoted* that is due only to a violation of the identification assumption. However, a test using *pastvoted* as a placebo outcome finds no effect of $DirectHarm$ on *pastvoted* ($\beta = 0.02$, $p = 0.62$) (estimated from entropy balancing followed by a fixed effects regression with covariates). Notably, using the same specification, exposure to direct harm increasing the probability of participants reporting they *would* vote in future elections (*wouldvote*) by 11% ($p < 0.01$). The fact that $DirectHarm$ has an effect on *wouldvote* but passes the placebo test on *pastvote* is thus reassuring evidence that no self-selection into violence or targeting of politically active individuals seems to be occurring conditional on village and gender.

Third, two types of sensitivity analyses are useful for examining the sensitivity of the results to violations of the identification assumption. On the matched data, using the approach of (Rosenbaum; 2002), I find that the p-value for the effect of $DirectHarm$ on $PeaceFactor$ is robust to a γ of 1.45. In brief, this implies that for some confounder to be strong enough that the actual treatment effect is no longer significant, it would have to make the unit with the higher value of y in each matched pair 1.45 times as likely to have received the treatment than the other unit in that pair. In our case, such a confounder would be one that, among two individuals from the same village and gender and matched on other covariates, it makes the individual with the more pro-peace attitude 1.45 times as likely to have experienced direct harm. This seems unlikely, especially given that we would expect the opposite – that those with higher pro-peace attitudes would have been less likely to be harmed.

For the regression approach, I conduct an analysis similar in spirit to Imbens (2003). Suppose the "true" model is $y = X\beta + Z\gamma + \epsilon$, where X contains the treatment, intercept, and covariates, β is the true (causal) effect of each variable in X on y, Z is an unobserved confounder, and γ is the effect of this confounder on y. If we estimate this model using OLS on only the observables (X), then $\beta = \hat{\beta} - \gamma(X'X)^{-1}X'Z$.Figure 5 shows the "true" treatment effect implied by varying the degree of confounding, indexed by the effect of Z on $PeaceFactor$ (γ) and the partial relationship between the confounder and the treatment controlling for the rest of X (i.e. $(X'X)^{-1}X'Z$, or equivalently $E[Z|directharm = 1, X] - E[Z|directharm = 0, X]$). For comparability, the plot shows the confounding effects of each covariate included, had it not been observed. This shows that in order for an omitted confounder to reduce the true effect so far that it cannot be distinguished from zero (the red dotted line), it would have to be a considerably stronger confounder than any observed covariate. For example, given a confounder Z, to produce a treatment effect statistically indistinguishable from zero it would have to be as strongly correlated with $DirectHarm$ as age, but would need to have an effect on $peacefactor$ more than 10 times larger than that of age. The covariate that would have been most confounding had it been left out is $female$. However for a confounder that is as strongly related to the outcome as $female$ (which is substantial) to reduce the treatment effect below significance, it would have to be three times more strongly related to the treatment.

Robustness to Other Violations

Two other threats to the validity of this analysis are spillover, and survivorship biases or other selection pressures on the population that we are able to sample.

Regarding spillover, in the present case one cannot reasonably assume that the "Stable Unit Treatment Value Assumption" (SUTVA) assumption of zero spillover is valid. Instead, it is more productive to examine possible violations of the no-spillover assumption, and determine how each would alter the meaning of the estimated effect. One possibility is that, when person j is exposed to violence, the effect of j's exposure on person i's non-treatment outcome (y_i^0) is, on average, in the same direction as the average treatment effect. This is the view of "partial treatment" – i.e. that those who are not exposed to the treatment directly but who know others who are exposed get effectively a smaller dose of the treatment. If this is the case, it ensure a bias towards zero on the estimated treatment effect. In this sense, it is not a problem for the analysis, but suggests that the true effect is stronger than the estimates here.

Alternativley, "negative" spillover may be possible: it could be that when person j experiences violence, it's effect on person i is on average opposite in direction to the average treatment effect. Similarly, it could even be that there is no substantial treatment effect on j directly, but a spillover effect on i that is oppposite to the average treatment effect as it would be estimated. This is not entirely unreasonable: it could be that those who witness violence or have family members and friends injured get more angry, while those who are directly harmed instead desire peace. Such a scenario could be induced by "survivor's guilt". Another would be that due to cultural pressures, individuals are expected to be outraged and seek vengeance when their communities have been attacked, but those who are directly harmed are able to press for peace without fear of violating these norms. If such scenarios occur, they influence the interpretation of the treatment effect. For example, what appears to be a pro-peace effect of violence on those who are harmed, may instead be an anti-peace effect of violence specific to those who are not directly harmed, or some combination of the two.

These cannot be entirely ruled out and should be borne in mind as possible explanations for the observed effect. However he data do not show evidence for substantial spillover of either the partial-treatment type or the negative type. In addition to asking about $DirectHarm$, we also asked individuals how many family members were killed or maimed, whether they witnessed other family members being injured, or whether they witnessed non-family members being injured. Using the same specifications and models as above but with these measures of indirect exposure to violence as the treatments, no significant effects are found in either direction. Thus neither positive (partial-treatment) spillover nor negative spillover appear to be powerful explanations.

Another potential threat is that some respondents are of a "sophisticated" type, who seek to show the survey enumerators that (a) they have suffered and are thus in need of support from donors, and (b) are of a pacific, conciliatory nature, more likely to attract donors to continue supporting the camps. This is effectively a concern about non-classical measurement error: the error or mis-representation on the measurement of the treatment (part a) may be correlated with error on the outcome (part b). However this is unlikely to explain the observed effect here: if strategic mis-representation of this type was driving the effect, we would also expect to see a (false) effect for indirect forms of violence, such as the loss of family members. The same individuals would be expected to over-report there losses on these measures, again creating a confound. However we see no measurable effect of indirect forms of exposure.

Survivorship Biases

A final set of concerns is that the population from which we sample is censored in some way that may influence the results. It is certainly true that the population from which we sample is in no way representative of the population of Darfur: individuals only appear in the population studied here if they survive the initial attack, chose to come to refugee camps in Chad rather than seek refuge elsewhere or join the rebel movements to stay and fight in Darfur, and survive the trip to Chad.

To the degree that these selection pressures on who makes it to the camps are experienced by both those who are directly harmed and those who are not, it alters the population about which we make inferences, but causes no bias on the causal estimates. In contrast, selection pressures that occur differentially depending on whether one is directly harmed could cause a bias. The first concern of this type is that among those who are directly harmed during an attack, the chances of death are higher. It seems plausible, however, that among those who are directly harmed, whether they survive that injury or not is unrelated to their potential attitudes. Likewise among those who are not harmed, whether they survive is surely uncorrelated with their attitudes. As long as this reasonable assumption holds, then the higher death rate among those who are injured does not introduce a bias.

Another concern regards selective mobilization into to rebel groups in a way that depends on $DirectHarm$. Among those present during the attack but*not* directly harmed, the more "angry" ones may have joined the rebel movements rather than coming to the refugee camp, while among those directly harmed, even the angry ones may come to the camp for medical care, regardless of their attitudes. This would bias the results, however it would cause a bias opposite in direction to the observed effect. If the more "angry" individuals from the unharmed sub-population join the rebel movements rather than coming to the camps, it would make the resulting non-harmed sample in the camp appear less angry, but we observe the opposite.

More Weary than Angry?

The evidence found here is strikingly consistent with the "weary" hypothesis, and provides no support to the "angry" hypothesis. This holds across a range of variables that collectively appear to relate to a single construct regarding the preference for peace rather than vengeance or continued violence. So far every test of the identifying assumption have shown it to be supported, and the bias that one expects from a violation of this assumption is almost surely opposite in direction to the observed effect.

The predominance of the "weary" hypothesis may seem surprising at first. Two types of explanations seem most plausible. The first is a cultural category, in which all individuals in the group are expected to show anger and a desire for vengeance in response to attacks. One version of this argument is a "culture of honor" explanation: particularly among communities that live far from government protections and have easily lootable capital stocks, which certainly applies in Darfur, it becomes important for groups to develop and maintain a reputation for toughness and the willingness to use reciprocal violence (Nisbett and Cohen; 1996). However individuals who are injured earn an "exemption" from this obligation. They may have the credibility to speak out against the use of violence, without being branded as cowardly. Another variant on this explanation is the "survivors guilt" argument, in which individuals who were not harmed are obligated to be outraged, while those who are actually harmed are not. However these arguments that rely on "negative spillover" are somewhat less likely given that we see no effect of variation in exposure to indirect forms of violence on attitudes.

Alternatively, a simpler calculus of perceived cost and benefit may be at work. Symbolically important events, rhetoric, propaganda, stories about past losses and mis-treatments, and any other source of grievances will lead to an increased belief that it is impossible to live together peacefully with the enemy. Having conducted this survey in refugee camps, all participants have similar exposure to this environment. When assessing the causal effect of direct exposure to violence, we are examining the impact of such exposure to violence *above and beyond* the effects of these shared grievances. It is less surprising then, that those who share similar grievances but who additionally have been exposed to greater costs of conflict, may be more ready to end the conflict. In short, "anger" may be a prominent characteristic of the background condition shared by all participants in the study, but individual exposure to violence appears to increase the "weariness' of those who experience it.

V. Conclusion

Violence against non-combatant civilians is a common feature of many civil wars, and beyond the obvious human cost of such violence, it can shape the possible outcomes of a conflict. Yet measuring the effects of violence on individual attitudes has been difficult, since violence is often strategically distributed and cannot be randomized.

This paper contributes to a small but growing literature that exploits identification opportunities in which exposure to violence is argued to be (conditionally) exogenous. It is the first paper seeking to causally identify the effects of exposure to violence on attitudes towards the possibility of peace. In particular, I focus on arbitrating between two reasonable hypotheses that make important but opposite predictions: does exposure to violence make individuals more "angry" and likely to view peace as impossible, or more "weary" and ready for peace?

Using regression, regression on a balanced sample, and matching analyses, the consistent finding is that while some forms of violence had no detectable effect on attitudes, the effect of being directly injured or maimed is consistently in favor of the "weary" hypothesis. Those who report being injured or maimed show approximately a 10% or larger increase in the likelihood of saying it is possible to live in peace with former enemies, to live in peace with individual *Janjaweed*, or to live in peace with the tribes from which the *Janjaweed* were recruited. The simplest explanation for this is that while the population in general is exposed to a range of grievances and fears, those who have been directly harmed have experienced the costs of war more directly than those who have not.

While not directly comparable to any prior study, these findings are broadly consistent with other studies that find, perhaps surprisingly, that exposure to violence is associated with some improvements in social or political outcomes such (e.g. Bellows and Miguel; 2009, Blattman; 2009, Gilligan et al.; 2011and Bateson; 2012) or possible reductions in violence (e.g. Lyall; 2009). It conflicts, by contrast, with observational studies such as Vinck et al. (2007), who found associational evidence that those exposed to more violence are more likely to see violence as a way of achieving peace.

These findings are of theoretical relevance, as they bear on the relationship between exposure to violence and attitudes that may influence the course of conflict through perceptions of whether peace is possible or desirable. This also has implications for practitioners and policymakers involved in reconciliation and political settlement processes. While it may be widely expected that those directly exposed to violence would be more vengeful and thus harmful to include in these processes, we find the opposite. Individuals most directly exposed to violence instead appear to be among the most pacific, while having the legitimacy to press for peace.

VI. Tables and Figures

Table 1: Multivariate Balance Conditional on Village Fixed Effects (Male Only)

	β	SE	t	p-val
age	-0.003	0.002	-1.295	0.196
farmer_dar	-0.031	0.090	-0.341	0.733
herder_dar	0.201	0.100	2.013	0.045
pastvoted	-0.011	0.065	-0.175	0.861
hhsize_dar	-0.004	0.006	-0.612	0.541
merchant_dar	0.135	0.081	1.671	0.096
tradesman_dar	0.037	0.148	0.254	0.800
constant	0.539	0.090	5.999	0.000
Joint F(8,338)	1.10			
Joint p	0.362			
N individuals	640			

Table 2: Multivariate Balance Conditional on Village Fixed Effects (Female Only)

	β	SE	t	p-val
age	-0.001	0.002	-0.252	0.801
farmer_dar	0.027	0.078	0.342	0.732
herder_dar	0.144	0.100	1.441	0.151
pastvoted	0.093	0.087	1.066	0.287
hhsize_dar	0.006	0.006	0.932	0.352
constant	0.201	0.124	1.625	0.105
Joint F(6,321)	0.99			
Joint p	0.429			
N	588			

Table 3: *DirectHarm* Regression Estimates

	Enemies Peace	Enemies Peace	PeaceWith JJIndiv	PeaceWith JJIndiv	PeaceWith JJTribes	PeaceWith JJTribes	Execute GoS sold.	Execute GoS sold.	Peace Factor	Peace Factor
DirectHarm	0.087	0.083	0.075	0.082	0.097	0.099	-0.092	-0.11	0.13	0.13
	(0.042)	(0.043)	(0.034)	(0.035)	(0.039)	(0.040)	(0.043)	(0.044)	(0.044)	(0.044)
Intercept	0.49	0.61	0.51	0.22	0.40	0.51	0.56	0.50	0.32	0.46
	(0.030)	(0.078)	(0.024)	(0.064)	(0.028)	(0.076)	(0.032)	(0.082)	(0.032)	(0.078)
Female	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Village FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓		✓		✓		✓		✓	
N	1294	1270	1303	1279	1303	1278	1225	1223	1188	1168

Controls: age, farmer, herder, past vote, household size in Darfur
Robust SEs in parentheses

Table 4: *DirectHarm* Entropy Balanced Regression Estimates

	Enemies Peace	Enemies Peace	PeaceWith JJIndiv	PeaceWith JJIndiv	PeaceWith JJTribes	PeaceWith JJTribes	Execute GoS sold.	Execute GoS sold.	Peace Factor	Peace Factor
DirectHarm	0.097	0.083	0.094	0.088	0.11	0.10	-0.092	-0.11	0.16	0.14
	(0.046)	(0.044)	(0.036)	(0.036)	(0.042)	(0.041)	(0.045)	(0.046)	(0.047)	(0.045)
Intercept	0.36	0.59	0.13	0.22	0.28	0.51	0.68	0.40	0.15	0.44
	(0.026)	(0.081)	(0.019)	(0.067)	(0.023)	(0.078)	(0.025)	(0.085)	(0.025)	(0.082)
Female	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Village FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓		✓		✓		✓		✓	
N	1270	1270	1279	1279	1278	1278	1203	1203	1168	1168

Controls: age, farmer, herder, past vote, household size in Darfur
Robust SEs in parentheses

Table 5: *DirectHarm* Matching Estimates for Effect of Direct Harm

	Enemies Peace	PeaceWith JJIndiv	PeaceWith JJTribes	Execute GoS sold.	Peace Factor
All					
Estimate	0.13	0.11	0.14	-0.12	0.19
p-val	0.00	0.00	0.00	0.00	0.00
N_{pairs}	254	258	260	248	231
Male					
Estimate	0.20	0.14	0.21	0.09	0.26
p-val	0.00	0.00	0.00	0.01	0.00
N_{pairs}	118	119	118	108	101
Female					
Estimate	0.08	0.07	0.03	-0.04	0.10
p-val	0.01	0.00	0.35	0.20	0.00
N_{pairs}	94	97	99	91	85

Figure 1: Propensity Score Balance, no village adjustment

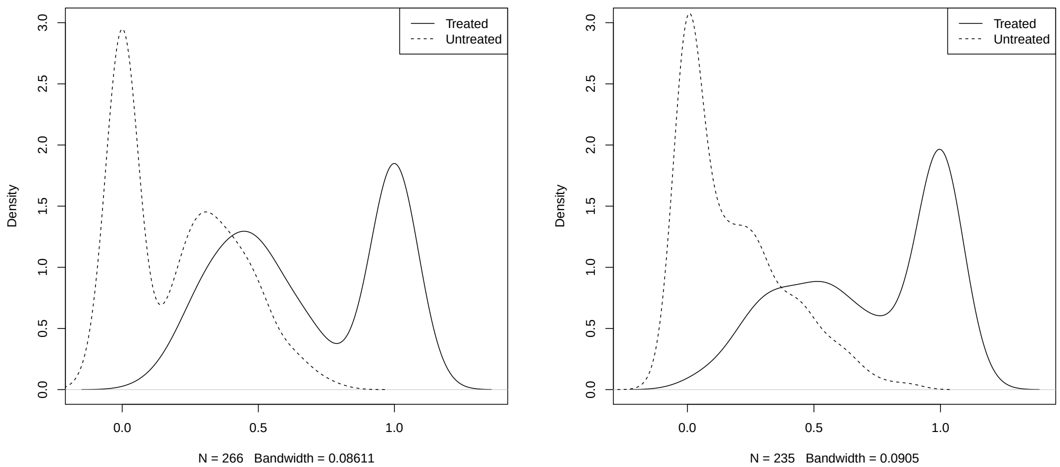

Propensity scores for treated and untreated units, using same linear model of pre-treatment covariates used in multivariate balance testing.*Left panel:* male only; *Right:* female only

Figure 2: Propensity Score Balance, with village adjustment

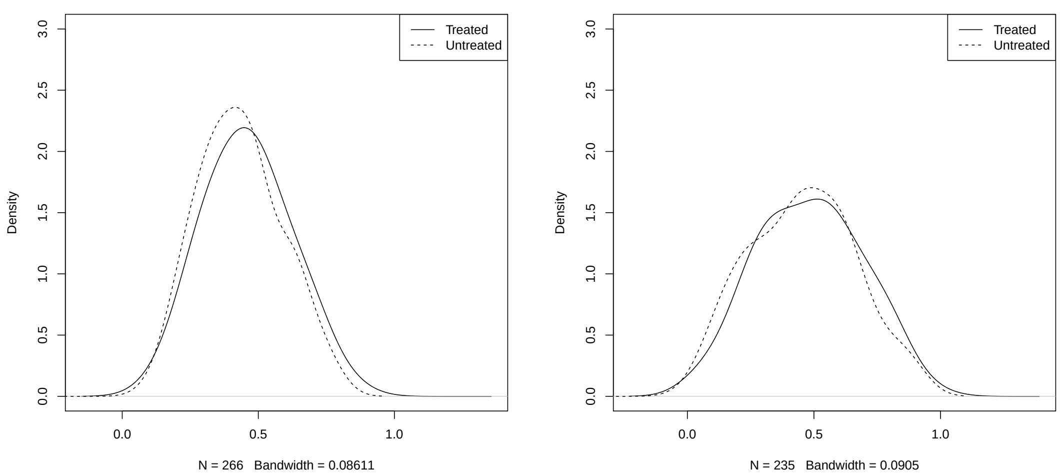

Propensity scores for treated and untreated units, using same linear model of pre-treatment covariates used in multivariate balance testing.*Left panel:* male only; *Right:* female only

Figure 3: True treatment effects given hypothetical confounders

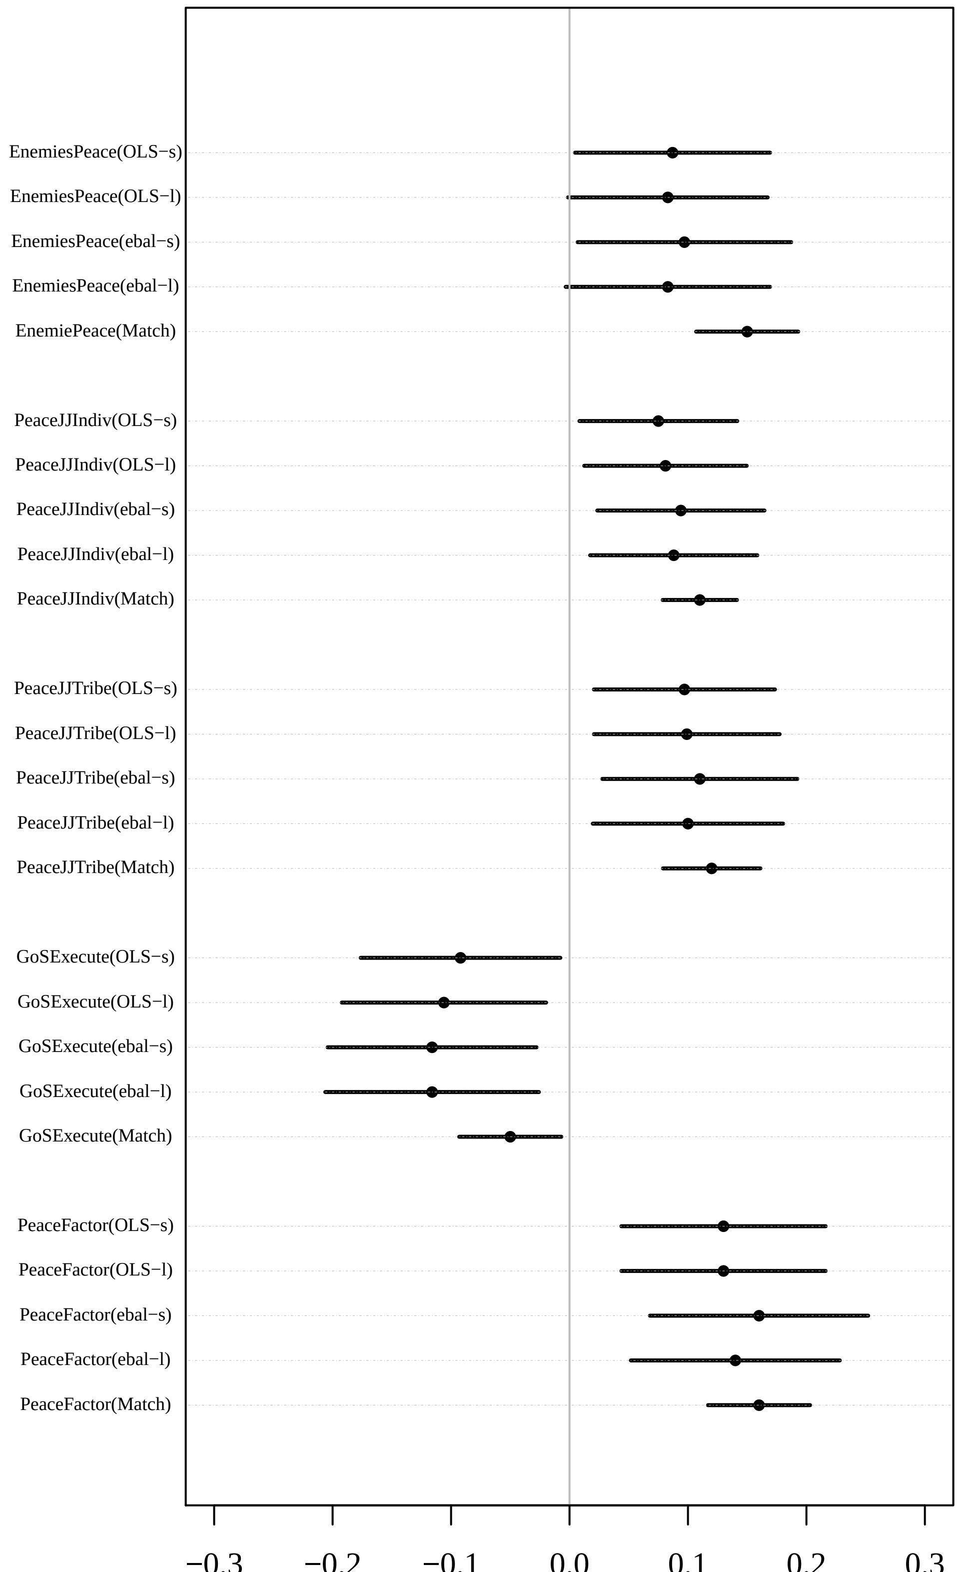

Summary of effect estimates on various outcome under five models: OLS with only minimal covariate (OLS-s), OLS with additional covariates (OLS-l), entropy balancing followed by weighted OLS with minimal covariates (ebal-s) and with additional covariates (ebal-l), and matching.

Figure 4: True treatment effects given hypothetical confounders

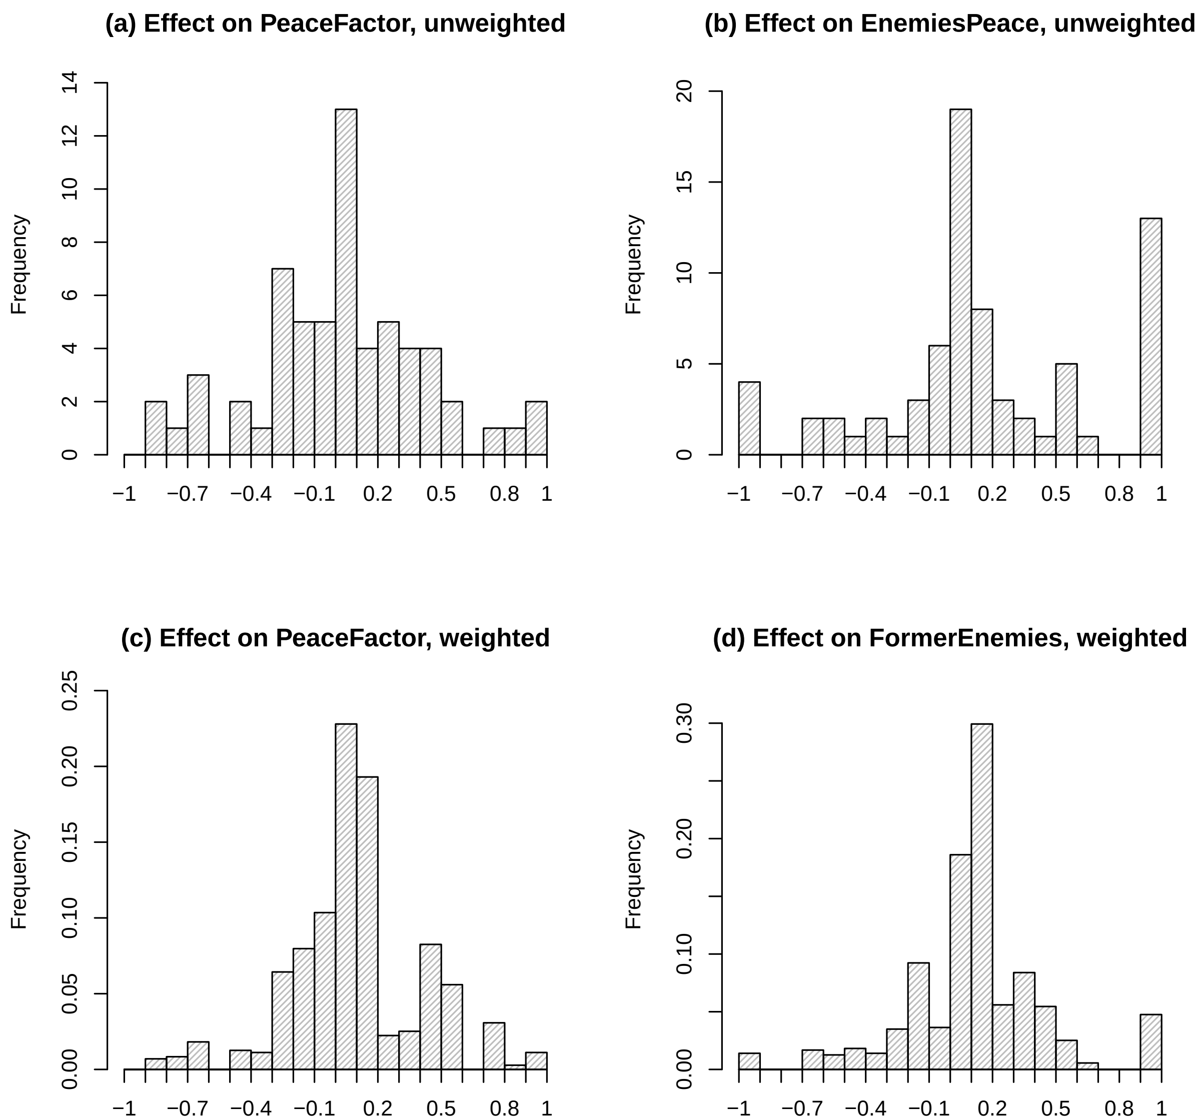

Top Row: Panels (a) and (b) are histograms showing the distribution of village-level effect estimates, weighting each village equally. *Botom Row*: Panels (c) and (d) are weighted histograms, with each village-level effect estimate now weighted by the number of times that village appears in the analysis. The results indicate that the average treatment effects found in various models(all near 0.10) are driven by a number of villages, including larger ones, with treatment effects in this range, rather than a few outliers or a combination of highly heterogenous village-level effects.

Figure 5: True treatment effects given hypothetical confounders

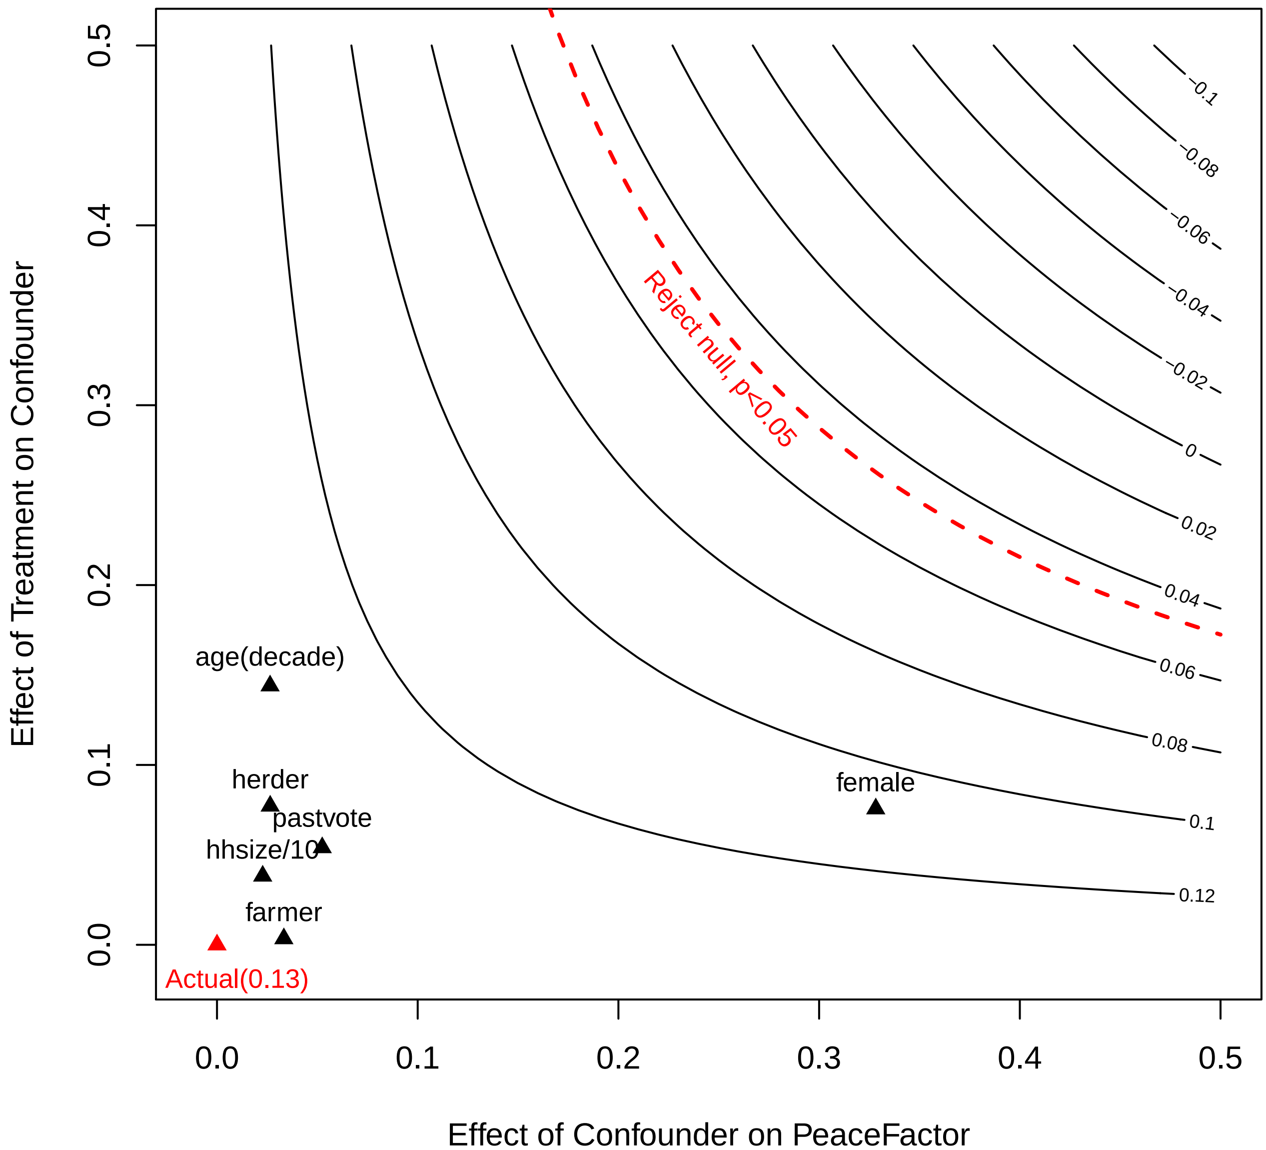

The "height" shown by contour lines gives the expected true size of the effect of *directharm* on *peacefactor*, given the biasing effects of confounders of varying relationship to the treatment (vertical axis) and the outcome (horizontal axis), compared to observed covariates. The "height" indicated by the contour lines shows

References

Akresh, R. and De Walque, D. (2008). Armed conflict and schooling: Evidence from the 1994 rwandan genocide.

Bateson, R. (2012). Crime victimization and political participation, *American Political Science Review* **106**.

Bellows, J. and Miguel, E. (2009). War and local collective action in sierra leone, *Journal of Public Economics* **93**(11): 1144–1157.

Blattman, C. (2009). From violence to voting: War and political participation in uganda, *American Political Science Review* **103**(02): 231–247.

Blattman, C. and Annan, J. (2010). The consequences of child soldiering, *The review of economics and statistics* **92**(4): 882–898.

Bullock, W., Imai, K. and Shapiro, J. (2011). Statistical analysis of endorsement experiments: Measuring support for militant groups in pakistan, *Political Analysis* **19**(4): 363–384.

Colaresi, M. and Carey, S. (2008). To kill or to protect, *Journal of Conflict Resolution* **52**(1): 39–67.

De Waal, A. (2004). Counter-insurgency on the cheap, *Review of African Political Economy* **31**(102): 716–725.

Foltz, J. and Opoku-Agyemang, K. (2010). Civil conflict and schooling outcomes: Evidence from uganda, *Technical report*, Mimeo, Department of Agricultural and Applied Economics, University of Wisconsin, Madison.

Gilligan, M., Pasquale, B. and Samii, C. (2011). Civil war and social capital: Behavioral-game evidence from nepal.

Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies, *Political Analysis* **20**(1): 25–46.

Imbens, G. (2003). Sensitivity to exogeneity assumptions in program evaluation, *The American Economic Review* **93**(2): 126–132.

Kahneman, D. (2011). *Thinking, fast and slow*, Farrar Straus & Giroux.

Kalyvas, S. (2006). *The logic of violence in civil war*, Cambridge Univ Press.

Lyall, J. (2009). Does indiscriminate violence incite insurgent attacks?, *Journal of Conflict Resolution* **53**(3): 331–362.

Lyall, J., Imai, K. and Blair, G. (2011). Explaining support for combatants during wartime: A survey experiment in afghanistan.

Nisbett, R. and Cohen, D. (1996). *Culture of honor: The psychology of violence in the South.*, Westview Press.

Pham, P., Vinck, P. and Stover, E. (2009). Returning home: forced conscription, reintegration, and mental health status of former abductees of the lord's resistance army in northern uganda, *BMC psychiatry* **9**(1): 23.

Pham, P., Weinstein, H. and Longman, T. (2004). Trauma and ptsd symptoms in rwanda, *JAMA: the journal of the American Medical Association* **292**(5): 602–612.

Posen, B. (1993). The security dilemma and ethnic conflict, *Survival* **35**(1): 27–47.

Rosenbaum, P. (2002). *Observational studies*, Springer Verlag.

Rubin, D. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies., *Journal of Educational Psychology; Journal of Educational Psychology* **66**(5): 688.

Valentino, B., Huth, P. and Balch-Lindsay, D. (2004). draining the sea: mass killing and guerrilla warfare, *International Organization* **58**(02): 375–407.

Vinck, P., Pham, P., Stover, E. and Weinstein, H. (2007). Exposure to war crimes and implications for peace building in northern uganda, *JAMA: the journal of the American Medical Association* **298**(5): 543–554.