International Studies Quarterly

International Studies Quarterly (2013) **57**, 219–232

Information Politics Versus Organizational Incentives: When Are Amnesty International's "Naming and Shaming" Reports Biased?[1]

Daniel W. Hill Jr
University of Georgia

Will H. Moore
Florida State University

and

Bumba Mukherjee
Penn State University

''Information politics'' INGOs such as Amnesty International have incentives to maintain their credibility by carefully vetting information about rights abuses committed by governments. But they are also strategic actors that may inflate allegations of abuse to fulfill organizational imperatives. This raises an intriguing question: When are INGOs more likely to exaggerate their allegations? In answer to this question, we argue that news media reporting pressures INGOs to comment for organizational reasons, even if the information available to them is poor. On the other hand, higher numbers of domestic human rights NGOs increase the quality of available information, and INGOs will find more credible information provided about states as the winning coalition to the selectorate rises. Yet, an incentive to exaggerate allegations under certain conditions does not imply that INGOs will always do so. Indeed, there exists significant observed variation in INGO reports about government abuse. To test our hypotheses, we employ a zero-inflated ordered probit model with correlated errors that permits us to model an unobservable probability (the probability that the INGO exaggerates its allegations) and correct for potential bias. Results provide support for our hypotheses, and suggest that Amnesty International adheres to its credibility criterion, rarely succumbing to incentives to exaggerate abuse.

It has long been recognized that to achieve their normative goals, human rights international non-governmental organizations (hereafter INGOs) such as Amnesty International (AI) and Human Rights Watch (HRW) have strong incentives to vet information carefully and evaluate human rights practices by governments as accurately as possible (Korey 1968; Cmiel 1999; Clark 2001; Wong 2008:135; Lake and Wong 2009:140,144). More recently, scholars have argued that though INGOs have normative goals, they are also strategic actors (Finnemore and Sikkink 1998; Keck and Sikkink 1998; Hudock 1999; Cooley and Ron 2002; Ron et al. 2005; Kelley 2009; Murdie 2009b; Hafner-Burton and Ron 2009; Hyde 2010). That ''information politics'' INGOs are at times strategic helps explain why they recognize that their influence rests largely on their credibility. Yet only recently have researchers begun to explore the extent to which other organizational incentives might work at cross-purposes (Finnemore and Sikkink 1998:899; Kelley 2009:767). For example, Ron, et al. show that state power, US military assistance, and media coverage influence the ''naming and shaming'' activity of AI (Ron et al. 2005:558). In addition, Hafner-Burton and Ron observe that rights INGOs such as AI or HRW might continually paint a poor picture of states' human rights performance even if that performance had consistently improved (Hafner-Burton and Ron 2009:383). This study pursues some new questions within that domain: When the quality of the information upon which they rely degrades, do ''information politics'' human rights INGOs such as AI and HRW compromise their credibility demands? More generally, when are INGOs more (or less) likely to exaggerate their allegations? This inquiry adds another brick to the foundation of the nascent literature on the ethics versus bureaucratic politics literature on INGOs (Gibelman and Gelman 2004; Gugerty 2009; Hyde 2010; Gourevitch and Lake 2011; Hyde 2011).

To answer these questions, we develop a theoretical explanation that treats INGOs like AI as strategic actors that interact with domestic human rights NGOs as well as the states they seek to influence.[2] More substantively, we follow the lead of scholars who have drawn attention to the impact of domestic political institutions upon state's respect for human rights (Hathaway 2005; Davenport 2007; Powell and Mitchell 2007; Vreeland 2008; Powell and Staton 2009; Simmons 2009; Conrad and Moore 2010) and explore how states' ratio of the winning coalition to the selectorate (Bueno de Mesquita et al. 2003), the strategic role of domestic human rights NGOs (here-

[1] Previous versions were presented at the 2010 Annual Meeting of the International Political Economy Society, Boston, and the 2009 Annual Meeting of the American Political Science Association, Toronto. We would like to thank Clair Apodaca, Cullen Hendrix, and Wendy Wong for their helpful comments. Replication files are available at http://myweb.fsu.edu/dwh06c, and at URL FOR ISQ SITE

[2] On INGOs as strategic actors, see Hudock (1999:20–21), Cooley and Ron (2002:5–6,36), Murdie (2009b), Hyde (2010). On interactions between INGOs and domestic NGOs, see Tsutsui and Wotipka (2004), Wong (2008), Murdie (2009a).

Hill, Daniel W., Will H. Moore, and Bumba Mukherjee. (2013) Information Politics Versus Organizational Incentives: When Are Amnesty International's ''Naming and Shaming'' Reports Biased? *International Studies Quarterly*, doi: 10.1111/isqu.12022
© 2013 International Studies Association

after domestic NGOs) and information about torture committed by states influences the likelihood with which INGOs exaggerate their allegations of state-sponsored rights abuse such as torture.

The theory first suggests that an INGO is more likely to inflate allegations of torture against the government when information flows from the media about rights abuses (including torture) committed by the government increase. In particular, dissident–government violence in a country both reduces the *quality* of information INGOs can collect from domestic NGOs[3] about specific allegations *and* increases the amount of international media coverage of rights abuses. Increasing media coverage creates fertile ground for INGOs to recruit labor and donations (Simon 2006; Brown and Minty 2008), which are the primary resources upon which these INGOs rely to operate (Cmiel 1999; Wong 2008; Lake and Wong 2009). Thus, greater media coverage of rights abuses encourage the INGO to inflate allegations of torture conducted by the government, even though the quality of the information that they obtain from domestic NGOs may be below their ''credibility standard.'' While several scholars have called attention to INGOs' reliance on domestic NGOs for information (Keck and Sikkink 1998; Clark 2001; Stewart 2004; Hafner-Burton and Tsutsui 2005; Wong 2008), less explored are the consequences of this reliance for the quality of information disseminated by INGOs about government practices. Our theoretical account provides a first step toward exploring the impact of information provided by domestic NGOs about rights practices by states. More broadly, by examining how the behavior of international organizations varies across domestic contexts, our theory also gives heed to scholars who encourage exploration of the nexus between international organizations and domestic politics (Martin and Simmons 1998; Milner 2005; Lake 2010).

Second, the theory also explains that the political incentives created by domestic institutions affect not only the human rights behavior of governments, but also the belief held by INGOs about such behavior. Focusing specifically on the ratio of the winning coalition to the selectorate (the W/S ratio), our analysis suggests that INGOs issuing reports about abuses are less likely to exaggerate their reports when the W/S ratio is high. This is because a high W/S ratio makes it costlier for the government to abuse its citizens, and knowledge of these constraints causes INGOs to readily believe that the government does not engage in widespread abuse.

Yet, a theoretical incentive to exaggerate allegations (which is *unobservable*) on some occasions does not necessarily translate into behavior. Do INGOs such as AI succumb to this incentive with regularity when reporting rights abuses by governments including torture? We approach this question by first examining a plot of the allegations—which are *observable*—that INGOs such as AI have made.[4] The CIRI project conducts content analysis on such allegations, and we use their coding of torture allegations to produce Figure 1.[5] The figure summarizes 191 countries' ''torture report cards'' for each year from 1981 to 2007. The CIRI scale ranges from 0 to 2, where a value of 0 represents a country accused of engaging in widespread use of torture, a value of 1 indicates a country that sometimes engages in torture, and a value of 2 identifies a country against which no allegations of torture were made.

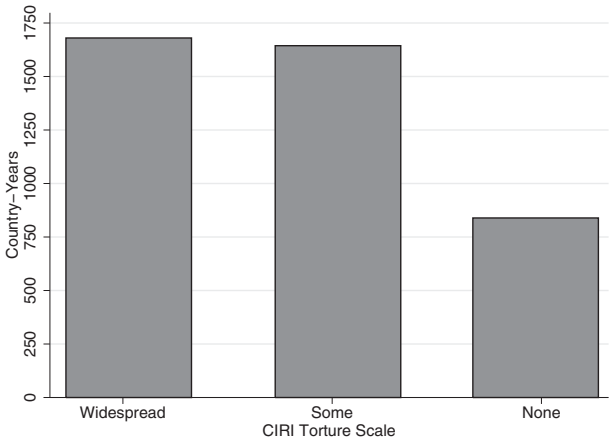

Fig 1. State's Use of Torture, 1981–2007

Figure 1 reveals some substantively interesting variation since it shows that roughly 40% of the observations reported by INGOs indicate widespread use of torture, while the remaining set of reported observations indicate greater respect by governments for the right to freedom from torture.[6] What explains this *observed* variation in reports issued by INGOs about the extent to which governments respect the right to freedom from torture? We argue that the need to raise resources, the amount and quality of information available, and the executive's dependence upon a large winning coalition each play important roles.

An appropriate test of the theoretical predictions summarized above must address a substantial econometric challenge: We can only observe the INGO's report; however, whether the INGO issued a report despite being uncertain about it is unobserved. Because of this, when we estimate the parameters for our statistical hypothesis tests we need to account explicitly for the factors that determine whether or not the INGO exaggerated a given allegation. In other words, to conduct a test of our predictions, we require a two-stage model that allows us to jointly test when INGOs are more (or less) likely to exaggerate their allegations (the first stage) *and* when they are likely to allege a particular level of abuse, given that their allegations are not exaggerated. To address this

[3] Reading AI reports suggests that there is a relationship between AI access to a country and the level of violence, and a personal interview with an AI official confirmed that access to information is generally an inverse function of violence in the country.

[4] Unlike most of the ''naming and shaming'' literature (Ron, Ramos, and Rodgers 2005; Franklin 2008; Hafner-Burton 2008), which focuses upon background reports and press releases, we focus on the Annual Reports issued by AI. We do so because the background reports and press releases are issued on an as-needed basis and are less likely than Annual Reports to be influenced by contradictory incentives. AI has staked out a reputation for issuing a brief report about all countries in its Annual Report, and thus has greater pressure to take a position in that publication than in the background reports or press releases that AI issues irregularly. Indeed, AI does not necessarily issue either a background report or a press release for all countries about which it writes a brief in its Annual Report. As such, the pressure to take a position despite lacking information that meets the quality standard it has set will be strongest when drafting Annual Reports.

[5] Cingranelli and Richards (2010). To code their Physical Integrity Rights variables, the CIRI project codes the annual reports issued by both AI and the US Department of State. When the two sources disagree, AI's allegations trump. As such, for Physical Integrity variables like torture, the CIRI data are effectively a coding of AI's allegations.

[6] Similar, but less stark, patterns exist for CIRI's variables that measure states' respect for other rights to the physical integrity of the person. We explain below why we focus on torture, but observe here that in ancillary materials to be made available online, we report analyses conducted using these other variables.

econometric challenge, we estimate a zero-inflated ordered probit model with correlated errors (ZiOPC model)—described below—that to our knowledge has not been employed in political science, nor in the human rights literature. The results from the ZiOPC model provide robust statistical and substantive support for our theoretical claims. The estimates from the two equations in the ZiOPC model also suggest that at least one human rights INGO, Amnesty International, largely adheres to its credibility criterion and rarely succumbs to incentives to exaggerate its allegations.[7]

In section *The Theoretical Framework*, we present a theoretical account that produces three testable hypotheses. Section *The Statistical Model* describes the ZiOPC model, and in section *Data and Measurement*, we discuss the data we employ. We present the results of the statistical analyses in section *Results*, and then provide a conclusion.

The Theoretical Framework

We present a theoretical account that explores how the following three factors affect the human rights reports of ''information politics'' INGOs (e.g., AI): Information flows from the media on the rights abuses (including torture) committed by states; the number of domestic NGOs within states; and domestic selectorate institutions. We first describe our theoretical arguments that explain how each of the three factors mentioned above determines when INGOs are more (or less) likely to exaggerate their allegations of torture committed by states. Our theoretical account not only produces interesting hypotheses, but also illuminates that a statistical analysis of human rights reports issued by INGOs such as AI must account for a two-step process, only one of which is observable: the INGO's report on the states' practices (which one can observe) conditioned by the probability that the INGO exaggerated that report. We then present a statistical model capable of accounting for such a process.

We begin our theoretical framework by focusing on the main goals and strategies of INGOs such as AI and then examine when INGOs have incentives to inflate their allegations of torture carried out by states. First and foremost, the key objective for human rights INGOs is to influence state behavior and pressure states to improve their human rights practices (Keck and Sikkink 1998; Risse and Ropp 1999). To this end, INGOs use a strategy of careful targeting of rights abuses committed by state actors via ''shaming and blaming'' or ''naming and shaming'' (Clark 2000; Schmitz 2002; Ron et al. 2005; Franklin 2008; Hafner-Burton 2008). For example, AI often targets repressive states with allegations of rights abuses committed by these states, which are issued from either its New York or London-based office in the form of background reports, press releases, or Urgent Action statements. These allegations are issued when AI's decision makers believe that the information is sufficiently trustworthy and that AI's grass-roots network and/or the press will respond to the information by pressuring the relevant government(s) (Hopgood 2006; Murdie 2009a; Murdie and Bhasin 2011). Furthermore, AI publishes an Annual Report on the degree to which governments around the world respect the right to freedom from torture (Keck and Sikkink 1998; Guzman 2008:99; Wong 2008; Lake and Wong 2009:139–41). The Annual Report differs significantly from AI's other publications in that it summarizes the respect for rights exhibited by every government. As such, there is greater pressure on the Annual Report in that AI cannot effectively offer ''no comment'' as it can with its irregular allegations (that is, background reports, press releases, and Urgent Action statements). To generate this set of reports, human rights INGOs first gather information about human rights practices of governments across all continents, and to do so, they rely upon contacts they develop with NGOs and reporters working in the area. They then process and evaluate this information before reporting it in their publications (Hopgood 2006; Wong 2008).

Note that processing and evaluating information on human rights practices by states is not a straightforward task, for at least two reasons. First, the quality of information about state-sponsored rights abuses that INGOs obtain from third-party sources may vary significantly which, in turn, could affect their ability to accurately assess such information. Indeed, we examine in more detail below how the quality of information of government-sponsored abuses (e.g., torture) that stems from two third-party sources—domestic human rights NGOs and reporters in the field—affects the strategic behavior of INGOs when evaluating the propensity with which states commit torture. Second, INGOs face *conflicting incentives* when assessing the information they obtain about human rights practices by governments.

On the one hand, they have a strong incentive to maintain their credibility by carefully vetting information committed by governments since their influence largely rests on their credibility. This is suggested by Keck and Sikkink (1998:19), who state that ''to be credible, the information provided by (human rights INGO) networks must be reliable and well documented.'' On the other hand, however, an INGO's incentives to raise labor and donations information might lead it to exaggerate allegations of torture particularly when the available information does not permit a firm conclusion about the state's use of torture. That is, the INGO may make a worst-case assumption about abuses committed by the state and thus exaggerate the extent of torture conducted by the government. In fact, Keck and Sikkink (1998:19) emphasize that ''both credibility and drama seem to be essential components of a strategy'' for INGOs that seek to induce ''policymakers to change their minds.''

Thus, given that human rights INGOs face ''conflicting incentives'' along the lines mentioned above, it is not surprising that INGOs, as suggested by extant research, may exaggerate rights abuses, such as torture, committed by governments in some cases, but not in other cases.[8] This variation therefore leads to the question raised above: When are INGOs more (or less) likely to exaggerate their allegations? To answer this question, we first propose that the likelihood with which an INGO exaggerates allegations of torture committed by states critically depends on the pressure to report on a country, specifically the volume of media coverage about violent conflict and its attendant rights abuses (including torture) committed by governments. In particular, when violence between governments and dissidents rises, it not only increases the

[7] Recent advances in split–population modeling make it possible to statistically model unobservable phenomena and assess how this may potentially influence observable outcomes (Harris and Zhao 2007; Svolik 2008; Xiang 2010).

[8] See, for example, Hafner-Burton and Ron (2009).

potential for governments to commit rights abuses but also attracts extensive media coverage. Indeed, the news maxim that ''if it bleeds, it leads'' suggests that the international media, on average, have incentives (for marketing reasons) to cover extensively violent conflicts and publicize rights abuses committed by governments around the world. Media coverage generates public attention and interest, and that in turn influences donors of both capital and labor: People are galvanized by media coverage of such events.[9] This in turn gives an INGO a strong incentive to report on such cases (Keck and Sikkink 1998:19). The opportunity to recruit labor and donations gives INGOs an incentive to comment, even though the quality of the information they are able to obtain may be below their ''credibility standard.'' The INGO is thus more likely to advance allegations, though it is unable to fully vet them.

Yet there is more to this story: The violence that stimulates the situation described above also degrades the quality of information that is available from both NGOs and field reporters. The extreme case can be seen in 1994 during the double genocide in Rwanda, when foreign NGOs and media literally fled the country and domestic NGOs were unable to maintain contact with INGOs. More generally, the quality of information is inversely related to the level of violence as the more abusive the state becomes the more likely it is to target those in civil society and the media who would report on its activities. Informal conversations the author(s) has had with AI staff establish that this conjecture fits their experience.

The outcome, then, is that the pressure to report for organizational reasons is greatest when the information is poorest. As noted briefly above, we expect that this conflict between organizational and reputational incentives will be greatest in Annual Reports as AI has established that it comments on all countries in that document, thus making it effectively impossible to remain silent even if the quality of information is insufficient to issue reports with confidence. What will AI allege in such a circumstance? We contend that they are much more likely to assume the worst—given greater media attention, which is highly correlated with civil violence, and hence state abuse—than assume the best. Put in technical language, we expect that the chances that AI's allegations are biased toward a claim of extensive abuse rise as both the incentive to issue an allegation rises and the quality of information declines. Based on the above theoretical argument, we thus posit the following hypothesis:

Hypothesis 1: *The (unobserved) probability that the INGO's allegations of torture committed by governments are inflated increases when stories of rights abuses (including torture) committed by governments provided by the international media increase.*

INGOs such as AI obtain their information about human practices of states from the domestic NGOs and field reporters that operate within these states. Though the argument we develop could be generalized to both sets of actors, we focus on the former both to simplify the presentation and because we are only able to secure data on human rights NGOs active in each country. The exchange of information from domestic NGOs to INGOs results from a symbiotic relationship between these two actors that has two key characteristics. First, as noted by Murdie and Bhasin (2011:3), ''INGOs provide connections, funds'' to domestic human rights NGOs to help them ''organize and pressure the state.'' Further,

> INGOs can also increase global awareness of the plight of these groups, encouraging the international community, including private foundations, the media in Western democracies, churches, intellectuals, intergovernmental organizations, and third-party state governments, to influence the repressive state as well (Murdie and Bhasin 2011:3).

In return for the connections, funds, and publicity, domestic NGOs that operate within countries provide information to INGOs about their government's respect (or lack thereof) for human rights and its (the government's) propensity to avoid (or commit) torture (Stewart 2004; Hafner-Burton and Tsutsui 2005; Murdie 2009a). Domestic NGOs provide such information to obtain funding *and* to persuade INGOs like AI to put pressure on their own government to reduce rights abuses (Keck and Sikkink 1998; Clark 2001; Hopgood 2006; Wong 2008). This particular behavior of domestic NGOs is emphasized by Hafner-Burton and Tsutsui (2005:1385):

> Domestic groups also reach out to external actors to publicize violations in their state. For example, Chilean activists, with the help of Amnesty International and other groups, sought to publicize forced disappearances committed by their government..., and human rights groups in Indonesia exchanged information with...international organizations to campaign for the release of political prisoners.

It is well known that information politics INGOs like AI rely upon local NGOs for information (Clark 2001; Murdie 2009a), and Hafner-Burton and Tsutsui's claim that domestic NGOs *publicize violations* in their state to external actors is important since it implies that domestic NGOs offer information to INGOs by providing *signals* of their government's human rights practices including the extent to which it uses torture to repress citizens. Building on their insight, we suggest that the higher the number of domestic NGOs in the state, the greater their information-gathering capacity to provide signals of rights abuses and torture committed by their government. Greater information-gathering capacity enhances the domestic NGOs' ability to be more precise and careful with respect to their analysis of the government's human rights practices. This in turn increases the credibility of the signals that they provide to the INGO.

This leads us to expect that an INGO is likely to recognize that the higher the number of domestic NGOs operating in the state, the greater their ability to effectively monitor the government and discourage it from using repressive policies such as torture. Because effective monitoring discourages the government from resorting to torture, the INGO is therefore more likely to logically conjecture that the government cannot resort to using torture against its citizens to protect its interests as the number of domestic NGOs increases. This conjecture decreases the likelihood that the INGO will inflate its allegations of torture against the government. Thus, we have the following hypothesis:

Hypothesis 2: *The (unobserved) probability that the INGO's allegations of torture committed by governments are inflated increases as the number of NGOs decreases.*

[9] For evidence of this relationship, see Simon (2006) and Brown and Minty (2008).

Apart from media stories and domestic NGOs, we further suggest that INGOs like AI also account for the domestic political institutions of countries when evaluating the human rights practices of governments in these countries. This suggestion is partly drawn from existing research that claims that certain types of domestic institutions can act as an effective constraint on human rights violations. The first large-N studies of governments' respect for human rights focused on democracy writ large (Poe and Neal Tate 1994), but more recent efforts have sought to examine specific institutions of both minimalist and liberal definitions of democracy (Bueno de Mesquita, Downs, Smith and Cherif 2005; Davenport 2007), and we work within the latter vein. For instance, some studies find that greater executive constraints induce governments to value respect for human rights, including the right to freedom from torture, more highly (Davenport and Armstrong 2004; Hafner-Burton and Ron 2009). Other studies, however, focus on how the level of freedom of expression, judicial independence, or electoral institutions and competition affect the degree to which governments respect human rights, which includes the right to freedom from torture (Powell and Staton 2009; Cingranelli and Filippov 2010; Conrad and Moore 2010). Our theoretical focus is the extent to which the size of the winning coalition, relative to the selectorate, an executive requires to retain power influences the extent to which she seeks to limit the practice of torture among those who interrogate and incarcerate on her behalf.[10] As discussed in the following paragraphs, we focus on the size of the winning coalition relative to the selectorate (that is, the W/S ratio) since we are interested in understanding how the extent of the marginal gains from patronage obtained by members from the winning coalition—which is determined by the W/S ratio—affects their *and* the executive's political incentive to use torture as a tool of repression. In fact, we suggest below that focusing on the W/S ratio, rather than broader factors such as accountability or the intensity of electoral competition (influenced by the level of democracy), provides a more nuanced understanding of when INGOs are less (or more) likely to inflate allegations of torture committed by governments. We advance two arguments to produce the implication that the probability that an INGO may inflate allegations of torture decreases when the W/S ratio increases.

First, we contend that as the size of the winning coalition—and therefore the W/S ratio—increases in states whose human rights violations are being assessed by the INGO, citizens of these states, including the members of the winning coalition, can more credibly threaten to remove their leaders from office. To survive in office, leaders therefore have an incentive not only to avoid using torture as a tool of repression, but also to increase the level of physical integrity rights enjoyed by their citizens.

Second, observe that leaders in each state have the potential to use repressive tactics such as torture to suppress challengers to their rule. But leaders cannot enact repressive policies alone: The size of the W/S ratio will shape the constraints and their ability to employ torture as a tool of repression. Specifically, in a system where the W/S ratio is low, leaders are relatively less constrained in their ability to adopt repressive tactics such as torture. This is because they can successfully rally their support base when faced with a threat, and owing to strong patron-client bonds, supporters are more likely to defend the ruler in order to maintain their benefits (Bueno de Mesquita et al. 2003). However, when the W/S ratio increases, the benefits of patronage will fall (Lake and Baum 2001; Bueno de Mesquita et al. 2003) and followers are less willing to pay the costs of engaging in repressive acts like torture. This will consequently place limitations on the extent to which leaders can resort to torture. Put differently, when the W/S ratio increases, leaders have less opportunity to engage in torture as they are subject to more constraints than their counterparts in systems with a low W/S ratio.

Since the W/S ratio is common knowledge and is thus well known to the INGO and the government of each state, the INGO will recognize that a government's political interest in respecting the rights of its citizens will at least partly be influenced by the W/S ratio. As a result, the INGO recognizes that when the W/S ratio increases, the government has greater incentives *not* to resort to torture as doing so will jeopardize its political survival. Hence, when the W/S ratio increases, the INGO is more likely to believe that the government does not habitually violate human rights via the use of torture given its incentive structure as elucidated above. This belief deters the INGO from inflating its allegations and thus decreases the probability that its accusations of torture are exaggerated. We state this expectation as a hypothesis:

Hypothesis 3: *The (unobserved) probability that the INGO's allegations of torture committed by governments are inflated decreases when the W/S ratio increases.*

The hypotheses posited above predict when INGOs are more or less likely to exaggerate and make a worst-case assumption about the torture practices of governments. In the following section, we describe the statistical model that we employ to test the three hypotheses posited in this section.

The Statistical Model

An appropriate test of the hypotheses posited above must address an important econometric challenge. As suggested by our hypotheses, a strategic INGO like AI may inflate its allegations of torture committed by a state particularly when increasing dissident–government violence in the country leads to more media coverage of this violence. To the extent that the INGO issues exaggerated allegations of widespread abuse under certain conditions, the data sets that scholars use to assess governments' respect for physical integrity rights will exhibit a biased inflation of the number of widespread abuses. Hence, when estimating how some variables may influence the government's respect for its citizens' physical integrity (as reported by the INGO), if we do not *statistically* account for the probability with which the INGO exaggerates its allegations, then the inferences that we draw about the effect of covariates on each government's use of torture will exhibit a specific bias due to the presence of more zeros in the data than should be there. To avoid bias, we require a two-stage model that allows us to test jointly our hypotheses about the circumstances that influence the likelihood that the INGO exaggerates its allegations (the first, or inflate, stage) *and* assess the effect of covariates that, according to extant studies, influence the degree of

[10] The W/S ratio concept has been introduced by Bueno de Mesquita, Smith, Siverson, and Marrow (2003).

torture committed by governments as reported by the INGO (the second, or outcome, stage).

It may be useful to examine this issue within the context of our observable dependent variable. As mentioned above, the CIRI project uses AI's annual country reports (in conjunction with US State Department's annual reports)[11] to develop a measure of the extent to which governments respect their citizens' right to freedom from torture per country-year.[12] This measure has three ordered categories (denoted as y_i where i = country and where the subscript t = time is dropped for notational convenience): A country-year with a value of y_i=0 was alleged to have engaged in widespread use of torture; those with a value of y_i=1 were alleged to have engaged in some use of torture; and a value of y_i=2 is assigned when the report does not allege that the state tortured any detainees. Our theory suggests that there are in effect two processes generating a country's observed value of respect for physical integrity rights in these ordinal data on state repression: one that determines whether the INGO issuing the report can reliably infer the ''true'' level of repression, and another that determines the level of repression reported by the INGO given that the INGO is able to infer reliably the ''true'' level. Put a bit more technically, the ordered measure (from high to low torture) described above, which constitutes the dependent variable for testing our hypotheses, has a critical characteristic that contribute to the number of 0s relative to 1s and 2s in Figure 1—specifically, the distribution of this variable might be ''zero-inflated'' such that it is skewed toward the y_i=0 value (high torture) relative to both y_i=1 (some torture) and y_i=2 (no torture).

Researchers who use these CIRI data implicitly assume that INGOs can infer or readily observe the actual level of repression in a country and subsequently report it with random measurement error. If that is a reasonable assumption, then one can safely ignore modeling the zero-inflation process. If that is not a safe assumption, and our theory suggests that it may not be, then one's inferences with respect to the factors affecting the second process will be biased. Equally important is the point that if inaccurate reporting is present and one ignores it, then it is impossible to draw inferences about the extent of these inaccuracies and the process producing such reporting. Given that our theory leads us to believe such a process may be at work in the data, a model that accounts for this seems desirable. Stated with greater precision, the zero inflation in the ordered measure of y_i is non-random because the strategic behavior of AI, according to the theory presented above, determines the likelihood with which it chooses to advance allegations of widespread abuse about which it is uncertain, thereby inflating some of its accusations of torture committed by a government. Thus, when estimating the impact of covariates that affect the ordered measure of the respect for freedom from torture in the outcome (second) stage, it is important to account explicitly for the zero-inflation process to avoid bias.

To address this econometric challenge, we estimate the zero-inflated ordered probit model with correlated errors (hereafter ZiOPC model) developed by Harris and Zhao (2007). This statistical model allows us to test explicitly the effect of variables that influence the probability with which INGOs like AI may exaggerate allegations of torture committed by governments and thus claim that abuse is widespread. The model does this by way of an ''inflation'' (binary) probit equation in the first stage that determines the impact of included covariates on the probability that an observation receives its ''true'' value—that is, the probability that the INGO does not make a worst-case assumption about government behavior —rather than always receiving a value of zero (which results when INGOs exaggerate accusations of rights abuses).[13] More formally, the probit equation in the first stage of the ZiOPC model is expressed as

$$a_i^* = \mathbf{z}_i'\boldsymbol{\gamma} + u_i \tag{1}$$

$$a_i = \left\{ \begin{array}{c} 1 \text{ if } a_i^* > 0 \\ 0 \text{ otherwise} \end{array} \right\}$$

where $\mathbf{z}_i'$ is the vector of covariates and u_i is a normally distributed error term. Though we cannot directly observe the probability that AI inflates its allegations of torture, we can use the first stage of the ZiOPC model to *estimate* it from observed data. Specifically, given the assumption of normality, the probability that INGOs inflate allegations of torture and say that abuse is widespread is $\Pr(a_i = 0) = 1 - \Phi(\mathbf{z}_i'\boldsymbol{\gamma})$ (this accounts for zero inflation in the data generating process), and the probability that it makes accurate allegations is $\Pr(a_i = 1) = \Phi(\mathbf{z}_i'\boldsymbol{\gamma})$ where $\Phi(.)$ is the standard normal c.d.f.

The outcome equation of the ZiOPC model allows us to test the effect of variables on the observed variation in the extent to which governments respect the right to freedom from torture. It does so via a standard ordered probit equation defined as

$$y_i^* = \mathbf{x}_i'\boldsymbol{\beta} + \varepsilon_i$$

$$y_i = \left\{ \begin{array}{c} 0 \text{ if } y_i^* \le 0 \\ 1 \text{ if } \mu_{j-1} < y_i^* \le \mu_j (j = 1) \\ 2 \text{ if } \mu_{J-1} \le y_i^* \end{array} \right\} \tag{2}$$

where $\mathbf{x}_i'$ is a vector of covariates, ε_i is a normally distributed error term, and μ_j is the boundary or cut-off parameter.[14] The error terms from the first-stage probit equation and the second-stage ordered probit outcome equation, that is u_i and ε_i, may be correlated since they correspond to the same unit of analysis. Hence, more formally, we assume that the error terms u_i and ε_i from the first- and second-stage equations of the ZiOPC model are correlated and follow a bivariate normal distribution with correlation coefficient ρ. If (u_i, ε_i) follows a bivariate normal distribution with correlation coefficient ρ and with the identifying assumption of unit variances, then

[11] As noted earlier, when there is a conflict between the AI and US State Department reports about a country's respect for a given physical integrity, the CIRI project assigns the value coded using the AI report.

[12] We limit the results reported in the study to the CIRI torture variable, but include in ancillary material that will be made available online and in the replication data set analyses that use other CIRI variables that measure states' respect for physical integrity rights.

[13] Note that when the model predicts that AI likely ''inflated'' an allegation (that is, alleged frequent torture when the amount of information available to them was below their credibility threshold), that does not necessarily mean that the government did *not* use torture frequently. In fact, it seems almost certain that AI's ''inflated'' guesses will sometimes be correct, and other times be wrong. That is, when we observe a value of 0 in the CIRI torture data, some of the ''inflated'' 0s will be cases in which the government did torture frequently, but AI lacked adequate information to know that with confidence, yet other cases where the government did not, and AI also lacked adequate information, but guessed wrong and reported widespread abuse.

[14] We assume without loss of generality that μ_0=0.

according to Harris and Zhao (2007:1077), the augmented ordered probit (outcome) equation of the ZiOPC model is given by

$$\Pr(y_i) = \left\{ \begin{array}{c} \Pr(y_i = 0|\mathbf{x}_i, \mathbf{z}_i) = [1 - \Phi(\mathbf{z}_i'\boldsymbol{\gamma})] \\ +\Phi_2(\mathbf{z}_i'\boldsymbol{\gamma}, -\mathbf{x}_i'\boldsymbol{\beta}; -\rho] \\ \Pr(y_i = 1|\mathbf{x}_i, \mathbf{z}_i) = \Phi_2(\mathbf{z}_i'\boldsymbol{\gamma}, \mu_j - \mathbf{x}_i'\boldsymbol{\beta}; -\rho) \\ -\Phi_2(\mathbf{z}_i'\boldsymbol{\gamma}, \mu_{j-1} - \mathbf{x}_i'\boldsymbol{\beta}; \rho] \\ \Pr(y_i = 2|\mathbf{x}_i, \mathbf{z}_i) = \Phi_2(\mathbf{z}_i'\boldsymbol{\gamma}, \mathbf{x}_i'\boldsymbol{\beta} - \mu_{J-1}; -\rho), \end{array} \right\} \quad (3)$$

where $\Phi_2(.)$ denotes the c.d.f. of the standardized bivariate normal distribution with correlation coefficient ρ.

Note that the probability of a zero observation in the augmented ordered probit outcome equation of the ZiOPC model accounts for ''inflation'' as it combines the probability that a country-year observation is assigned a value of zero (frequent use of torture) plus the probability that AI may inflate its allegations. Additionally, observe that the ordered categories in the expanded ordered probit outcome equation in (3) incorporate the probability that AI may not inflate its allegations in the first stage.[15] The fact that the ZiOPC model contains an inflation (first-stage) probit equation and an augmented ordered probit outcome equation is important because of several reasons. First, the presence of an inflation equation and an ordered probit outcome equation in the ZiOPC model permits one to draw inferences about (i) the covariates that influence the likelihood that AI's allegations of widespread abuse are inflated, (ii) the frequency with which AI inflates its allegations of widespread abuse, and (iii) correct for bias introduced by that process when estimating the impact of other covariates (suggested by extant studies) that influence the extent to which governments respect peoples' right to freedom from torture (as reported by AI).[16] Indeed, it is not possible to test hypotheses about the ''inflation'' process in the CIRI data without employing the ZiOPC model described above. Furthermore, as shown below, the estimates from the two equations in the ZiOPC model allow us to assess empirically the degree to which INGOs like AI adhere to high standards of credibility when reporting rights abuses rather than opting to exaggerate its allegations. We discuss the operationalization of the variables used for the tests and the results from the ZiOPC model in the following sections.

Data and Measurement

Our sample consists of 116 countries from the years 1990 to 1997, the period for which we have data on the full set of our variables.[17] The unit of observation is the country-year. The dependent variable used in the analysis is the CIRI measure of respect for the right to freedom from torture. As described above, a country-year with a value of 0 was alleged to widely used torture; those with a value of 1 were alleged to have tortured some detainees; and a value of two is assigned to a state that had zero torture allegations.

We feature torture for two reasons, but include in ancillary material to be made available online results using CIRI's other physical integrity variables. First, the CIRI data indicate that torture is the most widely violated of the physical integrity rights. Density graphs, like that in Figure 1, for the other physical integrity rights variables in CIRI exhibit smaller proportions of zeros. Second, torture is a signature AI issue and has been since 1975 (Clark 2001). We expect the pressure to comment in an Annual Report despite the low-quality information to be greatest for signature issues. AI was founded to support prisoners of conscience, and CIRI's political prisoners variable would thus be the primary competitor for a signature issue. Yet the density plots for these variables (only torture is shown here) suggest that torture is the most likely to provide evidence of AI compromising its principles due to organizational incentives.[18] The measures for the independent variables included in the model are discussed below.

The Inflation (γ) Equation

Our theoretical account suggests that the inflation stage of the ZiOPC should include variables that contribute to or detract from the information environment in which human rights NGOs operate, as well as a measure of political institutions. More specifically, we identify three concepts that we need to measure and include in our inflation (γ) equation: the amount of press coverage that discusses the country's observance of human rights; the size of the domestic NGO population; and the W/S ratio. As an indicator of the attention being paid to rights abuses by the media, we use data from Ron et al. (2005) that record the number of articles containing the term ''human rights'' in conjunction with a particular country published in *Newsweek* and the *Economist* (international edition) during the year in question. We use the average number of articles from these two publications.[19] As a measure of the number of domestic NGOs, we turn to the count of human rights NGOs collected by Murdie (2009b) and available in the replication data set to Murdie and Bhasin (2011), which covers the years 1990–1999. She produced a count of the number of INGOs with an office in a country that focus on human rights issue. To collect the data, she used the *Yearbook of International Organizations*. Though this measures the domestic presence of INGOs rather than NGOs that are strictly domestic in origin, it is the best available measure of human rights NGO presence that we are aware of.[20] We include

[15] Let $\widehat{\boldsymbol{\theta}} = (\boldsymbol{\gamma}', \boldsymbol{\beta}', \boldsymbol{\mu}', \rho)'$ for the full ZiOPC model. The log likelihood function of the full ZiOPC model where $\widehat{\boldsymbol{\theta}} = (\boldsymbol{\gamma}', \boldsymbol{\beta}', \boldsymbol{\mu}', \rho)'$ is given by $\ell(\boldsymbol{\theta}) = \sum_{i=1}^{N} \sum_{j=0}^{J=2} d_{ij} \ln[\Pr(y_i = j|\mathbf{x}_i, \mathbf{z}_i, \widehat{\boldsymbol{\theta}})]$ where $i \in \{1,2,...N\}$ is the number of observations, and where d_{ij}=1 if outcome j is realized in i, or is d_{ij}=0. The ZiOPC model is estimated by a Stata module written by the first author of this paper and described in Hill, Bagozzi, Moore, and Mukherjee (2011).

[16] Harris and Zhao (2007) perform Monte Carlo experiments to demonstrate that the ZiOP model outperforms (that is, its parameter estimates are less biased than) a standard ordered probit model in finite samples when the data are generated according to a ZiOP process (that is, when there are ''extra'' zeros in the data). Monte Carlo experiments conducted by the authors of this paper for a separate project also show that the bias of the ordered probit becomes larger as the proportion of false zeros in the sample becomes larger.

[17] Coverage for our measure of domestic human rights NGOs begins in 1990 and coverage for our measure of dissident violence ends in 1997.

[18] That said, as shown in the ancillary material to be made available online, the results when we use political prisoners in lieu of torture are only slightly different.

[19] Ron et al. (2005) examine the extent to which AI's background reports and press releases, which they use as a measure of naming and shaming activity, are influenced by factors such as press coverage as opposed to rights violations. As such, their study is similar to, but nonetheless distinct from, our own.

[20] In previous iterations of this study, we adopted the proxy measure used by, among others, Hafner-Burton and Tsutsui (2005), and collected by Wiik (2002). Wiik counts the number of INGOs registered with the Union of International Associations or listed in the *Yearbook of International Organizations*, Vol. 2, and so includes an INGO if citizens of that country claim membership in the INGO. The data collected by Murdie, in contrast, only include INGOs that have established an office in that country. In the ancillary materials that will be made available online, we describe the results obtained using that measure.

the natural log of the number of human rights NGOs in a country to reflect the fact that the impact of NGO presence on the availability of information about government abuse likely declines with each additional NGO.[21] Our measure of the W/S ratio comes from Bueno de Mesquita et al.'s (2003) measure of the size of the winning coalition relative to the selectorate. Larger values indicate increases in the proportion of the citizenry whose support is needed for the government to remain in power.[22]

We also include two important control variables in the inflation (γ) equation. First, as we argue above, one would expect that widespread violent conflict between the government and dissident groups will have an independent impact upon the quality of the information environment, thus inducing the credibility v revenue dilemma facing ''information politics'' INGOs. This variable is important to control for. To measure violent political conflict we employ data from the Global Terrorism Database.[23] From these data, we obtain a count of the number of terror attacks committed inside a country's borders within a given year. These data include attacks by both domestic and international dissident groups (with the overwhelming majority being carried out by domestic groups). Note that because this measure excludes acts of violence committed by the state, we avoid any endogeneity issues that might arise had we used other measures of violence.

Second, there is a potential endogeneity issue to address. When drafting its Annual Reports, AI makes use of the information it collected while writing the background reports, press releases, and Urgent Action notices that it produced during the preceding 12 months. Further, press releases are quite obviously intended to influence press coverage of human rights abuses and thus are extremely unlikely to be independent of the number of news reports published in *The Economist* and *Newsweek*. As a consequence, though our dependent variable is a trichotomous ordinal measure of the level of respect and the number of news reports is an event count, theory about the data generating process leads us to expect that the number of background reports and press releases will be strongly correlated with both the dependent variable and a key independent variable. To determine the impact of media reports, then, we must control for the number of AI reports (that is, background reports and press releases).[24] We lag the AI reports variable—the sum of background reports and press releases—by 1 year to ensure that it is exogenous.

The Outcome (β) Equation

The argument developed above provides the primary input to our specification of the inflation (γ) equation of the ZiOPC model, and we turn to the existing literature on state violence and repression to inform our specification of the outcome (β) equation. A considerable literature exists that uses either the CIRI measures of states' (lack of) respect for physical integrity rights or the older Political Terror Scale (PTS) (Cingranelli and Richards 2010; Wood and Gibney 2010). That literature identifies several covariates that are correlated with the CIRI and PTS measures, and we therefore include them in our outcome (β) equation. More specifically, scholars examining the determinants of state repression have generally focused on structural factors (that is, socioeconomic conditions) and political institutions (Henderson 1991; Poe and Neal Tate 1994; Poe, Neal Tate, and Keith 1999). Other studies have focused on the behavior of political dissidents as a cause of state violence (Davenport 1995; Moore 2000; Pierskalla 2010). Based on these studies, we include the following in the outcome (β) equation: the W/S ratio, size of macroeconomic output, population size, and acts of (violent) dissent.

As noted above, a number of studies have focused on the ability of domestic political institutions to curb repressive behavior, which leads us to include a measure of domestic institutions in the outcome (β) equation. Though early studies in this vein examined democratic institutions broadly defined, we follow more recent studies (Bueno de Mesquita et al. 2005; Davenport 2007; Cingranelli and Filippov 2010; Hafner-Burton and Ron 2009; Powell and Staton 2009; Conrad and Moore 2010) and focus on a more specific institution, the W/S ratio. We include in the outcome (β) equation the W/S ratio measure that appears in the inflation (γ) equation.

Relatively low levels of economic output are generally thought to lead to greater use of repression on the part of the government, since conditions of poverty and scarcity create grievances that the government is not adequately prepared to assuage through non-violent means. Governments in this situation are expected to be especially sensitive to challenges to their sovereignty, and thus, repressive acts become more likely. We include a measure of GDP per capita (logged) to account for this relationship. Another structural condition affecting the state's decision to engage in repression is population size. A large population is thought to create the same kinds of societal pressures as low levels of economic output, and also has the effect of increasing the number of opportunities a state has to repress its citizenry. We include a measure of population size (logged) accordingly. Both of these measures are drawn from the World Bank's World Development Indicators database (World Bank 2009).

A less frequently examined determinant of government-sponsored violence is the behavior of political dissidents. It has long been conjectured that governments generally respond to perceived threats to their sovereignty with force, and studies examining this proposition have generally found support for it (Davenport 2007). Though many studies of state repression ignore the behavior of dissidents, we argue that one cannot understand the government's decision to use force against its citizenry without examining the behavior of those who are opposed to its policies, particularly if that opposition manifests itself as a violent, direct challenge to the authority of the state. Consequently, we include the same measure of dissident violence featured in the inflation equation (that is, the GTD's event count of terror attacks committed by domestic and foreign dissident groups inside the country's borders during that year).

An emerging body of research has drawn attention to the effect of human rights treaties, such as the Convention Against Torture (CAT), on state violence. This

[21] In the online appendix, we report the results obtained using the raw count of human rights NGOs, and the results are nearly identical to those reported here.

[22] To avoid division by zero, Bueno de Mesquita et al implemented the following calculation: =W/(log((S + 1)*10)/3).

[23] Global Terrorism Database, http://www.start.umd.edu/gtd/.

[24] Ideally, we would also control for the number of Urgent Action announcements as well. It turns out that Ron et al. (2005) have counted the number of background reports and press releases that discuss each country, but we are unaware of a similar measure of Urgent Action calls. That said, background reports and press releases are produced to influence the media and other elite opinion makers, whereas the Urgent Action calls target AI's grass-roots membership. As such, if one of the three were to be left out, Urgent Action announcements would be the one to leave aside.

research produces the initially surprising finding that ratification is associated with *worse* state practices (Hathaway 2002; Hafner-Burton and Tsutsui 2005,2007; Hathaway 2007; Vreeland 2008; Powell and Staton 2009). This result is best explained as a self-selection process: Some states that have no intention of honoring the CAT hope to camouflage themselves among those who do, while other states that will generally respect the rights of their citizens will decline to sign the CAT.[25] Due to this research, we include in our model of the outcome (β) equation a variable that is coded 1 for states that have ratified the CAT and 0 for those that have not. We coded the data using the United Nations Treaty Collection Web site as a source.[26] Finally, we also include in the outcome (β) equation the natural log of the human rights NGO count collected by Murdie (2009b), which also appears in the inflation (γ) equation, as several recent studies have shown NGO/INGO activity to be related to state repression (Hafner-Burton and Tsutsui 2005; Landman 2005; Neumayer 2005).

Results

We report results from estimating both an ordered probit regression and a ZiOPC regression, both with standard errors clustered by country.[27] Yet before discussing the results, we provide an explanation of how parameters from the ZiOPC should be interpreted. First, the inflation (γ) equation estimates the probability that the INGO *does not* inflate its allegations of widespread torture. Because of this, the estimated parameters have the opposite sign one would expect given that our hypotheses are stated as the expected relationship between the variable and the probability that the INGO does inflate those allegations. It is thus important to keep in mind that the signs of the coefficients are opposite of what one would expect given the language in our hypotheses. More specifically, a positive, significant coefficient indicates that the variable has a positive effect on the probability that we observe the ''true'' value of repression rather than always observing zero regardless of the ''true'' value. Negative coefficients therefore indicate that as that variable increases, the probability becomes larger that a country will be assigned a value of zero regardless of the ''true'' value of repression.

Second, the parameters from the outcome (β) equation are interpreted in the same way as a standard ordered probit, with one caveat. The parameters indicate the effect of the variable in question on the probability that an observation will fall into the first or last category of the dependent variable given that the value of the dependent variable for that observation will not always be reported as zero.[28]

Inflate (γ) Results

The estimates from the γ equation have considerable substantive interest. Hypotheses 1–3 concern the factors that influence the probability that AI will *not* allege a level of violations that would lead CIRI project coders to assign an inaccurate score of ''widespread'' torture violations, and so we first turn our attention to the results from the inflate (γ) equation from the ZiOPC model. Note that by definition the ordered probit regression is mute with respect to these hypotheses: No information is available. Yet, with the estimates from γ in hand, we can produce a variety of quantities of interest. First, as is common, we can produce estimates of the substantive effects of a change in a given variable upon the probability that AI's Annual Report alleged widespread use of torture despite the fact that AI was not confident about the allegation. Second, we can create an interesting table of predicted probabilities that reports the estimated chance that allegations of No Respect for the rights enshrined in the CAT in AI's Annual Reports were based on a worst-case assumption about government treatment of citizens, and the chance that they were not. That this information is *unobservable but estimable* makes the ZiOPC model a remarkable and interesting statistical model.

Hypotheses 1–3 generate expectations between three variables and the probability that AI alleged a state broadly violated the right to freedom from torture despite being uncertain that this was the case: That the number of media stories about human rights has a positive impact, and that both the number of human rights NGOs and the ratio of the size of the winning coalition to the size of the selectorate (W/S) have a negative effect. The results reported in the ZiOPC column of Table 1 are consistent with all three expectations.[29]

Note that hypotheses 1–3 are supported by the estimates in Table 1. All three variables produce a statistically significant coefficient estimate with the expected sign (which is the opposite of that stated in the hypotheses).[30] The sign of the coefficient estimate for the number of media reports on human rights is negative and statistically significant. The negative sign indicates that the larger the number of such news reports, the lower the probability that AI does *not* allege that the state uses torture widely when, in fact, AI is unaware of the actual level of torture. Put plainly, as hypothesized, there is a positive association between the number of media reports on a country's rights record and the likelihood that AI's Annual Report includes an allegation about which it is uncertain, and thus at risk to exaggeration. Further, the control that we included to address the likely endogeneity among the level of respect reported in AI's Annual Report, the number of media stories, and the number of background reports and press releases—AI Report$_{t-1}$—is also statistically significant with the expected sign.

The second variable, the number of human rights INGOs with a local office, has a positive sign that is also statistically significant, providing support for the hypothesis that the greater the number of NGOs within a country, the more likely that AI would *not*, in the absence of

[25] For a description of a self-selection process of the first kind, see Simmons (2009:88–90). For descriptions of the second kind of process, see Hathaway (2005, 2007) and Powell and Staton (2009).

[26] Seehttp://treaties.un.org/Pages/ViewDetails.aspx?src=TREATYmtdsg_no=IV-9chapter=4lang=en, accessed June 2009.

[27] As mentioned above, the ZiOPC model that we estimate includes a parameter (ρ) that accounts for the possibility of a correlation between the error terms in the inflation and outcome equations.

[28] Though the effect of a variable on the probability of falling into categories between the first and last cannot be determined by looking at the coefficient alone, one can meaningfully interpret parameters from an ordered response model. For example, positive, statistically significant values may be interpreted as meaning that a one unit increase of the variable in question raises the probability that an observation will fall into the *last* category of the dependent variable and lowers the probability that it will fall into the *first* category of the dependent variable. Negative, significant values have the opposite interpretation.

[29] Recall that the signs in the table are the opposite of what one would expect given our hypotheses.

[30] The results when we use political prisoners in lieu of torture are only slightly different. We receive support for hypotheses 1 and 3 from the political prisoners model, but not hypothesis 2. That is, media reports and the W/S ratio still have a significant effect in the inflation equation, but human rights NGOs do not.

TABLE 1. Ordered Probit & ZiOP Estimates, 1990–1997

Dependent Variable: CIRI Torture Scale

Regressor	*O. Probit*	*ZiOPC*
Inflate (γ) Eq		
Media Reports	—	−0.327 (0.152)**
AI Reports$_{t-1}$	—	−0.044 (0.016)***
HR NGOs (logged)	—	0.452 (0.182)**
W/S	—	0.745 (0.405)*
Terror Attacks	—	−0.004 (0.002)*
Constant	—	0.679 (0.411)*
Outcome (β) Eq.		
HR NGOs (logged)	0.198 (0.090)***	0.051 (0.101)
W/S	1.186 (0.330)***	0.871 (0.451)*
CAT Ratification	−0.146 (0.134)	−0.258 (0.162)
Terror Attacks	−0.004 (0.002)**	−0.001 (0.003)
GDP per capita (logged)	0.238 (0.069)***	0.334 (0.089)***
Population (logged)	−0.291 (0.050)***	−0.274 (0.060)***
Constant	2.366 (0.896)***	2.111 (1.009)**
τ_2	1.436 (0.097)***	1.664 (0.148)***
ρ	—	−0.271 (0.186)
AIC	1408.969	1393.167
Observations	856	856

(*Notes.* *$p < .10$; **$p < .05$; ***$p < .01$ (two-tailed)
Standard errors are given in parentheses.)

information, label the government of that country as using widespread torture. That is, the probability that AI reports widespread abuse when it is uncertain about the actual level is negatively related to the number of NGOs in the country.

Observe third that the coefficient on the winning coalition to selectorate ratio is positive and significant, indicating that the greater the proportion of the population the government relies on to stay in power, the more likely it is that reports of rights abuses in that country are *not* an exaggerated allegation of widespread torture. Stated simply, larger W/S ratios are associated with lower probabilities that AI exaggerates its allegations of widespread torture.

We now turn to an examination of some substantive effects from the inflation (γ) equation. Estimates from the ZiOPC indicate that media reports, the number of NGOs, and the W/S ratio each significantly affect the probability that a human rights INGO will report violations as widespread regardless of their actual frequency. Figures 2–4 illustrate the effects of each of these variables in substantive terms. To generate the graphs below, three of the four covariates in the inflate (γ) equation were held constant, while the other was allowed to vary across its observed range. Where covariates are held constant, they take on the following values: The number of terror attacks is set to 3, the natural log of human rights NGOs is .76 (roughly equivalent to 2 human rights NGOs), human rights-related stories issued by popular media sources are set to 0, the winning coalition to selectorate ratio is 0.69, and the lagged value of the sum of Amnesty's background reports and press releases is 3.[31]

Figure 2 indicates that the effect of media reports about human rights violations on the probability that AI does *not* exaggerate its allegations of abuse is large and negative. While the probability of a human rights group not labeling a government as a frequent repressor regardless of its behavior is above .9 when no media reports are being issued, this probability falls quite quickly as media reports rise, dropping below .5 once the number of media reports is equal to 5. When media reports are at their in-sample maximum of 25.5, the probability that a government will be labeled a frequent torturer, regardless of how often it actually uses torture (that is, the probability that the country receives a value of zero for CIRI's torture variable even if its actual value is a 1 or a 2), is above .95.

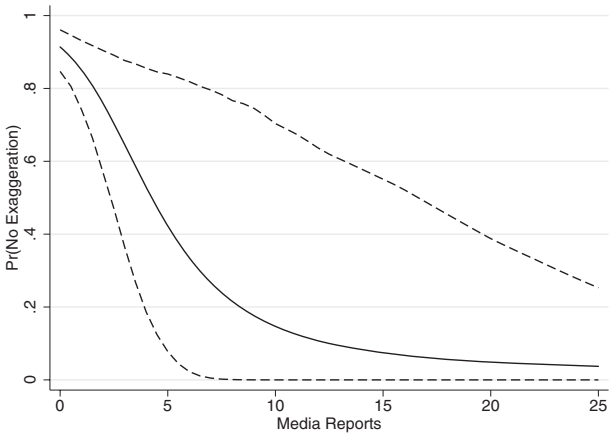

FIG 2. Effect of Media Reports on Probability of Accurate Report

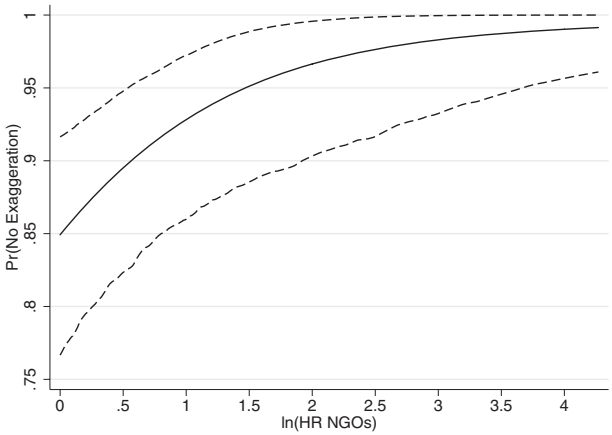

FIG 3. Effect of HRO Presence on Probability of Accurate Report

Figure 3 indicates that increases in the number of NGOs operating in a state are associated with an increase in the probability that a human rights group will *not* always allege frequent use of torture by the government of that state. When NGOs are non-existent, the probability of not always alleging frequent torture is roughly .85; when the natural log of human rights NGOs is at its in-sample maximum (roughly 4.27, or 71 NGOs), the probability that a report is not exaggerated is near certainty (.99). The W/S ratio also positively impacts the probability that INGOs do not always report widespread abuse. Figure 4 shows that in a completely closed autocracy (that is, where W/S = 0), the probability that a human rights group does not make the worst-case assumption of widespread torture is slightly below .8. As W/S increases to its maximum of 1, this probability is roughly .95.[32]

[31] These are the in-sample means of the natural log of human rights NGOs and the WS ratio, and the in-sample medians of terror attacks, human rights stories, and Amnesty reports.

[32] Recall that the baseline scenario is one in which popular media have not released any human rights-related stories, the W/S ratio is relatively high, and there are roughly two human rights NGOs present. This is a relatively favorable situation with respect to the accuracy of reporting, which is why the probability of not always alleging widespread abuse is quite high on the left-hand side of Figures 2–4.

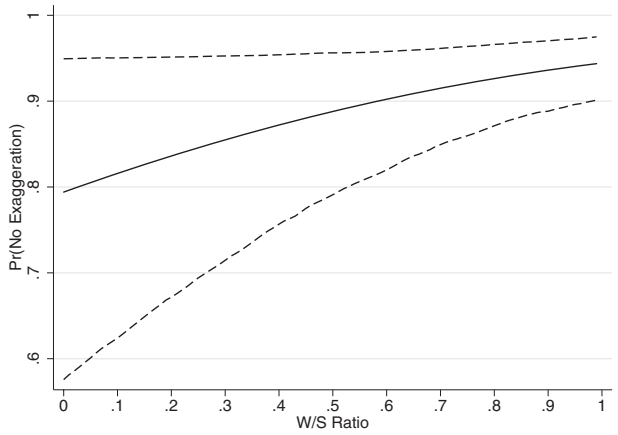

Fig 4. Effect of W/S on Probability of Accurate Report

We turn our attention to the table of predicted probabilities described above. Does AI frequently succumb to the incentive to make allegations of widespread abuse in its Annual Reports, despite the uncertainty about the level of abuse, or does the incentive to maintain its credibility consistently trump the organizational incentives to raise donations and labor? Table 2 provides estimates we can use to address that question.

The rows in Table 2 represent the predicted ''true'' value of a state's respect for the CAT.[33] The columns indicate our estimate of whether the information available to AI for a given country in a given year would lead AI to make an allegation of widespread torture regardless of actual torture practices. Thus, the cells in the table represent the intersection of a state's predicted respect for the right to freedom from torture and the predicted information condition AI faced, given the specification of the ZiOPC regression.[34] Each row of Table 2 shows the percentage of in-sample country-years predicted (based on estimates from the outcome (β) equation) to have a particular level of respect for the CAT that were likely (based on estimates from the inflate (γ) equation) to have been below AI's information threshold and thus labeled as frequent torturers regardless of the government's behavior, and the percentage for each predicted category that were likely above the threshold.[35] Confidence intervals for these predicted percentages are shown in parentheses.

The estimates indicate that during the 1990s, more than 90% of the information about states' respect for the CAT contained in AI's Annual Reports met or exceeded the threshold that AI sets for making allegations of abuse, and in more than 97% of the cases, AI's allegation matched the predicted ''true'' level of respect. That is, even when our model predicted that the information

Table 2. Predicted Percentages of Uncertain Allegations, 1990–1997

Predicted Respect for CAT	*Worst-Case Assumption*	*No Worst-Case Assumption*
None (0)	19.5% (7.6–42.6%)	80.5% (57.4–92.4%)
Some (1)	4.4% (0.2–13.6%)	95.6% (86.4–99.8%)
High (2)	0.8% (0–2.1%)	99.2% (97.9–100%)
% of Sample	8.3% (2.9–19.6%)	91.7% (80.4–97.1%)

(*Note.* 95% confidence intervals are given in parentheses.)

available to AI led to a worst-case assumption about torture, it predicts that AI only ''got it wrong'' in 2.6% of the country-years upon which it reported during the 1990s.[36] Put differently, there is a dramatic difference in the chance that the information in AI's Annual Reports was below the threshold given the predicted ''true'' value. That is, in <1% of the cases for which our model predicts that the ''true'' level of respect for the CAT is high did AI allege that it was none, and that figure rises to 4.4% for predicted ''true'' levels of low, but 19.5% for predicted ''true'' levels of none. Note that if AI does not have the level of information it desires, yet alleges in its Annual Report that a state has no respect for the CAT and the predicted ''true'' value is no respect, then AI did a good job balancing its conflicting incentives. Indeed, aside from finding evidence that an organization never compromises its standards for expediency—which is a standard organizations comprised by human beings are unlikely to meet—the best-case normative scenario for balancing incentives is the pattern we observe in Table 2. It suggests the AI almost never alleges that governments that exhibited high respect for the CAT were broad violators, and only very rarely alleges that governments that had medium respect for the CAT broadly violated it. In other words, the predicted probabilities from our ZiOPC model suggest that AI balances its credibility and organizational incentives strongly in favor of the former, and to the extent that it issues allegations that do not meet its threshold, it overwhelmingly does so against states that are in fact major violators of human rights.

Outcome (β) Results

We turn now to what would have been the sole focus of a standard analysis of the extent to which countries respect the CAT, which is to say an analysis that (implicitly) assumes that AI is an information collector, assembler, and distributor that is able to observe most, if not all, of what it needs to observe to do its job. The existing literature is populated with studies that implicitly make this assumption and utilize either CIRI or PTS data to study countries' (lack of) respect for human rights. Does this study suggest that the estimates reported in those studies suffer from bias due to a failure to model the conflicting incentives that INGOs like AI face?

The estimates in the O. Probit column of Table 1 are very similar to those in the Outcome (β) equation portion of the ZiOPC column. That is, there is little evidence of consequential bias should one use an ordered probit regression instead of the ZiOPC regression. That the bias in coefficients from an ordered probit would be small to non-existent is consistent with

[33] We put the word true in quotes to underscore that it is not possible to know a state's respect for the CAT. As Conrad and Moore (2010) have argued, the existence of a Principal–Agent problem inhibits even the executive's ability to know the true value of his state's respect for the CAT.

[34] As with any regression analysis, the results from a ZiOP regression are based on the assumption that the model is properly specified (that is, contains the correct set of X variables, which are measured well with only random error, and was produced by a data generating process that can be approximated by the ZiOP model). Statistical models were created to permit inference about unknown true values (typically population parameters), and in that sense, this use is very standard.

[35] Based on the ZiOP estimates, the observations in the second and third rows of column one would be observed as zeros, that is, would be governments reported to have used torture frequently.

[36] To calculate the 2.6% figure, we added the frequency of cases in the Below the Threshold column and the Medium and High rows, which are 21 and 1, respectively. 22/856 = 2.6%.

Monte Carlo evidence about the ZiOPC reported in Hill et al. (2011). They report that when the proportion of always zeros (represented by the Below the Threshold column in Table 2) is below 10% of the sample, the ZiOPC is unlikely to produce estimates that are superior to those produced by an ordered probit. Table 2 suggests that an estimated 8% of the sample is always zero, thus leading us to anticipate that the bias induced by organizational incentives is sufficiently small that it can be ignored.

Further, and importantly, the variables in the β equation produce estimates one would expect given the previous findings in the state repression literature. The only differences between estimates from the two models are that terror attacks shift from having a small, negative impact upon states' respect for the CAT to having no effect and that human rights NGOs changes from having a positive effect on respect for freedom from torture to having none.

The latter result is worth discussing since it is in mild tension with previous empirical findings that suggest that NGO/INGO presence is strongly (and negatively) related to government repression (Hafner-Burton and Tsutsui 2005; Landman 2005; Neumayer 2005). Prominent arguments as to why such a relationship would exist predict that strong connections to global civil society have an impact on repressive behavior because they make states vulnerable to bad publicity (Hafner-Burton and Tsutsui 2005), and the studies cited above use the count of registered INGOs collected by Wiik (2002) as a proxy for the extent to which a state is connected to global civil society and thus the extent to which it is susceptible to pressure from the international community. Our argument is that higher numbers of domestic NGOs increase the credibility of signals about government abuse received by INGOs such as AI, and increase the possibility of bad publicity such that INGOs are less likely to make a worst-case assumption about a government's behavior. Thus, our theoretical justification for the inclusion of human rights NGOs is related to, but distinct from, that of previous studies, and the measure we employ (human rights NGOs) is also different from the one used in these studies. Having said that, it bears mentioning that the argument and results presented here suggest a more nuanced mechanism linking NGO/INGO presence to government repression than is usually presented. Our model indicates that high numbers of domestic human rights NGOs increase the probability that reports about government abuse are not always dismal but, given the possibility of a ''good'' report (that is, one that does not indicate widespread abuse), NGOs have a negligible impact on reported levels of government abuse. In conjunction with our argument, this suggests that in addition to generating pressure for compliance with international human rights norms by creating bad publicity, it may be the case that a large domestic NGO presence affects the quality of information collected by groups like AI, and the beliefs such groups have about government abuse, so that relatively favorable reporting becomes more likely.

Conclusion

This study began with the following question: When are INGOs more (or less) likely to exaggerate allegations of torture committed by governments? To address this question, we construct a theoretical account of strategic interaction between an INGO, a government whose responsiveness to its citizens is influenced by the W/S ratio, and domestic NGOs that monitor the state. The argument suggests that INGOs are more likely to inflate their allegations about rights abuses conducted by the government when information flows from the media about rights abuses (including torture) committed by the government increases. It also predicts that INGOs issuing reports about abuses are less likely to exaggerate their reports when the W/S ratio or the number of domestic human rights NGOs is high. To both test our hypotheses and empirically evaluate the strategic behavior of an influential INGO that produces observable reports about rights abuses committed by governments, we estimate a zero-inflated ordered probit (ZiOPC) model with correlated errors. Results from the ZiOPC model provide robust support for our hypotheses, but suggest that AI largely maintains its credibility standard, responding infrequently to organizational incentives to exaggerate allegations.

The findings presented in this study have three main substantive implications for the literature. First, our analytical framework reflects the exhortation by Martin and Simmons (1998:749) to explore the effects of international organizations/institutions conditional upon domestic considerations; in particular, they argue that scholars should ''move toward genuinely interactive theories of domestic politics and international institutions.''[37] Our theoretical story moves a step forward in this direction as it shows how strategic interaction between an important international non-governmental organization (e.g., AI) and a state whose behavior is influenced by domestic institutions as well as domestic NGOs influences the state's respect for the right to freedom from torture. Second, recent studies have suggested that domestic human rights NGOs can successfully cajole governments to improve their human rights practices (Sundstrom 2005; Murdie 2009b). Our analysis provides a more nuanced result as it suggests that the relationship between domestic NGOs and reported levels of rights abuses may also be explained by the influence these groups have on the information available to, and beliefs held by, INGOs who report on government abuse.

Third, this study joins a small body of work that explores the incentives that ''information politics'' INGOs have to exploit strategically information asymmetries to fulfill certain organizational imperatives (Hudock 1999; Murdie 2009b; Hyde 2010). In fact, our theoretical account identifies conditions under which rights INGOs are more (or less) likely to exaggerate allegations of torture committed by governments. Part of what makes this work so challenging and interesting is that some of the implications are unobservable (indeed, the unobservability is the crux of the information asymmetry). In our theory, the probability that an INGO decides to issue an allegation that is a ''best guess'' is the unobservable quantity of interest. Exciting advances in statistical modeling make it quite feasible to test hypotheses about such unobservable quantities. This study introduces researchers to the ZiOPC model, which to our knowledge has not been employed in the literature. By employing this model, we are able to draw inferences about both the factors that influence an INGO's decision to exaggerate its allegations of torture and the frequency with which it has done so.

[37] See, also, Milner (2005) and Lake (2010).

Results from the ZiOPC model suggest that while certain organizational incentives do affect the probability that one INGO, AI, exaggerates its claims about state abuse with respect to torture, the probability of an exaggeration is quite low in most cases. From a methodological perspective, this is good news since it implies that standard ordered response models will not produce biased estimates when using data generated from AI reports to draw inferences about the relationships between state-level characteristics and human rights practices. This result is also normatively appealing because it suggests that AI rarely sacrifices credibility for drama, but rather maintains its credibility standards in the face of opportunities to marshal resources by issuing exaggerated reports. If one accepts that the political influence of human rights INGOs is determined by the perceived credibility of the information they distribute, then news that AI adheres to a relatively strict credibility standard should be good news to those who value AI's normative goals.

The research presented here can be extended in at least two directions. First, while we find that AI almost never trades its credibility for organizational reasons, we do not have information about the many other ''information politics'' INGOs out there. A natural extension of our study is similar inquiries into other rights INGOs as well as other ''information politics'' INGOs. Doing so is methodologically important for researchers who employ INGOs reports, and the data sets that analyze them, at face value. Second, valuable theoretical insights may be gained from analyzing how bargaining between domestic NGOs and governments under the shadow of reporting on rights abuses by INGOs may affect the prospects for better human rights practices. We hope that this study spurs future research in the issue-areas mentioned above in order to understand more closely the behavior of ''information politics'' INGOs and other organizations.

References

Brown, Philip H., and Jessica H. Minty. (2008) Media Coverage and Charitable Giving After the 2004 Tsunami. *Southern Economic Journal* 75 (1): 9–25.

Bueno de Mesquita, Bruce, Alastair Smith, Randolph M. Siverson, and James M. Morrow. (2003) *The Logic of Political Survival*. Cambridge: MIT Press.

Bueno de Mesquita, Bruce, George W. Downs, Alastair Smith, and Feryal M. Cherif. (2005) Thinking Inside the Box: A Closer Look at Democracy and Human Rights. *International Studies Quarterly* 49:439–457.

Cingranelli, David L., and Mikhail Filippov. (2010) Electoral Rules and Incentives to Protect Human Rights. *Journal of Politics* 72(1):243–257.

Cingranelli, David L., and David L. Richards. (2010) The Cingranelli and Richards (CIRI) Human Rights Data Project. *Human Rights Quarterly* 32(2):401–424.

Clark, Ann M.. (2001) *Diplomacy of Conscience: Amnesty International and Changing Human Rights Norms*. Princeton: Princeton University Press.

Clark, David H. (2000) Agreeing to Disagree: Domestic Institutional Congruence and US Dispute Behavior. *Political Research Quarterly* 53(2):375–400.

Cmiel, Kenneth. (1999) The Emergence of Human Rights Politics in the United States. *The Journal of American History* 86(3):1231–1250.

Conrad, Courtenay R., and Will H. Moore. (2010) What Stops the Torture? *American Journal of Political Science* 54(2):459–476.

Cooley, Alexander and James Ron. (2002) The NGO Scramble: Organizational Insecurity and the Political Economy of Transnational Action. *International Security* 27(1):5–39.

Davenport, Christian. (1995) Multi-Dimensional Threat Perception and State Repression: An Inquiry Into Why States Apply Negative Sanctions. *American Journal of Political Science* 39(3):683–713.

Davenport, Christian. (2007) *State Repression and the Domestic Democratic Peace*. New York: Cambridge University Press.

Davenport, Christian, and David Armstrong. (2004) Democracy and the Violation of Human Rights: A Statistical Analysis from 1976–1996. *American Journal of Political Science* 48(3):538–554.

Finnemore, Martha, and Katherine Sikkink. (1998) Norms and International Relations Theory. *International Organization* 52(4):887–917.

Franklin, James C. (2008) Shame on You: The Impact of Human Rights Criticism on Political Repression in Latin America. *International Studies Quarterly* 52(1):187–211.

Gibelman, Margaret, and Sheldon R. Gelman. (2004) A Loss of Credibility: Patterns of Wrongdoing Among Nongovernmental Organizations. *Voluntas: International Journal of Voluntary and Nonprofit Organizations* 15(4):355–381.

Gourevitch, Peter, and David A. Lake. (2011) Credible Ethical Action. In *TBA*, edited by Peter Gourevitch and David A. Lake. Toronto: University of Toronto Press.

Gugerty, Mary K.. (2009) Signaling Virtue: Voluntary Accountability Programs Among Nonprofit Organizations. *Policy Sciences* 42(3):243–273.

Guzman, Andrew T. (2008) *How International Law Works: A Rational Choice Theory*. New York: Oxford University Press.

Hafner-Burton, Emilie M. (2008) Sticks and Stones: Naming and Shaming and the Human Rights Enforcement Problem. *International Organization* 62(4):689–716.

Hafner-Burton, Emilie M., and James Ron. (2009) Seeing Double. *World Politics* 61(2):360–401.

Hafner-Burton, Emilie M., and Kiyoteru Tsutsui. (2005) Human Rights in a Globalizing World: The Paradox of Empty Promises. *American Journal of Sociology* 110(5):1373–1411.

Hafner-Burton, Emilie M., and Kiyoteru Tsutsui. (2007) Justice Lost! The Failure of International Human Rights Law to Matter Where Needed Most. *Journal of Peace Research* 44(4):407–425.

Harris, Mark N., and Xueyan Zhao. (2007) A Zero-Inflated Ordered Probit Model, With an Application to Modeling Tobacco Consumption. *Journal of Econometrics* 141(2):1073–1099.

Hathaway, Oona A. (2002) Do Human Rights Treaties Make a Difference? *Yale Law Journal* 111:1935–2042.

Hathaway, Oona A. (2005) Between Power and Principle: An Integrated Theory of International Law. *University of Chicago Law Review* 71:469–533.

Hathaway, Oona A. (2007) Why Do Countries Commit to Human Rights Treaties? *Journal of Conflict Resolution* 51(4):588–621.

Henderson, Conway. (1991) Conditions Affecting the Use of Political Repression. *Journal of Conflict Resolution* 35:120–142.

Hill, Jr. Daniel W, Benjamin E. Bagozzi, Will H. Moore, and Bumba Mukherjee. (2011) *Strategic Incentives and Modeling Bias in Ordinal Data: The Zero-inflated Ordered Probit (ZiOP) Model in Political Science*. Paper Presented at the New Faces in Political Methodology IV meeting, Penn State University, April 30. Available at http://qssi.psu.edu/_les/NF4Hill.pdf. (Accessed December 15, 2010.)

Hopgood, Stephen. (2006) *Keepers of the Flame: Understanding Amnesty International*. Ithaca: Cornell University Press.

Hudock, Ann C. (1999) *NGOs and Civil Society: Development by Proxy*. Cambridge: Polity Press.

Hyde, Susan. (2011) *The Pseudo-Democrats Dilemma: Why Election Observation Became an International Norm*. Ithaca, NY: Cornell University Press.

Hyde, Susan D. (2010) *Catch Us if You Can: Election Monitoring and International Norm Creation*. Available at http://ssrn.com/abstract=1150262. (Accessed December 15, 2010.)

Keck, Margaret E., and Katherine Sikkink. (1998) *Activists Beyond Borders: Advocacy Networks in International Politics*. Ithaca: Cornell University Press.

Kelley, Judith. (2009) D-minus elections: The Politics and Norms of International Election Observation. *International Organization* 63(4):765–787.

Korey, William. (1968) The Key to Human Rights: Implementation. *International Conciliation* 570:17–19.

Lake, David A. (2010) Rightful Rules: Authority, Order, and the Foundations of Global Governance. *International Studies Quarterly* 54(3):587–613.

Lake, David A., and Matthew A. Baum. (2001) The Invisible Hand of Democracy: Political Control and the Provision of Public Services. *Comparative Political Studies* 34(6):587–621.

Lake, David A., and Wendy H. Wong. (2009) The Politics of Networks: Interests, Power, and Human Rights Norms. In *Networked Politics: Agency, Power, and Governance*, edited by Miles Kahler. Ithaca: Cornell University Press.

Landman, Todd. (2005) *Protecting Human Rights: A Comparative Study*. Washington, DC: Georgetown University Press.

Martin, Lisa L., and Beth A. Simmons. (1998) Theories and Empirical Studies of International Institutions. *International Organization* 52(4):729–757.

Milner, Helen V. (2005) Rationalizing Politics: The Emerging Synthesis of International, American, and Comparative Politics. *International Organization* 52(4):759–786.

Moore, Will H. (2000) The Repression of Dissent: A Substitution Model of Government Coercion. *Journal of Conflict Resolution* 44(1):107–127.

Murdie, Amanda. (2009a) The Impact of Human Rights NGO Activity on Human Right Practices. *International NGO Journal* 4(10):421–440.

Murdie, Amanda. (2009b) *Signals Across Borders: The Conditional Impact of INGOs*. Unpublished Doctoral Dissertation, Emory University, Department of Political Science. Available at http://pid.emory.edu/ark:/25593/1dpxn. (Accessed December 15, 2010.)

Murdie, Amanda, and Tavishi Bhasin. (2011) Aiding and Abetting: Human Rights INGOs and Domestic Dissent. *Journal of Conflict Resolution* 55(2):163–191.

Neumayer, Eric. (2005) Do International Human Rights Treaties Improve Respect for Human Rights? *Journal of Conflict Resolution* 49(6):925–953.

Pierskalla, Jan H.. (2010) Protest, Deterrence, and Escalation: The Strategic Calculus of Government Repression. *Journal of Conflict Resolution* 54(1):117–145.

Poe, Steven, and C. Neal Tate. (1994) Repression of Personal Integrity Rights in the 1980's: A Global Analysis. *American Political Science Review* 88:853–872.

Poe, Steven, C. Neal Tate, and Linda C. Keith. (1999) Repression of the Human Right to Personal Integrity Revisited: A Global, Cross-National Study Covering the Years 1976–1993. *International Studies Quarterly* 43:291–313.

Powell, Emilia J., and Jeffrey K. Staton. (2009) Domestic Judicial Institutions and Human Rights Treaty Violation. *International Studies Quarterly* 53(1):149–174.

Powell, Emilia J., and Sara M. Mitchell. (2007) The International Court of Justice and the World's Three Legal Systems. *Journal of Politics* 69(2):397–415.

Risse, Thomas, and Stephen C. Ropp. (1999) International Human Rights Norms and Domestic Change. In *The Power of Human Rights: International Norms and Domestic Change*, edited by Stephen C. Ropp and Kathryn Sikkink. New York: Cambridge University Press.

Ron, James, Howard Ramos, and Kathleen Rodgers. (2005) Transnational Information Politics: NGO Human Rights Reporting, 1986–2000. *International Studies Quarterly* 49(3):557–588.

Schmitz, Hans P.. (2002) *From Lobbying to Shaming: The Evolution of Human Rights Activism since the 1940s*. Paper Presented at the annual meeting of the International Studies Association, New Orleans, LA, March 22–27.

Simmons, Beth A. (2009) *Mobilizing for Human Rights: International Law in Domestic Politics*. New York: Cambridge University Press.

Simon, Adam F. (2006) Television News and International Earthquake Relief. *Journal of Communication* 47(3):82–93.

Stewart, Susan. (2004) The Role of International and Local NGOs in the Transformation of the Georgian-Abkhazian Conflict. *The Global Review of Ethnopolitics* 3(3-4):3–22.

Sundstrom, Lisa M. I. (2005) Foreign Assistance, International Norms, and NGO Development: Lessons from the Russian Campaign. *International Organization* 59(2):419–449.

Svolik, Milan. (2008) Authoritarian Reversals and Democratic Consolidation. *American Political Science Review* 102(2):153–168.

Tsutsui, Kiyoteru, and Christine M. Wotipka. (2004) Global Civil Society and the International Human Rights Movement: Citizen Participation in Human Rights International Nongovernmental Organizations. *Social Forces* 83(2):587–620.

Vreeland, James R.. (2008) Political Institutions and Human Rights: Why Dictatorships Enter into the United Nations Convention Against Torture. *International Organization* 62(1):65–101.

Wiik, Robert. (2002) *Democratisation in the Name of Civil Society. A Quantitative Analysis of the Impact of Non-governmental Organisations on Democratisation*. Unpublished Doctoral Dissertation, Department of Sociology and Political Science, Norwegian University of Science and Technology.

Wong, Wendy H. (2008) *Centralizing Principles: How Amnesty International Shaped Human Rights Politics Through its Transnational Network*. Unpublished Doctoral Dissertation, San Diego, Department of Political Science, University of California. Available at http://tinyurl.com/WHWongDiss (Accessed December 15, 2010.)

Wood, Reed, and Mark Gibney. (2010) The Political Terror Scale (PTS): A Re-introduction and a Comparison to CIRI. *Human Rights Quarterly* 32(2):367–400.

World Bank. (2009) *World Bank Development Indicators*. Washington: World Bank.

Xiang, Jun. (2010) Relevance as a Latent Variable in Dyadic Analysis of Conflict. *Journal of Politics* 72(2):484–498.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Results from alternative specifications of the OP and ZiOP regression models.