

Torture allegations as events data: Introducing the Ill-Treatment and Torture (ITT) specific allegation data

Journal of Peace Research
2014, Vol. 51(3) 429–438
© The Author(s) 2014
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0022343314524427
jpr.sagepub.com



Courtenay R Conrad

University of California, Merced

Jillienne Haglund

Florida State University

Will H Moore

Florida State University

Abstract

The Ill-Treatment and Torture (ITT) Data Collection Project uses content analysis to measure a number of variables on more than 15,000 public allegations of government ill-treatment and torture made by Amnesty International (AI) from 1995 to 2005. The ITT specific allegation (SA) event data use the torture allegation as the unit of observation, thus permitting users to manipulate them for a wide variety of purposes. In this article, we introduce the ITT SA data. We first describe the key variables in the SA data and report a number of bivariate descriptive statistics to illustrate some of the research questions that might be usefully investigated with the data. We then discuss how we believe the ITT SA data can be used to study not only AI's naming and shaming behavior, but also states' (lack of) compliance with the United Nations Convention Against Torture (CAT). We conclude with an empirical analysis using the SA data that investigates the effect of domestic political institutions on formal complaints, investigations, and adjudication of torture allegations. We show that legislative checks are positively associated with complaints, investigations, and trials; elections and freedom of speech are positively associated with investigations and trials; and powerful judiciaries are associated only with investigations.

Keywords

data collection, human rights, torture

Introduction

Amnesty International (AI), an international nongovernmental organization (INGO) working for the protection of human rights, documented allegations of torture and ill-treatment against 98 countries in 2011 (Amnesty International, 2011). Although telling of the amount of government torture occurring worldwide, this figure masks a remarkable degree of heterogeneity in state torture practices. In January 2012, for example, AI criticized a Migrant Accommodation Center in Western Ukraine for ill-treating and forcing into isolation a group of Somalis and Eritreans detained at the center (Amnesty

International, 2012c). In the same month, security and military forces in Libya were accused of leaving 'visible marks' on pro-Gaddafi loyalists, including open wounds on the head and limbs (Amnesty International, 2012a). The scars left by the Libyan military stand in stark contrast to AI allegations of torture against the United States that focus on water boarding, sleep deprivation, the use of loud music, and exposure to cold temperatures (Amnesty International, 2012b).

Corresponding author:

cconrad@ucmerced.edu

As these vignettes suggest, torture is perpetrated by many government agencies against a variety of victims. Torture techniques also vary widely, and government responses to allegations of torture are similarly heterogeneous across states. Although human rights advocates and scholars acknowledge these distinctions, existing large-N, cross-national data (e.g. Hathaway, 2002; Cingranelli & Richards, 2010b) on torture focus almost exclusively on its incidence at the country-year level of observation – the level of torture reported to have occurred in a given country during a given year. While those data have greatly increased our understanding of the covariates of torture, they are unable to support inferences about agents and victims, methods of abuse, and state response to torture allegations.

In this article, we introduce new cross-national event data on Amnesty International allegations of government torture from 1995 to 2005. The Ill-Treatment and Torture (ITT) Data Collection Project specific allegation (SA) data are different from previous data on state torture in at least two important ways.¹ First, by coding allegations of torture as events, ITT is able to generate information on a multitude of characteristics of allegations not currently available for large-N, cross-national analysis. Aside from information on torture incidence, the SA data include information on the duration of a particular allegation, the location of the event, the type of torture, the government agency alleged to have committed a given abuse, the socio-political group of which the alleged victim is a member, and the government response to the torture allegation. Second, the ITT project explicitly recognizes that actual state torture and ill-treatment are unobservable (e.g. Spierer, 1990; Clark, 2001: 57; Cingranelli & Richards, 2001: 230–231; Goodman & Jinks, 2003) and instead focuses on what can be measured reliably and validly (e.g. Bollen, 1986): AI allegations of such violations. This is a significant conceptual departure from previous data collection efforts that have implicitly treated the allegations contained in reports by INGOs and governments as representative of state's (lack of) respect for human rights.

Overview of the ITT specific allegation (SA) data

To generate data on AI allegations of torture and ill-treatment,² the ITT Data Collection Project performed content analysis of all Amnesty International publications published between 1995 and 2005. An allegation of state torture or ill-treatment occurs when AI alleges that the perpetrator is an agent of the state, the victim(s) is a person detained under the state's control, and the alleged abuse meets the definition of torture and/or ill-treatment in the United Nations Convention Against Torture (CAT).

Unlike existing data on human rights violations, the ITT Data Collection Project measures allegations of torture at two units of observation: specific allegations (SA) and country-year (CY) allegations. The distinction between these units of analysis involves the breadth of their spatial-temporal domain. The CY data report allegations of abuse targeted at a particular government agency over the course of an entire year. In contrast, SA data include only precise allegations about abuse in a specific place that is smaller than the country itself or that occurred during a period of time less than a year in duration. AI allegations that the state tortured an individual victim, that ill-treatment was prevalent in a single prison, or that torture occurred during the three weeks following an election are included in the SA data.

Take, for example, the following allegation against Uruguay in 2005: 'Amnesty International has been concerned at reports of torture and ill-treatment, including cases involving minors, in prisons, detention centres, and police stations' (AMR 52/002/2005). Because this allegation is about country-wide torture and ill-treatment, it is coded in the ITT country-year (CY) data, not the SA data presented here.³ Contrast that allegation with one targeted at Azerbaijan in 1995: 'Rafiq Ismayilov, a barber from the village of Digah, was detained in December by police officers from Masalli district on suspicion of theft and taken to the Regional Police Department, where he later died. Unofficial sources, however, allege that he died as a result of injuries sustained when he was beaten by police officers.' (AI Annual Report 1996). This allegation constitutes one

¹ The ITT project has also released data at the country-year level of analysis. The difference between the ITT SA data described here and the country-year (CY) data described in Conrad, Haglund & Moore (2013) is discussed in more detail below. Additional information on the ITT Project is available at http://faculty.ucmerced.edu/cconrad2/Academic/ITT_Data_Collection.html.

² Twenty-three undergraduate and four graduate students were responsible for the content analysis of AI documents. For more information on the processes by which coders were selected and trained, please refer to our Online Appendix.

³ Coders were instructed to code this allegation as ill-treatment and unstated torture against (1) unknown victims and (2) marginalized individuals by (1) prison officials and (2) police.

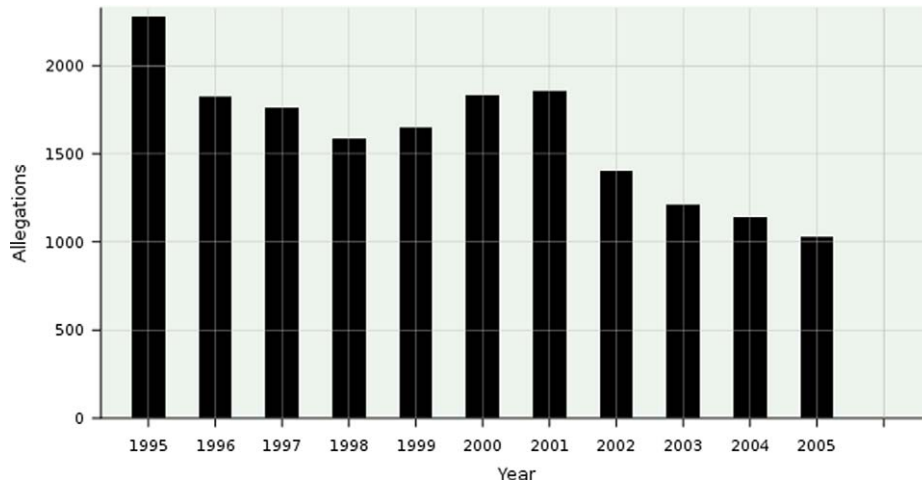


Figure 1. Number of AI allegations per year, 1995–2005

observation in the ITT SA data because it is about the abuse of one victim rather than patterns of violence across all of Azerbaijan in 1995. According to the ITT codebook used by our coders while generating the data,⁴ this AI statement is coded as an allegation of scarring torture by the police against one criminal victim that led to death. All of the government response variables for this observation are coded as Unknown.

Some bivariate relationships in the ITT SA data

There are over 15,000 events – allegations of torture and/or ill-treatment – included in the ITT SA Data.⁵ Figure 1 shows that the number of AI allegations per year tended to decline from 1995 to 2005. Although AI made fewer total allegations of torture over time from 1995 to 2005, there are over 1,000 allegations of torture and ill-treatment in 2005. While that change represents an over 50% decline from 1995, it is still a very large number of alleged violations of the UN Convention.⁶

⁴ This codebook is available on the ITT website (Conrad & Moore, 2010a).

⁵ The ITT SA data measure the population of AI allegations of ill-treatment and torture. Allegations are distinguishable from true rights violations, which are inherently unobservable (e.g. Spierer, 1990; Clark, 2001: 57; Cingranelli & Richards, 2001: 230–231).

⁶ Though it is beyond the scope of our effort, see Clark & Sikkink (2013) for an interesting analysis of change over time in the scores of human rights data projects.

What types of torture?

The ITT SA data distinguish among four ‘torture type’ allegations. Each of these is a binary variable indicating whether AI made an allegation of a particular type of abuse. *Scarring Torture* is coded when AI alleges torture that leaves marks on the human body (Conrad & Moore, 2011b: 11–12), and *Stealth Torture* or ‘clean’ torture is coded for allegations that do not leave marks on the body (Rejali, 2007). *Unstated Torture* distinguishes allegations of torture in which AI documents that torture occurred, but does not provide information regarding the type of torture alleged. The CAT not only prohibits torture, but also proscribes states from engaging in cruel, inhuman, or degrading treatment or punishment. *Ill-Treatment* is coded when AI alleges such behavior (Conrad & Moore, 2011b: 11–12). AI made roughly 8,000 allegations of scarring torture, over 3,000 allegations of stealth torture, 6,000 allegations of unstated torture, and over 6,000 allegations of ill-treatment between 1995 and 2005.⁷ The importance of the distinction between allegations and true rights violations raises an interesting question: Does AI allege scarring torture more frequently because it is the most utilized type of torture? Or are allegations of scarring torture more prevalent because it is the easiest type of torture to monitor? We do not yet know, but in the following section we discuss potential avenues for distinguishing between these possibilities.

⁷ Torture types are not mutually exclusive, and some of the individual events in the SA data include allegations of multiple torture types.

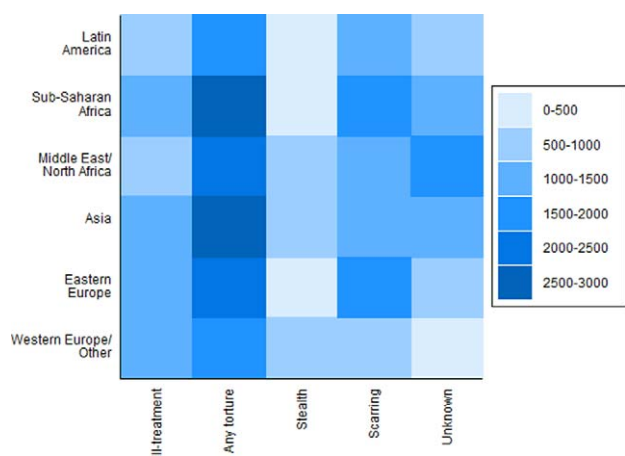


Figure 2. Number of AI allegations by region and torture type, 1995–2005

Cell values describe the number of AI allegations of each torture type occurring in each geographic region.

In the 11-year period from 1995 to 2005, Amnesty International issued over 2,000 allegations of torture and ill-treatment in each region of the world. We divided the world into the following geographic regions: Asia, Eastern Europe and the former Soviet Republics, Latin America, the Middle East and North Africa, sub-Saharan Africa, and Western Europe (including Australia, Canada, Japan, New Zealand, and the USA).⁸ Although AI is likely strategic in its decisions to publicize allegations, every region of the world is well-represented with regard to allegations of torture and ill-treatment from 1995 to 2005. Figure 2 displays the distribution of torture types by geographic region. All regions except Latin America and Western European were the subject of over 2,000 allegations of any torture, with the other two regions being named and shamed by AI more than 1,500 times each. With regard to ill-treatment, AI issued between 1,000 and 1,500 allegations for all regions except the Middle East and North Africa, which was the object of more than 500 allegations.

Figure 3 reports the percentage of each type of torture allegation occurring in the presence of several institutional variables: judicial power (Henisz, 2000); legislative checks (Henisz, 2000); freedom of speech (Cingranelli & Richards, 2010a); and competitive

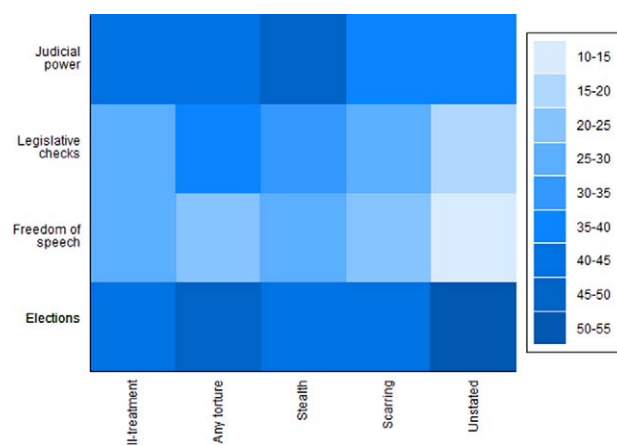


Figure 3. Prevalence of AI allegations by institution and torture type, 1995–2005

Cell values describe the percentage of AI allegations of each type of torture occurring in the presence of different political institutions.

elections (Cheibub, Gandhi & Vreeland, 2010). Each of these measures is a binary indicator, where 1 indicates the presence of the institution.⁹ The heat map depicts the percentage of each type of torture allegation occurring in states where each of these institutions is present, giving us a sense for how each type of torture is correlated with these institutions. AI made fewer than half of all ill-treatment, stealth, scarring, or unstated torture allegations in countries where these institutions are present. The prevalence of allegations of any type of torture or ill-treatment is also relatively lower in countries with full free speech protections. Further, approximately 15 to 20% of unstated torture allegations and 20 to 25% of all ill-treatment and scarring torture allegations occur in countries with an elected legislature. Contrast the relatively few allegations in the presence of these institutions with the number of stealth torture allegations in countries with a powerful judiciary. Half of all stealth torture allegations take place in countries with powerful judiciaries.

⁸ We split the Anglo democracies and Japan from their geographic regions because their wealth and democratic institutions generate a different global expectation of these governments, and they would be outliers in the Americas and Asia.

⁹ Henisz (2000) provides fractionalization measures of judicial power and legislative checks by coding the presence and party composition of the supreme (or constitutional) court and the legislature in each country. The freedom of speech variable is taken from the Cingranelli & Richards (2010a) dataset, which is coded on a 3-point scale. In Figures 3 and 4, we code freedom of speech as a binary variable, in which 1 represents states with full freedom of speech guarantees (2 on the CIRI coding scale) and 0 represents states with some censorship of the media (1 on the CIRI coding scale) or complete censorship of the media (0 on the CIRI coding scale).

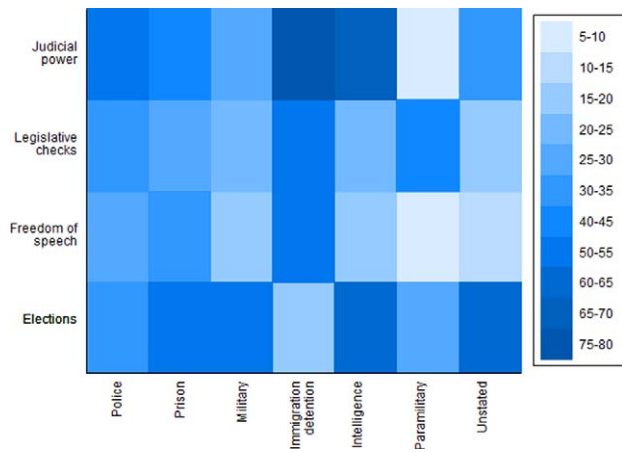


Figure 4. Prevalence of AI allegations by institution and agency of control, 1995–2005

Cell values describe the percentage of AI allegations of torture perpetrated by each agency of control occurring in the presence of different political institutions.

Which government agencies?

The SA data also include information on the government agency of control (AoC) alleged to be responsible for a given abuse. AI called out police for torture and ill-treatment considerably more frequently than any other AoC. AI published roughly 2,800 allegations of prison personnel violating the CAT; immigration detention centers, intelligence agencies, and paramilitary groups were each accused of fewer than 500 allegations.

Figure 4 reports the percentage of alleged responsibility by different AoCs for torture in the presence of the institutional variables described above. The figure shows the percentage of allegations of torture by police, prison officials, the military, immigration detention, intelligence, or unstated AoCs in states where there is a powerful judiciary, an elected legislature, free speech protections, or competitive elections. Approximately 50 to 65% of all allegations against police and 75 to 80% of allegations against immigration detention officials occur in countries where there is a powerful judiciary. Further, Figure 4 indicates that there is little correlation between competitive elections and the number of allegations made against prison officials and the military. Free speech protections are correlated with relatively fewer allegations against all agencies. Perhaps the boomerang model of NGO influence explains the relatively higher number of allegations made by AI against AoCs in countries with limited or no free speech protections (Keck & Sikkink, 1998). In countries where channels between domestic actors and the government are limited, as is the case in countries lacking free speech protections, transnational

networks, particularly NGOs such as AI, can bypass the state and amplify the demands of domestic groups.

The figure also shows that most allegations against immigration detention officials occur in countries where these domestic institutions are present. About 75 to 80% of all immigration detention allegations occur in countries with a powerful judiciary, and around half of all immigration detention allegations occur in countries with an elected legislature and free speech protections. There are two plausible explanations for this pattern. First, countries with a powerful judiciary and freedom of speech protections tend to have greater levels of macroeconomic productivity and higher average wages, making them more likely to not only have large immigrant populations (thus generating political demand for immigration restrictions), but also the state capacity to incarcerate. Second, AI tends to prioritize the publication of allegations of more grave violations relative to those of less grave violations.¹⁰ To the extent that political prisoners are more likely to be abused in states that lack powerful courts and/or a free press, AI may be less likely to report allegations of abuse that occur in immigration and detention centers.

Statistical models of state practice with ITT

The ITT Project provides data on the population of AI *allegations* of torture. As such, scholars interested in drawing inferences about the determinants of AI allegations – or naming and shaming – can use the SA data to do just that. But scholars who wish to use SA data to study the actual rights behavior of states (e.g. the extent to which the state complies with the CAT) must account for the strategic process by which AI generates allegations. Existing quantitative data on allegations of human rights violations are frequently employed as measures of violations of international and domestic law. As Bollen (1986), Spiner (1990), Clark (2001: 57), Cingranelli & Richards (2001: 230–231), and Goodman & Jinks (2003) point out, the implicit assumption that such data measure actual human rights practices is rather strong. In what follows, we elaborate on the distinction between allegations and violations and explain that researchers can utilize data on allegations to draw inferences about state violations of human rights by controlling for variables they believe will impact the likelihood that AI makes an allegation (i.e. learns about and then reports it).

¹⁰ Interview with senior AI staff, fall 2009.

It is standard practice in the quantitative study of human rights abuse to use content analytic data from AI and/or US State Department annual reports as a measure of states' performance with respect to their obligations under international treaties. Although allegations in AI documents are credible, there has been little discussion in the literature about whether the resulting data are representative of the actual level of state torture in a given country-year. Because AI is unlikely to report allegations with equal probability across countries, allegations are not an unbiased undercount of state violations of human rights. States (and their agents) face incentives to hide human rights abuse, and INGOs like AI face strategic and budgetary incentives that influence the extent to which it makes allegations of human rights violations (Hill, Moore & Mukherjee, 2013). Thus, researchers drawing inferences about human rights violations – rather than drawing inferences about AI allegations or naming and shaming – are typically interested in a latent concept measured using (1) a biased observable indicator produced by (2) a strategic actor.

In order to draw inferences about *actual* violations of the CAT, we show that a control variable approach can be successful – with a caveat. Researchers need to include as controls a set of variables that model the process by which AI generates allegations. The caveat is that some variables that will influence the likelihood that AI publishes an allegation will also impact the government's (lack of) respect for the CAT. Each variable that the researcher suspects will contribute to both the likelihood AI learns of, and reports, an allegation and the state's (lack of) respect can be included only once in the regression. As a result, the coefficient for that variable will be a combination of the two different impacts (upon both AI allegations and the state's performance), and there is no way to decompose the estimate into its two parts from that regression. There are more complex statistical models that permit one to decompose these effects, but they are not readily available in software packages such as Stata.¹¹ We illustrate this in the following empirical application, and in an Online Appendix demonstrate why that is so. We also provide a discussion in the Appendix about how researchers can directly model both AI's allegation production *and* states' (lack of) respect for the CAT.¹²

An example: Complaints, investigations, and adjudications

The ITT SA dataset considerably increases the number of research questions about government torture allegations and violations that scholars can test quantitatively. Unlike previous data on state torture, both the ITT SA and CY data include information about the victims alleged to have been abused and the agents alleged to be responsible for a given abuse. Such data allow researchers to investigate questions like: Does international treaty commitment affect torture by police officers and military officials differently? Does judicial effectiveness have the same (decreasing) effect on government torture by prison guards and immigration officers? Further, and we focus our empirical illustration here, the SA data include information about what happens *after* a torture allegation. In what follows, we illustrate a potential use of the ITT SA data by offering a preliminary inquiry into the question: What impact do domestic institutions have upon formal complaints, investigations, and adjudications of torture allegations?

Data and empirics

Because the SA data are the first to contain information about states' investigations into allegations of torture, the outcome of the investigation, and the outcome of any judicial proceedings that were reported, we have selected those three variables as interesting for an empirical illustration. To that end, we explore whether or not domestic institutions – contested elections, judicial effectiveness, freedom of speech, and institutional veto – influence formal complaints, investigations, and adjudications of torture allegations. Because the majority of data on our independent variables are coded at the country-year unit of analysis, we aggregate the ITT SA data to the country-year unit of observation. To do so, we calculate the raw sum, within each country-year in our data, of (1) formal complaints lodged against the state, (2) state investigations into those complaints, (3) investigations that produced an international court case, and (4) investigations that produced a domestic court case.¹³ Because our dependent variables are counts, we require count models.¹⁴ Judicial power captures the extent to which the

¹¹ Finite mixture models permit one to use two (or more) equations to estimate the impact of a variable on two different outcomes, such as the level of a state's (lack of) respect for rights and the likelihood that AI reports it.

¹² See Cameron & Trivedi (1998: Chapter 10) for a discussion of event count regression models that permit one to model measurement error.

¹³ In instances in which an allegation of torture is reported to have begun in year t and ended in year $t+N$, we code all years as having experienced an event.

¹⁴ The negative binomial model adds a parameter to account for dispersion and reduces to the Poisson when this parameter is equal to zero (Long & Freese, 2001).

Table I. Effect of covariates on complaints, investigations, and international and domestic trials

	<i>Complaints</i>	<i>Investigations</i>	<i>International trials</i>	<i>Domestic trials</i>
Elections	1.185 (0.129)	1.308* (0.163)	6.225* (3.925)	1.484* (0.234)
Judicial power	1.211 (0.269)	1.624 [†] (0.405)	0.290 (0.296)	1.571 (0.486)
Legislative checks	1.769* (0.446)	3.055** (0.904)	1.292 (1.828)	2.084* (0.774)
Freedom of speech	1.062 (0.075)	1.179* (0.097)	1.894 [†] (0.695)	1.099 (0.117)
GDP	1.184** (0.58)	1.199* (0.002)	0.898 (0.173)	1.248* (0.088)
Population	1.125 (0.333)	1.028 (0.356)	0.156 (0.352)	0.415* (0.180)
Restricted access	2.518** (0.003)	2.307** (0.290)	5.111** (2.308)	2.978** (0.442)
Constant	1.988** (0.162)	0.657** (0.066)	0.006** (0.003)	0.289** (0.037)
<i>N</i>	1,217	1,217	1,217	1,217

Incidence response rates reported, unadjusted standard errors in parentheses; [†] $p < 0.10$; * $p < 0.05$; ** $p < 0.01$ (two-tailed).

domestic judiciary is autonomous and effective. We utilize a measure of *Judicial Power* constructed by Linzer & Staton (2011), who employ a latent variable measurement model composed of eight indicators of de facto judicial independence, producing a continuous variable ranging from 0 to 1. To measure legislative checks we use the *Checks* variable from the Database of Political Institutions (Keefer, 2005), which is an integer count of the number of veto players.¹⁵ *Elections* is a binary variable, indicating the presence of competitive elections and comes from Cheibub, Gandhi & Vreeland (2010). *Freedom of Speech* is a trichotomous variable, coded 0 for complete government censorship of the media, 1 for some censorship, and 2 for no censorship (Cingranelli & Richards, 2010a). Data on country wealth (*GDP*, measured in billions of US dollars) and *Country Population* (measured in millions) come from the World Development Indicators (WDI).

To use SA data to draw inferences about the conditions under which states violate the CAT, we must include covariates that model the likelihood that AI reports an investigation, adjudication, etc. The extent to which AI is able to generate allegations of human rights violations and report about their investigation and adjudication is dependent on its ability to work within a given country. Although AI maintains local offices in many countries, some governments prevent NGOs from operating within their borders. In these cases, it is more difficult for AI to have access to victims, as well as local advocates, and thus more difficult for them to make allegations and report on investigations and adjudications – even when violations

occur. To account for such biases in the production of allegations and reports of investigations and adjudications, we include a variable from the ITT country-year (CY) data (Conrad & Moore, 2011a).¹⁶ *Restricted Access* is a binary measure that captures whether or not AI published a statement that it, or another INGO, had difficulty gaining access to detainees during a given country-year (Conrad & Moore, 2011a).

Empirical results

Table I reports the results from four negative binomial regression models. These results are intended to be suggestive and to encourage future work on these topics. Note that we have not conducted a battery of robustness checks, we have not sought to model any potentially conditional relationships,¹⁷ and we do not have a strong theoretical case for the model specification.

To provide a sense of the substantive impact of the variables, rather than report coefficient estimates and their standard errors, we report incidence response ratios (IRRs) in Table I. IRRs are the exponentiated value of the coefficient and can be interpreted as the expected change in the number of torture events, given that AI observes one event. An IRR value of 1 indicates no change, and a value less than 1 indicates a decrease by a factor of IRR in the expected count. An IRR value greater than 1 represents an expected increase of $100 \times (\text{IRR} - 100)$ percent.

¹⁵ These operational indicators differ from the binary indicators we used to produce Figures 3 and 4, which we selected because a binary indicator is preferable for the heatmaps. In a multiple regression we prefer the greater information available in the continuous and integer indicators (Linzer & Staton, 2011; Keefer, 2005).

¹⁶ In the analyses that follow, we are interesting in modeling the influence of covariates on the investigation and adjudication of torture. Other scholars may be interested in modeling torture as their dependent variable. For scholars interested in drawing inferences about torture using the ITT SA data, we encourage them to model AI's organizational incentive to report an allegation in addition to (lack of) access.

¹⁷ For example, it may very well be the case that AI's reporting about trials is restricted to cases where there was an investigation.

To begin, note that restricted access, the variable that we include as a control for AI's propensity to report, is not only highly statistically significant in all four regressions, but also has an estimated IRR of at least 2 in each model (100%+ increase). AI's reports of restricted access are associated with relatively higher numbers of AI allegations of complaints, investigations, and trials.¹⁸ This result highlights the importance of carefully modeling the data generating process by which AI produces allegations when using the ITT SA data to draw inferences about state behavior.¹⁹

Turning to the institutional variables, elections are positively related to the number of investigations and trials, having a particularly large impact upon international trials (over 500% increase compared to countries that do not hold elections). The independence and effectiveness (i.e. power) of courts increases the number of investigations countries conduct, but is not associated with greater numbers of complaints and domestic trials. Veto, the extent to which the legislature can check the executive, is positively associated with the number of complaints brought against the state, the number of investigations it conducts, and the number of domestic trials that occur. Given these relationships it would be interesting to explore whether executives misjudge the extent to which legislatures (1) might generate information about violations, thereby spurring complaints, (2) pressure the bureaucracy to investigate (or do so on their own), and (3) signal to rights advocates that domestic courts will be receptive to cases brought against the state. Finally, greater respect for freedom of speech is positively associated with both larger number of investigations and international trials. That freedom of speech has a positive impact upon investigations is unsurprising, but it is interesting that it does not similarly impact complaints against the state. Further, international trials are more likely as respect for freedom of speech rises, but domestic trials are unaffected.

¹⁸ Given that restricted access is coded only for the first year in which the ill-treatment or torture allegation occurred, it is plausible that AI (as well as other INGOs) gain access in subsequent years, or the years in which formal complaints are filed, and investigations and trials are conducted. This indicates that INGO reporting of restricted access in year t when the allegation occurred may be related to relatively higher numbers of AI allegations of complaints, investigations, and trials in following years.

¹⁹ See the discussion in the 'Statistical models of state practice with ITT' section and in Conrad, Haglund & Moore (2013) for further discussion of the distinction between AI allegations and state behavior.

Conclusion

Between 1995 and 2005 Amnesty International made public more than 15,000 allegations against states for violating the United Nations Convention Against Torture (CAT). Those allegations represent a jarring, though uncountable, number of human beings who, while under state control, were subjected to ill-treatment, stealth torture, and/or scarring torture. Yet AI allegations are an undercount of actual violations: because both perpetrators and states generally have an incentive to hide violations of the CAT, it is not possible to know how many violations actually occurred. As a consequence, the ITT project departs from previous practice, and rather than use AI allegations as a direct measure of state violations, uses them instead as a measure of AI activity. Nevertheless, we argue that it is possible to use the ITT data to draw inferences about state compliance by controlling for the factors that affect the generation of AI allegations.

We report a number of intriguing patterns that we hope stimulate interest in the ITT SA data. For example, Figure 1 records a near monotonic decline in the annual number of allegations made by AI between 1995 and 2005. The two years that deviate from the pattern of decline are 2000 and 2001, not 2002 and 2003 as one might anticipate in the wake of states' responses to Al Qaeda's attacks on 11 September 2001 and the Bush administration's roll-out of its enhanced interrogation and extraordinary rendition programs. That temporal decline is intriguing and warrants investigation. It might represent a decline in AI's capacity to observe abuse, but might also indicate a variety of alternative possibilities including the justice cascade (Meernik, Nichols & King, 2010; Sikkink, 2011). Similarly, we show interesting variation in allegations across regions of the world, and Figure 3 indicates that freedom of speech is associated with types of alleged torture. We learn that police are the agency of control most likely called out by AI, and that is true across all regions of the world. We hope that the ITT SA data will encourage researchers to investigate these and other patterns.

We pursued the ITT project to permit us to examine research questions that arose out of the research reported in Conrad & Moore (2010b). In particular, we are not satisfied with that study's inability to distinguish national security torture from criminal and social control torture (Rejali, 2007). The ITT SA data make it possible for us to conduct work that makes a much more satisfying distinction, and in so doing will

reveal new and interesting questions that otherwise would have likely gone unnoticed. We have plans to conduct a number of inquiries that will explore the impact of liberal democratic institutions on states' compliance with the CAT. Those studies but scratch the surface of what might be done with the ITT SA data, and we are eager to learn what others are able to illuminate.

Replication data

Data, replication files, and an online appendix to accompany this article can be found on the authors' websites and at <http://www.prio.no/jpr/datasets>.

Funding

The ITT Data Collection Project has received support from the US National Science Foundation (Grants #0921397 and #1123666), the Department of Political Science at Florida State University, the School of Social Sciences, Humanities, and Arts at the University of California, Merced, and the Department of Political Science and Public Administration at the University of North Carolina at Charlotte.

References

- Amnesty International (2011) *Human Rights Facts and Figures 2011*. New York: Amnesty International.
- Amnesty International (2012a) Libya: Deaths of detainees amid widespread torture. AI news report, January 26 (<https://www.amnesty.org/en/news/libya-deaths-detainees-amid-widespread-torture-2012-01-26>).
- Amnesty International (2012b) Torture and rendition: Guantánamo case profiles. AI news report, January 9 (<http://www.amnesty.org/en/news/guantanamo-case-profiles>).
- Amnesty International (2012c) Ukraine: Asylum seekers detained and ill-treated. AI Urgent Action document, January 31 (<http://www.amnesty.org/en/library/info/EUR50/001/2012/en>).
- Bollen, Kenneth A (1986) Political rights and political liberties in nations: An evaluation of human rights measures, 1950 to 1984. *Human Rights Quarterly* 8(4): 567–591.
- Cameron, A Colin & Pravin K Trivedi (1998) *Regression Analysis of Count Data*. New York: Cambridge University Press.
- Cheibub, José Antonio; Jennifer Gandhi & James Raymond Vreeland (2010) Democracy and dictatorship revisited. *Public Choice* 143(1–2): 67–101.
- Cingranelli, David L & David L Richards (2001) Measuring the impact of human rights organizations. In: Claude E Welch, Jr (ed.) *NGOs and Human Rights: Promise and Performance*. Philadelphia, PA: University of Pennsylvania Press, 225–237.
- Cingranelli, David L & David L Richards (2010a) The Cingranelli and Richards (CIRI) human rights data project. *Human Rights Quarterly* 32(2): 401–424.
- Cingranelli, David L & David L Richards (2010b) The Cingranelli-Richards (CIRI) human rights dataset (<http://humanrightsdata.blogspot.com/>).
- Clark, Ann M (2001) A calendar of abuses: Amnesty International's campaign on Guatemala. In: Claude E Welch, Jr (ed.) *NGOs and Human Rights: Promise and Performance*. Philadelphia, PA: University of Pennsylvania Press, 55–71.
- Clark, Ann M & Kathryn Sikkink (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.
- Conrad, Courtenay R; Jillienne Haglund & Will H Moore (2013) Disaggregating torture allegations: Introducing the Ill-Treatment and Torture (ITT) country-year data. *International Studies Perspectives* 14(2): 199–220.
- Conrad, Courtenay R & Will H Moore (2010a) The Ill-Treatment & Torture (ITT) data project coding rules and norms (http://faculty.ucmerced.edu/cconrad2/Academic/Under_the_Hood_files/ITTCodebookNorms6-May10.pdf).
- Conrad, Courtenay Ryals & Will H Moore (2010b) What stops the torture? *American Journal of Political Science* 54(2): 459–476.
- Conrad, Courtenay R & Will H Moore (2011a) The Ill-Treatment & Torture (ITT) data project: Country-year data user guide (http://faculty.ucmerced.edu/cconrad2/Academic/Data_files/ITT_CY_UsersGuide19July11.pdf).
- Conrad, Courtenay R & Will H Moore (2011b) The Ill-Treatment & Torture (ITT) data project: Specific allegation data user guide (<http://faculty.ucmerced.edu/cconrad2/Academic/Data.html>).
- Goodman, Ryan & Derek Jinks (2003) Measuring the effects of human rights treaties. *European Journal of International Law* 14(1): 171–184.
- Hathaway, Oona A (2002) Do human rights treaties make a difference? *Yale Law Journal* 111(8): 1935–2042.
- Henisz, Witold J (2000) The institutional environment for economic growth. *Economics and Politics* 12(1): 1–31.
- Hill, Daniel W; Will H Moore & Bumba Mukherjee (2013) Information politics versus organizational incentives: When are Amnesty International's 'naming and shaming' reports biased? *International Studies Quarterly* 57(2): 219–232.
- Keck, Margaret E & Kathryn Sikkink (1998) *Activists beyond Borders: Advocacy Networks in International Politics*. Ithaca, NY: Cornell University Press.
- Keefer, Philip (2005) DPI2004 Database of Political Institutions: Changes and variable definitions. Washington, DC: The World Bank (http://siteresources.worldbank.org/INTRES/Resources/DPI2004_variable-definitions.pdf).
- Linzer, Drew A & Jeffrey K Staton (2011) A measurement model for synthesizing multiple comparative indicators: The case of judicial independence. Unpublished manuscript

- (<http://polisci.emory.edu/faculty/jkstato/resources/WorkingPapers/LinzerStaton.pdf>).
- Long, J Scott & Jeremy Freese (2001) *Regression Models for Categorical Dependent Variables Using Stata*. College Station, TX: Stata.
- Meernik, James D; Angela Nichols & Kim L King (2010) The impact of international tribunals and domestic trials on peace and human rights after civil war. *International Studies Perspectives* 11(4): 309–334.
- Rejali, Darius (2007) *Torture and Democracy*. Princeton, NJ: Princeton University Press.
- Sikkink, Kathryn (2011) *The Justice Cascade: How Human Rights Prosecutions are Changing World Politics*. New York: WW Norton.
- Spirer, Herbert F (1990) Violations of human rights – how many? *American Journal of Economics and Sociology* 49(2): 199–210.
- COURTENAY CONRAD, b. 1983, PhD in Political Science (Florida State University, 2010); Assistant Professor, University of California, Merced; research interests: state repression, international organizations, domestic political institutions; articles have appeared in the *American Journal of Political Science*, *International Studies Quarterly*, and *Journal of Conflict Resolution*, among others.
- JILLIENNE HAGLUND, b. 1985, MA in Political Science (University of Memphis, 2009); PhD Candidate in Political Science, Florida State University; research interests: human rights, international law/courts, domestic political institutions; articles have appeared in *International Studies Perspectives*.
- WILL H MOORE, b. 1962, PhD in Political Science (University of Colorado, 1991); Professor of Political Science, Florida State University; research interests: dissent and repression, human rights, intrastate conflict, research methods; author of *A Mathematics Course for Political and Social Research* (with David A. Siegel, Princeton University Press, 2013); articles have appeared in the *American Journal of Political Science*, *International Studies Quarterly*, and *Journal of Conflict Resolution*, among others.