

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 have an incentive to conceal actions of their subordinates and destroy evidence of abuse. In this article, 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 repression levels, the model also estimates the probability that a state engaged in one-sided-killing and the predictive distribution of deaths for each country-year in the dataset. 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. Our modeling framework also provides a principled method for inferring unobserved count variables based on conceptually related categorical information.

*We thank James Lo, Zhanna Terechshenko, and three anonymous reviewers for many helpful comments and suggestions. The estimates from this article along with the code necessary to implement the models in STAN and R are publicly available at a dataverse archive: <https://doi.org/10.7910/DVN/7C7KPU>. Fariss acknowledges research support from the SSK (SocialScience Korea) Human Rights Forum, the Ministry of Education of the Republic of Korea, and the National Research Foundation of Korea (NRF- 2016S1A3A2925085).

Introduction

Recording repressive events is integral to the scientific study of peace and conflict. Doing so accurately, however, is complicated by the fact that state-leaders often have strong incentives to conceal these events from the international community and destroy evidence associated with abuse. Even when monitors, activists, and journalists have complete access, resource constraints may limit their ability to observe or record state violence. The lack of access, and constraints on monitoring resources, combine to potentially bias counts of repressive events (Brysk, 1994; Davenport and Ball, 2002).

Researchers recognize that differences in information sources may lead to divergent inferences and have spent considerable time seeking to resolve these problems by integrating data derived from multiple sources (e.g., Hendrix and Salehyan, 2015; Krüger et al., 2013). These approaches provide a promising means of validating inferences from the study of repressive events, but they are seldom applied to the time-series cross-sectional analyses that are central to much of the empirical human rights literature. Remaining concerns over the validity of repressive events count data contributed to a movement away from these data in human rights research (Jabine and Claude, 1992; Poe, 2004). Yet standards-based indicators are subject to other forms of bias (Fariss, 2014, 2019) and are less well suited to precisely track the patterns of repressive events. Much could be learned with counts of these events but only if the political and operational processes that make obtaining such data problematic are conceptually identified and empirically addressed.

We take up this challenge by developing and validating a solution to the problem of bias in count data within the context of estimating one-sided government killings (Eck and Hultman, 2007; Pettersson, Högladh and Öberg, 2019). We demonstrate how latent variable modeling techniques can be used to triangulate information from a variety data sources to improve and expand upon existing estimates of whether and how many individuals were killed by their governments. Our approach builds on existing latent variable models of human rights, which assume that a government's underlying level of repression is not observed directly, but can be estimated based on observable pieces of information captured through a wide variety of human rights monitoring sources. Specifically, we leverage a useful property of latent variable models: the ability to generate predictive distributions for each input variables, regardless of whether they are observed or missing in a particular country-year. The result is the creation of new

estimates of one-sided-killing that account for and reduce bias stemming from resource constraints and the incentive to conceal repression.

We use the [Eck and Hultman \(2007\)](#) data from the Uppsala Conflict Data Program (UCDP) as our benchmark for our model. These data are central to conflict research and have been deployed in hundreds of published articles.¹ We nevertheless endeavor to make several improvements. [Eck and Hultman \(2007\)](#) code one-sided-killings as absent in country-years where there is no direct and reliable evidence of at least 25 individuals deaths. The risk is that these data under count repressive events. We empirically assess this possibility by using information conveyed in other human rights indicators to identify instances where killings were likely to have occurred, but for which insufficient documentary evidence exists. Determining which regimes are within this category is itself of substantive interest to scholars interested in concealed or otherwise unobserved abuse. In addition, while existing data begin in 1989, our approach is able to generate predicted counts extending back as far as 1946, allowing researchers to expand the temporal range of empirical tests.² This is particularly useful for testing theories of slow-changing and/or system-level changes of repression. Finally, we improve estimates of uncertainty around these events by moving beyond the provision of ‘low’, ‘best’, and ‘high’, estimates and instead generating full probability distributions for the number of individuals killed in any particular country-year.

Our latent variable model allows us to expand the empirical testing ground for the scientific study of political violence by providing a principled means of combing biased, incomplete, or otherwise imperfect pieces of information. This approach is critical for scholars because of the increasing volume and breadth of information about repression. The transformation of information from analog to digital is allowing scholars of peace and conflict studies to generate counts of important events like killings but also other forms of political events and patterns of communication in and around conflict (e.g., [Steinert-Threlkeld, 2017](#)).

Beyond these substantive contributions, we also innovate methodologically by extending latent variable modeling techniques to account for zero-inflated count processes. Zero-inflation occurs in count processes when zeros are observed (perhaps excessively) for one of two reasons: (1) the event being

¹As of October 2019, the data have been cited 582 times according to Google Scholar.

²The human rights regime has been developing contemporaneously over the same period of time. For a review see [Fariss and Dancy \(2017\)](#).

counted did not, and could not have occurred; or (2) the event being counted was not observed but could have happened. Within our context, this corresponds to instances where: (a) no killings are observed because none occurred; and (b) killings occurred, but were not recorded due to reporting biases driven by a regime's attempt to conceal these events or because of capacity limitations on the monitoring organizations that collect information about repression. Though used within the context of human rights, we expect these techniques will be useful for a wider body of research linking political event-counts to unobserved concepts of interest.

Below, we introduce our modeling innovations in further detail and validate our estimates of one sided killing. In so doing, we highlight several instances where we believe one-sided killings occurred, despite their absence from existing data sources. Leveraging the expanded temporal domain of our data, we track a reduction in one-sided killing over time, corroborating evidence of more widespread reductions in other forms of political violence (Goldstein, 2011; Lacina, Gleditsch and Russett, 2006). We close with a discussion of the promise of latent variable models generally, with suggestions for addressing limitations of our model and improving the measurement of repressive events.

Existing measurement models of repression

All latent variable models have the assumption that an underlying concept of interest cannot be observed directly but can be approximated through observable manifestations. For latent variable models of human rights, the underlying concept is broadly defined as respect for physical integrity rights, which comprises a unidimensional spectrum ranging from complete and widespread abuse of these rights at one extreme, and complete respect of the rights at the other (Eck and Fariss, 2018; Schnakenberg and Fariss, 2014; Fariss, 2014, 2019). Where a state sits along this spectrum is determined by data collected from a variety of sources on different types of abuse and repression: the observed manifestations, which are broadly classified as either “events-based” indicators or “standards-based” indicators.

Events-based indicators capture extreme rights violations, identifying the scope and scale of repressive events. Typically, these indicators are collected and periodically updated through the continual evaluation of primary and secondary sources. By contrast, standards-based indicators capture both less extreme abuse and widespread abuses. These categorical variables rely on contemporaneous human

rights reports. Though we focus on reducing bias in event-counts, both types of data sources are subject to a variety of biases. A latent variable approach lessens these biases by combining and aggregating information through a principled and transparent measurement model.

The online appendix contains descriptions of the specific datasets and the sources for the standards-based indicators and event based indicators that we use for our updated model. The standards-based indicators are almost all derived from Amnesty International and US State Department reports (Cingranelli, Richards and Clay, 2015; Gibney et al., 2017; Hathaway, 2002). 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; 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 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, but excludes extra-judicial killings that occur inside a prison and combatant deaths that occur during civil conflicts (Eck and Hultman, 2007). Extra-judicial 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; Gibney et al., 2017).

The models constructed by Schnakenberg and Fariss (2014) and Fariss (2014) are outlined in Table 1. These models assume there is an unobserved latent trait, θ_{it} — the level of respect for physical integrity rights — from which we observe manifest indicators y_{itj} . The human rights variables are indexed with i , t and j , where $i = 1, \dots, N$ indexes cross-sectional units (countries), $t = 1, \dots, T$ indexes time periods (years), and $j = 1, \dots, J$ indexes indicators. We also use k_j to indicate the values that the 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” $(\alpha_{jk})_{k=1}^{K_j}$. These parameters connect the observed indicator to the latent variable and are analogous to a slope and intercept term in a logistic regression or the slope and

Table 1: Existing latent variable models of repression

Model and Description	Prior Distributions	
Schnakenberg & Fariss (2014) Dynamic ordinal IRT model	Latent Variable 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	$\beta_j \sim \text{Gamma}(4, 3)$
	Cut-points	$\alpha_{jk} \sim N(0, 4)$
Likelihood Function: $\mathcal{L} = \prod_{i=1}^N \prod_{t=1}^T \prod_{j=1}^J \underbrace{\left[F(\alpha_{jy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{jy_{it(j-1)}} - \theta_{it}\beta_j) \right]}_{\text{Ordinal Indicators}}$		
Fariss (2014) Dynamic ordinal IRT model with Changing standard of accountability	Latent Variable 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	$\beta_j \sim \text{Gamma}(4, 3)$
	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: $\mathcal{L} = \prod_{i=1}^N \prod_{t=1}^T \prod_{j=1}^J \underbrace{\left[F(\alpha_{tjy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{tjy_{it(j-1)}} - \theta_{it}\beta_j) \right]^{(v_j)}}_{\text{Ordinal Items (Standards-Based Indicators)}} * \underbrace{\left[F(\alpha_{jy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{jy_{it(j-1)}} - \theta_{it}\beta_j) \right]^{(1-v_j)}}_{\text{Ordinal Items (Events-Based Indicators)}}$		

cut-points in an ordered logistic regression. The likelihood function for this model is in Table 1, with $F(\cdot)$ denoting the logistic cumulative distribution function. The likelihood is akin to a logistic regression, but with multiple outcome variables for each observation.

Fariss (2014) extends this model by allowing the difficulty cut-points for some of the items to vary over time, changing $(\alpha_{jk})_{k=1}^{K_j}$ to $(\alpha_{tjk})_{k=1}^{K_j}$.³ Note the t subscript indicating that the cut-points for the standards based variables are estimated for each year of data.⁴ This parameterization is used for 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 (Fariss, 2019). Put

³In the likelihood models we replace this with $\alpha_{jy_{itj}}$ to account for the observed values of y_{it}

⁴This is in contrast to the cut-points or intercepts for the other indicators with a Constant standard and so have only one of these parameters estimates for all time periods. This is important generally because researchers may have a conceptual reason to suspect the relationship between a manifest indicator and the latent trait changes over time (e.g., Fariss, 2014, 2019; Kenwick, 2020; Terechshenko, 2020).

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: $(\alpha_{jk})_{k=1}^{K_j}$. The standard of accountability likely affects the documentation used to code event-based variables as well. 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. For the event-based variables, documentary evidence is taken from multiple sources 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 is updated. The 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 a set of different documents and are updated periodically. The standards-based coding process cannot directly account for this variation (Fariss, 2019).

A latent variable model for binary, ordered, and zero-inflated count processes

We can extend this latent variable model to take advantage of the event-counts from some of the event-based data sources that have been coarsened to binary indicators in existing models. This extension allows for the incorporation of more information into latent variable estimates.

The new model requires us to assume a parametric form for event-counts that fits the underlying data generating process. Specifically, we need to account for the fact that killings may be unobserved either because they did not occur, or because they did, but evidence of the event was unobserved. This second category likely features 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.

$$\begin{aligned} \mathcal{L}(\beta, \alpha, \theta, r|y) = & \left(p^* + (1 - p^*) \left[\left(\frac{r}{\exp(\alpha + \theta_{it}\beta) + r} \right)^r \right] \right)^{y_{it}=0} + \\ & \left((1 - p^*) \left[\frac{\Gamma(r + y_{it})}{\Gamma(r)y_{it}!} \left(\frac{r}{\exp(\alpha + \theta_{it}\beta) + r} \right)^r \left(\frac{\exp(\alpha + \theta_{it}\beta)}{\exp(\alpha + \theta_{it}\beta) + r} \right)^{y_{it}} \right] \right)^{y_{it}>0} \end{aligned} \quad (1)$$

As with the above models, we again are assuming that the observed indicators (here the amount of one-sided-killing) is a function of the underlying latent trait. We parameterize this with α , β , r , and p^* . The α and β have a similar interpretation as in other latent models, and we unpack the other parameters below. Observations enter into different portions of likelihood function depending on whether a zero was observed or not. The first line of equation 1 accounts for instances where no deaths were counted (i.e. $Y_{it} = 0$). In this instance the zero count could be structural—i.e. a zero was recorded because no killings occurred—with a probability of p^* or it might result from a situation where killings may have occurred, but none were observed. The second case has a probability of $(1 - p^*)$ multiplied by the probability of the negative-binomial producing a zero count. The second line is for cases where there are non-zero counts (i.e. $Y_{it} > 0$), and so is the probability of $(1 - p^*)$ multiplied by the probability of observing a non-zero count from the negative-binomial distribution. p^* is parametrized as:

$$p^* = F(\alpha^* - \theta_{it}\beta^*) \quad (2)$$

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 $\mu + \frac{\mu^2}{r}$ where μ is the expected count 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.

Extending the latent variable model of repression

Next we must integrate this negative binomial framework into the broader model of human rights so that UCDP Eck and Hultman (2007) data can be incorporated. Though our 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 has categorical estimates of uncertainty, providing researchers with ‘low’, ‘best’, and ‘high’ fatality estimates.⁵

To integrate count data, we construct a model with the following likelihood function:

$$\begin{aligned}
 \mathcal{L} = & \prod_{i=1}^N \prod_{t=1}^T \prod_{j=1}^J \underbrace{\left[F(\alpha_{tjy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{tjy_{itj}-1} - \theta_{it}\beta_j) \right]^{(v_j)*(1-c_j)}}_{\text{Ordinal Items (Changing standard)}} * \\
 & \underbrace{\left[F(\alpha_{tjy_{itj}} - \theta_{it}\beta_j) - F(\alpha_{tjy_{itj}-1} - \theta_{it}\beta_j) \right]^{(1-v_j)*(1-c_j)}}_{\text{Ordinal Items (Constant standard)}} * \\
 & \left[\left(p_{it}^* + (1 - p_{it}^*) \left[\left(\frac{r}{\exp(\alpha_J + \theta_{it}\beta_J) + r} \right)^r \right] \right)^{y_{itj}=0} + \right. \\
 & \left. \underbrace{\left((1 - p_{it}^*) \left[\frac{\Gamma(r + y_{itj})}{\Gamma(r)y_{itj}!} \left(\frac{r}{\exp(\alpha_J + \theta_{it}\beta_J) + r} \right)^r \left(\frac{\exp(\alpha_J + \theta_{it}\beta_J)}{\exp(\alpha_J + \theta_{it}\beta_J) + r} \right)^{y_{itj}} \right] \right)^{y_{itj}>0}}_{\text{Count Indicators (Constant standard)}} \right]^{(1-v_j)*(c_j)}
 \end{aligned} \tag{3}$$

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 parameters. 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

⁵Because 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 indicate no killings took place. We leave the ‘best’ and ‘high’ estimates missing for these cases (See online appendix C).

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 their interdependence 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 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 single α_j , β_j , α^* , β^* , and r parameter values are 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.

Along with 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 can be used to produce predictions for each manifest indicator y_{itj} . For the event-counts, the expected value of one-sided-killing is:

$$E(y_{it}) = (1 - p_{it}^*) \exp(\alpha_j + \beta_j \theta_{it}) \quad (4)$$

Often, these predictions are used as a form of model-checking (Gelman and Hill, 2007). Yet, relaxing 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. A 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*. We therefore generate predictive distributions for one-sided-killing both for years where UCDP did not find reliable documentary evidence of one-sided-killing resulting in at least 25 fatalities and for years that are outside the temporal domain of the UCDP dataset (1946-1988). With regard to the first set of cases, this allows researchers to weaken the assumption that zero killings took place for country-years not included in the UCDP data. Thus we can identify countries where killings were not observed, but were probabilistically likely to have occurred based on the high levels of other repression variables.

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 killing counts. 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. We leave this task for future research.

Priors and estimation

The parameters for the binary and ordered data are given the same distributions as in Fariss (2014) (Table 1) 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:

$$\theta_{i1} \sim T_{1,000}(0, 1) \quad \forall i \in [1, N] \quad (5)$$

$$\theta_{it} \sim T_4(\theta_{i(t-1)}, \sigma) \quad \forall i \in [1, N] \quad \text{and} \quad \forall t \in [2, T] \quad (6)$$

$$\sigma \sim N(0, 3)I(\sigma > 0) \quad (7)$$

The wider tails of the Student’s T distribution allows estimates of the latent trait to experience sudden changes within a given time-series. This is a desirable modeling innovation as repression levels may change rapidly because of regime change, military coups, or the onset of rebellion.

We now need to assign priors to the parameters for our count data. α_j , β_j , α^* and β^* are given the same priors as the α_j and β_j as the other binary manifest indicators. The rate parameter is given the following prior:

$$r \sim \text{gamma}(1, .5) \quad (8)$$

Results and validation

Validating estimates of respect for human rights and one-sided-killings is difficult because we cannot observe the “true” values with complete certainty (Adcock and Collier, 2001). In our application, we are particularly concerned with whether our model fits the data reasonably well and whether the estimates of one-sided-killing produced are valid.⁶

We assess the validity of our model using several types of criterion-related validity checks, which validate a measure based on its relationship with existing measures and determining whether each be-

⁶Though this overarching conceptualization of validity is sometimes called construct validity, we avoid making use of this term. As Jackman (2008) notes, “there are several species of measurement validity. But at least in the context of latent variables, the term ‘construct validity’ has lost much of the specificity it once had, and today is an umbrella term of sorts.” We note further that Adcock and Collier (2001) review how different fields vary in the usage of the term construct validity term, which leads to some confusion.

have in theoretically plausible ways (Trochim and Donnelly, 2008, 59). First, we compare the model's estimates to previous latent estimates of repression, which is one type of convergent validation check. For convergent validity, the latent variable estimates should closely relate to other measures that are known to be valid measures of the concept of interest. Second, we conduct a posterior predictive check of how well our models estimates correspond to the data used to generate the model (Gelman and Hill, 2007). Specifically, we examine the correlation between model predictions of one-sided-killing and the original UCDP variables. A strong correlation would be evidence that our model fit the data relatively well, which itself functions as a form of predicted validity check (Trochim and Donnelly, 2008, 57). Third, we take advantage of the expanded temporal scope of our data to track changes in one-sided-killing over time. As a convergent validity check, we determine whether the predicted number of one-sided-killings corroborates recent findings of a declines in other forms of political violence such as fatalities during war (Goldstein, 2011; Lacina, Gleditsch and Russett, 2006). Fourth, we assess the predicted count variables through the examination of the Democratic Republic of Congo and the 10 cases most likely to have experienced one-sided-killing that are not covered by UCDP. Such assessments of deviant cases is a concurrent validity check that can use a case with an unexpected value, a positive count that is missing from the UCDP dataset, to understand if the estimates from our model are in line with qualitative information about that case (Seawright, 2016; Trochim and Donnelly, 2008).⁷

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). The addition of count-based data into the model produces more variation between the latent traits, 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 uncovers more granular estimates for respect for physical integrity rights across country-years when observing one-sided-killings.

⁷Concurrent validity determines whether the measure can distinguish between groups it should theoretically be able to distinguish between—in this case, years where one-sided-killing did or did not occur. In particular, this concurrent validity assessment reveals a deviant case. A deviant case is an observation that is coded at a value along some theoretical concept that is unexpected (Eck and Fariss, 2018; Seawright, 2016). We discuss the issues in more detail in the online appendix.




Figure 1: Comparing estimates of latent respect for physical integrity rights. This figure compares the latent variable model estimates to a previous set of latent estimates, which is one type of convergent validation check.

Model predictions of one-sided-killing

Figure 2 displays predictions across values of θ_{it} along with the observed UCDP ‘low’, ‘best’, and ‘high’ counts. The red line in the main figure 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 the prediction. The curved red line in the bottom section of the panels is the probability of observing killings given the estimate of θ_{it} . These lines are the same across the different figures as we only estimate one set of parameters for the three different UCDP counts.

One-sided-killings reduce to zero 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.

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 predictions for this subset of observations. For most country-years, the model produces predictions tightly clustered around 0, which is consistent with the decision to exclude them from the UCDP data. For observations that are otherwise low on the latent trait, however, the model predicts non-zero values for one-sided-killings.

Next we focus in on the cases where UCDP reports no one-sided-killing while our model predicts high one-sided-killings. In Figure we show the top-10 country-year units for unobserved one-sided-killings. Afghanistan in 1989 is the most extreme case. We estimate a median of 1,098 killings taking place, with a 33% to 66% range of 649 to 1,778. In 1989 the USSR withdrew from Afghanistan starting a civil war that would last until 1992. Iraq in 1989 is the next with a median estimate of 705 and a range of 470 to 1,024. This was the last year of the Anfal genocide where the Iraqi government systematically massacred Kurdish Iraqis. Although the bulk of killings occurred just prior to this date, documentary evidence suggests active repression campaigns persisted, and our estimates suggest that one sided killings remained probable. Eight of ten of the cases we identify are between 1989 and 1991, and almost all of

these are cases were countries caught between the United States and the Soviet Union during the end of the Cold War. The only cases that are not from this period are Sudan in 2009 and 2010 where we have median one-sided-killing estimates of 190 and 179 respectively. Reports from Human Rights Watch mention civilian bombing in Sudan during these years in addition to reports of 200 individuals that were ‘disappeared’ between 2008 and 2009.

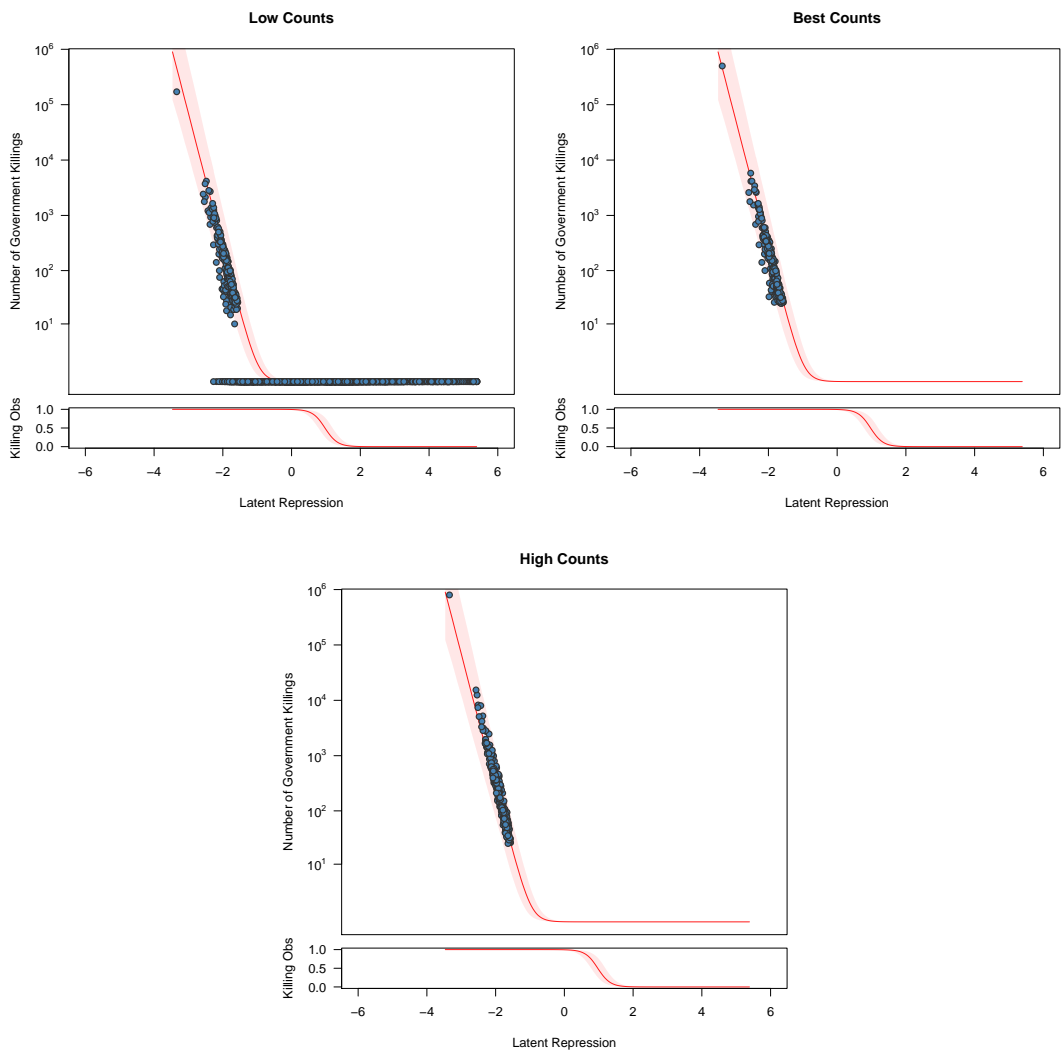


Figure 2: Predictions counts of one-sided-killing along values of the latent trait

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. In contrast to Figure 2, these posterior predictive checks allow us to demonstrate the ability of our model to predict the observed one sided government killings based on values of the latent trait.

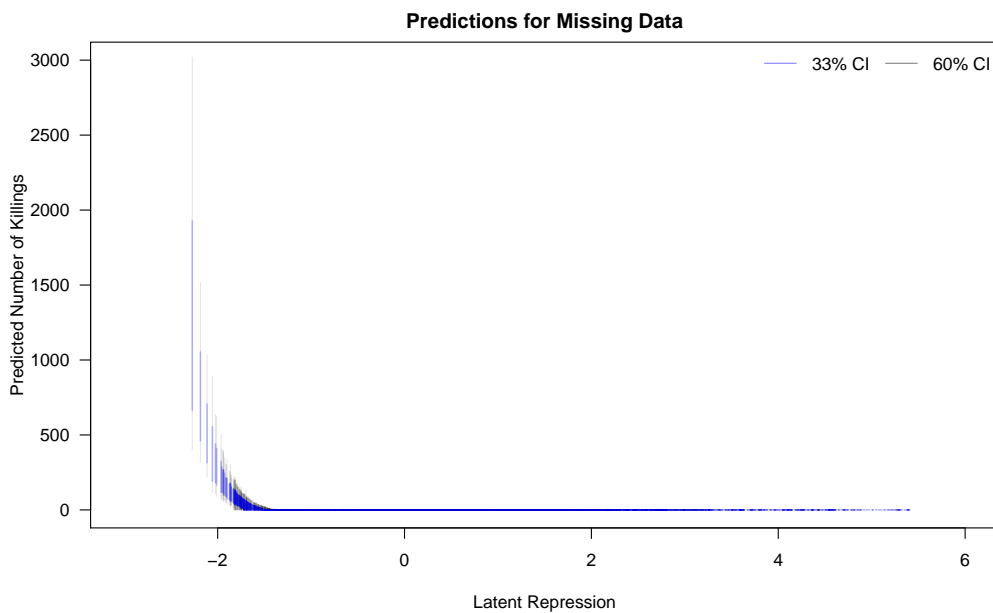


Figure 3: Predictions of one-sided-killing for country-years omitted from the UCDP data

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 33% and 60% credible intervals displayed in grey and blue lines, respectively. Plus signs correspond to the highest posterior density value. These posterior predictive checks demonstrate the ability of our model to predict un-observed one sided government killings based on values of the latent trait.




Figure 4: Predicted one-sided-killing for worst country-years with no reported killings in UCDP data

Note: Violin plots showing the estimated number of one-sided-killing in a country-year. The country-years with the highest estimated killings with no reported killing in the UCDP are displayed. The X-axis is log transformed after adding 1 to each value to preserve 0s. Like Figure 3, these posterior predictive checks demonstrate the ability of our model to predict the top-10 most likely un-observed one sided government killings based on values of the latent trait. This is a type of concurrent validity check.




Figure 5: 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. 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 observed and unobserved data.

Figure 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 observed and unobserved data.

Changes in government killing over time

Because of the temporal coverage of the other human rights variables, our model produces estimated counts of one-sided-killings beginning in 1946, allowing us to conduct a convergent validity test by examining whether our measure can corroborate existing findings that the level of violence has declined over time. Figure 6 displays the total number of one-sided-killings each year. Each annual count is created by taking draws from the posterior of each country’s predicted one-sided-killing for a given a year and then summing across all countries. Readers should keep in mind that our predictions of one-sided-killing are based primarily on our the country-year estimates of the latent respect for human rights (θ_{it}), which is in turn informed by the available manifest indicators. From 1946 to 1988, the UCDP data are missing and our predictions are therefore based on information conveyed in the categorical human rights indicators and this loss of data results in an increase in predicted uncertainty during that time.⁸

The model suggests that the total number of one-sided-killings was relatively low starting in 1946 before increasing 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-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 (these deaths do not included extra-judicial killings that occur in custody). These estimates corroborate the results from other studies that find a decline in fatalities during war (Goldstein, 2011; Lacina, Gleditsch and Russett, 2006), a decline in the level of violence more generally (Pinker, 2011), and improvements in respect for human rights (Fariss, 2014, 2019). All of these authors point out that the decline in violence has not been steady. A fact thrown into stark relief as recent conflicts in Ukraine,

⁸We assume that the relationship between the latent trait and observed one-sided-killing does not vary substantially over time. This is reflected in the fact that the item-specific parameters in the negative binomial link function are not time-varying.

Venezuela, and Syria presage heightened violence.




Figure 6: Model based estimates of the yearly number of one-sided government killings beginning in 1946 and ending in 2017. 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).

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 (e.g., Brysk, 1994). 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 armed groups of militia with even more varied affiliations than the state sponsored combatants (Schatzberg, 2012). Acquiring information in such an environment is, not surprisingly, challenging (Sundaram, 2014).

Though the best open-source information — much of which is provided by journalists and monitors on the ground — is used by the UCDP coders, the estimates given are still just that: *estimates*. Figure 7 displays distributions of the number of one-sided-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. 1994 is notable because it is the only year for which UCDP was unable to uncover direct evidence of 25 killings or more. Nevertheless, an absence of killing is unlikely, given qualitative evidence, the values on the other observed repression variables, and the UCDP estimates from 1993 and 1995.

According to the US State Department Human Rights reports, in the Democratic Republic of Congo, then Zaire, “[p]rovincial officials continued to incite ethnic strife leading to massive displacement and deaths in Shaba, although on a smaller scale than the unprecedented violence in 1993.” The report goes on to provide more detail about the scale of this type of repressive event, stating (1) that the “undisciplined security forces committed numerous extrajudicial killings”; (2) that “[h]uman rights observers, the press and eyewitnesses reported several dozen such fatal altercations, many committed by uniformed personnel”; and (3) that “[i]t is highly likely that additional incidents went unreported, especially in Zaire’s remote interior” (see more details about the information for this case in the online appendix).

Though the frequency of the events was less than prior years, security forces engaged in the extrajudicial killing of civilians in 1994. Because the information environment for this case was poor, sufficient, reliable information was not available for this case to enter the UCDP database. Our latent variable model produces a distribution of potential estimates for this case, with a median estimate of 61 deaths.

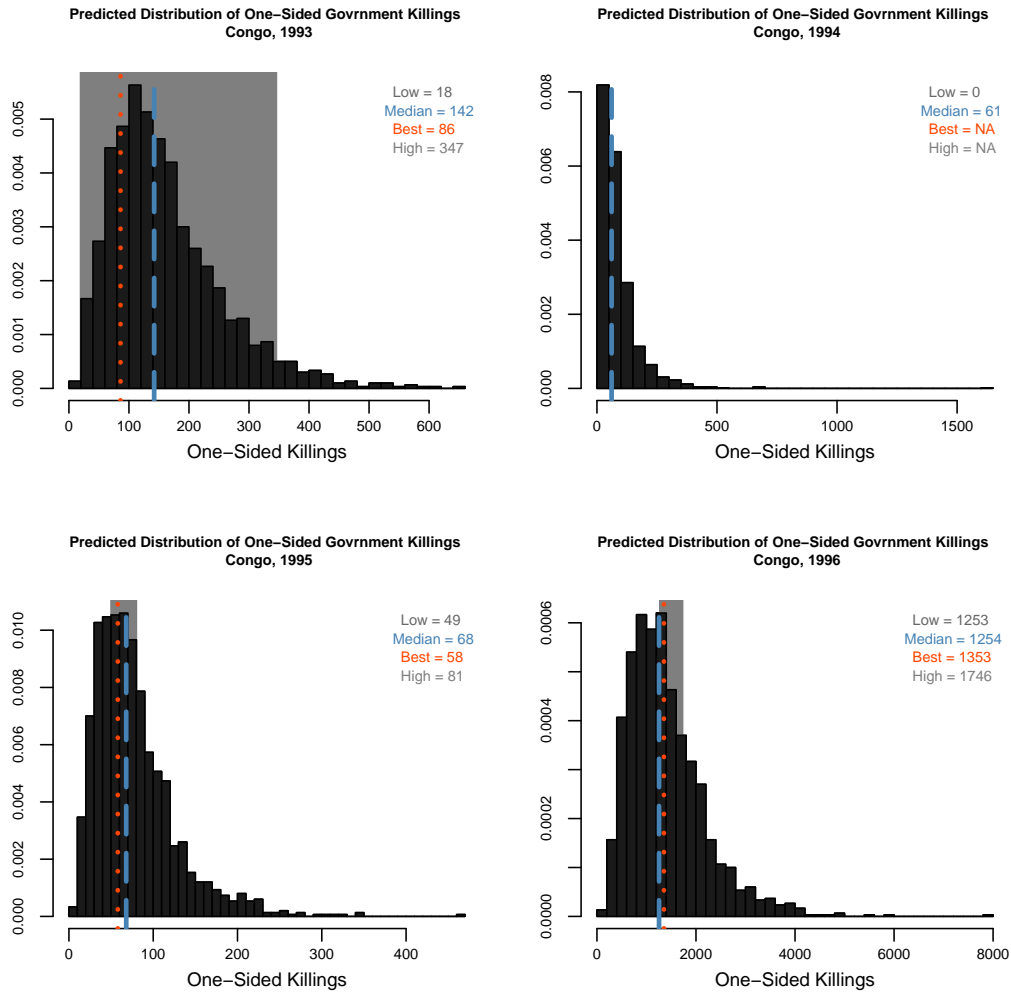


Figure 7: Country-year distribution of the number of one-sided government killings. The distributions each contain the median prediction from the model and the original observed UCDP ‘low’, ‘best’, and ‘high’ counts. In 1994, UCDP does not provide an estimate of the number of one-sided-killings but the model is able to estimate a distribution of potential values. This is a type of concurrent validity check that uses a deviant or unexpected case to assess the performance of our latent variable model relative to the original observed count from the UCDP dataset.

Conclusion

The measurement model and validation tests presented in this article provide insights into the validity of the new latent count measures developed in this article and contributes to a growing research area on using measurement models like to improve the validity of the variables used to study peace and conflict (e.g., [Anders, 2020](#); [Barnum and Lo, 2020](#); [Clay et al., 2020](#); [Cordell et al., 2020](#); [Huddleston, 2020](#); [Krüger and Nordås, 2020](#); [Marquardt, 2020](#); [Meserve and Pemstein, 2020](#); [Montal, Potz-Nielsen and Sumner, 2020](#); [Terechshenko, 2020](#)). Our research makes important improvements in the measurement and understanding of repressive events by linking together count data and categorical variables of repression. In particular, our modeling strategy leverages disagreements in event-counts within data sources, which allows us to generate country-year distributions of estimates of one-sided-killings as part of the latent variable model of repression. Disagreements between event-count estimates exist because of the reporting incentives of monitoring organizations and a lack of transparency or resources with which to completely observe all repressive events. This framework allows us to bring together different sources of information about repression and assess how well each piece of information works together and then, based on assumptions about the way the information was produced, modify, validate, and update the model. We conclude by noting three remaining threats to measurement validity. For each point, the expanding set of human rights data and increasing adoption of new measurement modeling techniques promise to yield additional insights about counting repression.

First, additional data collection for cases where UCDP does not report reliable counts above 25 deaths would allow us to further refine the latent variable model estimates. There is a general challenge in the study of repressive practices because scholars and activists tend to focus monitoring capacity and attention on the most violent cases ([Brysk, 1994](#); [Eck and Fariss, 2018](#)). Additional knowledge of the cases not likely to have committed one-sided-killing would improve the performance of the model. Relatedly, our model assumes that the standard of accountability for UCDP one-sided-killing data does not change. The UCDP data, like other event-based variables, are created using many sources and updated over time. These features of coding process help to account for bias from particular sources, which makes the event-based variables suitable to act as a baseline for comparison with the standards-based variables that do not share this feature.

Second, and building on point one, scholars are beginning to acknowledge and quantify disagreements between different sources of information. Such efforts should assuage concerns about models that use event-counts from disparate sources of information. Recent research has exploited multiple systems evaluation and capture-recapture models as a promising means of leveraging disagreements between sources to produce more accurate accounts of repressive events (e.g., [Hendrix and Salehyan, 2015](#); [Krüger et al., 2013](#)). These analyses are limited to a smaller number of spatial and temporal units. As source-specific information becomes increasingly available, our latent variable measurement strategy provides a principled, model-based approach for incorporating information from new count-based estimation procedures. Linking information from multiple-systems estimation and latent variable models is an important area of new methodological research that directly confronts this challenge.

Third, researchers and activists want to make inferences about more than just country-year units. 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, Haglund and Moore, 2013](#)), and the UCDP data project ([Eck and Hultman, 2007](#); [Pettersson, Högladh and Öberg, 2019](#)) 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 (e.g., [Salehyan et al., 2012](#)). The models presented in this article are capable of systematically linking diverse sources of information and multiple levels of information in one model (e.g., country-year, country-year-actors, country-year-victims, country-year-regions). Each of these measurement challenges represent opportunities for new theorizing, new data collection, and new measurement modeling.

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.
- Anders, Therese. 2020. "Territorial control in civil wars: Theory and measurement using machine learning." *Journal of Peace Research* .
- Barnum, Miriam and James Lo. 2020. "Is the NPT Unraveling? Evidence from Text Analysis of Review Conference Statements." *Journal of Peace Research* .
- Brysk, Alison. 1994. "The Politics of Measurement: The Contested Count of the Disappeared in Argentina." *Human Rights Quarterly* 16(4):676–692.
- Cingranelli, David L., David L. Richards and K. Chad Clay. 2015. "The Cingranelli-Richards Human Rights Dataset Version Version 2014.04.14."
URL: <http://www.humanrightsdata.com/p/data-documentation.html>
- Clay, K. Chad, Ryan Bakker, Anne-Marie Brook, Jr. Daniel W. Hill and Amanda. 2020. "Using Practitioner Surveys to Measure Human Rights: The Human Rights Measurement Initiative's Civil & Political Rights Metrics." *Journal of Peace Research* .
- 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.
- Cordell, Rebecca, Kristian Skrede Gleditsch, Florian G Kern and Laura Saavedra-Lux. 2020. "Measuring institutional variation across American Indian constitutions using automated content analysis." *Journal of Peace Research* .
- 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 Christopher J. Fariss. 2018. "Ill Treatment and Torture in Sweden: A Critique of Cross-Case Comparisons." *Human Rights Quarterly* 40(3):591–604.
- 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* 113(3):868–881.
- Fariss, Christopher J. and Geoff Dancy. 2017. "Measuring the Impact of Human Rights: Conceptual and Methodological Debates." *Annual of Law and Social Science* 13:273–294.
- 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* .
- Goldstein, Joshua S. 2011. *Winning the War on War: The Decline of Armed Conflict Worldwide*. New York, NY: Dutton.
- Harff, Barbara. 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.
- Huddleston, R. Joseph. 2020. "Continuous Recognition: A Latent Variable Approach to Measuring International Sovereignty of Self-Determination Movements." *Journal of Peace Research* .
- Jabine, Thomas B. and Richard P. Claude, eds. 1992. *Human Rights and Statistics: Getting the Record Straight*. Philadelphia, PA: University of Pennsylvania Press.
- Jackman, Simon. 2008. Measurement. In *The Oxford Handbook of Political Methodology*, ed. Janet M. Box-Steffensmeier, Henry E. Brady and David Collier. Oxford University Press.
- Kenwick, Michael R. 2020. "Self-Reinforcing Civilian Control: A Measurement-Based Analysis of Civil-Military Relations." *International Studies Quarterly* .
- Krüger, Jule, Patrick Ball, Megan E. Price and Amelia Hoover Green. 2013. It Doesn't Add Up: Methodological and Policy Implications of Conflicting Casualty Data. In *Counting Civilian Casualties: An Introduction to Recording and Estimating Nonmilitary Deaths in Conflict*, ed. Taylor Seybolt. Oxford University Press.
- Krüger, Jule and Ragnhild Nordås. 2020. "A latent variable approach to measuring wartime sexual violence." *Journal of Peace Research* .
- Lacina, Bethany, Nils Petter Gleditsch and Bruce M. Russett. 2006. "The Declining Risk of Death in Battle." *International Studies Quarterly* 50(3):673–680.
- Marquardt, Kyle L. 2020. "How and how much does expert error matter? Implications for quantitative peace research." *Journal of Peace Research* .
- 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* .
- Meserve, Stephen A. and Daniel Pemstein. 2020. "Terrorism and Internet Censorship." *Journal of Peace Research* .
- Montal, Florencia, Carly Potz-Nielsen and Jane Lawrence Sumner. 2020. "What States Want: Estimating Ideal Points from International Investment Treaty Content." *Journal of Peace Research* .

- Pettersson, Therese, Stina Högladh and Magnus Öberg. 2019. "Organized violence, 1989–2018 and peace agreements." *Journal of Peace Research* 56(4):589–603.
- 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. Aldershot: Ashgate pp. 16–42.
- Reuning, Kevin, Michael R. Kenwick and Christopher J. Fariss. 2019. "Exploring the Dynamics of Latent Variable Models." *Political Analysis* 27(4):503–517.
- Rummel, Rudolph J. 1994. "Power, Geocide and Mass Murder." *Journal of Peace Research* 31(1):1–10.
- 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.
- Seawright, Jason. 2016. "The Case for Selecting Cases That Are Deviant or Extreme on the Independent Variable." *Sociological Methods & Research* 45(3):493–525.
- Steinert-Threlkeld. 2017. "Spontaneous Collective Action." *American Political Science Review* 111(2):379–403.
- Sundaram, Anjan. 2014. *Stringer: A Reporter's Journey in the Congo*. New York, NY: Doubleday.
- 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.
- Terechshenko, Zhanna. 2020. "Hot under the collar: A latent measure of interstate hostility." *Journal of Peace Research* .
- Trochim, William M.K. and James P. Donnelly. 2008. *Research Methods Knowledge Base*. 3rd ed. Mason, OH: Atomic Dog.
- Wayman, Frank W. and Atsushi Tago. 2010. "Explaining the onset of mass killing, 1949–87." *Journal of Peace Research* 47(1):3–13.