

Examining the Effects of International Organizations Using Newspaper Coverage

Stephen Chaudoin*

March 27, 2026

Abstract

Many scholars have hypothesized about the impact of international organizations on the public. International organizations are generally unable to enforce decisions themselves and must rely on compliance partners among the general domestic public. But how do we know that the public actually learns about or from an international organization's actions? This chapter describes the creative and complex process of measuring the impact of organizations using newspaper coverage, machine learning, and topic models. The author highlights innovations in both conceptualization and measurement, as well as exciting quantitative models that are still relatively new to human rights work. The author demonstrates proof of concept through an application to coverage of the International Criminal Court in Philippine newspapers following investigations into the War on Drugs.

*Assistant Professor, Harvard University, Department of Government. Email: chaudoin@fas.harvard.edu

International organizations and non-governmental organizations (NGOs) have long played a critical role in providing information about human rights abuses. They convey details about the frequency and severity of abuses. International legal organizations give further information about whether abuses run afoul of international law.¹ This information reaches a variety of audiences whose reactions can potentially affect the decision-making calculus of the perpetrators of the abuse. Ideally, information about abuse triggers disapproval from internal and/or external audiences, which pushes current or would-be rights abusers to change their behavior.²

The core questions addressed here and in preceding research are: How does information about human rights reach broader audiences? What is the specific content of that information? How does that potentially affect audience reactions?

Earlier and recent research have shared an emphasis on information transmission via the media. As Tumber and Waisbord (2017) write: “It is not an exaggeration to say that the way societies come to understand and experience the situation of ‘human rights’ is largely shaped by mediated communication... communication and persuasion are central processes by which societies come to define, understand and implement rights-based policies” (3). Hafner-Burton (2014) makes clear that the media is a key conduit channel: “a key lesson is that the success of advocates’ efforts to spread human rights through the dissemination of information depends in part on uptake by institutions such as the media that operate according to incentives other than the promotion of rights” (p. 281). Dancy (2021) describe creative measures of human rights resonance based on Google search trends. Search data is important because it captures individual-level engagement with a topic or event. The top hits for many of those Google searches are media reports. When searching for “human rights” using Google’s current default settings, the browser results begin with a mosaic of “Top Stories” highlighting mainstream media reports.

This chapter continues this emphasis on information transmission via the media and describes how the research focus and approach for answering these questions has shifted over time in two

¹E.g. Chaudoin (2016). There is a wide body of related literature. Where possible, I give examples most directly related to the International Criminal Court and the context below.

²E.g. Chapman and Chaudoin (2020).

ways.

The media can transmit information to foreign/external audiences or to domestic/internal audiences. By “foreign,” I mean audiences – citizens or elites or other actors - beyond the country where a rights abuse takes place, and “domestic” refers to audiences within that country. For example, audiences in Washington D.C. and in Manila both read media articles about the Philippines’ so-called war on drugs. Researchers can focus on variation at the extensive margin of this coverage or its content. By extensive margin, I mean whether and how much content an audience receives. Do American outlets cover the WOD, and if so, how many articles cover this issue per year? Content refers to variation in how different outlets or articles cover a particular issue. Does an article about the WOD use a particular frame or emphasize a particular aspect of the issue?

Both dimensions matter. An emphasis on whether the media transmits information to external audiences gives support for key mechanisms in the boomerang model.³ Local NGOs or other sources provide information to foreign media outlets, in many cases, the United States. The U.S. media investigates and then transmits information to political elites and the general population, who can push for additional pressure on abusive governments. A general conclusion from this part of research on media is that this dynamic occurs and that it matters.

As text-as-data tools have become more commonplace in social science applications and media data from non-U.S. sources have become available in greater quantities, this has facilitated a shift to quantitative studies of the content of local coverage. In this chapter, I describe the approach used in Chaudoin (2023) to analyze the content of the Philippines’ media coverage of the WOD. I also use this to demonstrate two common challenges: corpus churn and contextuality, which must be dealt with to make these approaches their most persuasive. I conclude with a discussion of the different conclusions from research on the extensive margin of foreign coverage versus the content of local coverage.

Earlier studies gave persuasive evidence about how the media transmitted information to *foreign/external audiences*, and they emphasized variation at the *extensive margin* of information

³Keck and Sikkink (1998).

transmission. Hafner-Burton and Ron (2013) are an excellent example of scholarship paying close attention to the amount of information transmitted to potentially important external audiences. They analyzed data on the content of major U.S. media outlets from 1981-2000. They regress the number of articles mentioning the term “human rights” and a particular country on several independent variables and dummy variables for various global regions. They find evidence of a Latin bias, where abuses in Latin America are more likely to receive coverage, all else equal. Dancy and Fariss (2017) use a measure of human rights resonance – the degree to which the language of human rights is used to describe events – based on the number of New York Times articles using the term. They show that the proportion of articles using that term has increased over time. Completing the full circle described in the Boomerang Model, Murdie and Davis (2012) show that more shaming events, as detected in events data derived from Reuters media coverage, can improve human rights practices, under certain conditions.⁴ This brief, non-exhaustive list shows how scholarship has analyzed variation at the extensive margin of foreign media coverage to provide evidence of key mechanisms for broader theoretical models of human rights.

To motivate a discussion of work analyzing the content of coverage in local media, I describe in greater detail the arguments and approach in Chaudoin (2023). My goal was to assess whether and how the actions of a critical international organization, in this case the International Criminal Court (ICC), changed the extent and content of media coverage. I needed a single, underlying human rights abuse and media data from before and after that event, ideally covering as long a time frame as possible. The Philippines’ WOD was an excellent case. The WOD was a distinct issue with clear ties to human rights violations, but which could be covered with many possible media frames, from different angles. The media environment in the Philippines is complex, with multiple prominent outlets that operate to varying degrees and direction of political bias.

I found that a key ICC announcement shifted the content of WOD coverage towards a greater emphasis on human rights. Though, this predominantly consisted of coverage of contestation between international human rights advocates and domestic elites that supported the WOD. I paired

⁴For a recent survey of this research, see Brathwaite et al. (2025).

this with survey experimental showing that these changes mattered. Increased coverage of the WOD related to the ICC and human rights decreased support for the WOD, which is encouraging. Though, coverage of contestation partially blunted that effect.

My data collection efforts were certainly not the first to consider the importance of local media content. For example, Gordon and Berkovitch (2007) sampled issues of an Israeli newspaper from 1983 to 1996 and first identified whether the issue talked about social wrongs that could be classified as human rights violations. They then coded whether the articles actually frames the wrong in the language of human rights. This allows them to identify when human rights framing first took off in Israeli media, around the First Intifada. Sorensen (2011) established that human rights content can appear in local media, even when the local media ecology is highly constrained by the government. They show that some Chilean media outlets provide coverage of the legacy of human rights abuses under Pinochet, sometimes in highly encoded ways and sometimes via satirical outlets. Ayoub (2014) describes different media framing of LGBT issues in Eastern Europe, some emphasizing acceptance and others emphasizing the need to defend the nation from international influences.

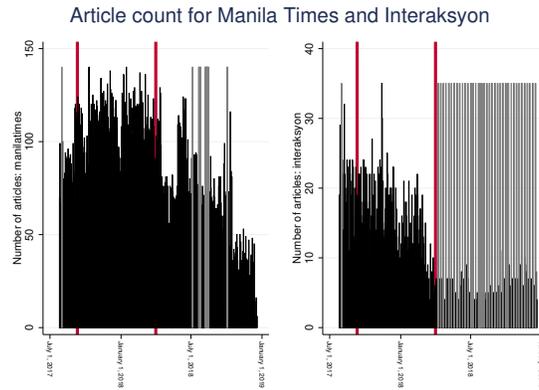
The first task in systematically analyzing the content of local coverage is constructing a corpus of articles. The Cline Center at the University of Illinois at Urbana-Champaign⁵ was ahead of the curve in text-as-data collection and analysis. They used an RSS-feed scraper to collect an extremely large set of articles from global media outlets over a long time frame. Several of the feeds were from the Philippines and covered time before and after February 2018.

I started by identifying the set of media outlets with consistent coverage for the widest possible time frame. I used all articles from the RSS feeds of The Manila Times, The SunStar, Interaksyon, and the Philippine Daily Inquirer, covering 9-10-2017 to 4-1-2018. The Cline Center's corpus also contained articles from other prominent news outlets, like the Manila Standard, ABS-CBN, and GBN. My initial searches collected all articles from January 2016 through December 2018.

Corpus consistency and its nemesis - corpus churn - are underappreciated aspects of text-as-

⁵<https://clinecenter.illinois.edu/>.

Figure 1: Corpus Churn in the Philippines Media Data



data collection efforts. Corpus churn refers to how the presence of certain outlets or content can change over time, sometimes in subtle ways. The most important type of corpus churn occurs when certain outlets are present for some dates but not for others. A subtler example of corpus churn occurs when a subset of an outlet’s articles are absent for certain dates. For example, many outlets have national and local coverage. A database may have both sections of coverage for some dates, but only one for others. This is harder to detect because the outlet is still present in the database, but the data generating process may have changed in unseen ways.

One way to diagnose issues in the underlying corpus is to simply plot the number of articles per day, by outlet. Figure 1 shows this for two of the outlets I used. Each of the black bars shows the number of articles for a particular day. The thick red bars delineate the analysis window that I chose.

There is variation across outlets, which is to be expected. Some outlets simply publish more articles than others. There is variation across time, which is also to be expected. Outlets might publish more articles on the weekend, as opposed to weekdays. Holidays, like Christmas, might also bring dips in the number of articles. The figure also shows zero articles for single days, which occurs when the RSS scraper was offline for technical maintenance.

Overall, within the window that I analyzed, the arrival rate of new articles is fairly constant over time for both outlets. However, outside of this window I can less credibly claim to have a

consistent corpus. The Manila Times shows a large decrease in the number of articles it published per day around May of 2018. It is unlikely that this major outlet produced roughly half as many articles per day for an extended period. More likely, they changed the location of some sections to a different RSS feed for a particular time period (eg if the metro section was converted to as separate feed). The grey bars also indicate periods of time with zero articles for more than one day. These likely indicate technical changes by the outlet, such as a prolonged RSS feed maintenance period. Interaksyon also shows a relatively stable corpus production rate within the analysis window. However, it then shifts to periodic gaps that occur every few days. This is likely caused by an outlet instituting article limits per time period, which would not be automatically detected by an RSS feed scraper.

Why does this matter? The target population for analysis of media content is every article produced by an outlet for a set period of time. Undetected corpus churn means that the researcher is not analyzing the target population, potentially in ways that bias their desired inferences. If corpus churn affects a particular news outlet, it may be “noise” for the purposes of any downstream analysis.

However - in practice – these issues are not trivial. “Noise” in a dependent variable is not benign (at least from the perspective of biasing estimates) if it is correlated with explanatory variables of interest. These variables themselves could vary across time or media outlet. In my context, media outlets have potential biases in favor of particular politicians or issues. The owner of the Manila Times is so closely associated with former President Rodrigo Duterte that he was named Duterte’s “special envoy for international public relations.” The potential for the Manila Times to have a slanted view of the WOD or the International Criminal Court is obvious. If the Manila Times became more or less prominent in the corpus, this could easily lead to over- or under-estimation of key features, like the prevalence of content that connected the WOD to human rights or the degree of domestic political contestation over the WOD.

Corpus churn is also present in commercially available datasets, like Lexis Nexis.⁶ As part

⁶A major new effort to identify shaming incidences also uses Lexis, Brathwaite et al. (2025).

of revisions to my article on the Philippines, reviewers pushed me to consider other countries. I therefore replicated my approach, focusing instead on Kenya and its experiences with the ICC. Instead of the WOD, I collected articles that were about post-electoral violence, whose events spurred the ICC's involvement. I no longer had institutional access to the Cline Center's resources, so I collected articles from Lexis Nexis' Nexis Uni database. They only had two major Kenyan broadsheets in their collection, The Nation and The Star. Even using this (very expensive) database from a university subscription, there was substantial corpus churn. I had to limit the window of my analysis because Nexis Uni only had three articles mentioning any variation of the string "post-electoral violence" (often shorthanded as "PEV") from The Nation for the time period before January 2010 and zero from The Star. That string is an obviously important topic in Kenya. It is unfathomable for the true level of media coverage of PEV to be so minimal during those time periods. And indeed, manual searches of The Nation and The Star's websites produced lots of articles using that term from throughout this time period. Those articles simply weren't part of Nexis Uni's corpus for reasons that are impossible to diagnose. These two outlets also go dark on Nexis Uni for roughly two months in 2012 for unknown reasons. These could be damning problems for any analysis that did not do due diligence about underlying corpus churn in data construction.

Corpus churn is also present in regional news aggregators. Kitagawa and Shen-Bayh (2024) use a word embeddings approach to assess the degree to which the ICC is discussed in conjunction with words indicating political polarization. Their outcome of interest is the degree to which the term "ICC" is associated with a set of words indicating political polarization. They constructed a much larger corpus that mine from Lexis. Their corpus includes articles from four Kenyan newspapers from 2007 to 2020. They use change point detection to find that the association increases then decrease in November 2010 and June 2016 respectively, which correspond roughly to the start and end of the ICC's actions.

They conducted searches for "ICC/International Criminal Court" in the legal affairs section of a database called AllAfrica which aggregates news from African outlets. AllAfrica's coverage of

Kenyan outlets is more extensive than Nexis Uni, but it too shows significant signs of corpus churn. Figure 2 shows the number of articles per day in the data set used for the changepoint analysis, with vertical bars marking the two change points. Three outlets make up 99.7% of the corpus and none have consistent coverage across the full time period of analysis. One outlet, The Nation, has zero articles per day for approximately 42 months, from February 2012 to July 2015. Though, The Nation produced many articles about the ICC during this time period. Mueller (2014) analyzes many in her analysis of the Kenyatta campaign against the ICC.

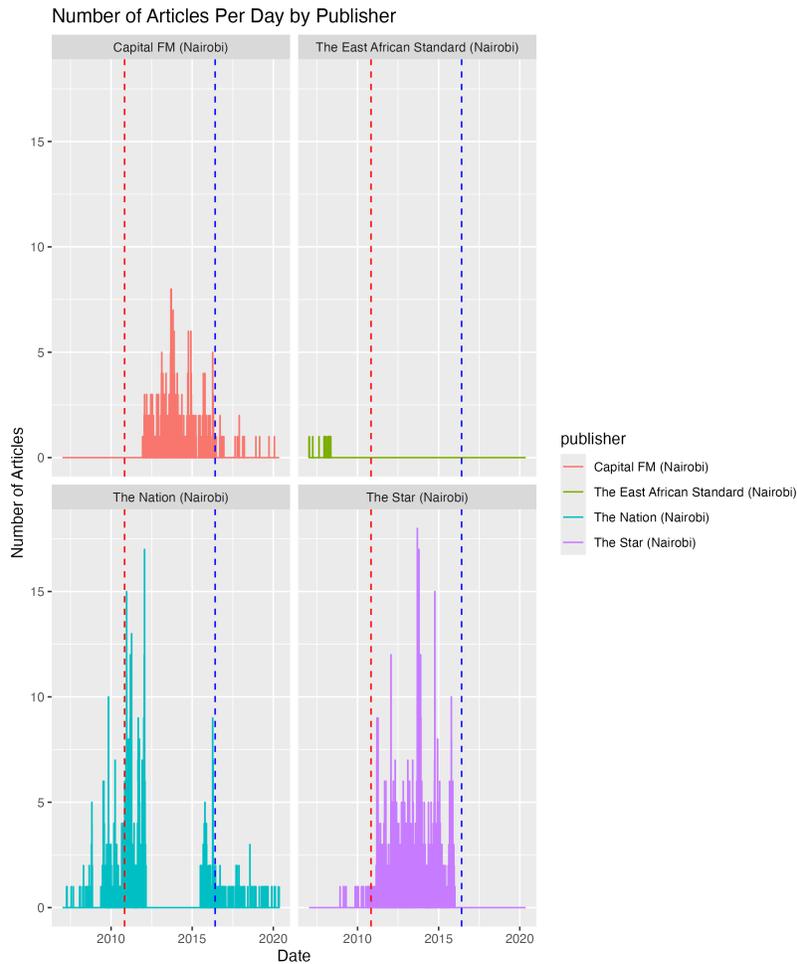
The arrival of articles for one of the other outlets nearly coincides with the departure of The Nation from the corpus. It seems likely that AllAfrica supplemented their coverage of Kenya by adding articles from Capital FM when they lost access to articles from The Nation. When AllAfrica's coverage of The Nation returns, it appears to be throttled by an article limit for an extended period after the end of the ICC's case against Kenyatta.

Whether this churn affects conclusions is a context-specific question that is beyond the scope of this chapter. My overall point though is that, among three commonly used datasets - from RSS feeds, Lexis, to regional news aggregators - all show signs of corpus churn.

There are many options for systematically describing media content. The appropriate choice depends on the particular question and research design. I used topic modeling for articles that were about the WOD. Topic modeling consists of identifying clusters of words that tend to co-occur. The prevalence of particular clusters can indicate a greater emphasis on certain themes or frames within an article. This is important variation that would not be detected in analysis at the extensive margin of coverage. For example, the number of articles about the WOD might stay the same, but there could be meaningful shifts in whether human rights frames were emphasized instead of criminality frames. The number of articles using the term "human rights" could also change at the extensive margin, but there could be changes in whether these are references to advocates pushing for accountability of WOD-supporters flaming foreign actors.

Gordon and Berkovitch (2007), writing before the more widespread availability of topic-modeling packages, described a similar intent. They wanted to analyze whether "issues that had in the past

Figure 2: Corpus Churn in AllAfrica Media Data



been depicted as mere ‘problems’ would be described as violations of human rights” (260). They achieved this by hand-coding a random sample of issues of Israeli news outlets. Part of the appeal of more recent approaches like topic modeling is their scalability. They are intended to automatically identify features of content from very large amounts of text that could be beyond what a human or team of coders could feasibly achieve with manual coding.⁷

These content analysis approaches a lot of manual, unglamorous, context-specific data clean-

⁷For other related applications, see Terman (2017) who show that the prevalence of “Women’s Rights and Gender Equality” topics is higher in U.S. news coverage of Muslim and MENA countries, even controlling for measures of equality. Sentiment analysis is another commonly used tool. Zhou, Kiyani and Crabtree (2023) use it to measure the intensity of negative sentiment in the particular reports underlying many shaming events, such as Amnesty International or U.S. State Department reports.

ing. The modal article (mine included) includes some variation on the phrase “I use machine learning to...” or “I use automated content analysis to...”. In reality, every step of this pipeline likely needs to be carefully analyzed with context-specific knowledge in mind. For example, some outlets contain standard phrases that they include in every article, like “also available on your mobile phones, laptops, and tablets.” Many text-as-data approaches take a bag of words approach, where syntax and word order are not accounted for. This means that commonly occurring phrases might not identify the prevalence of a topic, but rather an idiosyncrasy of a particular outlet. The Manila Times isn’t especially inclined to cover the topic of “laptops,” even if that word appears in every one of their article’s preambles.

Different articles may use different words to refer to the same person or thing. For example, a prominent event in the Philippines’ WOD was the murder of a boy named Kian delos Santos. Different articles used several variations of his name, so any bag of words approach might mistakenly conclude those articles were about different things. This is especially thorny for names in general, which might contain common components (like “delos Santos”) across individuals. For my application, I combined many of the most important strings into a single string to ensure that idiosyncrasies did not influence results. For example, I combined variations of Kian delos Santos’ name into a single, one word string: kiandelossantos. I had originally planned to use automated named-entity recognition approaches to help with this problem. In 2018, these processes were not advanced enough to make me confident in their output. It is possible (likely) that they have advanced substantially in the meantime.

One value of these new approaches is that they can provide systematic, quantitative evidence about key theoretical arguments. While a general description of work on the extent of coverage in foreign media would loosely say “it happens and it works,” analysis of the content of local coverage tends to emphasize intense heterogeneity in reactions and contestation between actors. In general, it tends to be more measured in the enthusiasm for the effects of international organizations or external shaming. In addition to support, defiance of foreign shaming frequently plays out on the pages of the media in the targeted country (Terman, 2023). The clear importance of contestation

has also led to research on which persuasive appeals and rhetorical strategies are most effective for international organizations.⁸

In terms of methods and measurement, my general advice would be that the ability to process vast quantities of text data creates interesting opportunities. Analysis only appears to be getting easier. Many have (or soon will) use AI tools for content analysis, and I do not doubt those tools will perform equally well at content measurement. This will place even greater emphasis on corpus construction and close attention to whether features of the input bias conclusions about the output. I only felt confident drawing conclusions from automated text-as-data processes after reading hundreds (probably more like thousands) of articles from the corpus, to anticipate and react to unexpected things in the data pipeline. Since part of the impetus for studying local coverage, outside of US outlets, is to better understand context-specific dynamics, that same concern will need to extend to data and measurement steps, as well.

References

- Ayoub, Phillip M. 2014. “With arms wide shut: Threat perception, norm reception, and mobilized resistance to LGBT rights.” *Journal of Human Rights* 13(3):337–362.
- Brathwaite, Robert, Shanshan Lian, Amanda Murdie and Baekkwon Park. 2025. “Tailoring the message: A new dataset on the dyadic nature of NGO shaming in the media.” *Journal of Peace Research* p. 00223433241276975.
- Chapman, Terrence L and Stephen Chaudoin. 2020. “Public reactions to international legal institutions: The international criminal court in a developing democracy.” *The Journal of Politics* 82(4):1305–1320.
- Chaudoin, Stephen. 2016. “How Contestation Moderates the Effect of International Institutions: the International Criminal Court and Kenya.” *Journal of Politics* 78(2):557–571.

⁸Zvobgo and Chaudoin (Forthcoming); Morse and Pratt (2022).

- Chaudoin, Stephen. 2023. "How international organizations change national media coverage of human rights." *International Organization* 77(1):238–261.
- Dancy, Geoff and Christopher J Fariss. 2017. "Rescuing human rights law from international legalism and its critics." *Human Rights Quarterly* 39(1):1–36.
- Dancy, Geoffrey Thomas. 2021. "The hidden impacts of the ICC: An innovative assessment using Google data." *Leiden Journal of International Law* 34(3):729–747.
- Gordon, Neve and Nitza Berkovitch. 2007. "Human rights discourse in domestic settings: How does it emerge?" *Political studies* 55(1):243–266.
- Hafner-Burton, Emilie and James Ron. 2013. "The Latin bias: Regions, the Anglo-American media, and human rights." *International Studies Quarterly* 57(3):474–491.
- Hafner-Burton, Emilie M. 2014. "A social science of human rights." *Journal of Peace Research* 51(2):273–286.
- Keck, Margaret E. and Kathryn Sikkink. 1998. *Activists beyond borders: advocacy networks in international politics*. Cornell.
- Kitagawa, Risa and Fiona Shen-Bayh. 2024. "Measuring Political Narratives in African News Media: A Word Embeddings Approach." *The Journal of Politics* 86(3):1087–1092.
- Morse, Julia C and Tyler Pratt. 2022. "Strategies of contestation: International law, domestic audiences, and image management." *The Journal of Politics* 84(4):2080–2093.
- Mueller, Susanne D. 2014. "Kenya and the International Criminal Court (ICC): politics, the election and the law." *Journal of Eastern African Studies* 8(1):25–42.
- Murdie, Amanda M and David R Davis. 2012. "Shaming and blaming: Using events data to assess the impact of human rights INGOs." *International Studies Quarterly* 56(1):1–16.

- Sorensen, Kristin. 2011. "Chilean Print Media and Human Rights: Mainstream Silence Versus Satirical Subversion." *Peace & Change* 36(3):400–426.
- Terman, Rochelle. 2017. "Islamophobia and media portrayals of Muslim women: A computational text analysis of US news coverage." *International Studies Quarterly* 61(3):489–502.
- Terman, Rochelle. 2023. *The Geopolitics of Shaming: When Human Rights Pressure Works—and When It Backfires*. Vol. 201 Princeton University Press.
- Tumber, Howard and Silvio Waisbord. 2017. media and human rights. In *The Routledge Companion to Media and Human Rights*. Routledge pp. 1–14.
- Zhou, Yuan, Ghashia Kiyani and Charles Crabtree. 2023. "New evidence that naming and shaming influences state human rights practices." *Journal of Human Rights* 22(4):451–468.
- Zvobgo, Kelebogile and Stephen Chaudoin. Forthcoming. "Complementarity and Public Views on Overlapping International and Domestic Courts." *Journal of Politics* .