

News Media and Violence Against Women: Understanding Framings of Stigma

Shravika Mittal¹, Jasmine C. Foriest¹, Benjamin D. Horne², Munmun De Choudhury¹

¹Georgia Institute of Technology

²University of Tennessee Knoxville

smittal87@gatech.edu, jay@gatech.edu, bhorne6@utk.edu, munmund@gatech.edu

Abstract

Discussions of Violence Against Women (VAW) in publicly accessible forums like online news media can influence the perceptions of people and organizations. Language reinforcing stigma around VAW can result in negative consequences such as unethical representation of survivors and trivialization of the act of violence. In this work, we study the presence of stigmatized framings in news media and how it differs based on media attributes like regionality, political leaning, veracity, and latent communities of news sources. We also investigate the interactions between VAW-based stigma and 14 issue-generic policies used to describe political communications. We found that articles from national, right-leaning, and conspiratorial news sources contain more stigma compared to their counterparts. Furthermore, alignment of articles to the issue-generic policies offers the highest explanation for the presence of stigma in news articles. We discuss implications for institutions to improve safe reporting guidelines on VAW.

Introduction

Violence Against Women (VAW) is defined by the United Nations (UN) as “any act of gender-based violence that results in, or is likely to result in, physical, sexual, or mental harm or suffering to women, including threats of such acts, coercion or arbitrary deprivation of liberty, whether occurring in public or in private life” (United Nations 1993). It is recognized as a global problem that can have adverse effects on the health, safety, productivity, and overall well-being of women. Despite designated efforts by associations such as World Health Organization (WHO) and UN Women to prevent and eliminate all forms of VAW, its prevalence remains high – “globally 1 in 3 women experience physical and/or sexual violence in their lifetime” (United Nations 1993).

Media, including mass and print, is considered a “key entry point for preventing VAW in the long-term because of their unique reach to broad sections of the population” (UN Women 2019). In addition, media has the ability to influence and shape public opinion about what is considered socially acceptable – Bandura (2001) identified media to “serve as socializing agents that aid in construction and perpetuation of perceptions and learned behaviors.” News media, in particular, can potentially raise public awareness to bring jus-

tice to survivors and shine light upon the extent, causes, and repercussions of VAW (Owusu-Addo et al. 2018). At the same time, through harmful or stigmatized representations, news media may reinforce problematic norms that promote gender inequality, trivialize the issue, or encourage reproduction of violent acts (“copycat effect” (Lira and Toledo 2014)). A contributing factor to explain these depictions of VAW stems from Entman’s sociological notion of “framing” (Entman 1993) – news media may “select some aspects of a perceived reality and make them more salient in a communicating text, to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation for the item described.”

Given the significant impact of news media portrayals on personal, political, and social justice responses, researchers have studied representations of VAW in mainstream news media. For instance, a body of work shows that journalists often use “episodic framing” (Lindsay-Brisbin, DePrince, and Welton-Mitchell 2014), eliciting individualistic details instead of posing VAW as a systemic societal problem to their audience. Research has also discovered that, unfortunately, news stories depict sexual assault in a “sensationalist” manner (Sacks, Ackerman, and Shlosberg 2018), reinforcing cultural stereotypes, stigma, and rape myths (Meyers 1996). In addition, journalistic reports can intensify stigma around VAW through “victim-blaming” – “an act of blaming victims for criminal acts committed against them” (Berns 2001) – such as attributing VAW incidents to physical appearances and descriptions. Usage of such stigmatizing language not only impacts public opinion, as highlighted earlier, but it also trivializes the act of violence, and discourages (*silences*) the survivors from reporting the crime (Eastal, Holland, and Judd 2015), emphasizing the need to understand (and eliminate) VAW-based stigma in news.

To the best of our knowledge, existing works uncovering stigma of VAW in news only examine articles from a few publishers. However, to inform change, it is essential to recognize and contextualize stigma around VAW in the larger news media landscape, factoring in the diversity of today’s media. In this work, we do so by looking at articles from 396 national, local, and conspiratorial news sources, within the United States. Furthermore, there is a gap in understanding the factors that may drive journalists’ usage of specific linguistic framings, in the context of VAW. For instance, Ele-

jalde, Ferres, and Schifanella (2019) observed that news media coverage is highly influenced by the geographic reach and the political leaning of outlets. Moreover, with mis- and dis-information threatening the American news ecosystem in recent years, it is imperative to also consider news quality as we seek to better understand framings of VAW. In this context, misinformation on VAW can contribute to harmful attitudes, making it more difficult to prevent and address gender-based violence. This misinformation may gain perceived credibility through inauthentic coordination amongst fringe media outlets (Horne, Nørregaard, and Adalı 2019). Consequently, we make the first attempt to study and explain stigma around VAW in news using outlet attributes such as *regionality*, *partisanship*, *veracity*, and *media community membership* (captured via content sharing networks) as well as article attributes, such as *issue-generic policies* to study political issues highlighted in news (Boydston et al. 2013). Our research pursues the following two research questions:

RQ1: How is stigma around VAW perpetuated across U.S. news media and how does it differ based on (a) regionality, (b) political leaning, (c) veracity of news sources, (d) media communities, and (e) issue-generic policies?

RQ2: What are the relative roles of the above mentioned factors in explaining stigma in VAW news reports?

To answer these two research questions, we filtered two publicly available news datasets – NELA-GT and NELA-Local – using relevant key phrases, identified by domain experts, to only contain VAW-specific articles in the United States. In total, we looked at 98,107 articles in NELA-GT and 68,892 articles in NELA-Local – published from April 2020 to March 2023. To explore the presence of stigmatized VAW framings in this data, we adopted a BERT-based (Devlin et al. 2019) representation learning framework proposed by Mittal and De Choudhury (2023). We then looked at differences in stigma-based framings, along the five factors mentioned in RQ1. Our analysis revealed that national, right-leaning, and low-veracity outlets tend to use more stigma compared to their counterparts. Our analysis of media communities supported these findings, showing that outlets in the fringe, right-leaning community used stigmatized framings more than the average outlet, but also added nuance to the findings. In particular, we found that small clusters of progressive outlets used stigmatized framings more than the average outlet despite left-leaning outlets more generally using less. Our article level analysis indicated that articles with mentions of certain issue-generic policies such as those highlighting (in)equality of penalty or rewards, policy outcomes across organizations, and effectiveness of existing systems contain more stigma. On the other hand, articles highlighting overall well-being, social responsibility, and existing social norms contain less stigma. Lastly, an article’s alignment (or lack thereof) with issue-generic policies was the most informative in determining whether an article uses stigmatized VAW framings or not.

Broader Perspectives, Ethics, and Competing Interests: Our findings bear implications for survivor-centered reporting guidelines that avoid perpetuation of political biases and

conspiracy theories, provide contextual information highlighting the underlying systemic issues, and use trauma-informed language instead of sensationalist and victim-blaming framings. We followed best practices in our analysis (Chancellor, Baumer, and De Choudhury 2019), such as working with deidentified data, and refraining from sharing raw or personally identifiable data in any form. All quotes are paraphrased to reduce traceability and potential harm to those referenced in the news articles. *We provide a content warning as the paper contains descriptions of VAW.*

Background and Related Work

Stigma, Stereotypes, and Violence Against Women

Stigma and stereotypes have a documented influence on incidence and prevalence of VAW. “This can take the form of promoting stereotypes and myths about ‘real’ and ‘deserving’ victims and ‘provoked’ or ‘tragic’ perpetrators and failing to convey the true scale of violence against women and girls” (Eaves and Equality Now 2012). But how existing research comes to identify presentations of stigma relating to VAW, and the resulting implications for individual help-seeking behaviors, are varied by individual contexts. As it relates to intimate partner violence, research suggests disparities for women of color are best understood when accounting for cultural messaging that communities of color are not victims of abuse (Plough 2000). However, in the context of sex work, VAW research offers that stigma is structural, and barriers to prevention are shaped by social stigmatization and criminalization of sex work – enabling violence against sex workers to persist.

Despite the richness of these perspectives, scholars have criticized that existing research on stigma has had an *individualistic* focus on characteristics. For instance, prior works concentrate on a single manifestation of stigma or micro-level interactions (e.g., cultural stereotype or victim-blaming), leading to “vague and varied conceptualizations of stigma” (Oliver 1990). To address this concern, we adopt a BERT-based framework to explore stigmatized VAW framings (Mittal and De Choudhury 2023), which is built on a *holistic* conceptualization of stigma – “stigma exists when elements of labeling, stereotyping, separation, status loss, and discrimination occur together in a power situation that allows them to unfold” (Link and Phelan 2001) and thus goes beyond an individualistic focus due to a reliance on a very large news dataset reporting on diverse VAW incidents.

News and Violence Against Women

Framing, defined as the “portray[al] of an issue from one perspective to the necessary exclusion of alternative perspectives of issues,” (Schaffner and Sellers 2010) of VAW may be correlated with harmful cultural and structural stigma perpetuated in mass communication. Studies attribute problematic portrayals of VAW to media’s proclivity toward frames (Bullock 2007) of objectivity (Meyers 2004), concealing broader social context (Bullock 2007), and directing the responsibility of prevention onto victims (Carline and Eastaer 2014). Research has also identified media’s occasional use of “social problem frame” in reports

of VAW (Berns 2004). Another examination investigated changes in journalistic framing of sexual violence following the #MeToo movement (Noetzel et al. 2022).

While all of these works contribute to the cognizance of problematic VAW framings, they each assess a narrow manifestation, restricted to a few news articles and non-standardized dimensions of framing. We harness the 14 issue-generic policies derived from political communication theory to understand framing of subject matter in mass communication (Boydston et al. 2013). Prior work has employed them to study framing of immigration (Mendelsohn, Budak, and Jurgens 2021), tobacco use, and same-sex marriage (Card et al. 2015). In the context of VAW, the policies provide a systematic way to comprehend discourse on distribution of rewards/punishments, capacity of existing judicial systems and public well-being, and stigma they may attract.

Additionally, existing studies exploring framing of VAW examine media in a limited number of outlets (Chagnon 2020), locale (Bullock 2007), and types of violence (Nettleton 2011). However, to inform change, it is essential to understand framings in the larger news media landscape. This paper contributes by studying media content from two publicly available datasets, expanding on the number (396), locality, and variety of news outlets.

Media Attributes: Regionality, Partisanship, Veracity, and Communities

Audience reach and targeting efforts of news are influenced by geographical factors and differ for local and national news outlets (Elejalde, Ferres, and Schifanella 2019). In tandem, the political bias of a news outlet is correlated with audience targeting efforts intended to appeal to those they reach and maximize revenue (Bahamonde et al. 2018). Barnett and Hilz (2018) found that individuals with politically conservative beliefs are more likely to hold rape supportive attitudes. In addition, republican men reported the highest levels of sexual assault myth acceptance, lowest levels of concern for sexual assault and were least likely to perceive the #MeToo movement as having a positive impact (Ortiz and Smith 2022). Hence, both the *regionality* and the *partisanship* of a news outlet may influence coverage of VAW.

The news consumed today is not always reliable or written by a professional journalist. Information may be filtered through fringe, alternative media outlets, where sensationalist framing and moral-emotional language are used to gain reader engagement (Brady et al. 2017). In the context of VAW, fringe media that blames victims of violence for their own victimization, or suggests that women should be more vigilant to avoid being assaulted, can shift the focus away from holding perpetrators accountable for their actions. This can discourage victims from coming forward and can lead to a culture where VAW is normalized and minimized. Hence, we also consider the *veracity* of an outlet.

Lastly, there are various sub-groups of media that impact VAW coverage, but those groups may be opaque to the information consumer and can be hard to define. For example, fringe media may coordinate to make certain narratives seem more widely reported (Horne, Nørregaard, and Adalı 2019), or the ownership of outlets may impact story fram-

ing and therefore the public's view of events (Levendusky 2022). Given that latent sub-groups like these can be difficult to directly capture in data, we estimate them through what we call "media communities" – leveraging prior work (Horne, Nørregaard, and Adalı 2019; Starbird et al. 2018).

We contribute by studying the relationship between stigmatized VAW coverage and four outlet factors: regionality, partisanship, veracity, and latent communities.

Data

We used articles published by U.S. national, local, and fringe news media outlets, present within the publicly available NELA-GT (Nørregaard, Horne, and Adalı 2019) and NELA-Local (Horne et al. 2022) datasets. To have an up-to-date news coverage, we used an extended version of both NELA-GT and NELA-Local, collected from April 2020 to March 2023. Table 1 provides an overview of these.

We chose these datasets as they contain nearly every news article published within the specified timeline. Apart from the textual content of news articles, the datasets provide information on attributes such as source of publication, date of publication, estimates of political leaning, estimates of veracity (for NELA-GT), and the state and the county the news outlet is headquartered in (for NELA-Local).

Filtering and Data Cleaning

We filtered news articles, which were relevant to study discourse on VAW, in both the NELA-GT and NELA-Local datasets. To do so, we used key phrases identified by UNFPA domain experts (Purohit et al. 2016), and further fine-tuned by ElSherief, Belding, and Nguyen (2017), to define three categories of Gender Based Violence (GBV): (1) Physical Violence (e.g., *'intimate partner violence'*, *'woman attacked'*, and *'domestic violence'*), (2) Sexual Violence (e.g., *'sexual assault'*, *'rape victim'*, and *'stalking woman'*), and (3) Harmful Practices (e.g., *'woman trafficking'*, *'sex slave'*, and *'woman abuse'*). We carefully selected a subset of these GBV key phrases such that they were applicable to VAW. For instance, we dropped gender neutral phrases like *'child abuse'* and *'physical violence'* from the original list. After this processing and filtering out sources not headquartered in the U.S., we were able to get 115,197 VAW-specific news articles in NELA-GT and 68,892 in NELA-Local.

Lastly, following P. Aires, G. Nakamura, and F. Nakamura (2019) and Nakov and Da San Martino (2021), we used source-level veracity labels – already available in the NELA-GT dataset, obtained through the MBFC¹ resource, to remove articles originating from low-veracity conspiracy-pseudoscience, and questionable sources. Finally, we were left with 98,107 articles in NELA-GT and 68,892 in NELA-Local. We grouped the filtered out 17,090 low-veracity articles, to form a separate dataset containing articles from conspiratorial news sources, to answer RQ1(c).

Data Validation

To validate our filtered VAW-specific NELA-GT and NELA-Local datasets, we randomly sampled 50 news ar-

¹<https://mediabiasfactcheck.com/>

Dataset	#Articles	#N	Mean (#words)	SD	Example Sources
NELA-GT	7,246,352	361	3702.47	108.33	CNBC, Washington Post, The New York Times, The Duran, US News, The Blaze
NELA-Local	2,089,785	313	4440.32	773.01	Daily News, The Daily Gazette, Philadelphia Weekly, Daily Herald, The Republic

Table 1: Overview of the news datasets, before filtering. $N = \#sources$. Mean, SD (Std dev) represent article-length statistics.

ticles each from the two data sources. The first two authors hand annotated them as relevant or irrelevant (to VAW) based on their prior experience and expertise with studying news related content. Throughout this process, the annotations were discussed amongst the two authors to resolve conflicts, and reach agreement. The subsets contained a high percentage of relevant articles (84% for NELA-GT and 92% for NELA-Local). The annotations were reliable in terms of a high inter-rater agreement or Cohen’s Kappa of 0.89 for NELA-GT and 0.91 for NELA-Local.

Methods

Measuring Stigma

In this study, our goal was to understand how news media describes VAW, through a stigma lens. Hence we adopted the Approval/Stigma frame characterization, proposed by Mittal and De Choudhury (2023), to quantify levels of stigma in news media reports. These authors curated well-validated dictionaries, referring to existing literature on conceptualization of stigma (Link and Phelan 2001; Goffman 2014), to represent *approval* and *stigma* dimensions. Link and Phelan (2001) provide a comprehensive understanding of stigma – “stigma exists when elements of labeling, stereotyping, separation, status loss, and discrimination occur together in a power situation that allows the components of stigma to unfold.” As supported by prior work, stigma around VAW contains the elements mentioned above:

- Labeling:** Survivors are labeled as false accusers (Larson 2018); e.g., “her account did not add up.”
- Stereotyping:** Dominant beliefs link survivors to undesirable characteristics (Plough 2000); e.g., “engage in behaviors such as thumbsucking.”
- Separation:** Survivors are “othered”, receive “out-group treatment” (Pickup, Williams, and Sweetman 2001).
- Status loss and Discrimination:** Survivors experience disadvantaged outcomes (Jina and Thomas 2013); e.g., “dealing with PTSD, scarred and may never recover.”

Some exemplar terms present in the Approval/Stigma dictionaries are ‘support’, ‘accept’, ‘regard’ for approval and ‘disgrace’, ‘exclusion’, ‘flaw’ for stigma. We carefully studied the dictionaries to check for relevancy. Using these dictionaries the authors then generated vector representations for *approval* and *stigma* dimensions via BERT (Devlin et al. 2019), a pre-trained large language model. They averaged the word embeddings, extracted using BERT, for all the keywords present in the two dictionaries, to create the representations for *approval* and *stigma*. Finally, following the *semantic axes* approach (An, Kwak, and Ahn 2018), a single linguistic representation for Approval/Stigma was obtained by subtracting the embedding for stigma from approval.

Next, to quantify stigmatized framings of VAW in news media, we extracted sentence-level BERT embeddings to get a linguistic representation of the news article content similar to Mittal and De Choudhury (2023). Lastly, following prior work (Kwak et al. 2021), we compared the vectors representing the news article and the Approval/Stigma frame using the cosine similarity metric. This comparison resulted in scores ranging from -1 to 1 such that a score closer to -1 (alternatively, 1) indicates an alignment to stigma (alternatively, approval). Following Mittal and De Choudhury (2023) we then set a threshold on 0 cosine similarity to say that news articles with a score lesser (greater) than 0 align more with the stigma (approval) dimension.

Validating Approval/Stigma Frame Extraction Two authors hand-annotated a random sample of 10 news articles identifying the presence of either stigma or approval-based framings. As a demonstration of annotator credibility, we note that one annotator has expertise in social computing research and the other is proficient in violence-related public health work. They agreed on 70% of the samples and discussed their annotations to reach 100% consensus. Following this, they sampled and annotated 50 articles (25 each from NELA-GT and NELA-Local datasets) to compare human-assigned ground truth labels against those generated by the Approval/Stigma frame characterization. The framework achieved high precision (82%) and recall (78%).

Error analyses: The framework mislabeled presence of stigma in 4 articles when they contained stigmatizing quotes from subjects, but the overall reporting argued against the said ideas. E.g., “[name] said ‘[...] claims were corroborated.’ This allegation stands void for multiple reasons.”

Media Communities

Next, we generated broad, latent sub-groups of news outlets, extending beyond regionality, political leaning, and veracity of sources captured in the NELA datasets – our motivation was that this fine-grained grouping could provide insights into different communities of VAW framings. Accordingly, we constructed a content sharing network (CSN) of news outlets, where nodes represent outlets and edges represent the proportion of articles copied between outlets (directed as information flow, $A \rightarrow B$ means B copies from A). As demonstrated in prior studies (Gruppi et al. 2022; Horne, Nørregaard, and Adalı 2019; Starbird et al. 2018), this structural conceptualization of news networks consistently yields clusters that represent meaningful and distinct parts of the news landscape. These clusters may capture inauthentic coordination, ownership, partisanship, and regionality.

This network was constructed using the algorithm described by Horne, Nørregaard, and Adalı (2019). Specifically, we constructed a Term Frequency Inverse Document

Issue-generic Policy Dimension	Description
Economic	Costs, benefits, and other monetary implications
Capacity and resources	Availability of physical, human, or financial resources, and capacity of existing systems
Morality	Religious, ethical, or social implications
Fairness and equality	Equality or inequality with which laws, punishments, or rewards are distributed
Constitutionality and jurisprudence	Rights, freedoms and authority of individuals, government, and corporations
Policy prescription and evaluation	Discussion of policies proposed for addressing an identified problem
Crime and punishment	Effectiveness and implications of existing laws and their reinforcement
Security and defense	Security or threats to security of an individual, group or nation
Health and safety	Healthcare access and effectiveness, public safety
Quality of life	Effects of policy on individuals' wealth, well-being, and happiness
Cultural identity	Social norms, values, and customs constituting culture(s), in relevance to a policy issue
Public opinion	Attitudes and opinions of the general public
Political	Considerations related to politics and politicians including elections and lobbyists
External regulation and reputation	Comparison of policy outcomes across nations, states, and groups

Table 2: Issue-generic policy dimensions devised by Boydston et al. (2013).

Frequency (TF-IDF) matrix for each 5 day window of the dataset and created article pairs by keeping articles that had a greater than 0.85 cosine similarity (near verbatim copies of each other). We then ordered these edges by time and aggregated individual articles into outlets. Importantly, to ensure we were appropriately representing the news landscape, we constructed this network using all of the outlets and articles in *NELA-GT* and *NELA-Local* (9,336,137 articles from 674 outlets) rather than only using those in our final filtered VAW dataset (184,089 articles from 396 outlets).

A visualization of this network is given in Figure 2. The community membership (computed using modularity) of a news outlet is used in our subsequent analyses for RQs 1, 2.

Operationalizing Issue-generic Policies

Issue-generic policies (refer to Table 2) – proposed by Boydston et al. (2013) – summarize aspects that provide high-level insights into publicly available political communications and their influence on public attitudes.

To answer RQ1(e), i.e. to explore differences in stigmatized VAW framings across mentions of issue-generic policies, we first identified presence of the 14 issue-generic policy dimensions, as listed in Table 2, in our news datasets. We used two publicly available datasets, manually annotated for the presence of the 14 policy dimensions, to extract mentions of issue-generic policies within our news articles:

1. **Media Frames Corpus** (Card et al. 2015): This consists of annotated news articles on three political issues: (1) immigration (5, 549 articles), (2) smoking (4, 077 articles), and (3) same-sex marriage (6, 298 articles).
2. **Immigration Discourse on Social Media** (Mendelsohn, Budak, and Jurgens 2021): This consists of 4, 497 immigration-related annotated tweets.

We used a BERT-based representation learning framework (Devlin et al. 2019) to get linguistic representations of the 14 policy dimensions. For this, we used all the news article excerpts and tweets – contained in the two datasets mentioned above – annotated for the presence of each issue-generic policy. For each policy, the corresponding annotated

article excerpts and tweets were then transformed into vector representations using sentence-level BERT. Finally, we averaged the sentence-level embeddings of all the articles and tweets annotated for an issue-generic policy to get its linguistic representation. This process resulted in 14 embeddings representing the 14 issue-generic policies.

Similar to measuring stigma, to identify the presence of the 14 policy dimensions in our news datasets we represented the news articles using sentence-level BERT embeddings. Again, per Kwak et al. (2021), we compared embeddings representing the news article and the 14 policy dimensions using cosine similarity. This generated a 14-length vector for each news article, enumerating the presence of the 14 issue-generic policies within the news article.

Validating Issue-generic Policy Extraction Two authors first had an in-depth discussion to reach a common understanding of the 14 policies. Following this, they labeled the presence of issue-generic policies for 10 randomly sampled news articles and agreed on 80% of the sample. The authors then resolved their differences to reach 100% co-author agreement. Finally, they sampled and annotated 150 news articles, spanning *NELA-GT* and *NELA-Local* datasets, to compare hand annotations against machine labels for the policies. Our issue-generic policy extraction system achieved high precision (79%) and recall (74%) against the human-assigned ground truth.

Error analyses: The system mislabeled 8 articles as ‘Constitutionality and jurisprudence’. This happened when there were mentions of justice institutions and law enforcement but no direct references to individuals’ rights/freedom (e.g., “supreme court justice was in question”). The framework also mislabeled 4 articles as ‘Economic’ when there were numerical comparisons with no monetary implications (e.g., “43 times more likely to commit assault”).

Explanatory Model

In response to RQ2, we developed a logistic regression model (*all-factors*), using the default *statsmodels* Python module, to understand the relative roles of differ-

ent factors, namely, (a) regionality, (b) political leaning, (c) veracity of news sources, (d) media communities, and (e) issue-generic policies, in explaining stigmatized VAW framings in news media. Our choice of using logistic regression stems from prior work that adopt this methodology to investigate the association of various independent variables with particular outcomes of interest in understanding online media behaviors (Sivaraman et al. 2023). The `all-factors` model was trained using the five factors as independent variables, with an 80-20 train-test split, to perceive their explanatory power in determining whether an article uses stigmatized framings or not (dependent variable).

Results

Difference in Stigma: Regionality of News Sources

Corresponding to RQ1(a), to analyze differences in stigmatized VAW framings based on regionality of news sources, we compared cosine similarity score distributions (indicating articles' alignment to Approval/Stigma frame) for two clusters of news articles: (1) published by *national* media outlets, i.e., articles within `NELA-GT` and (2) published by *local* media outlets, i.e., articles within `NELA-Local`.

We found that 82,986 (15,121) news articles published by national news outlets skewed towards the stigma (approval) dimension of the Approval/Stigma frame. This shows that a larger proportion of articles by national sources aligned more with the stigma dimension compared to the approval dimension, indicating their VAW discourse to use more stigmatizing linguistic cues than approving or accepting framings. Along similar lines we observed that 51,471 (17,421) articles published by local outlets aligned with the stigma (approval) dimension. Articles from national outlets 'US News', 'CNN' and local outlets 'Atlanta Daily World', 'The Daily Republic' contained VAW-based stigma.

On comparing prevalence of the Approval/Stigma frame across the national and local outlets we found that the average cosine similarity score was -0.0428 (-0.0332) for articles within the national (local) clusters, indicating that local media sources contain less stigmatizing VAW framings compared to national sources. This finding is also captured in the cosine similarity distribution plot shown in Figure 1a. The histogram plot for local media outlets is skewed towards the right, compared to that of national outlets, suggesting that local outlets use more approval-based framings than national. Effect size measurement revealed meaningful differences in score distributions of national and local sources, [Cohen's d]: 0.513. These differences were also statistically significant under Mann-Whitney U-tests (U-stat: 29.903; $p < 0.01$).

Elucidating these quantitative observations further, in the paraphrased excerpts below, we observe the presence of supportive, positive, and progressive language (e.g., "communicate priorities as one voice", "make the experience comfortable") around survivors and structures for both national and local articles falling towards the approval dimension:

"WRAP helps victims of rape and domestic violence to bounce back emotionally, physically and mentally. 'We want to make the experience as comfortable and relaxing as possible [...] the conversation you are having is going to

be one of the most difficult of your life as you open up about the experience,' [...] said." (*National (CBS): approval*)

"This year's legislative session brought a more broad and organized effort to communicate priorities as one voice. [...] also filed HB 1906, dubbed 'Athena's Law,' which would give judges the authority to delay release of those accused of sexual assault." (*Local (Kennebec Journal): approval*)

In contrast, we observe critical and condemning language directed at social processes and inequities for articles aligned with the stigma dimension. For instance, the paraphrased excerpt below blames the systemic processes and culture of a country for the prevalence of sexual assault ("conviction rate in trials is low", "blame increasing number of assaults on a culture of impunity"). Additionally, the other excerpt uses derogatory language to describe sexual assault survivors ("engage in behaviors such as bedwetting").

"When survivors file complaint for sexual assault in [country name] prosecution is very rare, and the conviction rate in trials that do go to court is very low. [...] 88% of respondents said they faced no legal consequences. [...] blame assaults on a culture of impunity." (*National (US News): stigma*)

"A few common signs of sexual assault include: changes in self-care, such as paying less attention to hygiene, or appearance [...] may engage in behaviors, such as thumbsucking or bedwetting." (*Local (Atlanta Daily World): stigma*)

Difference in Stigma: Political Leaning of News

In response to RQ1(b), i.e., to study differences in stigmatized VAW framings based on political leaning of news sources, we looked at cosine similarity distributions for articles grouped by political leaning of outlets analyzed separately within the national (`NELA-GT`) and local (`NELA-Local`) source clusters. Again, > 0 cosine similarity indicates alignment to approval, while < 0 to stigma.

National-left and -right leaning sources On segregating articles based on political leaning of sources within `NELA-GT`, i.e., articles published by left/right leaning sources, we observed that 32,123 (14,371) articles published by national-left leaning outlets aligned with the stigma (approval) dimension (Figure 1b). In addition, 43,874 (7,739) articles published by national-right leaning outlets aligned with stigma (approval). These raw statistics suggest that for both left and right leaning sources within `NELA-GT`, more articles tended to use a stigmatizing framing compared to an approval-based framing. However, the ratio of number of articles leaning more towards approval than to stigma is higher, 0.447, for national-left leaning outlets in comparison to that for national-right outlets (0.176), indicating that national-left sources use lesser stigmatizing linguistic framings than national-right. This is also summarized in the cosine similarity distribution plots shown in Figure 1b where the histogram distribution for national-left sources is shifted towards the right to that of national-right sources. Effect size measurement revealed meaningful differences in score distributions of national-left and -right leaning sources, [Cohen's d]: 0.751. These differences were also statistically significant under Mann-Whitney U-tests

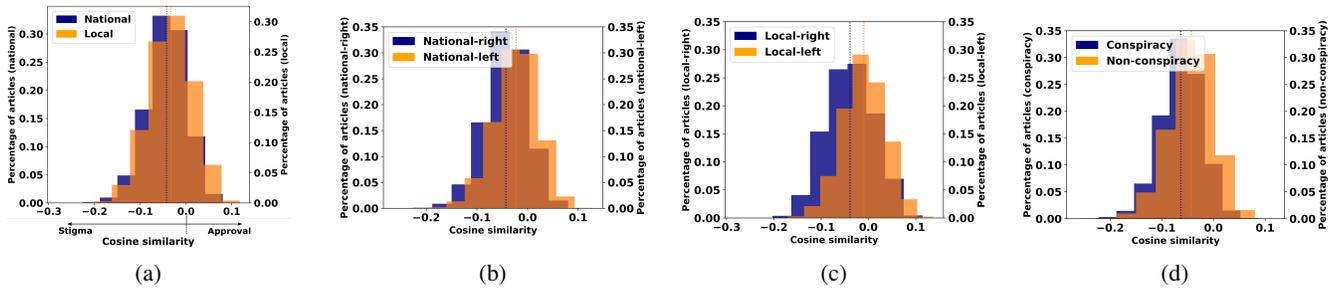


Figure 1: Histogram of cosine similarity scores for articles published by (a) national outlets (left Y-axis) and local outlets (right Y-axis), (b) national-right leaning (left Y-axis) and -left leaning (right Y-axis) outlets, (c) local-right leaning (left Y-axis) and -left leaning (right Y-axis) outlets, and (d) conspiracy (left Y-axis) and non-conspiracy outlets (right Y-axis). (All) Dotted lines represent the average cosine similarity for articles from corresponding groups of news sources.

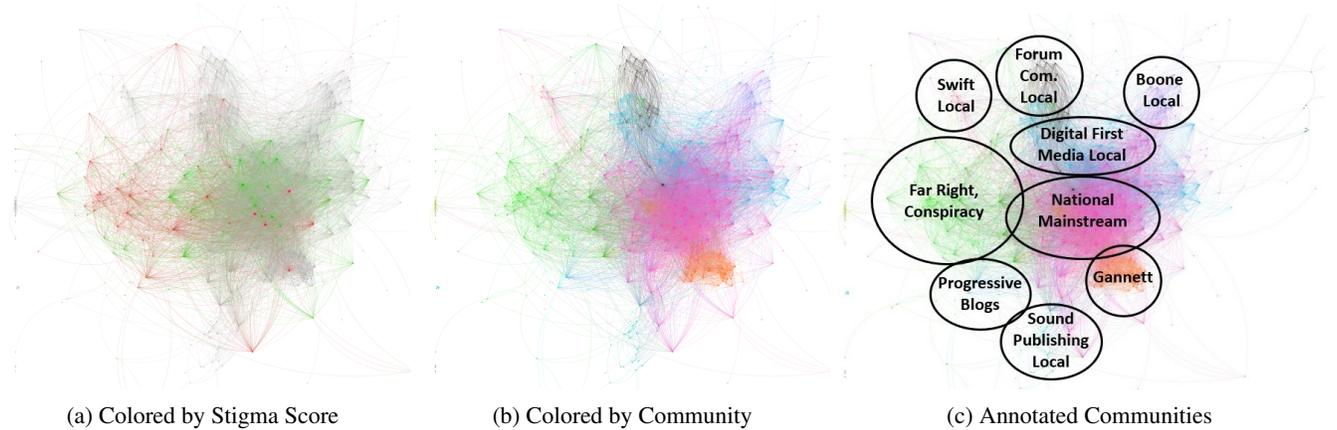


Figure 2: Content sharing networks. In (a), outlets are colored green if they use less stigma framing than the average outlet, red if they use more stigma framing than the average outlet, and grey if they are not included in our VAW sample. In (b), outlets are colored by community membership, determined by modularity. In (c), to better explain the types of media communities captured in the network, we provide high-level annotations of the largest communities.

(U-stat: 15.196; $p < 0.05$). Quotes below from national-left and national-right articles show comparable degrees of stigma – left portrays structural systems with stigma while the right portrays survivors and supporters with stigma:

“News sparked outrage in [...], where allegations of SA often go under reported due to a slow legal system. [...] brutal cases of assault are often dealt with poorly under the country’s justice system.” (*National-left (CNN): stigma*)

“Supporters falsely claim that survivors of sexual assault are denied care [...] advocating for procedures that put women at risk.” (*National-right (RealClearPolitics): stigma*)

Local-left and -right leaning sources Next, looking at cosine similarity scores for articles in left and right leaning sources within NELA-Local (Figure 1c), 24,197 (17,439) articles released by local-left leaning sources aligned with the stigma (approval) dimension of the Approval/Stigma frame. Furthermore, 19,995 (6,176) articles published by local-right sources skewed towards the stigma (approval) dimension. Again, similar to what was observed for NELA-GT, for both left and right leaning sources within NELA-Local, a larger proportion of articles used stigma-

tizing language than positive approval-based framings. On comparing the ratio of number of articles aligning with approval to those with stigma, we observed that local-left sources achieve a higher value (0.721) compared to local-right (0.309), suggesting that left-leaning sources use more approval-based framings than right-leaning sources within NELA-Local. This finding is also featured in the score distribution plots shown in Figure 1c such that the histogram for local-left sources appears towards the right to that of local-right sources. Effect size measurement revealed meaningful and large differences in score distributions of local-left and -right leaning sources, $|Cohen's d|$: 0.812. These differences were also statistically significant under Mann-Whitney U-tests (U-stat: 72.48; $p < 0.01$).

To summarize, within both NELA-GT and NELA-Local datasets, left-leaning sources use more approval-based framings compared to right. This is also indicated by a higher average similarity score for articles from national-left (-0.0228) and local-left (-0.0097) when compared to those from national-right (-0.0427) and local-right (-0.0390).

Difference in Stigma: Veracity of News Sources

Next, for RQ1(c), i.e., to study differences in stigmatized VAW framings based on veracity of news sources, we looked at cosine similarity distributions across two clusters of news articles published by: (1) low-veracity or conspiracy outlets and (2) high-veracity or non-conspiracy outlets (NELA-GT).

Per Figure 1d, 14,760 (2,330) articles from conspiracy news sources (e.g., Breitbart, Infowars, The Gateway Pundit) skewed towards the stigma (approval) dimension, showing that conspiracy sources use more negative stigma-based language than approval-based framings. Similarly, for non-conspiracy news outlets, like NPR, PBS, ABC News, 82,986 (15,121) articles aligned with the stigma (approval) dimension of the Approval/Stigma frame. However, comparing the stigmatized framings across conspiracy and non-conspiracy outlets, we found that non-conspiracy sources use more approval-based framings than conspiratorial sources. E.g., non-conspiracy outlets achieved a higher average cosine similarity, -0.0428 , as compared to conspiracy outlets (-0.0632). In addition, the ratio of number of news articles aligning with approval to those with stigma is higher for non-conspiracy (0.182) than for conspiracy outlets (0.053). This finding is also supported by the cosine similarity score distributions shown in Figure 1d, where the histogram plot for non-conspiracy outlets is shifted towards the right to that of conspiracy outlets. Effect size measurement revealed meaningful and large differences in distributions of mainstream and conspiratorial sources, |Cohen's d |: 0.803. Differences were statistically significant under Mann-Whitney U-tests (U-stat: 26.891; $p < 0.05$).

Difference in Stigma: Media Communities

Media communities captured by the CSN supported the above findings. In Figure 2a, we colored nodes green if they used stigmatized framings less than the average outlet overall, and red if they used stigmatized framings more than the average outlet overall, where the average outlet had a cosine similarity score of -0.044 . As expected, outlets with similar levels of stigmatized language were clustered together.

Then, overlaying this network with community memberships (computed using modularity), we found that the community that used the most stigmatizing language on average is a community of far-right and conspiracy-peddling sources (labeled *Far Right, Conspiracy* in Figure 2c). As a whole, outlets within this community had a lower cosine similarity score (-0.057) than the average outlet (-0.044). In particular, several outlets central to the community used highly stigmatizing language across their articles: The Blaze (-0.104), Free Beacon (-0.075), and Breitbart (-0.053).

Next, when comparing this community to the largest community in the network, which contains mostly U.S. mainstream news outlets (labeled *National Mainstream* in Figure 2c), we found similar results to our previous findings. That is, the U.S. mainstream community used less stigmatizing language than the far right, conspiracy community (-0.045 vs. -0.057) but slightly more than the average outlet overall (-0.045 vs. -0.044). While much of the community uses less stigmatized language than the average outlet (as indicated by the green nodes in Figure 2a), several outlets that

are highly central to the community use more stigmatized language than average (CBS (-0.056) and NPR (-0.056)) and multiple outlets on the periphery of the community also used more stigmatized language than the average outlet (Shareblue (-0.076) and Business Insider (-0.072)).

The media communities also supported the finding that national outlets used more stigmatizing language than local outlets, particularly more than left-leaning local news outlets. E.g., the community with the highest average cosine similarity score (lowest stigma), -0.033 , is a group of local news outlets owned by Digital First Media (labeled *Digital First Media Local* in Figure 2c, e.g., The Denver Post, The Mercury News, Daily Democrat, St. Paul Pioneer Press) that are headquartered in left-leaning counties.

Nonetheless, there are some nuances to this support. For example, despite left-leaning outlets, both nationally (e.g., The New York Times) and locally (e.g., The Columbus Dispatch), producing more articles leaning towards the approval dimension rather than the stigma dimension, a small community of progressive outlets (labeled *Progressive Blogs* in Figure 2c) used stigmatized framings more than the average (-0.051 vs. -0.044). This community included outlets such as The Raw Story (-0.048), and Crooks and Liars (-0.049).

Difference in Stigma: Issue-generic Policies

Per Section “Operationalizing Issue-generic Policies”, for RQ1(e), we generated a 14-length vector representation for each news article to quantify the presence of the 14 dimensions, listed in Table 2. Articles published by national, local, and conspiracy sources were then separately clustered into 14 groups – representing the 14 issue-generic policies. We performed this clustering by assigning an article to the issue-generic policy cluster with which the article aligned the most, i.e., for which it attained the highest numerical value in the 14-length vector. After obtaining these 14 clusters, for each of the three types of news sources (national, local, and conspiracy), we averaged the cosine similarity to the Approval/Stigma frame for all articles in a cluster. This analysis gave the heatmap shown in Figure 3, representing average Approval/Stigma alignment of articles with pronounced presence of an issue-generic policy.

From Figure 3 we see that articles mentioning issue-generic policies ‘Fairness and equality’, ‘External regulation and reputation’, ‘Health and safety’, and ‘Capacity and resources’ use stigmatizing language the most across all three categories of news sources, with conspiratorial news sources demonstrating 26.19% more stigma compared to national for these 4 issue-generic policies, and national outlets with 43.73% more stigma than local ones, also for the same policies. Table 3 lists top uni-, bi-, and tri-grams, based on TF-IDF, in articles belonging to the policy clusters, across the 3 source types. Phrases such as ‘justice system’, ‘cancel culture’, and ‘war crimes’ indicate themes that attract VAW-based stigma. Consider the following two paraphrased excerpts. In the first, instead of providing references to support services or ways to combat the issue, the article sensationalizes the judicial system of the mentioned country. In addition, the next excerpt not only exaggerates the act of violence unnecessarily but uses a visual representation (“forced her”)

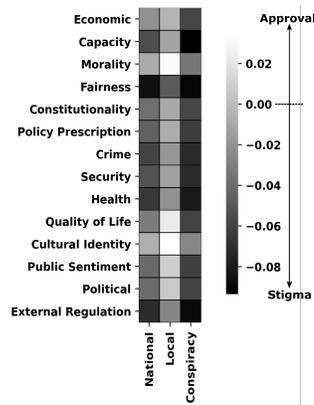


Figure 3: Heatmap showing alignment to Approval/Stigma, for articles with an issue-generic policy (Y-axis); published by national, local or conspiracy outlets (X-axis).

to reinforce stereotypes on female subordination.

“Sexual crimes against women are widespread in [...], brutal cases of rape are often dealt with poorly under the country’s justice system.” (*Capacity and resources: stigma*)

“[...] forced her into the apartment, where he assaulted her. [...] Though the perpetrator had a prior record on property crimes, the judge ruled against punishment using a ‘benefit of the doubt’ argument.” (*Fairness and equality: stigma*)

On the other hand, articles containing issue-generic policies such as ‘Morality’, ‘Cultural identity’, and ‘Quality of life’ tend to mostly use approval-based framings, again, across all three types of news sources, with local outlets having the most approval by 31.14% compared to national and conspiratorial outlets. This finding is evident through the following two exemplars. The first paraphrased article excerpt raises awareness about the increased domestic and child abuse during the COVID-19 pandemic (Piquero et al. 2021), acknowledges its impact on the youth’s well-being, and advocates for appropriate support. Terms in Table 3 like ‘better life’, ‘health care’ also indicate this approval-based usage. The excerpt follows existing media guidelines to responsibly report and represent VAW, without exacerbating stigma, such as by using a “survivor-centered approach” and highlighting that “situations (like COVID-19) can lead to an increase in VAW, but should not be used as an excuse for violence” (UNICEF and UN Women 2020). Similarly, the second paraphrased excerpt not only normalizes the impact of domestic violence across different identities and demographics but also affirms the need for the audience to collectively care for the survivors (“all of us should care”).

“During the pandemic, for many teenagers, home was not always a safe place. [...] having therapists in schools provides a way of supporting youth, particularly female students, who may be experiencing physical abuse in the home during COVID-19.” (*Quality of life: approval*)

“DV is a crime that affects both women and men, of every race, religion, culture, status and income level. [...] all of us should care, and should understand the resources available to those in danger.” (*Cultural identity, Morality: approval*)

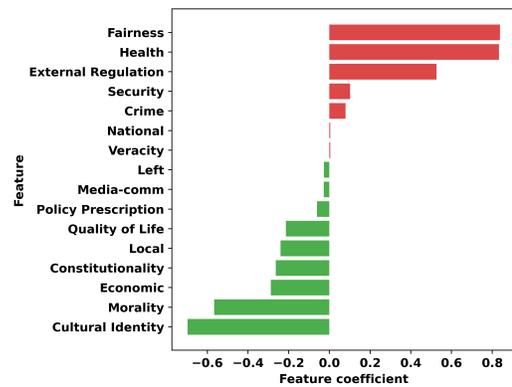


Figure 4: Standardized regression coefficients (β) of independent variables with the highest positive and negative values ($p < 0.01$) in the all-factors model.

Importance of Factors that Explain VAW Stigma

For RQ2, here we present our evaluation of the all-factors logistic regression model, used to explore the relative roles of the five factors in explaining stigma around VAW in news media. To do so, we provide the positive and negative standardized coefficients (β) of the most salient independent variables in Figure 4. Larger positive (negative) β value of an independent variable represents its greater ability in explaining the likelihood of an article containing stigma-based framings (or not). Amongst the five factors, issue-generic policies (or article alignment to the 14 issue-generic dimensions) offer the highest explanation in our dependent variable, i.e., presence/absence of stigma – issue-generic policy dimensions obtain large absolute β values as shown in Figure 4. In addition, independent variables such as veracity, national regionality, and left political leaning are relatively weak in identifying stigma- or approval-based framings since they achieve small absolute β values.

Lastly, variables with the highest importance in identifying stigmatized framings include issue-generic policies like Fairness and equality (β : 0.839), Health and safety (β : 0.834), and External regulation and reputation (β : 0.527), which were also shown to contain pronounced stigmatized framings in response to RQ1(e). In contrast, independent variables with the highest importance in identifying approval-based framings include local regionality (β : -0.240) and issue-generic policies like Morality (β : -0.566) and Cultural identity (β : -0.697), which could be attributed to some level of adherence to media guidelines.

Discussion

Why News Sources Differ in their VAW Framings

Journalists are usually supported by a set of comprehensive guidelines that they are recommended to follow while reporting accounts of VAW (UNICEF and UN Women 2020; UNDP 2022). Most of these resources expand on steps that journalists can take to ensure their reporting is part of the solution to VAW and does not cause additional harm to the survivors reinforcing stigmatizing attitudes and behaviors.

Issue-generic Policy Dimension	Top uni-, bi-, tri-grams based on TF-IDF
Economic	fsa, layoffs, 19 million, lifetime warranty, labor, money allocated, abortion money, compensation pay
Capacity and resources	justice system, reserve court, school staff, abortion resource access, hospital admin, white house
Morality	fair, civil rights, taking care, acted unethically, responsible activists, sexual morality, social ethics
Fairness and equality	against punishment, black lives matter, cancel culture, right wing, resource allocation, racism
Constitutionality and jurisprudence	us citizen, green card, voter id, reform, forbidden, joe biden, 2021 vawa passed, district attorney
Policy prescription and evaluation	abortion ban, we need, president trump, the bill, anti abortion, reauthorizing vawa, 14th amendment
Crime and punishment	violence, sex trafficking, law enforcement, los angeles, johnston county, police said, sheriff office
Security and defense	president trump, brief survey, security upgrade, defense department, prison camps, amendment needs
Health and safety	make sure, health experts, emergency, safety guidelines, precaution, public health, volunteer
Quality of life	mental health, health care, public health, better life, live work, hostile environment, feel like, covid 19
Cultural identity	country built, black lives, citizen, white women, culture wars, meanwhile indians, high school, racism
Public opinion	petition, please sign, public health, rights, pro life, liberal media scream, abortion rights, lives matter
Political	far right, crisis, united states, president trump, right wing, hearing, joe biden, attorney general
External regulation and reputation	united nations, biden administration, white house, war crimes, human rights, rape incest

Table 3: Top uni-, bi-, and tri-grams based on TF-IDF for articles contained in the 14 issue-generic policy clusters.

Still, our findings for RQ1(a-c) revealed pronounced presence of stigma in articles from national, right-leaning, and conspiratorial news outlets. There could be many factors at play behind these observations and we highlight them below.

First, news outlets tend to tailor and optimize their content using audience targeting techniques (Elejalde, Ferres, and Schifanella 2019). Right-leaning news media may be more likely to hold conservative or traditional views on gender roles and sexuality. This can lead to stigmatizing attitudes. The implications of this are particularly acute in local media systems, where communities may only receive local VAW related news from one particular angle. As our analysis indicated, the largest difference in the usage of stigmatized framings was between local outlets embedded in right-leaning counties versus those embedded in left-leaning counties (Figure 1c). Further, conspiratorial outlets, often driven by covert or overt political agendas, may attempt to downplay the incidents of VAW if they conflict with their beliefs or priorities. In fact, such outlets are known to promote conspiracy theories that are not based on facts or evidence, which may result in sharing beliefs that VAW is a hoax. Since they also tend to use fear-mongering tactics to attract viewers or readers, these outlets may portray women as vulnerable and helpless. Importantly, conspiratorial news outlets may not be held accountable to the same standards of accuracy or fairness as mainstream news, resulting in reporting that contains sensationalist language or imagery, biased, exploitative, or lacking in context. Indeed, we found more stigma-based framings in articles from conspiratorial outlets compared to mainstream news (Figure 1d). Such framings can preclude victimized women from reporting incidents for fear of shame, judgment, and minimization.

Next, despite the differences between local left- and right-leaning outlets, local news as a whole used stigmatizing frames less than national news (Figure 1a). This finding expands upon prior work, which found that local newspapers in Utah challenged patriarchal institutions on domestic violence (Bullock 2007). Distance from the local community may drive national outlets' stigmatizing VAW frames – these outlets may be less connected to the communities where the

violence occurred. As a result, they may not have access to local resources or community perspectives that can provide context and nuance to the reporting. Given their scope, they may also tend to focus on high-profile cases of VAW, overlooking or under reporting on the prevalence of VAW in everyday life, which can perpetuate the belief that these incidents are rare or exceptional. In contrast, local news sources may be more likely to report on VAW in a way that is sensitive to the survivors involved. They may have closer relationships with local resources and organizations that provide support and advocacy for survivors, and they may be more attuned to the nuances of the local context.

Our study also found that news articles adopting certain issue-generic policies displayed different levels of stigma towards victims of VAW. Specifically, policy dimensions related to fair and equal treatment and health and safety were found to perpetuate stigma in different ways. News media focusing on the 'Fairness and equality' policy dimension may demonize women who have experienced violence by framing them as liars and manipulators; thus exploit their trauma for political gain. Meanwhile, media focusing on the 'Health and safety' dimension may perpetuate the stereotype that women are most at risk from unknown men and ignore the fact that women are more likely to experience violence from someone they know. This narrative may blame victims for not taking steps to protect themselves. Additionally, when news media frames VAW as a policy issue of capacity and resource allocation, it may portray survivors as burdens on the system and ignore the root causes of VAW, such as systemic sexism and gender inequality. It perpetuates stigma by suggesting that VAW is solely a problem of resource allocation rather than a societal issue requiring systemic change.

Finally, our RQ2 findings suggest that media attributes such as regionality, political leaning, and veracity alone are not enough to identify stigma in news articles about VAW. The relationship between these attributes and stigma is more complex and context-dependent. For instance, left-leaning articles showed stigma when criticizing policies or actions that contribute to VAW, while right-leaning articles stigmatized survivors, advocates, and political rivals. Thus, these

attributes alone are not enough to capture the underlying context. In contrast, systematically categorizing news outlets based on media communities and aligning articles to issue-generic policy dimensions offered a better explanation for the presence/absence of stigma, as they can capture contextualized information on the outlet's position in the news media landscape and the political themes used in the articles.

Survivor-Centered Safe Reporting Guidelines

Based on our findings noting the presence of stigma in VAW news reporting in different types of outlets, we suggest the following considerations for VAW safe reporting guidelines.

Use trauma-informed language and avoid victim-blaming: Articles focusing on the 'Fairness and equality' policy dimension perpetuate stigma by portraying survivors as manipulators (RQ1(e)). News outlets should avoid language that downplays the severity of the problem. Aside from prioritizing survivors' voices and experiences, trauma-informed language should be used when reporting on VAW, which means avoiding language and narratives that suggest survivors are responsible for their own victimization or that minimizes the severity of the harm they have experienced.

Hold perpetrators accountable: The 'Capacity and resources' policy dimensions propagates stigma by justifying that VAW merely results from incapacity of existing systems (RQ1(e)). News outlets should avoid language that excuses or justifies perpetrators' behavior and instead focus on the harm caused to survivors and consequences of their actions. This includes providing information on legal consequences.

Provide context and educate readers: We found that paraphrased excerpts contextualizing VAW in RQ1(e), e.g., by advocating for students experiencing increased domestic abuse during COVID-19 or normalizing domestic violence across different cultural identities, contained approval-based framings. News outlets should provide context for incidents of VAW and educate readers on the broader societal issues that contribute to the problem. This can include information on systemic sexism, gender inequality, the impact of cultural norms and beliefs on VAW, and local/community information where a particular VAW incident may have occurred.

Avoid perpetuating conspiracy theories: News outlets should adopt evidence-based reporting and thus avoid perpetuating false mis- and disinformation that undermine the credibility of survivor accounts or that survivors could have prevented VAW by simply taking different actions.

Avoid perpetuating political biases: Our results for RQ1(b) show that right-leaning sources use more stigmatized framings than their left counterparts. Right-leaning outlets may hold conservative views on VAW to appeal to their target audiences (Elejalde, Ferres, and Schifanella 2019). News outlets should avoid using VAW as a way to advance underlying political agendas that undermine the credibility of survivor accounts or that suggest VAW as a strategy to exploit existing institutions for personal advantage.

Conclusion, Limitations, and Future Work

In this paper, we sought to understand the presence of stigmatized framings in news media reports of VAW. Using a BERT-based framework, we quantified approval- and

stigma-based language in news articles from 396 national, local, and conspiratorial news outlets within the U.S., covering a broad news media landscape. Articles from local, left-leaning, and high-veracity sources adopted more acceptable and approval-based framings, in comparison to their counterparts. In addition, articles with mentions of specific issue-generic policies such as those around the distribution of laws and rewards, those speaking to the capacity of existing societal infrastructures, and those with references to the public safety threats of VAW had pronounced stigma, and it also showed the best ability to explain the presence of stigma.

We note some limitations which provide excellent directions for future research. We did not examine how specific stigmatized VAW framings in news influence consumers. Next, VAW saw an increase in reported incidents during the COVID-19 pandemic (UN Women 2021). Researchers could study temporal trends of stigma to quantify the impact of unforeseen crises. Finally, this work was executed under the premise that framing is a journalistic strategy influencing partisan action. Future work might consider evolution of stigma around the implementation of national VAW policy. All code and data are available at <https://github.com/mittalshravika/VAW-Framing>.

Acknowledgements

Mittal, Foriest, and De Choudhury were partly supported through NIH grants P50MH115838 and R01MH117172. Foriest was also supported through an ORISE Fellowship from the U.S. Centers for Disease Control and Prevention.

References

- An, J.; Kwak, H.; and Ahn, Y.-Y. 2018. SemAxis: A Lightweight Framework to Characterize Domain-Specific Word Semantics Beyond Sentiment.
- Bahamonde, J.; Bollen, J.; Elejalde, E.; Ferres, L.; and Poblete, B. 2018. Power structure in Chilean news media. *PLOS ONE*.
- Bandura, A. 2001. Social cognitive theory of mass communication. *Media psychology*.
- Barnett, M. D.; and Hilz, E. N. 2018. The Psychology of the Politics of Rape: Political Ideology, Moral Foundations, and Attitudes Toward Rape. *Violence Against Women*.
- Berns, N. 2001. Degendering the problem and gendering the blame: Political discourse on women and violence. *Gender & society*.
- Berns, N. 2004. *Framing the victim: Domestic violence, media, and social problems*. Aldine Transaction.
- Boydston, A. E.; Gross, J. H.; Resnik, P.; and Smith, N. A. 2013. Identifying media frames and frame dynamics within and across policy issues. *New Directions in Analyzing Text as Data Workshop*.
- Brady, W. J.; Wills, J. A.; Jost, J. T.; Tucker, J. A.; and Van Bavel, J. J. 2017. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*.
- Bullock, C. F. 2007. Framing Domestic Violence Fatalities: Coverage by Utah Newspapers. *Women's Studies in Communication*.
- Card, D.; Boydston, A. E.; Gross, J. H.; Resnik, P.; and Smith, N. A. 2015. The Media Frames Corpus: Annotations of Frames Across Issues. In *Proc. ACL*.
- Carline, A.; and Easteal, P. 2014. *Shades of grey – Domestic and sexual violence against women: Law reform and society*.

- Chagnon, N. 2020. It's a Problem of Culture (for Them): Orientalist Framing in News on Violence Against Women. *Race & Justice*.
- Chancellor, S.; Baumer, E. P.; and De Choudhury, M. 2019. Who is the "human" in human-centered machine learning: The case of predicting mental health from social media. *Proc. CSCW*.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proc. NAACL-HLT*.
- Easteal, P.; Holland, K.; and Judd, K. 2015. Enduring themes and silences in media portrayals of violence against women. In *Women's Studies International Forum*.
- Eaves; and Equality Now. 2012. End Violence Against Women.
- Elejalde, E.; Ferres, L.; and Schifanella, R. 2019. Understanding news outlets' audience-targeting patterns. *EPJ Data Science*.
- ElSherief, M.; Belding, E.; and Nguyen, D. 2017. #NotOkay: Understanding Gender-Based Violence in Social Media. *ICWSM*.
- Entman, R. M. 1993. Framing: Towards Clarification of Fractured Paradigm (Journal of Communication). *PP/S*.
- FORCE11. 2020. The FAIR Data principles. <https://force11.org/info/the-fair-data-principles/>.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*.
- Goffman, E. 2014. Stigma. In *Classic and Contemporary Readings in Sociology*.
- Gruppi, M.; Smeros, P.; Adali, S.; Castillo, C.; and Aberer, K. 2022. SciLander: Mapping the Scientific News Landscape. *arXiv:2205.07970*.
- Horne, B. D.; Gruppi, M.; Joseph, K.; Green, J.; Wihbey, J. P.; and Adali, S. 2022. NELA-Local: A Dataset of U.S. Local News Articles for the Study of County-Level News Ecosystems. *ICWSM*.
- Horne, B. D.; Nørregaard, J.; and Adali, S. 2019. Different spirals of sameness: A study of content sharing in mainstream and alternative media. *ICWSM*.
- Jina, R.; and Thomas, L. S. 2013. Health consequences of sexual violence against women. *Best Practice & Research Clinical Obstetrics & Gynaecology*.
- Kwak, H.; An, J.; Jing, E.; and Ahn, Y.-Y. 2021. FrameAxis: characterizing microframe bias and intensity with word embedding. *PeerJ Computer Science*.
- Larson, S. R. 2018. Survivors, Liars, and Unfit Minds: Rhetorical Impossibility and Rape Trauma Disclosure. *Hypatia*.
- Levendusky, M. S. 2022. How does local TV news change viewers' attitudes? The case of Sinclair broadcasting. *Political Comm*.
- Lindsay-Brisbin, J.; DePrince, A. P.; and Welton-Mitchell, C. 2014. Missed Opportunities: Newspaper Reports of Domestic Violence. *Journal of Aggression, Maltreatment & Trauma*.
- Link, B. G.; and Phelan, J. C. 2001. Conceptualizing Stigma. *Annual Review of Sociology*.
- Lira, C. P. L.; and Toledo, P. 2014. The Media and Gender-Based Murders of Women: Notes on the Cases in Europe and Latin America. *Heinrich Boll Stiftung*.
- Mendelsohn, J.; Budak, C.; and Jurgens, D. 2021. Modeling Framing in Immigration Discourse on Social Media. In *NAACL-HLT*.
- Meyers, M. 1996. *News coverage of violence against women: Engendering blame*. Sage Publications.
- Meyers, M. 2004. African American women and violence: gender, race, and class in the news. *Critical Studies in Media Comm*.
- Mittal, S.; and De Choudhury, M. 2023. Moral Framing of Mental Health Discourse and Its Relationship to Stigma: A Comparison of Social Media and News. In *Proc. CHI*.
- Nakov, P.; and Da San Martino, G. 2021. Fake News, Disinformation, Propaganda, Media Bias, and Flattening the Curve of the COVID-19 Infodemic. In *Proc. ACM SIGKDD*.
- Nettleton, P. H. 2011. Domestic Violence in Men's and Women's Magazines: Women Are Guilty of Choosing the Wrong Men, Men Are Not Guilty of Hitting Women. *Women's Studies in Comm*.
- Noetzel, S.; Gentile, M. F. M.; Lowery, G.; Zemanova, S.; Lecheler, S.; and Peter, C. 2022. Social campaigns to social change? Sexual violence framing in U.S. news before and after metoo. *Journalism*.
- Nørregaard, J.; Horne, B. D.; and Adali, S. 2019. NELA-GT-2018: A Large Multi-Labelled News Dataset for the Study of Misinformation in News Articles. In *ICWSM*.
- Oliver, M. 1990. *The Politics of Disablement*. Macmillan Ed.
- Ortiz, R. R.; and Smith, A. M. 2022. A Social Identity Threat Perspective on Why Partisans May Engage in Greater Victim Blaming and Sexual Assault Myth Acceptance in the #MeToo Era. *VAW*.
- Owusu-Addo, E.; Owusu-Addo, S. B.; Antoh, E. F.; Sarpong, Y. A.; Obeng-Okrah, K.; and Annan, G. K. 2018. Ghanaian media coverage of violence against women and girls: implications for health promotion. *BMC Women's Health*.
- P. Aires, V.; G. Nakamura, F.; and F. Nakamura, E. 2019. A Link-Based Approach to Detect Media Bias in News Websites. *WWW*.
- Pickup, F.; Williams, S.; and Sweetman, C. 2001. *Ending violence against women: A challenge for development and humanitarian work*. Oxfam.
- Piquero, A. R.; Jennings, W. G.; Jemison, E.; Kaukinen, C.; and Knaul, F. M. 2021. Domestic violence during the COVID-19 pandemic - Evidence from a systematic review and meta-analysis. *Journal of Criminal Justice*.
- Plough, A. 2000. Domestic violence in the African American community: A local public health perspective. In *Forum Proc.: Health and mental health consequences of domestic violence in the African American community*.
- Purohit, H.; Banerjee, T.; Hampton, A.; Shalin, V. L.; Bhandutia, N.; and Sheth, A. 2016. Gender-based violence in 140 characters or fewer: A #BigData case study of Twitter. *First Monday*.
- Sacks, M.; Ackerman, A. R.; and Shlosberg, A. 2018. Rape Myths in the Media: A Content Analysis of Local Newspaper Reporting in the United States. *Deviant Behavior*.
- Schaffner, B. F.; and Sellers, P. J. 2010. Winning with words. *The Origins and Impact of Political Framing, New York-London*.
- Sivaraman, V.; Bukowski, L. A.; Levin, J.; Kahn, J. M.; and Perer, A. 2023. Ignore, Trust, or Negotiate: Understanding Clinician Acceptance of AI-Based Treatment Recommendations in Health Care. In *Proc. CHI*.
- Starbird, K.; Arif, A.; Wilson, T.; Van Koevering, K.; Yefimova, K.; and Scarnecchia, D. 2018. Ecosystem or echo-system? Exploring content sharing across alternative media domains. *ICWSM*.
- UN Women. 2019. The big conversation: Handbook to address violence against women in and through the media.
- UN Women. 2021. Measuring the Shadow Pandemic: Violence Against Women During COVID-19.
- UNDP. 2022. Guidelines on Media Reporting on VAW.
- UNICEF; and UN Women. 2020. Responsible representation and reporting of violence against women and violence against children; Guidelines for Media Professionals.
- United Nations. 1993. Violence against women.

Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes, as discussed in Sections “Introduction” and “Discussion,” our work advances the space of computational social science in general and computational journalism in particular, in the context of VAW. We offer survivor-centered safe reporting guidelines as a contribution that follows from our findings. We provide considerations to use trauma-informed language, state contextual information, and avoid the spread of political biases or conspiracy theories, without violating any social contracts.**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Yes. We have carefully reviewed that the claims made in abstract and Introduction accurately reflect our contributions.**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes, please refer to Section “Methods.” We justify the appropriateness of our methods by referring to prior works that use them for similar tasks/objectives. We also qualitatively validate the BERT-based frameworks to extract approval/stigma frame and issue-generic policy dimensions in our context (refer to Sections “Validating Approval/Stigma Frame Extraction” and “Validating Issue-generic Policy Extraction”). Further, we contextualize our findings on VAW framings in the existing public health and violence prevention literatures.**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes. Please refer to Section “Data” and Table 1 where we describe the characteristics of our data such as, number of articles, publication sources, and meta information. We also elaborate on our data cleaning process in Section “Filtering and Data Cleaning”.**
- (e) Did you describe the limitations of your work? **Yes, we discuss limitations of our work in Section “Conclusion, Limitations, and Future Work.”**
- (f) Did you discuss any potential negative societal impacts of your work? **No, because our work does not promote any societal violations. For instance, we do not build prediction frameworks to infer sensitive attributes. In fact, our findings provide ways to combat existing negative perceptions on VAW in news media (refer to Section “Survivor-Centered Safe Reporting Guidelines” under “Discussion”).**
- (g) Did you discuss any potential misuse of your work? **No. We do not curate/release any resources that promote harm or bear potential for misuse. We perform an observational study of already published datasets (refer to Section “Data”) that are available for public use.**

- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes, refer to Section “Broader Perspectives, Ethics, and Competing Interests.” We followed best practices in social computing research and worked with deidentified data. We also provide carefully paraphrased excerpts in the paper to reduce traceability.**
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes, we carefully read the ethics guidelines and ensured that the paper conforms to them. For instance, in Section “Broader Perspectives, Ethics, and Competing Interests,” we elaborate on the societal impact of our findings, data anonymization, and responsible release of paraphrased excerpts.**

2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? **NA**
- (b) Have you provided justifications for all theoretical results? **NA**
- (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
- (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**
- (e) Did you address potential biases or limitations in your theoretical framework? **NA**
- (f) Have you related your theoretical results to the existing literature in social science? **NA**
- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA**

3. Additionally, if you are including theoretical proofs...

- (a) Did you state the full set of assumptions of all theoretical results? **NA**
- (b) Did you include complete proofs of all theoretical results? **NA**

4. Additionally, if you ran machine learning experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes. We provide the URL to an anonymized repository containing all our code and instructions to reproduce the results in Section “Conclusion”. We used publicly available news datasets (refer to Section “Data” for their citations).**
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes, in Section “Explanatory Model” we specify training details such as data split and Python package reference for the logistic regression model implementation.**

- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA.
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? No, because our work was not compute intensive and we did not use any external storage/compute resources.
- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? Yes. Please refer to Section “Methods” where we relate our proposed evaluation procedures to the research questions. Also, refer to Section “Discussion” in which we justify that the evaluation/findings are appropriate with respect to our claims, research goals, and prior work.
- (f) Do you discuss what is “the cost” of misclassification and fault (in)tolerance? Yes. Although we do not perform any classification here, in the paper, we did use BERT-based representation learning frameworks to automatically extract approval/stigma frame characterization and the 14 issue-generic policy dimensions. In Sections “Measuring Stigma” and “Operationalizing Issue-generic Policies” under “Methods”, we provide a validation of the representation learning frameworks via qualitative evaluation and error analyses on a small sample of our dataset to shine a light on cases that the representation learning framework missed.
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
- (a) If your work uses existing assets, did you cite the creators? Yes, refer to Section “Data” where we cite the creators of NELA-GT and NELA-Local datasets.
- (b) Did you mention the license of the assets? Yes, we mention the license of the datasets used in our research in Section “Data.” As specified, the NELA-GT and NELA-Local datasets are available for public use, licensed under CC0 1.0 Universal Public Domain Dedication.
- (c) Did you include any new assets in the supplemental material or as a URL? NA
- (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? No, because we utilized publicly available datasets licensed under CC0 1.0 Universal Public Domain Dedication (refer to Section “Data”). Under this license “[one] can copy, modify, distribute and perform the work, even for commercial purposes, all without asking permission².” Analyzing this retrospective data did not constitute human subjects research and thus informed consent was not required from the authors of the news articles. However, we followed best practices, worked with deidentified data and refrained from revealing any identifiable data to avoid potential harm to those referenced in the news datasets (see Section “Broader Perspectives, Ethics, and Competing Interests”).
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, we discuss best practices followed in our work in Section “Broader Perspectives, Ethics, and Competing Interests”. These include working with deidentified data and paraphrasing quotes throughout the paper to reduce traceability. We also provide a content warning in the same section.
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? NA
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? NA
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
- (a) Did you include the full text of instructions given to participants and screenshots? NA
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? No. We utilize publicly available news articles (refer to Section “Data”) without any interaction with the authors of those articles or the people referenced in them. As an observational study of retrospectively gathered data, our research did not qualify as human subjects research, per our Institutional Review Board guidelines. This IRB approval was not required. However, ethical considerations extend beyond IRBs, when it comes to web research. Therefore, we used deidentified data and provided paraphrased quotes in the paper to avoid traceability and potential harm to those mentioned in our dataset (refer to Section “Broader Perspectives, Ethics, and Competing Interests”).
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
- (d) Did you discuss how data is stored, shared, and deidentified? Yes. As described in Section “Data”, we make use of already published publicly available datasets. Further, as mentioned in Section “Broader Perspectives, Ethics, and Competing Interests”, we follow best practices proposed by Chancellor, Baumer, and De Choudhury (2019) to work with deidentified data, to use secure machines for our analyses with only authorized access to the paper’s authors granted through the principle of least privilege, and to avoid sharing personally identifiable data in any form.

²<https://creativecommons.org/publicdomain/zero/1.0/>