

# Recording Abuse: How the Editing Process Shapes our Understanding of Human Rights Abuses

Carly Potz-Nielsen, Robert Ralston, and Thomas R. Vargas

February 20, 2018

# 1 Introduction

Amnesty International and the U.S. State Department produce annual country-level reports on human rights abuses that are widely used by scholars, policymakers, and organizations to gauge the status of human rights within countries across time. For instance, State Department country-year reports are utilized by judges to adjudicate asylum requests in the U.S., as principal documents in the legal systems of Canada, Australia, and New Zealand, and by foreign ministries, press, and business interests as a source of information. Similarly, Amnesty International reports are frequently used in court decisions and as a trusted source for media and government agencies, including the U.S. State Department.

Previous research has provided evidence that the content and structure of these reports—for example, their length and focus, as well as the quality of the information therein—often changes across countries and time.<sup>1</sup> These documents are products of distinct organizations that research, write, and publish their reports using different standards, norms, and processes over their history. Under such conditions, it becomes difficult to distinguish between real changes in the level of human rights abuses world-wide, shifting meanings with respect to what counts as a ‘human right,’ or the interests, preferences, and politics of the organizations that produce the reports. This raises concerns about how useful the information relayed by the reports is for tracking *temporal* trends in human rights violations. Indices derived from these reports treat the information as comparable from year-to-year, yet the emphasis within the organizations is on the documents being comparable over space, not time. The question then, is whether dynamics at the organizational level render coding of PTS and CIRI inappropriate when used in temporal analyses.

Other scholars have considered the organizational constraints on collecting human rights information contained within the reports, and how an organization’s information-gathering techniques and political agendas affect how they report on human rights abuses.<sup>2</sup>

---

<sup>1</sup>(Clark and Sikkink, 2013; Fariss, 2014)

<sup>2</sup>(Fariss, 2014; Richards, 2016; Clark and Sikkink, 2013)

Our contribution builds upon these efforts to “look under the hood” of the reports to understand the information generating processes creating them. It is important, however, to note what we are not doing. First, we do not offer a new measure of human rights abuses. Our task is to examine the human rights reports themselves—not the scores that are derived from them. Second, and as a result, we do not argue that existing measures of human rights are unusable for inference. Rather, we urge researchers using such measures to think carefully about what sort of information the underlying source material is capable of offering and what it is not.

In this research note, we examine whether these reports can serve as good indicators of *over-time* variation in human rights abuses. We do this in two steps. First, we interview individuals familiar with the drafting and editing process of the reports at both Amnesty International and the State Department. Second, we use two computer-assisted text analysis techniques—event data coding and Latent Semantic Analysis—to analyze how these documents change over time.

We find that the editing process matters. While the reports may provide indication of relative abuses across countries within a year, the reports are not comparable over time. Our examination of the process of compiling the annual reports within Amnesty International and the State Department and interviews with people who write and edit the country-year reports indicates that using the documents to measure abuses over time is inappropriate. This evidence supports the claim that scholars using these reports as an indication of trends of human rights abuses are actually capturing changes in organizational procedures, not necessarily changes in the underlying levels of abuse. The information within the reports that is used to create human rights indices is a product of multiple data-generating processes that vary over time. This implies using these reports as a source of over time patterns of human rights abuses, whether as a key variable or a control variable within an analysis, may bias the results.

The remainder of this research note is divided into six sections. The first section pro-

vides some background on the country-year reports created by Amnesty International and the State Department, highlighting differences in the organizational infrastructure that informs how the reports are created. The second section considers previous work that drew attention to identifying and measuring global trends of bias in the reports.

The third section presents evidence from interviews with current and former State Department and Amnesty officials detailing the impact of editorial choices on the production of yearly reports. The fourth section outlines the two methods – event coding via the TABARI software<sup>3</sup> and Latent Semantic Analysis – we use to computationally analyze the available country-year reports for our three cases of Mexico, Argentina, and Guatemala. The fifth section interprets the variation observed in the content of information over time contained within the State Department and Amnesty reports. We conclude by suggesting that using the country-year reports as indicators of *temporal* change in the occurrence of human rights abuse is inappropriate and speculate alternative ways to account for over time trends in country-level human rights abuses.

## 2 Generating Human Rights Reports

Amnesty International and the State Department have produced country-year reports since the late 1970s. Both organizations use these reports to provide a snapshot of pertinent human rights abuses for a given country and year for their respective audiences. The State Department and Amnesty international produce these reports to convey *particular* information to a *particular* audience within *particular* time and space constraints. The determinants of information relayed within the reports exist at the global<sup>4</sup> as well as the organizational level,<sup>5</sup> with both political and practical motivations. However, at the most basic level, what is contained within the reports is a function of what the writers and

---

<sup>3</sup>(Schrodt, 2014)

<sup>4</sup>(Fariss, 2014; Bagozzi and Berliner, 2016)

<sup>5</sup>(Clark and Sikkink, 2013; Fariss, 2014)

editors decide to include. The editing guidelines followed by the people compiling the country-year reports provide the most insight into what they contain.

We focus on the process beginning before the initial draft is completed, with the circulation of guidelines for the report until the final document is published. Both organizations follow the same basic procedure to generate the reports. Country/region ‘experts’ collect information and use it to compile the reports according to the standardized specifications made by upper-level management. Reports then go through rounds of editing for substance, sources, formatting, spelling, and grammar. However, there are interesting differences in the procedures of Amnesty International and the State Department that suggest organizational differences in the information intentionally relayed in the reports. Table 1 summarizes the key differences in the processes used to generate the annual country reports, providing information on who determines the information of interest, the audience of the reports, and the constraints faced by the authors.

The differences we see in the way the organizations generate the yearly reports provides insight into the different amount and type of information contained within each report. Amnesty uses their human rights reports to inform their members about the most egregious human rights abuses in different countries in a given year, or to raise awareness of a particular type of abuse. The reports are written by individuals that may have a regional expertise, but are not necessarily stationed within the country.<sup>6</sup> Authors usually write reports for several countries in their region, and usually visit each country at least once year. Amnesty places a priority on primary sources of information but also uses local media reports and secondary sources as evidence in their reports. As a nongovernmental organization funded by fees and donations of members worldwide, the scope of countries covered in their annual reports has been limited by both its agenda and its capacity. Put another way, these reports, are the product of an NGO for a particular audience: their members. Ron et al. note that Amnesty “produces more written work on some countries

---

<sup>6</sup>Amnesty has offices within London, Brussels, Moscow, Geneva, Hong Kong, Jakarta, Bangkok, Dakar, Nairobi, Johannesburg, Mexico City, Lima, New York, Washington, Beirut, Tunis, and East Jerusalem.

Table 1: Procedural Differences in Organizations

	Amnesty International	U.S. State Department
<i>Who determines the focus of the report?</i>	Secretary General and Regional Offices	Congress/Senior Officials in Washington
<i>How are violations organized?</i>	In rare instances, country-specific violations may be included separately if they do not fit the specified structure.	Violations must be fit within the structure defined yearly by instruction cable from D.C.
<i>How are the reports drafted?</i>	Initial draft → Thematic Experts and Lawyers → Editorial Staff → Fact Verification → Regional Research Manager → Copy Editorial Team → Translation → Year End Country Developments Incorporated → Publication	Embassies → Editors in D.C. → Country Desk Officers and DRL Regional Editors → Senior Staff* → Embassies
<i>What sources do the reports usually use?</i>	Personal network and secondary sources within country	Embassy: personal network and NGOs Editors: Internet and D.C.-based NGOs
<i>Is there a word count?</i>	Word count (500 - 2500)	No, but each country report must address every section.
<i>Where does the organization receive its funding?</i>	Donations	Congressional budget
<i>When are the reports drafted and released?</i>	Sept - January; Released late February	July - February; Usually released within first half of the year.

\* Depending on whether country report is of interest or controversial. For instance, reports on China and Russia are viewed by senior management every year.

than others to maximize advocacy opportunities, shape international standards, promote greater awareness, and raise its profile.”<sup>7</sup> Amnesty also faces particular space constraint; because the reports must fit into a booklet that is under 25,000 words total, each individual report can only be a maximum 2500 words. This is, arguably, the most direct determinant on information contained within the reports; writers and editors must intentionally omit information to meet the constraint on length.

The country-year reports produced by the U.S. State Department were originally in-

<sup>7</sup>(2005, p. 557)

tended to inform Congress on governments receiving aid from the U.S.<sup>8</sup> The reports rely on the expertise of foreign service officers stationed in the relevant country, information from domestic and international NGOs, specifically Amnesty International or Human Rights Watch, and international news sources with a focus on abuses perpetrated by governments. However, because the State Department is an executive agency, it is especially sensitive to the political considerations of each administration, meaning that political and editorial attention varies over countries and time for strategic reasons. For example, our interviews suggest that certain “high interest” countries (like Russia, China, and the Middle East) consistently get more attention than others at higher levels of the State Department. Furthermore, previous analyses of the State Department reports have found that the State Department’s reports have often been marked by harsher coverage of leftist regimes and better coverage of U.S. trading partners and allies than reports produced by Amnesty.<sup>9</sup> Despite their popularity among academics and organizations, Congress remains the principal audience for the reports. Each country-year report is usually authored by foreign service officers<sup>10</sup> stationed within the country and edited both in the embassies and in Washington, D.C.<sup>11</sup> While, since moving to digital reports, there is no strict word limit on the document, each report must address every standardized section.<sup>12</sup>

Importantly, the focus and issue content of the reports is determined at the upper levels of both organizations. Prior to the first draft, Amnesty International and State Department personnel writing the reports are directed to focus on specific issue areas that are of interest to the organization for that year. In the case of the State Department, this manifests through the standardized specification of exclusive section headings and subheadings

---

<sup>8</sup>Mandated by the Foreign Assistance Act of 1961 and the Trade Act of 1974.

<sup>9</sup>(Poe et al., 2001). Our interviews with individuals familiar with the process through which the State Department produces reports confirm these dynamics are at play.

<sup>10</sup>See de Neufville’s insightful piece examining the role of these individuals as creators of the reports.(de Neufville, 1986)

<sup>11</sup>Staffing shortages may also impact the amount and quality of reporting in a given country; some embassies may choose to change information at the margins, while embassies with more staff may choose to completely update a report.

<sup>12</sup>This also means that reports can be updated, without notice, to change inaccuracies in the reports, something scholars should be aware of moving forward.

for every report for that year. Considering the variation in the information relayed by the reports, under the constraints of each organization’s audience and report format, it is clear that there is a process separate from the occurrence of human rights abuses that determines the information contained within each report.<sup>13</sup>

Crucially, neither Amnesty International or State Department reports are written with an eye on future coding by human rights scholars.<sup>14</sup> Scholars using these reports should understand the differences by which Amnesty International and the State Department produce these reports. This includes paying attention to the reports’ diverging aim, audience, and the particular organizational incentives and priorities in place when reporting on particular countries.<sup>15</sup> The interest of most scholars using these reports is to gain insight to the human rights situation within a country, which makes separating the actual change in human rights abuses within a country from changes in the way information is presented within a report essential. In the next section, we discuss previous efforts by scholars to identify and differentiate the processes that lead to the inclusion, or exclusion, of information within the annual country-level human rights reports.

### 3 The Information Generation Process

The key concern plaguing the use of the country-year reports produced by Amnesty International and the State Department is how to determine actual levels of human rights abuses from reported levels of human rights abuses. In particular, there are two different sources of variation that would influence the reporting of human rights abuses, but not the actual level of abuses within a country: (i) the scope and depth of investigations and the

---

<sup>13</sup>In other words, any scholar using these documents should move beyond the assumption that these reports are simply relaying “objective” or “ground truth” information. See, for example, the Political Terror Scale’s words of caution with respect to this aspect of the reports: <http://www.politicalterroryscale.org/archive/SDvsAI/>

<sup>14</sup>This was echoed across all of our interviews, with one interview voicing surprise that this was occurring (Interviewee 2).

<sup>15</sup>(Bagozzi and Berliner, 2016)

reporting of abuses; and (ii) changes in what is considered a “human rights abuse.” If either, or both, of these sources are active, then the reports will vary in the level of specificity, accuracy, and completeness in their recording of human rights abuses over time. In other words, the divergence between “on the ground” truth and what is relayed by the reports is a function of the ability and interest of the organizations generating the reports, and the audience they are generating them for, as well as changes in what is considered within the scope of the reports. These two dynamics map onto two broad processes identified within the literature: (1) information effects; and (2) changing standards of accountability. This section details these processes before arguing for the consideration of a third, additional, source of variation: the editing process.

Clark and Sikink (2013) detail the possible sources, and consequences, of failure to distinguish between changes in what is recorded by the country-year reports and what is actually occurring within the country.<sup>16</sup> The point to two data generating processes: (1) the process by which human rights violations occur and (2) the process by which information is contained within the reports. In particular, they point to the uneven reporting over time as increases in capability to collect information and interest in the reporting of abuses within a country varies over time and space. These issues of data quality in human rights reporting are termed “information effects,” defined as “patterns in the data stemming from the process of information collection and interpretation, rather than the process that actually gives rise to human rights violations or their mitigation.”<sup>17</sup>

If these two processes operate on the data, then coding human rights abuses from the reports is likely to provide invalid over-time abuse measures. More specifically, this critique questions the validity of human rights indices coded from these reports like the Political Terror Scale (PTS) and the Cingranelli-Richards Human Rights Data Set (CIRI).<sup>18</sup>

---

<sup>16</sup>This is a concern that is not unique to human rights data, and that has long been central to the literature on political violence, similarly concerned with how violent incidents or casualties, are reported (Danzger, 1975; Caliendo et al., 1999; Davenport and Ball, 2002; Davenport, 2010; Weidmann, 2014, 2016; Vargas et al., 2016).

<sup>17</sup>(2013, p. 540)

<sup>18</sup>The PTS provides a measure of threats to physical integrity rights, used to measure annual state-based

Clark and Sikkink use word count of the documents and brief case studies to illustrate how the quality and amount of information varies over time. They argue that changes in the information relayed by the reports due to these information effects implies that violations may become more observable over time as access increases and the cost of gathering information decreases. Violations may be consistently occurring over time, but we only observe the most recent occurrences, thereby leading to the erroneous conclusion that the violation is a recent development. In other words, yearly reports are not comparable over time because they are measuring human rights violations based on systematically differing amounts of information.<sup>19</sup>

Fariss (2014) builds from Clark and Sikkink's critique, but focuses on a separate process masking true changes in the level of human rights abuses over time: changes in the definition of a human right. While Clark and Sikkink highlight the changes in our ability and interest in human rights abuses, Fariss argues that the very nature of what constitutes a "human rights abuse" may be changing over time. He defines these 'changing standard of accountability' as product of three different points of variation: information, access, and classification.<sup>20</sup> What counts as a human rights violation, what is seen as worthy of reporting on, and the classification of human rights violations are not constant over time, particularly when comparing across decades.

The task is to separate actual trends in human rights abuses from trends in the reporting on human rights abuses. Fariss derives a measurement model used to extract the latent level of human rights abuses. Comparing the latent scores to CIRI scores coded by human coders, he argues that the scores used by most scholars are biased by the context in which the reports were written. By separating out trends in the reporting process from trends observed in the human rights scores, he identifies a decrease in human rights abuses over

---

political violence against a country's population. CIRI scores create four separate indicators of physical integrity abuses that takes into account the frequency at which violations occur.

<sup>19</sup>Richards (2016) responds to the criticism of Clark and Sikkink, arguing that human coders used to produce CIRI are sensitive to changes in information, and that CIRI was regularly updated to account for changes in information, so that these information effects did not result in bias within the CIRI coding.

<sup>20</sup>(Fariss, 2014)

time that has been masked by changes in the standards with which the reports are written. In other words, the downward trend in the actual occurrence of human rights abuses, measured by the latent human rights scores, was unobserved because the scope of information, access, and consideration of human rights abuses was increasing.

This research note contributes to this literature on two fronts. First, we bring attention to a process active, but undiscussed, in current research: the editing process. From our interviews with authors and editors of the country-year reports, we find evidence that at the most basic level, the amount and type of information contained in the reports is a direct function of the editing process. This is complementary to the points raised by others, such as Clark and Sikkink and Fariss, but operates differently from information effects or changing standards of accountability. As any writer knows, the final draft often looks very differently from the first and the process that decides what is reformatted, cut, included, and reworded is crucial to the final product. By understanding the country-year reports produced by Amnesty International and the U.S. State Department as documents, the role of this process in shaping the information ultimately contained within the reports is crucial.

The second contribution of this research note is to consider the information actually contained within the documents. We improve upon Clark and Sikkink's use of word count as a measure of the amount of information in a document, to consider how much information is being relayed by those words. By looking at the type of information contained within the reports, and how it changes over time, we address the concerns raised by Fariss regarding trends in what is considered a (reportable) human rights abuses. The variation we identify within the documents over time highlights the consequences of the divergence between what is occurring within a country and what is recorded within the reports. By examining the documents, instead of the scores coded from them, we are able to focus on the process by which the information is consolidated, instead of how it is translated into human rights scores. This allows us to examine the consequences of changes in the

amount and type of information relayed by the country-year reports, supporting our claim that the information contained within them is not comparable over time.

In the next section, we review the information gained through our interviews with Amnesty International and State Department employees and former-employees. These interviews focused specifically on the procedure of gathering the information, writing the report, editing the report, and finalizing the document. Our interviews support the claim that the one process that is most crucial to all writing, the editing process, makes it inappropriate to understand these reports as tracking over time change in human rights abuses.

## 4 The Effects of the Editing Process: Interviews with Authors and Editors

*“Most people in the State Department would be surprised that this [using reports to create quantitative indices and comparisons over time] is happening [...] the process and content of reports is always changing”*

- Interview: Former State Department FSO and Contributor to Reports

In order to understand how the State Department and Amnesty International research, write, and publish their reports, we interviewed members of each organization. Our semi-structured interviews included a researcher who has written country reports for Amnesty International, a former State Department embassy-based researcher, and the editor-in-chief/coordinator of the State Department’s annual reports.<sup>21</sup> We view the information provided by these sources as credible because it comes from individuals who have worked directly in the production of these reports, from the level of the country desk to the upper echelons of the organization and from different time periods, including the 1990s, 2000s,

---

<sup>21</sup>Interviews were conducted via Skype in the Spring of 2016. Our interviewees requested anonymity. Interviews were recorded for accuracy with informant consent. Information regarding the interviews available upon request.

and 2010s. Moreover, we do not expect them to have any particular agenda regarding the production of human rights scores from their reports. This is particularly so for State Department officials, who are simply complying with a Congressional request. We find it most compelling to believe that these individuals are interested in clearly delineating the limitations of their work in the interest of its integrity. Furthermore, since we were interested in matters of standardized procedure, we do not expect that the information gathered from the individuals was dependent on their personal experience. In other words, we would not expect that our understanding of the process of creating the reports would vary greatly if we were to have interviewed a separate set of individuals.

The common theme across the interviews was that the reports produced by both Amnesty International and the State Department should be used carefully, especially when compared over time. Interviewees from both Amnesty International and the State Department expressed unease at the notion that these reports should be used to track temporal variations of human rights abuses. Rather, as both our text analyses and our interviews suggest, these documents are the artifacts of the editorial process of the organizations producing the reports.

It is a simple, but important, point to highlight the shared concern interviewees had with coding documents. The country-year reports are produced by each organization for a specific purpose: to inform their audience of human rights abuses of interest to the organization for a snapshot in time. Our interviews support the arguments made by Clark and Sikkink (2013) and Fariss (2014). Both the interest and capability of the organization changes over time. However, there was also an emphasis that the very manner in which the reports are put together, the editing process, makes them inappropriate to measure temporal variation.

The reports are never meant to represent “ground truth” in each country, nor are they exhaustive accounts. This is especially important when thinking about the *lack* of events in a particular context—this is not necessarily indicative of zero human rights abuses of that

kind that year, but rather it could be a function of the organizational decisions that privilege some information over other information. Not all types of human rights abuses make it into the reports each year. Our source within Amnesty International informed us that if a particular violation is well outside the boundaries of the general theme that year, then it is omitted from the reports. Moreover, interviewees suggested that if there is no change in a particular situation over time, then the abuses often get dropped from subsequent reports, despite the abuse continuing. In the case of Amnesty, an interviewee provided an example of a case in which Amnesty wanted to draw attention to police brutality in a particular country. However, because Amnesty did not see themselves making any headway on the issue, and, considering the length constraints each report is under, the issue was dropped from their reports. Additionally, because of the sensitive and current nature of some abuses, like cases of torture, the primary sources may report to Amnesty, but under the condition that their claims are not made public to protect against retribution. In both examples, Amnesty would intentionally not mention the abuse within their annual reports despite having knowledge of its occurrence. This suggests a different filtering process than shifts in interest or in capability of gathering information; Amnesty is both aware and interested in these abuses. However, because of the purpose and constraints of the documents, the abuses are not represented within the country-year reports. This may lead to an under-reporting of particular types of abuses, or an over-reporting of others.

On the other hand, it may be the case that an abuse consistently occurring over time only enters the reports once it is considered relevant to the designated areas of interest. Because of the structure of the reports, both organizations must make decisions about what is and is not included. The strict sections mandated to appear in each of the State Department country-year reports mean that they are sensitive to changes in the focus of the administration in Washington D.C. The reports are compiled across embassies under the guidelines laid out by a uniform document sent each spring, the Instruction Cable, detailing the format and content expected within the reports for that year. Changes in

the Instruction Cable, which are determined by the State Department in D.C. and applied uniformly to all embassies, directly impact what abuses are reported or emphasized each year. A clear example is the focus on abuses targeted at women in the 1993 country reports.<sup>22</sup> If one were to read through the reports chronologically, it would appear that in 1993 there was a surge in gender-based violence across countries. However, in truth, the Instruction Cable for that year directed embassies to pay special attention to human rights abuses experienced by women. Instances of rape by government officials were to be included in the section on torture, issues regarding voting rights to be covered in the section on the right to choose one's own government, and there was a specific request to focus on the trafficking of women within the reports. This is not to say that these abuses were not occurring within countries before 1993. Instead it demonstrates how the document requirements meant they were omitted from reports prior to 1993, and, therefore, were omitted from any indicators based on these reports.

Both Amnesty and the State Department are constrained by the structure and use of their reports by their respective organizations following their publication. For example, in addition to being under strict length constraints, like a page count or word count, Amnesty reports must also be translated into numerous languages to be consumed by their global audience. This influences the word used, the structure of the sentences, and the depth of information provided in the reports. One interviewee emphasized the importance of the process of constructing the reports: “the annual report itself is a really carefully timed and crafted piece of work that has to adhere to a very specific page count, word count, in order to be translated into all of the different languages, to be edited, to be revised, etc.”<sup>23</sup>

Likewise, the country reports produced by officers within the State Department are heavily constrained by the structure of the document. The pre-determined and uniform section headings decide what information is contained within the reports. If there is a specific abuse occurring within a country that does not fit into the section headings, then

---

<sup>22</sup>(Shattuck, 1993)

<sup>23</sup>Personal Interview, Interviewee #1

it is omitted from the reports. Crucially, these section headings are not constant over time. Figure 1 presents the section headings used within a country-year reports in 1980 (left panel) compared to those used in 2013 (right panel). Not only did the amount of sections change with the inclusion of a section on workers rights and discrimination, but the descriptions of certain sections changed. For example, in 1980, Section 1.f reported abuses related to the “Invasion of the Home.” This same section in 2013 includes any abuses related to “Arbitrary Interference with Privacy, Family, Home, or Correspondence.” The abuses that could be included within this section in 2013 cover a much wider range than those that could be included in this section in 1980.

<b>Section Headings, 1980</b>	<b>Section Headings, 2013</b>
<p>Section 1. Respect for the Integrity of the Person, Including Freedom from:</p> <ul style="list-style-type: none"> <li>a. Torture</li> <li>b. Cruel, Inhuman or Degrading Treatment or Punishment</li> <li>c. Disappearances</li> <li>d. Arbitrary Arrest and-Imprisonment</li> <li>e. Denial of Fair Public Trial</li> <li>f. Invasion of the Home</li> </ul> <p>Section 2. Government Policies Relating to the Fulfillment of Such Vital Needs as Food, Shelter, Health Care and Education</p> <p>Section 3. Respect for Civil and Political Liberties, Including:</p> <ul style="list-style-type: none"> <li>a. Freedom of Speech, Press, Religion and Assembly</li> <li>b. Freedom of Movement Within the Country, Foreign Travel, Emigration and Repatriation</li> <li>c. Freedom to Participate in the Political Process</li> </ul> <p>Section 4. Government Attitude and Record Regarding International and Non-governmental Investigation of Alleged Violations of Human Rights</p>	<p>Section 1. Respect for the Integrity of the Person, Including Freedom from:</p> <ul style="list-style-type: none"> <li>a. Political and Extrajudicial Killing.</li> <li>b. Disappearance.</li> <li>c. Torture and Other Cruel, Inhuman, or Degrading Treatment or Punishment.</li> <li>d. Arbitrary Arrest, Detention, or Exile.</li> <li>e. Denial of Fair Public Trial.</li> <li>f. Arbitrary Interference with Privacy, Family, Home, or Correspondence.</li> </ul> <p>Section 2. Respect for Civil Liberties, Including:</p> <ul style="list-style-type: none"> <li>a. Freedom of Speech and Press.</li> <li>b. Freedom of Peaceful Assembly and Association.</li> <li>c. Freedom of Religion.</li> <li>d. Freedom of Movement Within the Country, Foreign Travel, Emigration, and Repatriation.</li> </ul> <p>Section 3. Respect for Political Rights: The Right of Citizens to Change Their Government</p> <p>Section 4. Governmental Attitude Regarding International and Nongovernmental Investigation of Alleged Violations Human Rights</p> <p>Section 5. Discrimination Based on Sex, Race, Religion, Disability, Language, or Social Status</p> <ul style="list-style-type: none"> <li>Women.</li> <li>Children.</li> <li>Indigenous People.</li> <li>Religious Minorities.</li> <li>People with Disabilities.</li> </ul> <p>Section 6. Worker Rights</p> <ul style="list-style-type: none"> <li>a. The Right of Association.</li> <li>b. The Right to Organize and Bargain Collectively.</li> <li>c. Prohibition of Forced or Compulsory Labor.</li> <li>d. Minimum Age for Employment of Children.</li> <li>e. Acceptable Conditions of Work.</li> </ul>

Figure 1: State Department Section Headings

Even though most scholars would agree that the country-year reports produced by Amnesty International and the State Department are not representative of the actual number of abuses on the ground, there is still a reliance on them, and the scores coded from

them, to illustrate change over time. Underlying this understanding is an assumption of consistency in the way that these reports are produced. Both Clark and Sikkink (2013), and Fariss (2014) highlight the dangers of this assumption by drawing attention to temporal change in the interest and ability to collect information on human right abuses, as well as change in what is defined as a human rights abuse. We present another, more basic, process that determines what information is included or excluded from the country-year reports. Changes in procedures that determine how the reports are constructed and the process by which they are edited causes variation in the comparability of the reports over time.

This is not to argue against the claims of previous scholarship, but rather to suggest that there is an *additional* source of systematic bias within these reports. The restructuring of the State Department's reports in 2010 provides a perfect example of our argument. In order to reduce the time and space dedicated to State Department reports, embassies were instructed to only report one pertinent event per type of abuse, considered to be representative of a pattern, instead of documenting different instances.<sup>24</sup> The amount and type of information relayed in the reports directly changed due to adjustments in the procedural instructions. As one interviewee noted, the final report we utilize as scholars changes with variation in "constraints, time, budget, staff turnover, organizational mission change, changes in editorial standards, changes in types of outputs, [...] aimed at different audiences, changes in multimedia platforms that impact the reports."<sup>25</sup>

The editing process informs what is contained in the final report available to researchers and used to code popular indices of human rights abuses. This means that validation exercises comparing the results of analyses using both the PTS and CIRI score, or the CIRI score and Fariss' corrected latent measure, are still susceptible to the bias introduced in the editing process. Comparing indicators created from the same reports allows for researchers to account for possible biases in the coding of the reports. However, because the scores are

---

<sup>24</sup>Personal Interview, Interviewee #3

<sup>25</sup>Personal Interview, Interviewee #1

coded from the same source documents, the comparison does not address the prior source of bias that enters in the writing and editing stages of the reports. In other words, they may all be drinking from the same “poisoned” well.

Further, while the editing process differs between the organizations, the reports are not created independent from each other. Members of the State Department and Amnesty International explained how the process of gathering and confirming abuses to be entered into a report relies on communicating and swapping information between the organizations. Amnesty International and the State Department country-year reports are not produced independent of each other, but instead use each other as sources of information. Amnesty International may set up meetings with State Department officers to bring their attention to particular abuses. Alternatively, State Department officers may contact Amnesty International to confirm reports of abuses before writing the reports. Even under the best circumstance where the editing process biases the reports is random (i.e. the case of classical measurement error), considering the reports from both organizations will still result in biased temporal understandings.

We argue that it is important for human rights scholars to think carefully about using these reports as indicators of human rights violations over time. The country-year reports produced by Amnesty International and the U.S. State Department are appealing precisely because they are among the few sources that cover a wide range of countries over a wide range of time. However, scholars should not take this to mean that the information contained within the reports is comparable across time. In fact, our interviewees suggested that the reports are not meant to be read in relation to each other, but rather are meant to give a succinct summary of what the organization deemed important for that year. The reports are partial snapshots of a larger picture of the level of human rights violations. Importantly, those involved in working with the reports understand this and were typically distressed to learn that others were trying to use these reports to account for variation in human rights violations over time. One former Amnesty official said: “pretty much the

last thing we think about when writing is how well it's going to code."<sup>26</sup>

The editing dynamics we outline in this research note point to the difficulty in comparing the snapshots provided within the reports over time. The scope and intensity of an organization's reporting on a given abuse in a given year for a given country varies considerably over time. The variation produced by the editing process directly impacts the information contained within the reports. This was confirmed through the interviews we conducted with senior level editors and writers of the reports. For researchers, this means that we cannot separate out changes in violations due to on the ground changes versus changes in the what the editors are looking for within a given year. As a result, we would expect that the trends in human rights abuses as they appear in the reports to be an invalid representation of on the ground developments, especially over time. In the next section, we examine the content of the country-year reports to highlight the consequences of our claim that these reports are unreliable indicators of over time variation in human rights abuses.

## **5 The Effects of the Editing Process: Examining the Amount and Type of Information in Reports Over Time**

We expect that the annual reports produced by Amnesty International and the U.S. State Department vary both in terms of the amount of information they contain as well as the type of information they contain. Using two methods of text analysis, we look at the country-year reports for Argentina, Guatemala, and Mexico produced by Amnesty International and the State Department. From our interviews, we know that the information contained within the reports is a function of the editing process, in addition to information effects and changes in the standards of accountability. This section is concerned with whether these processes impact the reliability of the reports as over time indicators.

---

<sup>26</sup>Personal Interview, Interview #1

It may be the case that despite these processes, the reports are consistent in the amount and type of information they contain. This would suggest that either there is less variation in the identified sources of bias or that the reports are more resilient to their impacts. In either case, the reports themselves could be considered stable sources for coding purposes. On the other hand, if it is the case that the reports vary greatly in the amount and type of information over time, then this raises concerns for the reliability of the indices coded from them. This could be a result of fluctuation in the processes that determine what information is contained, which would bias our understanding of human rights abuses over time. It could also be due to actual variation in the amount and type of human rights abuses that are being accurately reported within the documents. Either explanation would mean that the use of a single score to represent the extent human rights abuses within a country would be biased.

To identify and consider variation in how much and what type of information is contained within a given country-year report, we analyze the documents using two separate methods. First, we use the event-coding software TABARI to leverage the structure of the text in order capture instances of abuse.<sup>27</sup> This provides us with an measure of how much information a report relays. TABARI uses shallow parsing to identify ‘who did what to whom’ from the syntactic structure of relevant sentences and to code these actions into relational events.<sup>28</sup> The second method we use, Latent Semantic Analysis (LSA), allows us to examine broader changes in content within the reports over time. While TABARI maintains and uses the structure of the document to identify events, LSA identifies the latent meaning of the reports by treating them as a “bag of words.” This allows us to understand variation in the type of information relayed by the country-year reports.

The country-year reports relay information on the presence of particular human rights abuse within a country. Scholars are interested in the occurrence of abuses, specifically the extent to which an abuse is present within a country. Importantly, the indices coded

---

<sup>27</sup>(Schrodt, 2014)

<sup>28</sup>More information regarding TABARI and the process by which we coded events can be found in Appendix

from these reports use the frequency of a particular abuse to assign a score indicating the state of abuse within a country.<sup>29</sup> In order to capture this emphasis on the occurrence of a human rights abuse, we treat abuses as events within a reports. An event is understood as *something* happening to *someone*. TABARI identifies events within the reports through the grammatical structure of a sentence; it looks for a noun → verb → noun combination to identify the actor, the action, and the target.<sup>30</sup> For example, TABARI reads the sentence:

“In December 1980 prisoners transferred in handcuffs from Unidad 6, Rawson, to Unidad 9, La Plata were severely beaten by military personnel”<sup>31</sup>

and identifies the actor, “military personnel,” the action, “beaten,” and the target of the action, “prisoners.”<sup>32</sup> In the context of these reports, TABARI identifies an event as an instance where there is an actor, a verb describing the type of abuse, and a target of abuse. In understanding human rights abuses as events, we can get an indication of the information used to code the overall state of abuses within a country.

Beyond change in the amount of information contained within the reports over time (see Appendix for full discussion of the TABARI event output), we are interested in how the content of the reports varies from year to year. Understanding how the type of information relayed by the reports changes over time is crucial to evaluating our ability to treat these reports, and the scores coded from them, as data points belonging to the same series. Variation in the type of abuses included in the reports, especially the addition of new abuses over time, indicates that the country-year reports are qualitatively different measures of human rights abuses over time.

We examine the content of the country-year reports along two dimensions. First, we categorize the type of events extracted by TABARI, defined in terms of the verb or action of

---

<sup>29</sup>Specifically, the CIRI Physical Integrity Score is coded based on how often a particular right is violated. More information in the Appendix.

<sup>30</sup>TABARI relies on dictionaries of verbs and actors specified by the researcher. More information on our creation of the verb dictionaries and the process of reading text into TABARI can be found in the Appendix.

<sup>31</sup>Amnesty International Country-Year Report, Argentina 1980.

<sup>32</sup>To see how this sentence is parsed by TABARI, see Figure ?? the Appendix.

the event, to compare the type of abuses that each organization records in each country-year report. By categorizing the abuses classified by TABARI, we can understand how the type of events varies over time. However, while TABARI maintains the structure of the sentence, it reads each sentence as separate record independent of the next or previous sentence. This means that it is unable to consider the document as a whole. To consider the information relayed by the document, as opposed to each sentence, we use Latent Semantic Analysis (LSA). LSA captures the latent meaning of the reports. To get a sense of how this latent meaning changes over time, we evaluate the similarity between the country-year reports. As opposed to TABARI, LSA evaluates the documents independent from their textual structure through a “bag of words” approach. This allows us to examine changes in the content of the reports not just in the words TABARI uses to identify an abuse, but in the meaning of the documents as a whole.

Together these measures of variation in the information contained within the documents gives us a sense of how reliable the country-year reports are over time. If the amount and type of information varies from year to year, or trends in a distinct direction, we know that this variation is not solely the product of variation in on the ground abuses. Instead, the information relayed by the reports is a function of the actual abuses *and* the processes that form the reports, like the editing process, information effects, and changing standards of accountability. When we quantify the reports and use them to track temporal trends of abuses within a country, we treat each year as being equally informative. However, the presence of these processes suggests that the scores are a product of different levels and type of information, that is not solely due to actual changes of the amount or type of abuses occurring within a country.

## **5.1 Case Selection**

In this research note, we are interested in examining how TABARI and LSA are able to identify temporal variation in the amount and type of information within the country-year

reports. To that end, we examine three countries – Mexico, Argentina, and Guatemala – to illustrate the potential of these methods and provide an initial look into what these reports contain and how they are changing over time.<sup>33</sup> We focus on these three countries over the entirety of time period available in the data, from 1980 to 2012.

We choose to focus on Mexico, Argentina, and Guatemala as an initial “plausibility probe” to examine these reports at this level of detail for several reasons. First, we focus on countries in the same region to counteract differences in the regional focus of these reports. In terms of media reporting, Hafner-Burton and Ron (2013) demonstrate that there is a “Latin Human Rights Bias” in newspaper coverage. By initially investigating these three Latin American countries, we hold constant the potential for this sort of spatial bias, and follow Hafner-Burton and Ron’s suggestion for more “regionally attuned” human rights scholarship.<sup>34</sup> Our focus on Latin America, as opposed to Southeast Asia or Europe, is also due to the region’s geographic location, specifically it’s proximity to the U.S. In particular, we expect that due to the geographic and historical ‘closeness’ between the U.S. and Latin America, there is a consistent baseline of attention given to this region within State Department reports. We know that country-year reports change over time in terms of their focus on types of abuses<sup>35</sup> as well as the location.<sup>36</sup> However, we expect that attentiveness to abuses within these countries to be less volatile than other regions of the world. The interest in Latin America is unlikely to wane over time due to its geopolitical location.

Mexico represents a case in which we would expect the focus of the State Department, over time, to be fairly constant and robust; it’s shared border and relationship with the U.S. almost ensures consistently detailed reports. Likewise, Amnesty International has one of their 17 offices in Mexico City. Each organization has opposing interests in Argentina.

---

<sup>33</sup>More information regarding the process through which the documents are cleaned and read into TABARI and LSA is provided in the appendix.

<sup>34</sup>(Hafner-Burton and Ron, 2013)

<sup>35</sup>(Bagozzi and Berliner, 2016)

<sup>36</sup>(Fariss, 2014)

Argentina has had high-profile human rights abuses during certain periods of its history. For example, the period known as the “Dirty War”, spanning from 1974 to 1983, where the military junta targeted political dissidents, students, and trade unionists, disappearing some 30,000 people and imprisoning those considered a threat to the regime.<sup>37</sup> The State Department, specifically Henry Kissinger, approved of the measures used to protect the neoliberal regime, while Amnesty International voiced concern over the human rights abuses taking place.<sup>38</sup> Finally, Guatemala does not enjoy the same international position as Mexico or Argentina. Because of its size and location, neither organization has constant access to the human rights abuses on the ground. The variation across countries allows us to compare across countries of different sizes and potential focus by the U.S. State Department and Amnesty International, but that have a well-documented history of human rights abuses.

## 5.2 Variation in Information and Content over Time

In this section, we consider the second dimension by which information may vary over time: change in the type of information contained within the reports. We use both TABARI and LSA to examine the substance of the country-year reports produced by Amnesty International and the State Department and consider how the content of the reports may be changing over time. If the reports vary greatly in the type of information presented over time, this raises concerns about the reliability of indices coded from the reports. While human coders may be able to distill abuses from changes in the amount of information, changes in content produces a different hurdle; coders cannot go back and code for an abuse that was not recorded at the time. To illustrate change in the content of reports, we leverage the type of events coded by TABARI, classifying the verbs indicating an instance of an abuse into broad categories of human rights violations. In addition, we use LSA to

---

<sup>37</sup>(Campbell, 2003)

<sup>38</sup>(Taylor, 1997)

extract the latent meaning of the reports. This allows us to inductively examine how the reports change in their meaning over time.

First, to illustrate how the broad composition of abuses relayed by the reports change over time, we classify the events coded by TABARI into four illustrative categories. We extract the operative verb identifying an event from TABARI and categorize it as one of four verb categories: physical integrity rights, civil/political rights, economic/social rights, and ‘other’ verbs related to human rights.<sup>39</sup> Figure 2 provides the count of events falling in each of the four verb categories.<sup>40</sup> Since we are primarily interested in how the number of events contained within each category changes over time, we care more for the consistency of classification over accuracy. These four categories provide a simple and consistent way to illustrate variation in how the focus of the documents differs between organizations as well as how the focus changes over time within a the organizations.

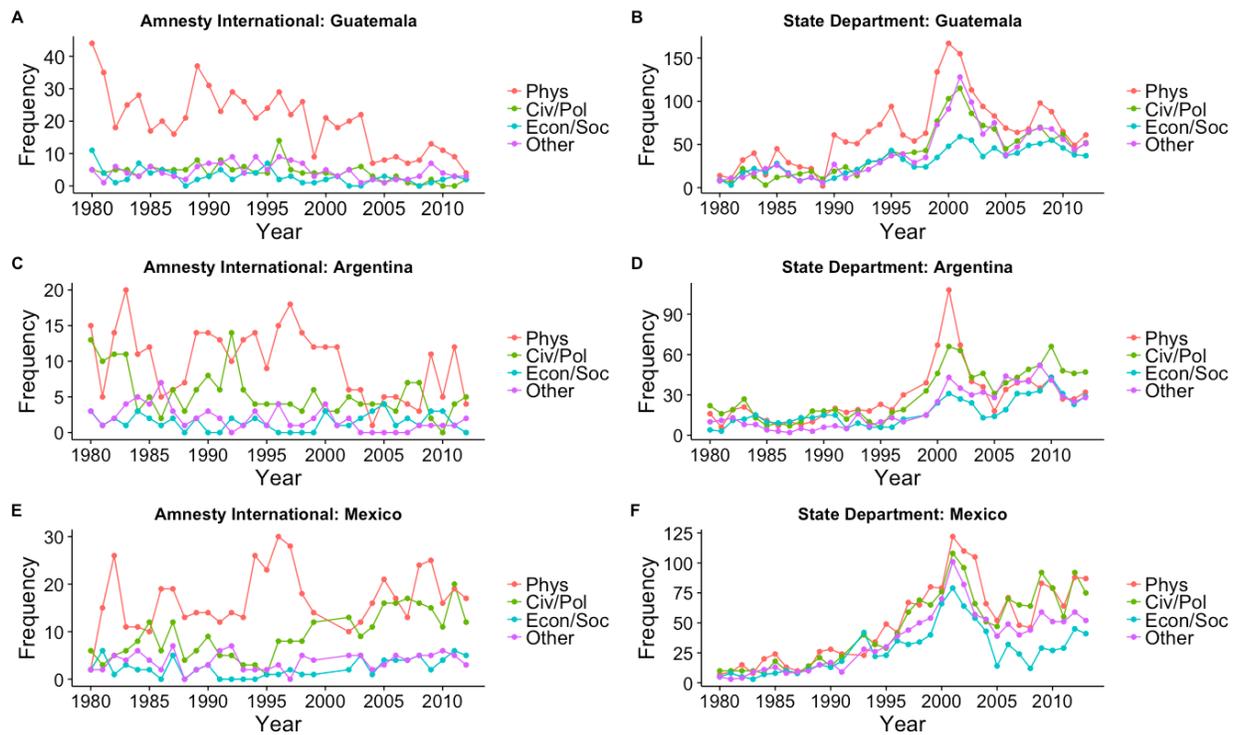


Figure 2: Verb Category Use Over Time

<sup>39</sup>Table of verbs provided in Appendix

<sup>40</sup>Graphs presenting proportion of events in each verb category appear in the Appendix. It should be noted that the verb categories are chosen specifically for ease of illustration.

As illustrated by the composition of the events extracted from the annual reports, Amnesty International clearly privileges physical integrity rights in their reporting. While this attention decreases overtime, their priority remains on reporting instances of physical abuse, aligning with their origin as an advocacy network against torture. By comparison, the State Department not only spreads its attention more evenly across the categories, but also exhibits a remarkable consistency over time and across countries. This aligns with the structure and process of State Department reports. The organization places a high degree of importance on the comparability of reports across countries within a year through the mandated structured format specified by the Instruction Cable sent to all embassies at the beginning of the compiling process each year.

However, since Figure 2 is based off individual events extracted from the sentences within the reports, it may tell us little about the content of the reports themselves. To inductively identify the type of information within the reports, we use LSA to map the change in content over time. We do this in two steps. First, we examine how similar State Department and Amnesty International reports are to each other. If the documents were identical from year-to-year, in terms of their latent content, they would occupy the same position in the “semantic space” and the correlation between them would be 1. If the documents were entirely distinct, they would occupy diametrically opposed positions in the latent semantic space, and the correlation between them would approach zero. By mapping the similarity of documents in a given year across the State Department and Amnesty International reports, we establish the level of similarity between the reports produced by each organization. Second, we can get a sense of what accounts for any difference between the reports by extracting the “top loading” words from each country-year report over time. The “top loading” words are give a sense of what the latent meaning of a document is for a given year. By comparing the differences in loading words, we can get a sense of how the latent meaning of the documents as a whole diverges across organizations.

Figure 3 plots the correlation between State Department and Amnesty country-reports for a given country-year. Since the reports cover similar subject matter, human rights, we would expect that their latent meaning be very similar. The correlation between the reports produced by LSA reflects this; reports are fairly consistent in their similarity (75% correlation) from year-to-year. Most interestingly, however, is how this otherwise high similarity varies over time. For all three countries, up until 1995, there is a substantial amount of variation in the similarity between State Department and Amnesty country-reports. Thereafter, the variation tightens. This is reflected across all three countries, suggesting that either both organizations are more consistently borrowing from each other in developing these reports, which our interviewees suggested happens regularly, or that the changing standards of accountability, identified by Fariss, are shaping the content of the reports across organizations.

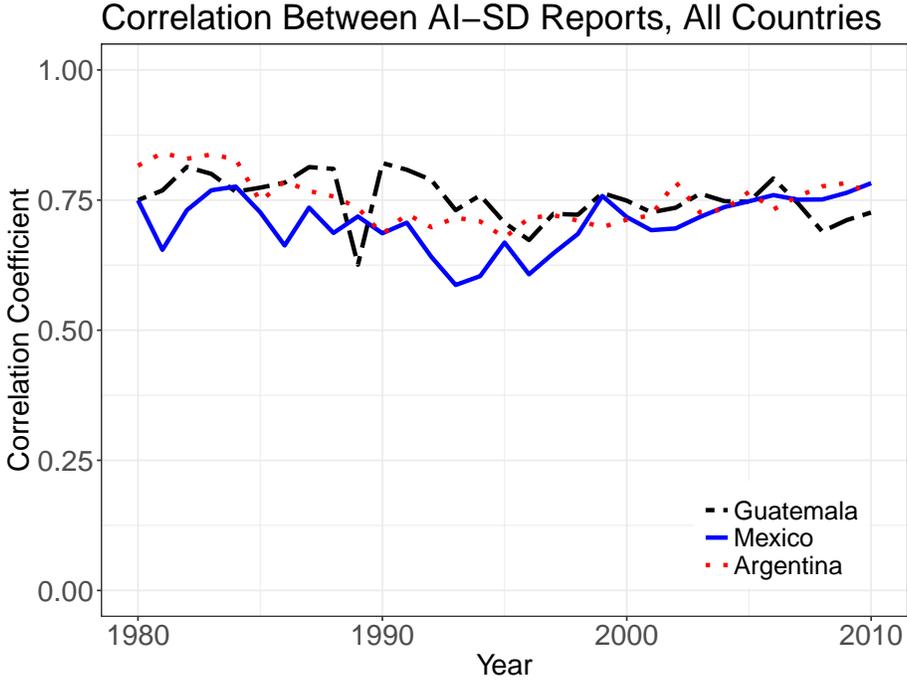


Figure 3: Correlation between Amnesty and State Department Country-Year Reports

Next, we examine the top 10 loading words for Amnesty and State Department reports for 1980, 1990, 2000, and 2010 (Table provided in Appendix). As a general trend, we

note that the State Department reports concentrate more heavily on institutions, such as courts, the military, and law, whereas Amnesty International reports focus most heavily on individuals and particular violations, such as killings, torture, and disappearances. More broadly, we observe shifts in the focus of the documents, particularly for Amnesty reports. The most recent Amnesty International reports on Argentina, place words such as ‘indigenous,’ ‘communities,’ and ‘provinces’ at the heart of the reports, suggesting a renewed interest in typically marginalized groups. A similar trend can be picked up in Guatemala where the 2010 report includes ‘women’ among the top 10 loading words. This is not to suggest, however, that the AI reports from 1980 and 2010 are diametrically opposed; it is likely that the reports are quite similar, but that the focus has changed slightly to reflect changes in human rights concerns over time. Taken along with the patterns of verbs across the organizations, this supports the claim that the type of information contained within the reports is not consistent over time.

## **6 Conclusion**

The research note argues that while the country-year reports produced by Amnesty International and the U.S. State Department may be useful for comparing human rights abuses across countries, they are inappropriate for over time comparisons. Through our interviews, we identified the editing process as playing a role in how much and what type of information is included in the published country-year reports. We then used machine-assisted text analysis techniques to evaluate our claim that the reports are unreliable over time indicators of human rights abuses. Our results suggest that scholars using quantitative indices based from these reports should think carefully about what process is driving the variation they are interested and whether it makes sense to use the snapshots provided by the reports to construct over time trends.

The interviews we conducted emphasized the danger of conceptualizing country-year

reports as accurate representations of “ground-truth”, particularly over time. Both Amnesty and the State Department decide upon yearly issues of focus that then determine what is included and emphasized within the country-reports for that year. The need for internal consistency and comparability across space limits what events and types of abuses are covered and in what way. While we are not the first to suggest that using country-reports as ground-truth is misleading, we seek to address not how accurate the information is, but how consistent it is over time.

Future work should also seek to address the issue of comparability by either supplementing analyses with case studies on both the countries of interest and the organization producing the reports, or by statistically accounting for changes within the data generating process in their models. Fixed effects may be one way to account for specific sources of variation originating at the organizational level.<sup>41</sup> In any case, researchers should take the role of organizations as actors producing this data seriously within their work. Whether it is using data-driven methods to identify changes in the scores, then cross-checking these changes with both qualitative information about violations within a country and possible changes within the organization, or by moving away from using the reports as measures of over-time variation, organizational dynamics should be accounted for when using these reports.

## References

Bagozzi, B. E. and Berliner, D. (2016). The Politics of Scrutiny in Human Rights Monitoring: Evidence from Structural Topic Models of US State Department Human Rights

---

<sup>41</sup>There are two primary issues for using fixed effects in this manner. First, to account for this level of variation, two-way fixed effects would need to be included accounting for organization-year variation. This puts a massive burden on most statistical models used by human rights scholars, with the potential for washing out interesting variation. In addition, two-way fixed effects rely heavily on the assumption of linearity meaning there is no dynamic causal relationship between treatment and outcome. (Imai and Kim, 2017). For human right scholars, this is problematic as it means assuming that the reports have no impact on the level of human rights violations within a country from one period to the next, and vice versa.

- Reports. *Political Science Research and Methods*, page 1–17.
- Caliendo, S. M., Gibney, M. P., and Payne, A. (1999). All the News that's Fit to Print. New York Times Coverage of Human-Rights Violations. *The International Journal of Press/Politics*, 4(4).
- Campbell, D. (Dec. 5, 2003). Kissinger approved argentinian 'dirty war'.
- Clark, A. M. and Sikkink, K. (2013). Information Effects and Human Rights Data: Is the Good News about Increased Human Rights Information Bad News for Human Rights Measures? *Human Rights Quarterly*, 35(3):539–568.
- Danzger, M. H. (1975). Validating conflict data. *American Sociological Review*, 40(5):570–584.
- Davenport, C. (2010). *Media Bias, Perspective, and State Repression: The Black Panther Party*. Cambridge University Press, New York City.
- Davenport, C. and Ball, P. (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.
- de Neufville, J. I. (1986). Human rights reporting as a policy tool: An examination of the state department country reports. *Human Rights Quarterly*, 8(4):681–699.
- Fariss, C. J. (2014). Respect for Human Rights has Improved Over Time: Modeling the Changing Standard of Accountability. *American Political Science Review*, 108(02):297–318.
- Hafner-Burton, E. and Ron, J. (2013). The Latin Bias: Regions, the Anglo-American Media, and Human Rights. *International Studies Quarterly*, 57(3):474–491.
- Imai, K. and Kim, I. S. (2017). When should we use linear fixed effects regression models for causal inference with longitudinal data? Working paper.

- Poe, S. C., Carey, S. C., and Vazquez, T. C. (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.
- Richards, D. L. (2016). The Myth of Information Effects in Human Rights Data: Response to Ann Marie Clark and Kathryn Sikkink. *Human Rights Quarterly*, 38(2):477–492.
- Ron, J., Ramos, H., and Rodgers, K. (2005). Transnational information politics: Ngo human rights reporting, 1986–2000. *International Studies Quarterly*, 49(3):557–588.
- Schrodt, P. A. (2014). Tabari version 0.8.4. <http://eventdata.parusanalytics.com/software.dir/tabari.html>.
- Shattuck, J. (Oct. 11, 1993). Article 8: Violations of Women’s Human Rights. *US DEPARTMENT OF STATE DISPATCH*, 4(41).
- Taylor, D. (1997). *Disappearing Acts: Spectacles of Gender and Nationalism in Argentina’s “Dirty War”*. Duke University Press.
- Vargas, T. R., Parkinson, S. E., and Bagozzi, B. E. (2016). Bleeding But Not Leading?: US Political Dynamics and Bias in International Event Data.
- Weidmann, N. B. (2014). On the Accuracy of Media-based Conflict Event Data. *Journal of Conflict Resolution*, 59(6).
- Weidmann, N. B. (2016). A Closer Look at Reporting Bias in Conflict Event Data. *American Journal of Political Science*, 60(1):206–218.