Estimating One-Sided Killings from a Robust Measurement Model of Human Rights

CHRISTOPHER J. FARISS*
The University of Michigan

MICHAEL R. KENWICK
Rutgers University

KEVIN REUNING
Miami University of Ohio

Abstract

Counting repressive events is difficult because state leaders often have an incentive to conceal the actions of their subordinates and destroy evidence of abuse. In this paper, we extend existing latent variable modeling techniques in the study of repression to account for the uncertainty inherent in count data generated for this type of difficult-to-observe event. We demonstrate the utility of the model by focusing on a dataset that defines one-sided killing as government caused deaths of non-combatants. In addition to generating more precise estimates of latent levels of repression, the model also allows researchers to estimate the probability that a state engaged in one-sided killing and the predictive distribution of deaths for each country-year in the data set. These new event-based count estimates will be useful for researchers interested in this type of data but skeptical of the comparability of such events across countries and over time.

*We thank Jesse Driscoll, James Fowler, Miles Kahler, David Lake, Brad LeVeck, Keith Schnakenberg, and Brice Semmens for many helpful comments and suggestions. The estimates from this paper along with the code necessary to implement the models in STAN and R will be made publicly available at:.

1 Introduction

Counting repressive events is difficult because state leaders often have an incentive to conceal the actions of their subordinates and destroy evidence associated with abuse. To make matters more complicated, even when monitors, activists and journalists have complete access, resource constraints will limit their ability to observe or record all types of violence by the state. The lack of access, and constraints on resources, combine to create potentially biased estimates of repressive events.

Political science scholars recognize that differences in the narrative accounts of primary source documentation may lead to divergent estimates and inferences (e.g., Althaus et al., 2011; Davenport, 2010; Davenport and Ball, 2002; Poe, Carey and Vazquez, 2001). Scholars sought to resolve these problems by integrating data derived from multiple sources (e.g., Davenport and Ball, 2002; Poe, Carey and Vazquez, 2001; Hendrix and Salehyan, 2015). These approaches provide a promising means of inferential validation in the study of human rights, but they have seldom been applied to the expansive time-series cross-sectional designs that are central to so much of the human rights literature. Remaining skepticism over the comparability of data that counts the number of repressive events in country-year observations was one of the main reasons for the movement away from event data in cross-national human rights research (Jabine and Claude, 1992; Poe, 2004).

In this paper, we introduce a new solution to this problem, and demonstrate its effectiveness for linking multiple estimates of event counts with other forms of data in the context of a specific type of repression: one-sided government killings (Eck and Hultman, 2007; Sundberg, 2009). Our approach builds on existing dynamic latent variable models for cross-sectional, time-series data (Fariss, 2014; Schnakenberg and Fariss, 2014; Reuning, Kenwick and Fariss, 2019) while explicitly incorporating uncertainty into our estimates. Through this we are able to simultaneously improve estimates of human rights abuse while generating new estimates of whether any state has engaged in one-sided killing and probabilistic distributions of how many individuals died as the result of these actions. This modeling strategy leverages information in other human rights indicators to generate estimates for country year units that are missing in existing one-sided killing data sets, either due to limited temporal coverage or because of information scarcity stemming from the different strategic incentives of the actors attempting to reveal or conceal this type of repressive event.

We make two primary methodological contributions. The first is a modeling extension that allows researchers to incorporate zero-inflated count processes into canonical latent variable modeling structures. For government one-sided killings, zero-inflation is driven by two distinct types of cases where no killings are observed: (a) no killings are observed because none occurred; and (b) killings may have occurred, but may not have been recorded due to reporting biases driven by a regime's attempt to conceal these events. The second contribution leverages this modeling structure to generate new estimates of one-sided killing. Specifically, we present a method for using this latent variable modeling structure and existing data to: (1) move beyond commonly used "low," "best," and "high" estimates for killing events from Eck and Hultman (2007) toward full predictive distributions of killings, and (2) estimate both whether and how many killings occurred for all country-years. This augments existing data by providing estimates for mass killings in countries where data are either missing, or where such events are assumed to have not occurred.

All of our new estimates quantify the uncertainty inherent in the construction of event-based count data by providing information about the underlying distribution from which such values are drawn for each country-year observation. We also demonstrate how the latent variable estimates confirm the inferences generated in a recent analysis by Fariss (2014) of increasing respect for human rights and how the new count estimates corroborate findings of a decline in the number of fatalities during war time (Goldstein, 2011; Lacina, Gleditsch and Russett, 2006) and a decline in the level of violence more generally (Pinker, 2011). We close with a discussion of the promise of latent variable models for improving the measurement and study of repressive events and political violence.

2 Existing Measurement Models of Repression

To measure the concepts of repression and respect for human rights, scholars have adopted a latent variable modeling approach (Schnakenberg and Fariss, 2014; Fariss, 2014). This approach leverages information from both standards-based and events-based indicators to generate country-year estimates for human rights respect or abuse. Tables 1 and 2 contain descriptions of the datasets and the sources for the standards-based indicators and event based indicators, respectively. The standards-based indicators are almost all derived from Amnesty International and US State Department reports (Cingranelli, Richards and Clay, 2015*a*,*b*; Gibney et al., 2017; Hathaway, 2002).[1]

The event-based indicators are: massive repressive events (Harff and Gurr, 1988); genocide and politicide (Harff, 2003; Marshall, Gurr and Harff, 2009); genocide and democide (Rummel, 1994, 1995; Wayman and Tago, 2010), one sided government killing (Eck and Hultman, 2007); and political executions (Taylor and Jodice, 1983). Fariss (2014, 2019) treats each of these variables as dichotomous indicators that identify whether each type of event occurred. The definitions of genocide, politicide, and massive repression variables each capture human rights violations at the extreme end of the repression spectrum. The measurement of one sided government killing captures instances in which more than 25 individuals (non-combatants) are killed, though this variable excludes extrajudical killings that occur inside a prison and combatant deaths that occur during civil conflicts (Eck and Hultman, 2007). This is a count variable that is left censored at 25. Extrajudicial killing more generally is captured by both the political execution data (Taylor and Jodice, 1983) in addition to several of the variables derived from the human rights reports described above (Cingranelli, Richards and Clay, 2015*a*,*b*; Gibney et al., 2017).

The models constructed by Schnakenberg and Fariss (2014) and Fariss (2014) are outlined in Table 3. The goal of these latent models is to estimate θ_{it}, the latent level of respect for physical integrity rights, from each indicator y_{itj}, where $i = 1, \ldots, N$ indexes cross-sectional units (countries), $t = 1, \ldots, T$ indexes time periods (years), and $j = 1, \ldots, J$ indexes indicators. Let k_j indicate the values that manifest indicator j can take on. In the original models, k_j is either ordinal or binary, such that for the binary indicators $K_j = 2$ while for the ordinal indicators $K_j > 2$

For each physical integrity item, the model estimates an "item discrimination" parameter β_j and a set of $K_j - 1$ "item difficulty cut-points" $\left(\alpha_{jk}\right)_{k=1}^{K_j}$. These parameters are analogous to a slope and intercept term in a logistic regression or the slope and cut-points in an ordered logistic regression. The likelihood function for this model is displayed in Table 3, with $F(\cdot)$ denoting the logistic cumulative distribution function.

Fariss (2014) extends this model by allowing the difficulty cut-points for some of the items to vary

[1]One indicator, Ill-Treatment and Torture, uses "Urgent Action Reports" published throughout the year by Amnesty International to create their index of country-year torture (Conrad and Moore, 2011; Conrad, Haglund and Moore, 2013). We treat this variable as standards-based because the operationalization is based on reports created in a specific historical context just like the other standards-based variables.

Table 1: Standards-Based Repression Data Sources

Dataset Name and Variable Description	**Dataset Citation** and Primary Source Information
CIRI Physical Integrity Data, 1981-2010 - political imprisonment (ordered scale, 0-2) - torture (ordered scale, 0-2) - extrajudicial killing (ordered scale, 0-2) - disappearance (ordered scale, 0-2)	Cingranelli and Richards (1999) Cingranelli, Richards and Clay (2015*a*) Cingranelli, Richards and Clay (2015*b*) Amnesty International Reports[1] and State Department Reports[2] *Information in Amnesty reports takes precedence over information in State Department reports*
Hathaway Torture Data, 1985-1999 - torture (ordered scale, 1-5)	Hathaway (2002) State Department Reports[1]
Ill-Treatment and Torture (ITT), 1995-2005 - torture (ordered scale, 0-5)	Conrad and Moore (2011), Conrad, Haglund and Moore (2013), Amnesty International (2006) Annual Reports[1], press releases[1], and Urgent Action Alerts[1]
PTS Political Terror Scale, 1976-2010 - Amnesty International scale (ordered scale, 1-5) - State Department scale (ordered scale, 1-5)	Gibney et al. (2017), Gibney and Dalton (1996) Amnesty International Reports[1] State Department Reports[1]

1. Primary Source; 2. Secondary Source

over time such that $\left(\alpha_{tjk}\right)_{k=1}^{K_j}$.[2] Note the t subscript here. This parameterization includes each of the standards-based indicators. This is done to account for the possibility that over time human rights monitoring agencies have applied increasingly stringent assessments of state behavior.[3] Put differently, this model accommodates the possibility that states have been subject to a changing standard of accountability regarding repressive behavior. The event-based indicators retain the constant item difficulty cut-point parameterization: $\left(\alpha_{jk}\right)_{k=1}^{K_j}$. The lack of a t subscript here reflects the assumption that the recording of these events is not subject to a changing standard of accountability.

The standard of accountability likely affects the documentation used to code these variables as well.

[2] In the likelihood models we replace this with $\alpha_{jy_{itj}}$ to account for the observed values of y_{it}

[3] See Fariss (2019) for additional discussion of this assumption.

Table 2: Event-Based Repression Data Sources

Dataset Name and Variable Description	**Dataset Citation** and Primary Source Information
Harff and Gurr Dataset, 1946-1988 - massive repressive events (1 if country-year experienced event 0 otherwise)	Harff and Gurr (1988) historical sources (see article references)[1]
Political Instability Task Force (PITF), 1956-2010 - genocide and politicide (1 if country-year experienced event 0 otherwise)	Harff (2003), Marshall, Gurr and Harff (2009) historical sources (see article references)[1] State Department Reports[2] Amnesty International Reports[2]
Rummel Dataset, 1949-1987 - genocide and democide (1 if country-year experienced event 0 otherwise)	Rummel (1994, 1995), Wayman and Tago (2010) New York Times[1], New International Yearbook[2], Facts on File[2], Britannica Book of the Year[2], Deadline Data on World Affairs[2], Kessing's Contemporary Archives[2]
UCDP One-sided Violence Dataset, 1989-2010 - government killing (event count estimate) (1 if country-year experienced event 0 otherwise) (3 death count estimates: best, low, high)	Eck and Hultman (2007), Sundberg (2009) Reuters News[1], BBC World Monitoring[1] Agence France Presse[1], Xinhua News Agency[1], Dow Jones International News[1], UN Reports[2], Amnesty International Reports[2], Human Rights Watch Reports[2], local level NGO reports (not listed)[2]
World Handbook of Political and Social Indicators WHPSI, 1948-1982 - political executions (event count estimate) (1 if country-year experienced event 0 otherwise)	Taylor and Jodice (1983) New York Times[1], Middle East Journal[2], Asian Recorder[2], Archiv der Genenwart[2] African Diary[2], Current Digest of Soviet Press[2]

1. Primary Source; 2. Secondary Source

However, unlike the CIRI, PTS, Hathaway, and ITT data projects, the event-based variables are not direct categorizations of documents but rather, are binary indicators that are coded 1 if sufficient documentary information exists in the historical record to support such a categorization. For the standards-based variables, the documents are directly categorized. Because the documents are never updated or revised, the standards-based variables are rarely updated. This is because, for the event-based variables, documentary

Table 3: Existing Latent Variable Models of Repression

Model and Description	Prior Distributions	
Schnakenberg & Fariss (2014)	**Latent Variable**	
Dynamic ordinal IRT model	Country-year latent variable	$\theta_{i1} \sim N(0,1)$, $\theta_{it} \sim N(\theta_{it-1},\sigma)$
	Innovation Parameter	$\sigma \sim U(0,1)$
	Categorical Indicators	
	Slope	$\beta_j \sim Gamma(4,3)$
	Cut-points	$\alpha_{jk} \sim N(0,4)$
Likelihood Function: $\mathscr{L} = \prod_{i=1}^{N} \prod_{t=1}^{T} \prod_{j=1}^{J} \underbrace{\left[F(\alpha_{jy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{jy_{it(j-1)}} - \theta_{it}\beta_j)\right]}_{\text{Ordinal Indicators}}$		
Fariss (2014)	**Latent Variable**	
Dynamic ordinal IRT model with changing standard of accountability	Country-year latent variable	$\theta_{i1} \sim N(0,1)$, $\theta_{it} \sim N(\theta_{it-1},\sigma)$
	Innovation Parameter	$\sigma \sim U(0,1)$
	Categorical Indicators	
	Slope	$\beta_j \sim Gamma(4,3)$
	Cut-points (event-based indicators)	$\alpha_{jk} \sim N(0,4)$
	Cut-points (standards-based indicators)	$\alpha_{1jk} \sim N(0,4)$, $\alpha_{tjk} \sim N(\alpha_{t-1,jk},4)$
Likelihood Function: $\mathscr{L} = \prod_{i=1}^{N} \prod_{t=1}^{T} \prod_{j=1}^{J} \underbrace{\left[F(\alpha_{tjy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{tjy_{it(j-1)}} - \theta_{it}\beta_j)\right]^{(v_j)}}_{\text{Ordinal Items (Standards-Based Indicators)}} * \underbrace{\left[F(\alpha_{jy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{jy_{it(j-1)}} - \theta_{it}\beta_j)\right]^{(1-v_j)}}_{\text{Ordinal Items (Events-Based Indicators)}}$		

evidence is taken from multiple sources and used to look for evidence that a particular type of repressive event occurred. If new documentary evidence emerges about a specific type of repressive event, the categorized value for the country-year unit is updated. event-based categorization process is therefore able to address variation in the underlying documentation processes that generates information because these variables are each based on set of different documents and are updated periodically. The standards-based coding process cannot directly account for this variation (Fariss, 2019).

3 A Latent Variable Model for Binary, Ordered, and Zero-Inflated Count Processes

The model can be further extended to take advantage of the event counts from some of the event-based data sources that have been reduced to binary indicators in existing models. Two modeling extensions are necessary. First, which we focus on in this section, we allow some indicators take on any positive integer value such that $k = 0, 1, 2, \ldots \infty$. In the next section we take up the second extension which is the need to account for the uncertainty inherent in recording repressive events. To do so, we assume that repressive events counts are generated by a zero-inflated process. Specifically, we account for the incentive leaders or regimes often have to conceal repressive events by assuming that recorded zeros for these cases (i.e. country-years for which no repressive events are recorded in existing data) can be subdivided into two categories: (1) cases where no repressive events were recorded because none occurred and (2) cases where killings were not recorded, but may have nevertheless occurred. This second category is likely to feature cases where killings were concealed due to information scarcity, potentially driven by deliberate concealment. We therefore use a zero-inflated, negative binomial probability distribution to link the latent repression variable with the event count data.

$$\mathscr{L}(\beta, \alpha, \theta, r|y) = \left(p^* + (1-p^*) \left[\left(\frac{r_j}{exp(\alpha_j + \theta_{it}\beta_j) + r_j} \right)^{r_j} \right] \right)^{y_{itj}=0} + \left((1-p^*) \left[\frac{\Gamma(r_j + y_{itj})}{\Gamma(r_j) y_{itj}!} \left(\frac{r_j}{exp(\alpha_j + \theta_{it}\beta_j) + r_j} \right)^{r_j} \left(\frac{exp(\alpha_j + \theta_{it}\beta_j)}{exp(\alpha_j + \theta_{it}\beta_j) + r_j} \right)^{y_{itj}} \right] \right)^{y_{itj}>0}$$

As with the above models, the negative binomial likelihood is parameterized with an α and a β, which have a similar interpretation as in other latent models. Note that observations enter into different portions of likelihood function depending on whether a zero was observed (i.e. $Y_{it} = 0$). For each country year, there is an unobserved probability of being in each of the groups described above: countries that did not engage in massive repressive events in a given year such that only zero counts can be observed; and country-years where killings are possible, but may or may not have been observed. In the above model, this probability is:

$$p^* = F(\alpha_{jy_{itj}} - \theta_{it}\beta_j)$$

where $F(\cdot)$ again denotes the logistic cumulative distribution function.

The negative binomial likelihood also incorporates a rate parameter, r. This accounts for the degree of 'over-dispersion' in the count data by allowing the variance to increase. The variance is equal to $\mu + \frac{\mu^2}{r}$ and μ is the expected value, and is equal to $\alpha + \beta * \theta$. r is assumed to be strictly greater than 0 and as it approaches 0 the negative binomial distribution converges to the poisson distribution.

3.1 Extending the Latent Variable Model of Repression

We build upon the modeling framework above to extend existing latent variable models of repression and to generate predictive distributions of one-sided killings. To do so, we incorporate the event count data on one-sided violence produced by Eck and Hultman (2007). These data identify the yearly number of non-combatants killed by government forces in years where the (observed) number killed exceeds 25. Though our modeling framework could be used to accommodate all count data, we chose to use the UCDP data as our primary data source for two reasons. First, these data have been widely used, widely scrutinized, and have relatively expansive spatial and temporal coverage. Second, as we discuss in more detail below, this is one of the only data sources that provide categorical estimates of uncertainty, providing researchers with "low," "best," and "high" fatality estimates.[4]

[4]Because the UCDP data only contain entries for country-years where more than 25 fatalities were found to have occurred, we set the "low" indicator to 0 for country-years not contained within the UCDP data, reflecting the fact that the UCDP data

To integrate these count data into the latent variable model of repression, we construct a model with the following likelihood function:

$$
\begin{aligned}
\mathscr{L} = \prod_{i=1}^{N}\prod_{t=1}^{T}\prod_{j=1}^{J} & \underbrace{\left[F(\alpha_{tjy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{tjy_{itj}-1} - \theta_{it}\beta_j)\right]^{(v_j)*(1-c_j)}}_{\text{Ordinal Items (Changing Standard)}} * \\
& \underbrace{\left[F(\alpha_{jy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{jy_{itj}-1} - \theta_{it}\beta_j)\right]^{(1-v_j)*(1-c_j)}}_{\text{Ordinal Items (Constant Standard)}} * \\
& \underbrace{\begin{aligned}\Bigg[& \left(p^*_{it} + (1-p^*_{it})\left[\left(\frac{r_J}{exp(\alpha_J + \theta_{it}\beta_J) + r_J}\right)^{r_J}\right]\right)^{y_{itj}=0} + \\
& \left((1-p^*_{it})\left[\frac{\Gamma(r_J + y_{itj})}{\Gamma(r)y_{itj}!}\left(\frac{r}{exp(\alpha_J + \theta_{it}\beta_J) + r}\right)^{r}\left(\frac{exp(\alpha_J + \theta_{it}\beta_J)}{exp(\alpha_J + \theta_{it}\beta_J) + r}\right)^{y_{itj}}\right]\right)^{y_{itj}>0}\Bigg]^{(1-v_j)*(c_j)}\end{aligned}}_{\text{Count Indicators (Constant Standard)}}
\end{aligned}
$$

where v_j and c_j are indicator variables that determine which portion of the likelihood function a particular manifest variable should be passed through. For standards-based indicators $v_j = 1$ and $c_j = 0$; for events-based indicators $v_j = 0$ and $c_j = 0$, and for the UCDP count data $v_j = 0$ and $c_j = 1$.

When constructing the model, one important choice was how to treat the "low," "best," and "high" variables. One option would have been to treat these as three independent indicators, and assign each their own difficulty and discrimination parameter. That is, we would assume that they are conditionally independent and only a function of the latent variable θ_{it}. A useful analogy for this would be three different coders, one liberal (high), one conservative (low), and one moderate (best). We did not adopt this approach for two reasons. First, treating the three estimates as independent of one another ignores the interdependence of these indicators and instead assumes that each reflects a distinct manifestation of the latent trait. Second, as we detail below, this would inhibit our ability to generate a single predicted distribution of one sided killing for all country-years.

We therefore used an alternative, and potentially more realistic, model parameterization that considers these values as the result of one coder or set of coders deliberately attempting to generate estimates of an unknown, true count of one sided killing, y^*_{it}. Because this quantity is not observed, coders provide an estimate of this quantity itself, $y_{it-best}$, and two additional estimates y_{it-low} and $y_{it-high}$ to produce

indicate no killings took place. We leave the "best" and "high" estimates missing for these cases.

a simple distribution around this mean to reflect uncertainty in the estimate of y^*_{it}. In other words, this approach removes the assumption that the low, best, and high estimates are observed independently and instead assumes that the variation across these three indicators reflect measurement error around the unobserved, true number of killings.

This assumption is reflected in the notation for count-indicators, where a single α_J, β_J, and r parameter values is estimated for all three one-sided government killing outcomes: $\{best, low, high\}$. The subscript on these item-specific parameters is J to denote that these parameters are assumed to be the last value in the j vector of α and β parameters and therefore the same for each the three government killing count variables.

In addition to generating improved estimates of the latent trait, this model specification generates two additional substantive quantities of interest. First, as detailed above, the model can be used to estimate p^*_{it}, which captures uncertainty related to whether it was possible to observe one-sided killing in a given country-year. Second, to generate country-year predictive distributions for one-sided killings, we leverage a useful property of latent variable models – that estimates of the latent trait (θ_{it}) and item-specific parameters (α_j, β_j) can be used to produce predictions for each manifest indicator y_{itj}. For the event count indicators, the expected value of one sided killing is:

$$E(y_{it}) = (1 - p^*_{it}) \exp(\alpha_J + \beta_J \theta_{it})$$

Often, these predictions are used as a form of model-checking (Gelman and Hill, 2007). Yet, if we relax the assumption that the manifest indicators are measured without error, these posterior predictions are also a useful means of approximating uncertainty around the indicators themselves. An additional desirable feature of this modeling framework is that predictions for y_{itj} can be generated *regardless of whether this indicator was observed within a particular year*. Put differently, we are able to use these parameter estimates to generate predictive distributions for one sided killing both for years that are outside the temporal domain of the UCDP data set (1949-1988) and for years where UCDP did not find evidence of one-sided killing resulting in at least 25 fatalities. With regard to the latter set of cases, this allows researchers to weaken the assumption that zero killings took place for country-years not included in the

UCDP data. This allows us to identify countries where killings were not observed, but were nevertheless likely based on otherwise high levels of repression.

The uncertainty around the number of killings is also quantifiable because the prior distribution of each of the model parameters and the latent variable allows for the approximation of the posterior distribution of each country-year distribution of one-sided government killings. Country-year heterogeneity is driven by either increased uncertainty in θ_{it}, which captures the latent degree of repression in a country year and is a function of variation between the "low," "best," and "high" estimates and the other manifest indicators, and uncertainty in p^*_{it}. Conversely, when the human rights indicators all point in a similar direction and there is less variation in the "low," "best," and "high" indicators, we expect more precise estimates of one-sided killing.

While this modeling structure offers meaningful extensions to conventional techniques, broader challenges to estimating count data nevertheless remain. Most importantly the number of primary sources available for each country varies and the quality and reliability of the information contained in each document varies as well. The model parameterizes each of these variables, which will eventually allow researchers to make probabilistic statements about the relative quality of the information used in the estimation itself.

3.2 Priors and Estimation

The parameters for the binary and ordered data are given the same distributions as in Fariss (2014) (see Table 3) with one exception. Recent work has applied robust-modeling techniques as a means of improving latent variable model estimates (Reuning, Kenwick and Fariss, 2019). Specifically, the conventional assignment of a standard normal prior to the latent trait is substituted with a Student's T distribution using the following prior specification on the latent trait and innovation parameter:

$$\theta_{i1} \sim \mathrm{T}_{1,000}(0,1) \quad \forall i \in [1,N]$$
$$\theta_{it} \sim \mathrm{T}_4(\theta_{i(t-1)},\sigma) \quad \forall i \in [1,N] \quad \text{and} \quad \forall t \in [2,T]$$
$$\sigma \sim \mathrm{N}(0,3)\mathrm{I}(\sigma > 0)$$

The wider tails of the Student's T distribution allows estimates of the latent trait to experience sudden changes within a given time-series. Within the context of human rights, this modeling innovation allows for sudden fluctuations in the the latent level of repression. This is a desirable modeling innovation, since repression levels may change rapidly as the result of regime change, military coups, or the onset of rebellion.

In addition, we now need to assign priors to the parameters in our count data. α_J and β_J are given the same priors as the α_j and β_j in the binary data. The rate parameter is given the following prior:

$$r_j \sim \text{gamma}(1,.5)$$

Computation is implemented in R using Stan, a program for Bayesian analytics (Carpenter et al., 2017). Sufficient effective sample sizes were obtained using six chains, run for 2,000 iterations each, with a 1,000 iteration burn-in period. Conventional diagnostics suggested convergence $\hat{\mathrm{R}}$ from Gelman and Rubin (1992), and standard graphical analysis.

4 Results and Validation

Validating estimates of respect for human rights and one-sided killings is difficult because we cannot observe the "true" values for these concepts with complete certainty. We therefore follow the framework recommended by Adcock and Collier (2001) whereby alternative measures, such as those produced by our model are treated as hypotheses, which must be evaluated using a variety of tests and procedures.

To determine whether this is the case we compare each the model's estimates to previous estimates of the latent trait. We then examine the correlation between predictions from the model and the original UCDP estimates. Next we examine general trends in the new count data over time, which corroborate recent findings of a decline in the number of fatalities during war time (Goldstein, 2011; Lacina, Gleditsch and Russett, 2006) and a decline in the level of violence more generally (Pinker, 2011). Finally, we present the utility of our new predicted count variables through the examination of one particular case involving the Democratic Republic of Congo.

4.1 Latent Variable Estimates of The Respect for Physical Integrity Rights

Figure 1 displays the mean estimates of the latent trait (θ_{it}), respect for physical integrity rights, across different model specifications. Comparisons are made between the model including counts of killings and those originally produced by (Fariss, 2014). While these estimates comport with those of Fariss (2014), the addition of count-based data into the model produces considerable more variation in the latent trait, as is reflected in the dispersion of estimates along the diagonal line, which would otherwise indicate perfect agreement in model estimates. Substantively, these patterns suggest that introducing count-based indicators appears allows us to uncover significantly more granular estimates for respect for physical integrity rights across country-years.

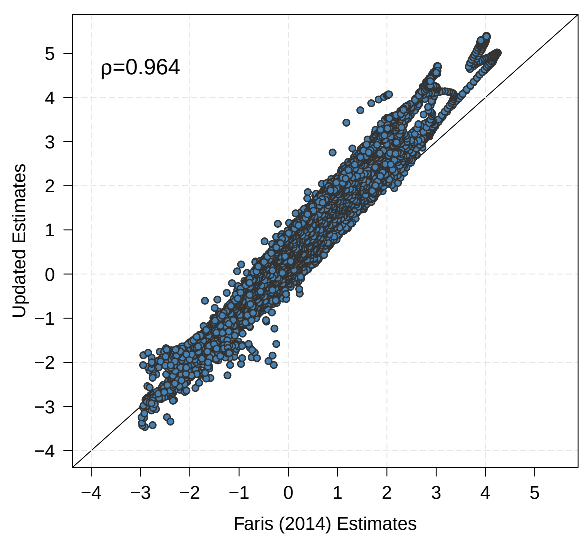

Figure 1: Comparing Estimates of Latent Respect for Physical Integrity Rights

4.2 Model Predictions of One-Sided Killing

Using the estimates of the latent trait (θ_{it}) along with the item-specific parameter estimates (α_j,β_j, and r_j), we generate predictions of one-sided killing for all country-years observed in our data. Figure 2 displays predictions across values of θ_{it}. Each panel plots these predictions in reference to the UCDP the "low," "best," and "high" counts. The red line in the main section of each panel corresponds to the mean posterior prediction of one sided killings given that killing is observed, while the shaded region corresponds to the 95 percent credible interval around this prediction. The curved red line in the bottom section of most of the panels is the probability of observing killings given the estimate of θ_{it}.

Our model suggests that one-sided government killing stops at approximately the mean value of latent repression estimates. The UCDP data begins recording frequent observed instances of one-sided killing at approximately one standard deviation (-1.0) below the mean value of the "true" level of repression. The magnitude of the predictions increases as the latent variable decreases. Though, only Rwanda (1994) nears the maximum observed value, the model makes predictions that accord with earlier episodes of domestic political violence that occurred prior to 1989 when the coverage of the UCDP conflict dataset begins. We discuss these patterns in more detail in the next section.

There are several observations where the UCDP data do not identify one-sided killings–reflected as zeros for the low count in these country years–but our model generates non-zero predictions. Figure 3 reports our predictions for this subset of observations. For the majority of country-years our model produces predictions tightly clustered around 0, which is consistent with their omission from the UCDP data. For observations that are otherwise low on the latent trait, however, the model consistently predicts non-zero values for one-sided killings.

This aids in identifying the subset of observations for which no killings were observed, but killings were likely. For example, UCDP does not record instances of one-sided killings for North Korea. Our estimates, however, predict that that killings occurred every year in North Korea, with median annual predictions ranging from dozens in the early 1990s, to over 100 at various points in the 1950s and mid 1980s. Recall that these estimates do not include extrajudicial killings that occur in custody. Though these predictions may be low relative to the true number of killings, it demonstrates the utility of the model in generating predictions for regimes where killings are concealed or difficult to observe, perhaps corresponding to regimes that effectively conceal abuses.

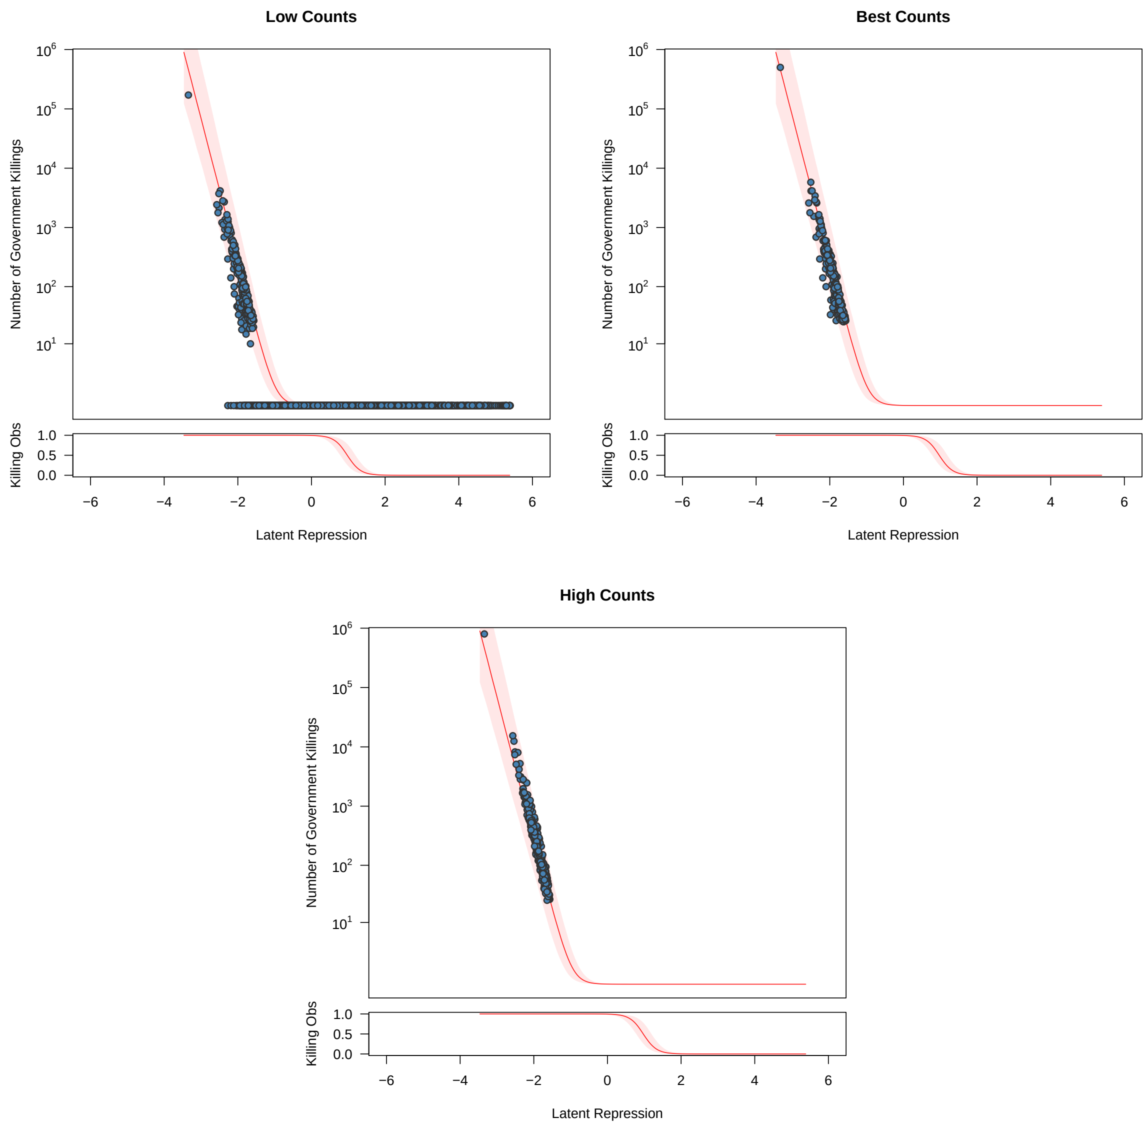

Figure 2: Predictions Counts of One-Sided Killing Along Values of the Latent Trait

Note: Plot reports mean predictions and 95% credible interval for one-sided killing from the model in red along values of the latent trait (respect for physical integrity rights). The "low", "best", and "high" estimates from the UCDP data are displayed with blue points in the top, center, and bottom panels, respectively.

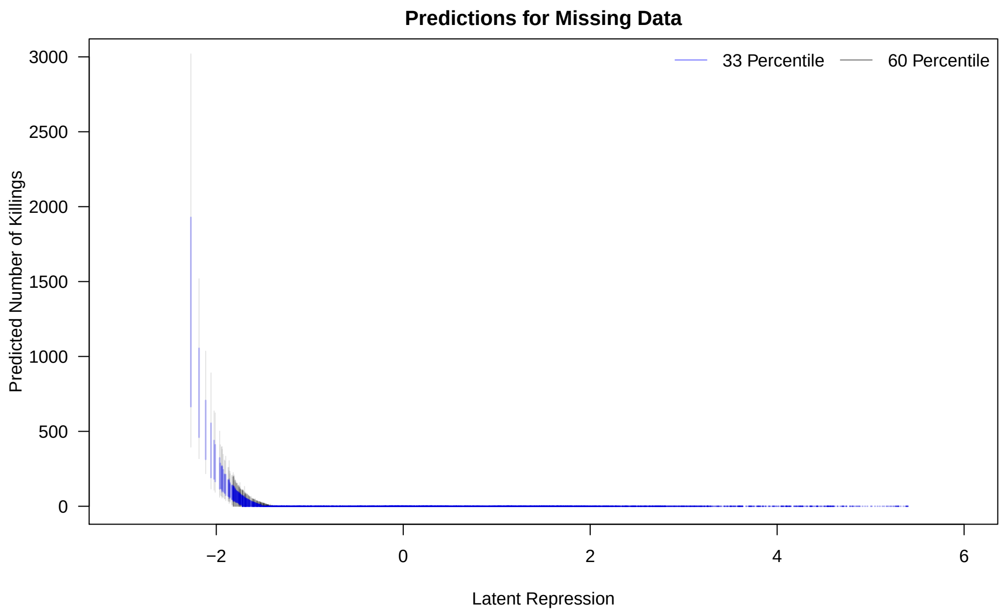

Figure 3: Predictions of One-Sided Killing For Country-Years Omitted from the UCDP Data

Note: Plot reports predictions for the number of one-sided killings among observations that are not contained in the UCDP and would otherwise be assumed to be zero. Triangles correspond to the mean value of these predictions, with the 88 and 33 percentile values reported in grey and blue lines, respectively. Plus signs correspond to the highest posterior density value

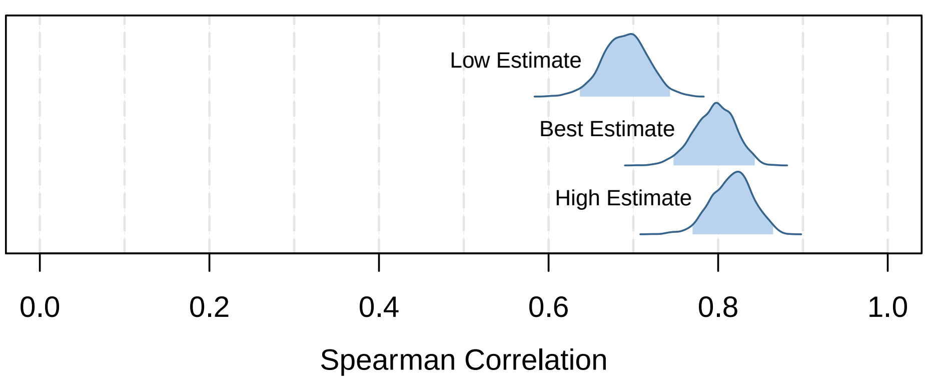

Figure 4: Correlations Between Model Predictions of One Sided Killings with UCDP Low, Best, and High Estimates

Note: Spearman correlations between the the model's predictions of one-sided killing and the original UCDP the "low," "best," and "high" counts.

Figure 4 displays the Spearman correlations between the the model's predictions of one-sided killing and the original UCDP the "low," "best," and "high" counts. The median spearman rank correlation for the low estimates is 0.69, while for the best estimates it is 0.80 and for the high it is 0.82. The correlations along with the other examinations of the predictions of one sided killing telling us that the model does a good job of fitting the data. We now move onto empirical validation by examining aggregate patterns across time as well as a specific case study of the Democratic Republic of Congo.

4.3 Changes in Government Killing Over Time

The latent variable model provides estimated totals that go back to the beginning of the series in 1949. Figure 5 displays the total number of one-sided killings each year for the entire period. Readers should keep in mind that these estimate are not based on count data from 1949 to 1988. Event-based information is, however, contained in several of the binary indicators included in the model. These estimated totals are an approximation of the overall level of one-sided government killing directed at civilians. Each annual count is created by taking one draw from the posterior of each country's predicted one-sided killing for a given a year and then summing across all countries. These are summations from independent zero-inflated negative binomials and so have a heavily skewed distribution. Because we are interested in describing what the *likely* number of total deaths we provide a posterior density that show the 66% credible intervals.

The model suggests that the total number of one-sided killings was relatively low between the mid 1940s before increasing beginning in the mid 1950s. This increase is driven in part by the independence of states like Sudan, who had violent entries into the international system. Estimates remain high throughout the Cold War; more than a million one-sided government killings occurred each year. The number dropped into the high 1000s during the 1990s (other than during the Rwandan genocide in 1994) and most recently to just below 1000 (recall that these deaths do not included extra-judicial killings that occur in custody). These estimates corroborate the results from other studies that find a similar decline in the number of fatalities during war time (Goldstein, 2011; Lacina, Gleditsch and Russett, 2006) and a decline in the level of violence more generally (Pinker, 2011). This is also consistent with improvements in respect for human rights over time (Fariss, 2014, 2019). All of these authors point out that the decline in violence has not been steady. This fact is thrown into stark relief as recent conflicts in Ukraine, Venezuela, and Syra presage heightened violence in coming years.

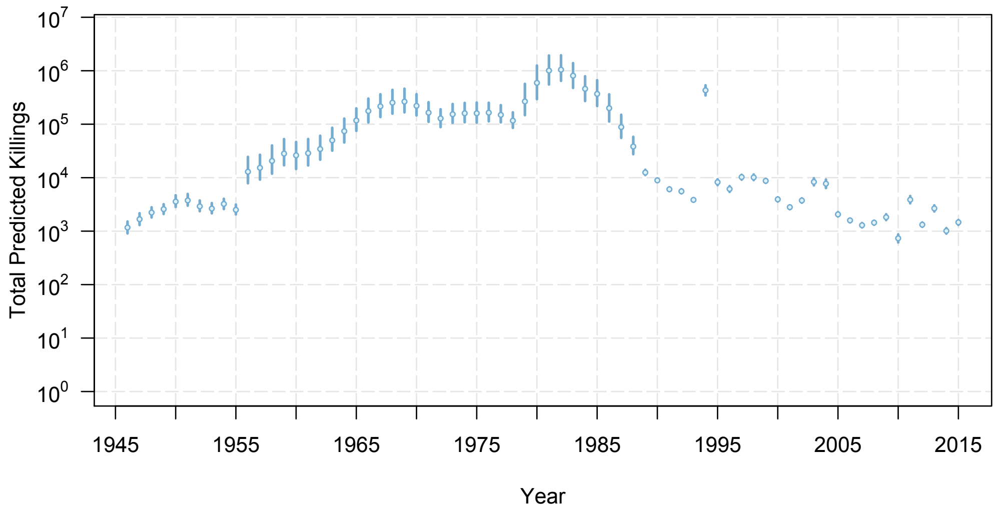

Figure 5: Model based estimates of the yearly number of one-sided government killings beginning in 1949 and ending in 2010. These estimated totals corroborate the results from other studies that find a similar decline in the number of fatalities during war time (Goldstein, 2011; Lacina, Gleditsch and Russett, 2006), a decline in the level of violence more generally (Pinker, 2011), and improvement in the respect for human rights (Fariss, 2014).

4.4 Country Example: The Democratic Republic of Congo (1993-1996)

Documenting repressive events in any country is difficult because of limited resources and limited access to areas in which repressive acts take place. This is especially true in places such as the Democratic Republic of Congo, which, over the last two decades, has experienced two large scale internationalized conflicts with armed participants from multiple countries as well as internecine violence between many armed groups of militia with even more varied affiliations than the state sponsored combatants (Schatzberg, 2012). Thus, acquiring accurate information in such an environment is, not surprisingly, a challenge (Sundaram, 2014).

Though the best open-source information — much of which is provided by journalists and monitors on the ground in places like the Democratic Republic of Congo — is used by the UCDP coders, the estimates given are still just that: *estimates* (Sundberg, 2009). Figure 6 displays distributions of the number of one-sided government killings for the Democratic Republic of Congo (1993-1996). Each plot contains the simulated distribution of potential values, the median prediction from our model, and the original UCDP "low," "best," and "high" counts. Though the number of one-sided killings dropped below 25 in 1994, it is unlikely that none of these events occurred in the Democratic Republic of Congo given the values on the other repression variables included in the model and the UCDP estimates from 1993 and 1995. The latent variable model provides a valid distribution of estimates for the number of one-sided government killings that occurred in a country that is well-known to be difficult to enter by monitors and journalists (Sundaram, 2014).

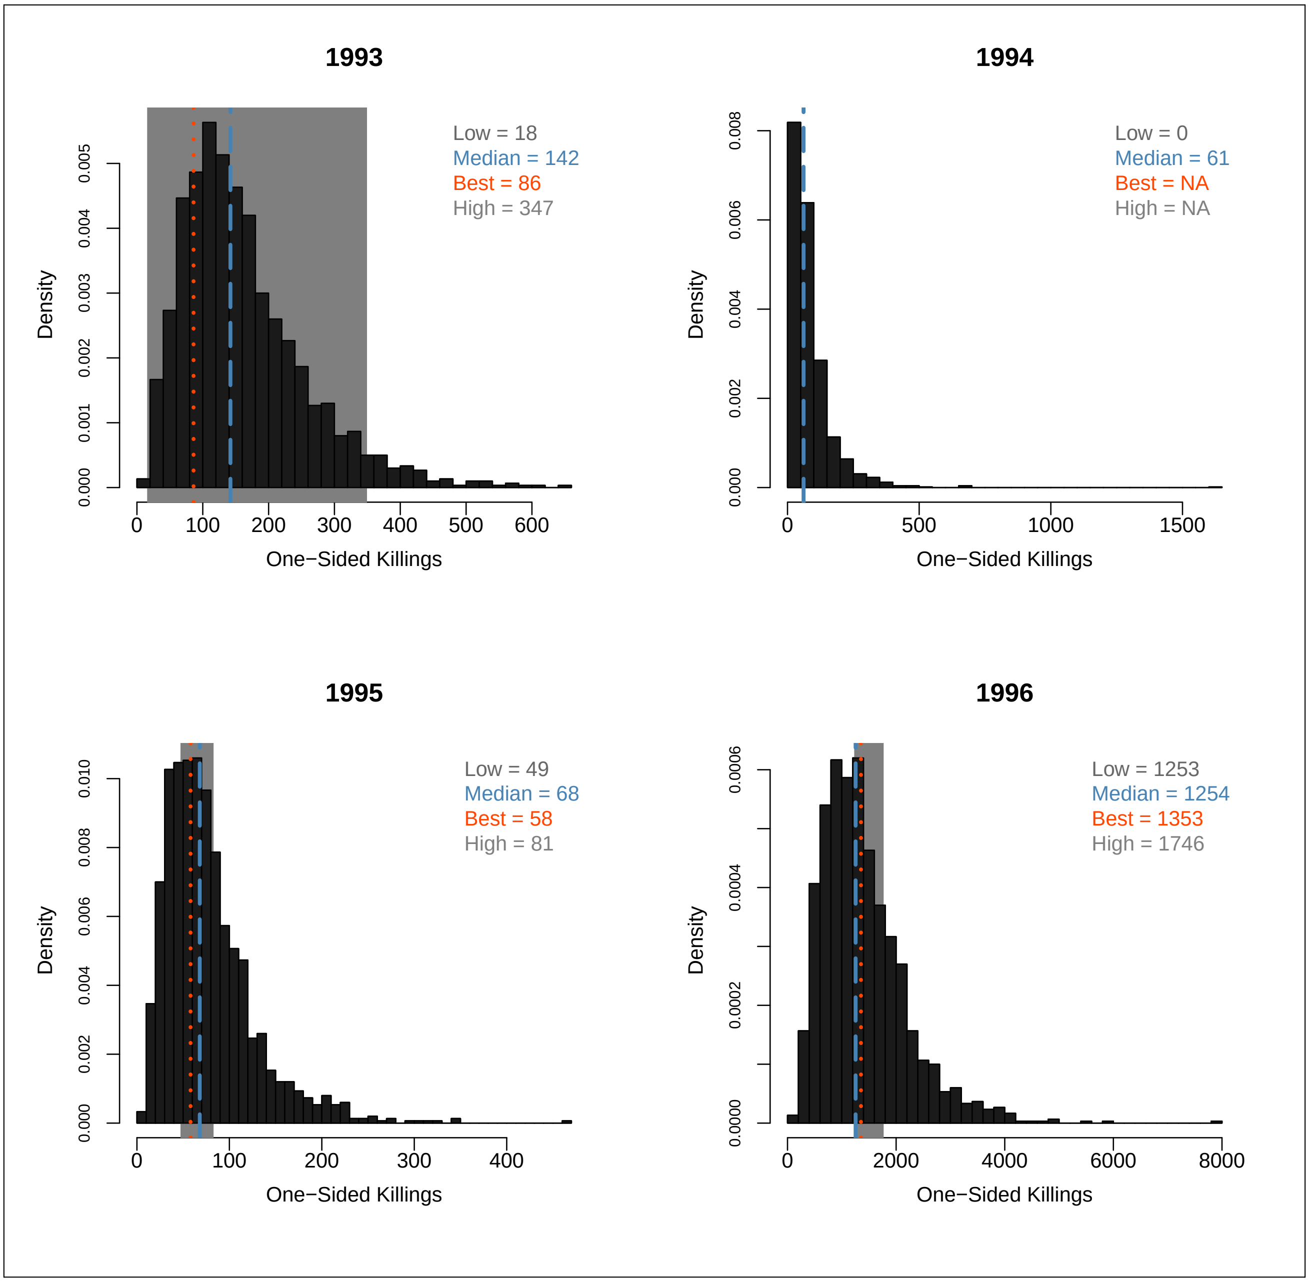

Figure 6: Country-year distribution of the number of one-sided government killings. The distributions each contain the median prediction from the model and the original observed UCDP "low," "best," and "high" counts. Note that in 1994, UCDP does not provide an estimate of the number of one-sided killings but the model is able to estimate a distribution of potential values.

5 Conclusion

The model developed in this paper allows for event-based count data to help improve the estimation of the latent repression estimates (Fariss, 2014; Schnakenberg and Fariss, 2014). The model goes beyond this however, by providing new count estimates of the expected number of one-sided government killings. The new estimates enhance the existing event-based government killing data by providing estimates for country-years that are otherwise assumed to be zero in applied research.

Our modeling strategy improves researchers' ability to leverage disparities in event counts within data sources to provide more valid estimates. Yet, the indicators used to construct our estimates are themselves subject to biases that may remain in primary and secondary sources due to the reporting incentives of news agencies and a lack of transparency in viewing repressive events. Nevertheless, the model-based approach we rely on can be adapted to address these remaining threats to measurement validity. We conclude by noting two domains of research where the expanding base of human rights data promises to yield additional insights with the modeling tools developed here.

First, researchers and activists want to make inferences about more than just country-year units of analysis. New data collection efforts are beginning to acknowledge and understand the role of different state actors who commit human rights violations and the different groups that are targeted. To date, the ITT data project (Conrad and Moore, 2011; Conrad, Haglund and Moore, 2013), and the UCDP data project (Eck and Hultman, 2007; Sundberg, 2009) are the only data efforts that systematically collect repression data about targets, agents, or non-state actors for all states. Other event-based data collection efforts exist and are also beginning to provide some of this information for specific regions (i.e., Saleyhan et al., 2012). The models presented in this paper are capable of systematically linking these new and diverse sources of information with existing categorical data to model multiple levels of information in one model (e.g., country-year, country-year-actors, country-year-victims, country-year-regions).

Second, scholars are beginning to acknowledge and quantify disagreements between different sources of information, which should begin to assuage concerns from some researchers still skeptical of models that compare event counts from disparate sources of information. Recent research has exploited multiple systems evaluation and capture-recapture as a promising means of leveraging disagreements between sources to produce more accurate accounts of repressive events (e.g., Hendrix and Salehyan, 2015). These analyses are often limited to a smaller number of spatial and temporal domains (e.g., Kosovo for four months in 1998). As source-specific information becomes increasingly available, the latent variable measurement strategy produced here will offer a principled, model-based approach for incorporating such information.

References

Adcock, Robert and David Collier. 2001. "Measurement Validity: A Shared Standard for Qualitative and Quantitative Research." *American Political Science Review* 95(3):529–546.

Althaus, Scott L., Nathaniel Swigger, Svitlana Chernykh, David J. Hendry, Sergio C. Wals and Christopher Tiwald. 2011. "Assumed Transmission in Political Science: A Call for Bringing Description Back In." *Journal of Politics* 73(4):1065–1080.

Amnesty International. 2006. *Amnesty International's Country Dossiers and Publications, 1962-2005*. Leiden: IDC Publishers.
URL: *http://www.idcpublishers.com/ead/ead.php?faid=127faid.xml*

Carpenter, Bob, Andrew Gelman, Matt Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Michael A. Brubaker, Jiqiang Guo, Peter Li and Allen Riddell. 2017. "Stan: A probabilistic programming language." *Journal of Statistical Software* 76(1).

Cingranelli, David L. and David L. Richards. 1999. "Measuring the Level, Pattern, and Sequence of Government Respect for Physical Integrity Rights." *International Studies Quarterly* 43(2):407–417.

Cingranelli, David L., David L. Richards and K. Chad Clay. 2015*a*. "The Cingranelli-Richards (CIRI) Human Rights Data Project Coding Manual Version 2014.04.14.".
URL: *http://www.humanrightsdata.com/p/data-documentation.html*

Cingranelli, David L., David L. Richards and K. Chad Clay. 2015*b*. "The Cingranelli-Richards Human Rights Dataset Version Version 2014.04.14.".
URL: *http://www.humanrightsdata.com/p/data-documentation.html*

Conrad, Courtenay R., Jillienne Haglund and Will H. Moore. 2013. "Disaggregating Torture Allegations: Introducing the Ill-Treatment and Torture (ITT) Country-Year Data." *International Studies Perspectives* 14(2):199–220.

Conrad, Courtenay R. and Will H. Moore. 2011. "The Ill-Treatment & Torture (ITT) Data Project (Beta) Country–Year Data User's Guide." *Ill Treatment and Torture Data Project* .
URL: *http://www.politicalscience.uncc.edu/cconra16/UNCC/Under_the_Hood.html*

Davenport, Christian. 2010. *Media Bias, Perspective, and State Repression: The Black Panther Party*. Cambridge, MA: Cambridge University Press.

Davenport, Christian and Patrick Ball. 2002. "Views to a kill - Exploring the implications of source selection in the case of Guatemalan state terror, 1977-1995." *Journal of Conflict Resolution* 46(3):427–450.

Eck, Kristine and Lisa Hultman. 2007. "Violence Against Civilians in War." *Journal of Peace Research* 44(2):233–246.

Fariss, Christopher J. 2014. "Respect for Human Rights has Improved Over Time: Modeling the Changing Standard of Accountability in Human Rights Documents." *American Political Science Review* 108(2):297–318.

Fariss, Christopher J. 2019. "Yes, Human Rights Practices Are Improving Over Time." *American Political Science Review* Forthcoming.

Gelman, Andrew and Donald B. Rubin. 1992. "Inference from iterative simulation using multiple sequences." *Statistical Science* 7:457–511.

Gelman, Andrew and Jennifer Hill. 2007. *Data analysis using regression and multilevel/hierarchical models*. Cambridge, MA: Cambridge University Press.

Gibney, Mark, Linda Cornett, Reed Wood, Peter Haschke, Daniel Arnon and Attilio Pisanò. 2017. "The Political Terror Scale 1976-2016." *Political Terror Scale* .

Gibney, Mark and Matthew Dalton. 1996. The Political Terror Scale. In *Human Rights and Developing Countries*, ed. D. L. Cingranelli. Vol. 4 of *Policy Studies and Developing Nations* Greenwich, CT: JAI Press pp. 73–84.

Goldstein, Joshua S. 2011. *Winning the War on War: The Decline of Armed Conflict Worldwide*. New York, NY: Dutton.

Harff, Barabara. 2003. "No Lessons Learned from the Holocaust? Assessing Risks of Genocide and Political Mass Murder since 1955." *American Political Science Review* 97(1):57–73.

Harff, Barbara and Ted R. Gurr. 1988. "Toward Empirical Theory of Genocides and Politicides: Identification and Measurement of Cases Since 1945." *International Studies Quarterly* 32(3):359–371.

Hathaway, Oona A. 2002. "Do human rights treaties make a difference?" *Yale Law Journal* 111(8):1935–2042.

Hendrix, Cullen S. and Idean Salehyan. 2015. "No News is Good News? Mark and Recapture for Event Data when Reporting Probabilities are Less than One." *International Interactions* 41(2):392–406.

Jabine, Thomas B. and Richard P. Claude, eds. 1992. *Human Rights and Statistics: Getting the Record Straight*. Philadelphia, PA: University of Pennsylvania Press.

Lacina, Bethany, Nils Petter Gleditsch and Bruce M. Russett. 2006. "The Declining Risk of Death in Battle." *International Studies Quarterly* 50(3):673–680.

Marshall, Monty G., Ted R. Gurr and Barbara Harff. 2009. "PITF - STATE FAILURE PROBLEM SET: Internal Wars and Failures of Governance, 1955-2009." *Dataset and Coding Guidelines* .

Pinker, Steven. 2011. *The Better Angels of Our Nature: Why Violence Has Declined*. Viking.

Poe, Steven C. 2004. The Decision to Repress: An Integrative Theoretical Approach to the Research on Human Rights and Repression. In *Understanding Human Rights Violations: New Systematic Studies*, ed. S. Carey and S. Poe. Aldershott: Ashgate pp. 16–42.

Poe, Steven C., Sabine C. Carey and Tanya C. Vazquez. 2001. "How are These Pictures Different? A Quantitative Comparison of the US State Department and Amnesty International Human Rights Reports, 1976–1995." *Human Rights Quarterly* 23(3):650–677.

Reuning, Kevin, Michael R. Kenwick and Christopher J. Fariss. 2019. "Exploring the Dynamics of Latent Variable Models." *Political Analysis* Forthcoming.

Rummel, Rudolph J. 1994. "Power, Geocide and Mass Murder." *Journal of Peace Research* 31(1):1–10.

Rummel, Rudolph J. 1995. "Democracy, power, genocide, and mass murder." *Journal of Conflict Resolution* 39(1):3–26.

Saleyhan, Idean, Cullen S. Hendrix, Jesse Hamner, Christina Case, Christopher Linebarger, Emily Stull and Jennifer Williams. 2012. "Social Conflict in Africa: A New Database." *International Interactions* 38(4):503–511.

Schatzberg, Michael G. 2012. "The Structural Roots of the DRC's Current Disasters: Deep Dilemmas." *African Studies Review* 55(1):117–121.

Schnakenberg, Keith E. and Christopher J. Fariss. 2014. "Dynamic Patterns of Human Rights Practices." *Political Science Research and Methods* 2(1):1–31.

Sundaram, Anjan. 2014. *Stringer: A Reporter's Journey in the Congo*. New York, NY: Doubleday.

Sundberg, Ralph. 2009. Revisiting One-sided Violence: A Global and Regional Analysis. In *States in Armed Conflict*, ed. Lotta Harbom and Ralph Sundberg. Uppsala: Universitetstryckeriet.

Taylor, Charles Lewis and David A. Jodice. 1983. *World Handbook of Political and Social Indicators Third Edition*. Vol. 2, Political Protest and Government Change. New Haven: Yale University Press.

Wayman, Frank W. and Atsushi Tago. 2010. "Explaining the onset of mass killing, 1949–87." *Journal of Peace Research* 47(1):3–13.