How Selective Reporting Shapes Inferences about Conflict

Yuri M. Zhukov* and Matthew A. Baum†

October 7, 2016

Abstract

By systematically under- or over-reporting violence by different actors, media organizations convey potentially contradictory information about how a conflict is likely to unfold, and whether outside intervention is necessary to stop it. These reporting biases affect not only statistical inference, but also public knowledge and policy preferences. Using new event data on the ongoing armed conflict in Eastern Ukraine, we perform parallel analyses of data from Ukrainian, rebel, Russian and third party sources. We show that actor-specific reporting bias can yield estimates with vastly different implications for conflict resolution: Ukrainian sources predict frequent unilateral escalation by rebels, pro-Russian rebel sources predict unilateral escalation by government troops, while outside sources predict that transgressions by either side should be quite rare. Experimental evidence suggests that news consumers tend to support intervention against whichever side is shown to be committing the violence. We argue that these kinds of reporting biases can potentially make conflicts more difficult to resolve – hardening attitudes against negotiated settlement, and in favor of military action.

*zhukov@umich.edu, Assistant Professor of Political Science, University of Michigan

†matthew_baum@harvard.edu, Marvin Kalb Professor of Global Communications, Harvard Kennedy School

How we respond to a civil conflict depends on what we know about it. That, in turn, depends on where we get our information. Not every event is observable, and not every observed event is publicly reported. Information providers diverge in the events and actors that attract their attention. One source may focus disproportionately on violence by rebels, another may emphasize government operations, while a third may not attribute violence to any armed group at all. Selective reporting may happen for commercial or partisan reasons, or because the government controls the press and requires it. As a result, different sources offer different perspectives on a conflict, and how violence begins, perpetuates and stops. This variation constitutes reporting bias – the systematic under- or over-reporting of events, or particular aspects of events.

Conflict scholars have sought to explain the directions and magnitudes of these biases (Davenport and Ball, 2002, Davenport, 2009), including how the intensity (Weidmann, 2014), location and timing of violence (Hammond and Weidmann, 2014) affect its visibility in the press, and its inclusion in social science datasets (Eck, 2012). These studies have uncovered systematic differences between media- and government-generated conflict data (Weidmann, 2014), between different types of media sources (Earl et al., 2004), and countries of publication (Drakos and Gofas, 2006, Baum and Zhukov, 2015). This research has highlighted the problems reporting bias can produce for statistical inference, but has mostly overlooked its effect on public opinion.

Reporting bias is more than a statistical nuisance. Open sources – like news articles, social media posts, and press briefings – produce the bulk of what the public knows about armed conflict. If it introduces systematic bias into the public record, selective reporting can not only contaminate the datasets upon which social scientists depend, but potentially skew the policy preferences of news consumers, and manipulate public opinion about the actors involved. Under the right conditions, warring parties can 'weaponize' reporting bias into a form of information warfare – using the media to shape the preferences of an intended audience to their advantage (Libicki, 1995).

Research on the effects of selective exposure to partisan media in the United States (Stroud, 2011, Arceneaux, Johnson and Murphy, 2012, Iyengar and Hahn, 2009) shows that one-sided arguments can drive opinions apart and make compromise more difficult. This tendency should be especially powerful in coverage of war, where consumers typically have little direct personal knowledge beyond what they consume in the media, and the international press may have limited direct access to the conflict zone (Baum and Groeling, 2009, DellaVigna et al., 2011). Despite this evidence, conflict research has so far failed to adequately measure the relative magnitude of reporting bias, how it shapes public understanding of conflicts and resulting support for, or opposition to, taking action to end them.

This study makes several unique contributions to the emerging literature on reporting bias during war. First, while previous studies have focused on the aggregate level (i.e. reporting bias in overall coverage of a conflict) we focus on *actor-specific reporting bias* – that is, systematic differences in media coverage of rebels vs. government forces. It is relatively straightforward to identify a unique conflict event: actors, casualties, locations and dates. Yet sources differ in which of these 'hard facts' they report. For instance, a news outlet may describe violence by actor A simply as 'violence,' while describing violence by actor B as 'violence by actor B.' News consumers may then, perhaps wrongly, perceive actor B as more responsible for violence than actor A.

Second, in light of this potential for misperception, we explore the implications of these biases not only for statistical inference, but also for public knowledge and public policy. Differential coverage of rebel and government attacks by the local press reveals information about how violence by one actor affects violence by the other, which side is more likely to cooperate, which side is more likely to escalate, and which is systematically more restrained. This information, in turn, shapes expectations both internationally and locally – given that local reporting may constitute the primary source material for the international press corps – about how a conflict is likely to unfold, and whether and what type of outside intervention is necessary or appropriate to stop it.

Using new event data from the ongoing armed conflict in Eastern Ukraine, we perform parallel analyses of media-generated event data from pro-government, pro-rebel and third party sources, to examine how reporting bias affects the empirical study of armed conflict. We investigate the extent to which different sources suggest different patterns of strategic interaction between warring sides, and advance different conclusions about the causes, location and timing of violence.

The Ukrainian conflict presents an opportune test case for several reasons. Due to its location in an economically developed and densely populated part of Eastern Europe, the conflict has received extensive coverage in local and foreign press, producing an abundant supply of event data. It is also a conflict where reporting biases are likely to have significant effects on public knowledge. Political authorities in Ukraine, Russia and rebel-held territories have imposed tight restrictions on news coverage, while limiting alternative sources of information for media consumers. As a consequence, consumers both within and outside the region are disproportionately dependent on local press reports for information about the conflict.[1] Given the highly politicized nature of the war's coverage, scholarly efforts to explain its causes and predict its future course

[1] For instance, between 18 and 39 percent of the 11,040 BBC News stories on the conflict in Eastern Ukraine between April 1, 2014 and August 16, 2016 cited local news sources. BBC News cited one such outlet – the Donetsk News Agency (DAN), which is the official mouthpiece of pro-Russian rebel authorities in Eastern Ukraine – in 406 stories since 2014, or 7.8 percent of all BBC stories on the conflict. These estimates are based on Lexis-Nexis search queries of BBC news transcripts that mentioned at least one of the 15 local news sources we used in our dataset (see below).

will depend on our ability to understand and account for these biases.

We find that actor-specific reporting bias can profoundly affect both statistical inference and public opinion. According to data from Ukrainian sources, rebels are more likely than the government to unilaterally escalate violence. According to rebel sources, the opposite is true. Both Ukrainian and rebel sources predict more violence in equilibrium than do Russian and outside, third-party sources – like Wikipedia and the Organization for Security and Cooperation in Europe (OSCE). Each perspective has its own implications for how different actors behave in war, the need for third-party intervention, and whether intervention should be neutral or one-sided.

To investigate the effect of reporting bias on policy preferences, we conduct a survey experiment, in which subjects read otherwise identical news stories, featuring different actors and tactics. We find that respondents tend to support intervention against whichever side is shown to be committing the violence. In addition to confounding the statistical analysis of conflict, these results suggest, reporting bias can mobilize support for and opposition to specific armed groups.

The broader implication of our research is that reporting biases could potentially make conflicts more difficult to resolve. For audiences inside a conflict zone, selective exposure and systematic over-reporting of unilateral violence by the 'other side' is likely to harden attitudes against negotiated settlement. For audiences outside a conflict zone, the reiteration of these biases in the foreign press can increase support for military intervention and escalation.

Our findings contribute to political science and communications research on media bias (Davenport and Ball, 2002, Davenport, 2009, Weidmann, 2016), and particularly to the growing literature on information and communications technology (ICT) and violence (Dafoe and Lyall, 2015). Due to the pervasive nature of reporting bias, these findings potentially apply to all empirical conflict research that relies on event data – whether the sources of the data are media firms (Raleigh et al., 2010), NGOs (Lyall, 2010) or government archives (Berman, Shapiro and Felter, 2011).

This paper proceeds as follows. We begin with an overview of recent research on media bias in the study of armed conflict. We then offer background on the Ukrainian case, and summarize main differences in coverage between Ukrainian, rebel, Russian, and international information providers. Next, we consider the consequences of these biases for data analysis and theory testing. Following this discussion, we consider the consequences for public opinion, using a survey experiment. The final section summarizes our results and conclusions.

Reporting bias and the study of armed conflict

Information providers differ in how they describe an event, and whether they choose to report it at all. Due to the growing reliance of armed conflict research on event data, these differences in reporting have never been more important to the study of violence than they are today. Recent years have seen several notable efforts by scholars of armed conflict to identify the sources of reporting bias, and their consequences for statistical inference.

Causes of reporting bias

For an event to become news, someone must observe and report it, and an information provider (e.g. media firm, government agency, or non-governmental organization) must publish the 'hard facts' (i.e. actors, casualties, location, date). For news to become data, social scientists need to detect the event report, classify it into a distinct category (e.g. rebel attack, government operation), and convert it into a format suitable for more rigorous analysis. Although selection issues abound in both processes, this project's focus is on the first component – why some events become news but others do not. In particular, the project addresses the 'whodunit' problem: why information providers report events perpetrated by some actors more than others, and how the resulting reporting biases shape what citizens and scholars know about conflict.

One of the most basic sources of reporting bias is lack of information: not all events are equally visible to observers. Events in densely populated urban areas tend to have more eyewitnesses than events in rural areas (Danzger, 1975). The likelihood that eyewitnesses report the observed event may depend on the proximity of event locations to reporting agencies (Moeller, 1999, Gans, 1980, Davenport, 2009), or the availability of communications infrastructure, like cell phone towers (Weidmann, 2016). Some event locations – like those with ongoing battles – may simply be too dangerous for reporters to access (Weidmann, 2014).

Once information providers learn of an event, they decide whether to report it – internally or publicly. Here, the incentives of reporting agencies vary greatly. Profit-oriented media firms tend to publish information that maximizes their audience. 'Newsworthy' events tend to be large-scale (Woolley, 2000), rare (Snyder and Kelly, 1977), new (Davenport and Stam, 2006), located in close proximity to an outlet's home bureau (Morton and Warren, 1992, Rosengren, 1974), or otherwise salient to the intended audience (Galtung and Ruge, 1965). Journalists and media consumers also tend to lose interest in a conflict over time, with 'coverage fatigue' generating a secular downward trend in the volume of reporting (Davenport and Stam, 2006, Baum and Groeling, 2010).

Where the opportunity costs of event coverage are high, as in print journalism or other media with limited space or time to feature competing stories (Snyder and Kelly, 1977, Davenport and Ball, 2002), the relative 'newsworthiness' of an event is a far stronger predictor of coverage than it is for less physically constrained digital media, like newswires, blogs or social media platforms (Wu, 1998, Shoemaker and Cohen, 2012). These market incentives may compound or offset other potential sources of media bias, like ownership structure (Djankov et al., 2001, Gehlbach and Sonin, 2014) and ideology (Davenport, 2009).

Ironically, another newsworthiness criteria, the norm of balanced reporting (Baum and Groeling, 2009) – that is, including the perspectives of both sides – may ultimately be at least as consequential for public opinion as any of these other factors. Balance implies neutrality. Neutrality in a conflict where one side bears the bulk of responsibility, in turn, may be quite different from "truth" or "objectivity." Borrowing an example from American politics, much of the mainstream American media in the mid-2000s pursued a policy of balance, or neutrality, in its coverage of climate change. When a scientist appeared on a news outlet, like CNN, discussing the scientific evidence in support of human caused climate change, the network would feature a climate skeptic arguing the other side. By treating the views of the skeptic, who represented a small fraction of scientific opinion, as of equal consequence, the network created a false equivalence between the two, making it difficult for viewers to understand which side represented the dominant scientific view (Mayer, 2012). In the context of civil conflict, if a media outlet provides equal time to the perspectives of both sides regarding some violent act or fails to attribute blame, even when one side is primarily responsible, consumers will lack the information they need to appropriately attribute responsibility. This, in turn, may depress or misdirect support for external intervention.

Government and NGO sources face somewhat different, yet also in some ways overlapping incentives. Government records are not constrained by market pressures, and tend to report a higher proportion of observed events than media sources (McCarthy, McPhail and Smith, 1996, Weidmann, 2014). However, the specific mission of the government agency (e.g. internal vs. public reporting), secrecy, and lag time to archival declassification can still limit the scope of this reporting. NGOs face similar problems of specialization – focusing, for instance, on particularly egregious human rights violations, rather than the day-by-day dynamics of armed conflict (Davenport and Ball, 2002).

Beyond source-specific variation in coverage, recent research has highlighted the importance of the political environment in which information providers are based. The extent to which media firms are able to act in accordance with 'newsworthiness' considerations depends on the level of press freedom in their media market (Baum and Zhukov, 2015). Even where they do not directly

Table 1: **Types of reporting bias**.

Type of bias	Example	Causes
1. Location-specific	'report violence in location A, not B'	reporters lack access to B; more witnesses in A ('urban bias'); location A is more 'newsworthy' (e.g. capital city, holy site)
2. Time-specific	'report violence at time A, not B'	A is earlier in conflict ('coverage fatigue'); fewer competing stories at time A ('slow news day'); time A is more 'newsworthy' (holiday, symbolic date)
3. Casualty-specific	'report high-casualty event A, not low-casualty event B'	A has more witnesses; deadlier events more 'newsworthy'('bad news bias')
4. Actor-specific (✓)	'report violence by actor A, not B'	political bias/pressure to focus on A; norms of balanced reporting; lack of access to B

own the media, ruling regimes can impose regulations on what media can and cannot report (Whitten-Woodring and James, 2012) or create norms of self-censorship (Djankov et al., 2001), producing cross-national variation in coverage of certain categories of events (Drakos and Gofas, 2006). Even democratic regimes may impose wartime restrictions on coverage of sensitive topics, particularly those that may compromise ongoing operations or discredit government policy (Sweeney, 2001, Norris, Kern and Just, 2003, Allan and Zelizer, 2004, Hallin, 1989).

Any one of these potential sources of bias may affect the relative likelihood that government or rebel violence will receive coverage. Table 1 summarizes the types of reporting bias most common to conflict event data, and their most widely-cited causes. Although existing research has examined reporting biases with respect to three of the four 'hard facts' of conflict events – casualties (Gohdes and Price, 2012), location and timing (Hammond and Weidmann, 2014, Weidmann, 2014) – actor-specific reporting bias has, with a handful of exceptions (Davenport 2009, Baum and Zhukov 2015), mostly eluded rigorous study.

Consequences of reporting bias

If violent events by some actors are more likely to receive coverage than violence by others, what impact will these biases have on public knowledge and opinion, and on the empirical study of conflict? Research on the effects of reporting bias is less voluminous than that on its causes, but several findings have emerged.

The impact of reporting bias on statistical inference depends on two primary considerations: whether the direction and magnitude of the bias is correlated with the explanatory variable of theoretical interest, and whether the bias is common to all sources. If reporting bias is uncorrelated with the explanatory variable (e.g. random disruptions in a communication network), then selective reporting represents measurement error rather than selection bias (Weidmann, 2014). The potentially large number of false negatives may favor models that under-predict levels of violence, and under-estimate the strength of causal relationships (Type II error). If reporting bias is correlated with the explanatory variable, the risk of detecting a false causal relationship is much greater (Type I error). For example, if there is more reporting in locations with more cell phone towers, the estimated 'cell tower effect' on violence will be biased upward (Weidmann, 2016).

If sources vary in their direction and magnitude of bias, then such problems are, in one sense, easier to empirically address. Recent research has explored methods to offset gaps in coverage with information from other sources, including multiple systems estimation (Ball et al., 2003), capture-recapture techniques (Nichols, 1992, Hendrix and Salehyan, 2015), and pooling with a one-a-day filter (Leetaru and Schrodt, 2013). If the bias is common to all sources (e.g. more reporting of higher-casualty events) it becomes more difficult to correct. Recent studies have proposed diagnostic procedures to detect some of these problems, such as the reanalysis of event subsamples with varying levels of severity (Weidmann, 2016). To our knowledge, none of this previous research has specifically addressed the problem of actor-specific bias.

Turning to the consequences for citizens and public policy, while social scientists have sought to correct or at least detect reporting bias, most news consumers have neither the time nor interest to seek out multiple alternative sources of information (Popkin, 1994), and tend to consume content that already aligns with their worldview (Stroud, 2011, Iyengar and Hahn, 2009). Political communication scholars have long been interested in the effect of news coverage on public opinion and knowledge (Zaller, 1992, Prior, 2007, Baum and Kernell, 1999). Conflict scholars have generally avoided this topic.

This gap is surprising, since political actors often seek to alter the information environment to their own advantage (Davenport, 2009), promoting reporting favorable to their cause, and restricting information that could be damaging (Kumar, 2006, Rampton and Stauber, 2003, Taylor, 1992). The U.S. government, for instance, does not report civilian casualties from counterinsurgency operations. Protest movements tend to deny or under-emphasize violent elements within their own ranks, while calling attention to the brutality of the police response. Such biases are particularly acute for information providers whose audience has a direct stake in the conflict – like agencies and media outlets located in close proximity to a conflict zone.

Depending on one's source of information, a news consumer will likely see only one side of a multifaceted story. One-sided information streams can have important effects on public opinion, polarizing the attitudes of individuals exposed to conflicting narratives (Pariser, 2012, Stroud, 2011, Levendusky, 2013). This polarization, in turn, makes political compromise and conflict resolution more difficult.

Ukraine's information war

One of the defining features of the ongoing armed conflict in Ukraine has been an 'absence of transparent, agreed-upon truth' (Darden, 2014). After the Euromaidan protest movement swept President Viktor Yanukovych from power in February 2014 – and Russia annexed the Crimean peninsula – residents of Ukraine's eastern and southern provinces launched a series of demonstrations against the new authorities in Kyiv. These demonstrations escalated into a Russian-backed separatist rebellion in Ukraine's eastern Donbas region, comprising the heavily-industrialized and densely populated provinces of Donetsk and Luhansk.

Before the revolution in Kyiv, Russian media had a heavy presence in Ukraine, particularly in Crimea and other parts of the country's south-east (Broadcasting Board of Governors, 2014). In contrast to Western media portrayals of the Euromaidan as a largely peaceful protest movement confronting riot police and hired thugs, mainstream Russian media devoted their coverage to nationalist militants storming parliament and hurling Molotov cocktails. Both images were in a narrow sense true, but neither represented the complete picture. The Russian perspective on events seemed to leave an impression on crowds rallying in Crimea and the Donbas, who condemned the Euromaidan movement as a 'Western-backed coup' and 'fascist junta.'

Concerned over the mobilizational potential of Russian media, Ukraine's post-revolutionary authorities took a series of steps to create an 'hermetically sealed information environment' (Vikhrov, 2014). In March 2014, before the first shots were fired in east, Kyiv banned Russian federal broadcasters from Ukrainian cable TV, followed several months later by bans on some Russian films and serials, and travel bans on Russian journalists. In December, Ukraine established a Ministry of Information Policy to protect Ukrainians from 'unreliable information,' register media outlets and define professional journalistic standards. To spread government-approved content in social media, the Ministry launched an 'Information Army' of patriotic volunteers.

Ukrainian authorities also exerted direct pressure on some information providers. In September 2014, Ukraine's Security Service (SBU) raided the offices of the newspaper Vesti, accusing it of violating Ukraine's territorial integrity through its coverage of the Donbas conflict. In February 2015, Ukrainian authorities arrested a blogger on charges of treason, for posting a YouTube video criticizing the government's military mobilization campaign. The same month, Ukraine's Television and Radio Council accused popular TV host Savik Shuster of violating a law on 'incitement of hatred' after a Russian journalist criticized the government's military operations on his show.

In the rebel-held territories of the Donbas, separatists moved to create a similar zone of 'informational sovereigny' (Pomerantsev, 2014). After seizing the Donetsk regional administration building in April 2014, one of the rebels' next steps was to take control of TV towers in the region, take Ukrainian channels off the air, and put Russian ones back on. Later that year, the self-proclaimed Donetsk People's Republic established an official News Agency (DAN), while multiple privately-owned pro-rebel outlets emerged to fill the regional media vacuum. Wary of journalists from outside Russia and the region, rebels detained several reporters on suspicions of espionage, including an American journalist with Vice News.

In 2014, across rebel- and government-controlled territories of Ukraine, there were 7 documented killings of journalists, 286 physical assaults, 78 abductions, multiple physical attacks on offices and cyberattacks on websites (Freedom House, 2015). Many of these developments have predictably raised concerns over freedom of speech (Gorodnichenko and Mylovanov, 2015). Some analysts have worried that an informational firewall between dueling and contradictory media narratives will only deepen existing divisions (Darden, 2014).

How has Ukraine's information war affected public attitudes toward the conflict? Survey evidence suggests that very few Ukrainians outside of the Donbas see Russian state media as a reliable or truthful source – which may be evidence either of the success of Ukraine's counter-propaganda efforts, or ineffectiveness on Russia's part (Snegovaya, 2015). Residents of rebel-held areas appear to have a similarly skeptical view of Ukrainian media, particularly due to the latter's unwillingness to report on civilians killed by pro-government troops – incidents which Kyiv routinely denies (The Economist, 2015).

Despite much anecdotal speculation over who is 'winning' Ukraine's information war, there have been no systematic empirical studies on variation in coverage across information providers, or the impact of this variation on statistical results and public opinion.

Quantifying the conflict in Ukraine

To take stock of reporting biases in the Ukrainian conflict, we examine new violent event data based on human-assisted machine coding of news reports, press releases and blog posts from Ukrainian, rebel, Russian and external, third-party sources. These sources include official newswires, television channels, Internet news sites, and blogs. We also included the Russian-language edition of Wikipedia, and daily briefings from the OSCE Special Monitoring Mission to

Ukraine. Following previous quantitative studies of this conflict (Zhukov, 2016), we created a separate electronic text corpus for each data source, which contained all incident reports published on the Donbas between February 2014 and May 2016. Altogether, our data include 72,010 violent events reported by 17 information providers, between February 23, 2014 and May 2, 2016.

To determine the geographic locations of events mentioned in the reports, we ran an automated geocoding script that identified populated place names referenced in the text, and matched them against the U.S. National Geospatial Intelligence Agency's GeoNames database.[2] Table 2 shows the resulting spatial distribution of events, along with a description of each source.

We used a supervised learning algorithm (Support Vector Machine) to classify each event into a series of pre-defined categories, by event type, initiator, target, tactic, and casualties. The events of primary interest are *rebel violence* and *government violence*.[3] We define incidents of rebel violence as specific acts of organized violence initiated by any anti-Kyiv armed group or regular Russian Armed Forces.[4] Incidents of government violence involve organized violence by any pro-Kyiv armed group.[5] For each source shown in Table 2, the authors and a team of research assistants read and classified a randomly-selected training set of 130-600 reports (depending on the size of the corpus), in Russian, Ukrainian and/or English.[6] We used these manually-coded training data to train the SVM classifier to construct 17 separate datasets of violence in eastern Ukraine.

[2] We used a one-to-many mapping algorithm, to account for multiple events mentioned in the same report. To identify and correct geocoding errors and double-counts, we referenced each list of geocoded locations against a lookup table of regular errors (e.g. to ensure that 'Donetsk oblast' isn't mis-coded as 'Donetsk city,' and that references to the 'Shakhtar battalion' are not mis-coded as 'Shakhtarsk city'). We also performed manual inspection.

[3] The SVM classifies documents by fitting a maximally-separating hyperplane to a feature space, examining combinations of features that best yield separable categories. Formally, the SVM separates data points from each other according to their labels ($y_{it} \in \{-1, 1\}$), and finds maximum marginal distance Δ between the points labeled $y_{it} = 1$ and $y_{it} = -1$, solving the optimization problem

$$\arg\max_{\Delta,\alpha,\phi} \Delta \text{ s.t. } y_{it}(\alpha + \phi(X_{it})) > \Delta$$

where $y_{it}(\alpha + \phi(X_{it})\beta)$ is a functional margin, $\phi()$ is a function that maps the training data X to a high-dimensional space, and $\mathbf{K}(x_i, x_j) = \phi(x_i)'\phi(x_j)$ is a kernel function. The advantage of the SVM is that it is well-suited to sparse, high-dimensional data, is highly robust, and can handle a low training-to-test data ratio.

[4] A specific act of violence is a reference to a single ongoing or recent military operation, act of terrorism, targeted killing, detention, other violent event. Not included are general summaries of war statistics or press statements. Anti-Kyiv groups include any forces explicitly labeled as 'insurgents,' 'rebels,' 'terrorists,' as well as specific formations like the Novorossiya Armed Forces, Donetsk People's Republic (DNR), Lugansk People's Republic (LNR), Vostok Battalion, Oplot, Kal'mus battalion, Bezler band, Zarya battalion, Russian Orthodox Army (RPA), People's Militia of Donbass (NOD), Prizrak battalion, Army of the South East, Don Cossacks, Russian National Unity, Eurasian Youth Union, Yovan Sevic.

[5] Pro-Kyiv groups include Ukraine's regular Army, Air Force, Airborne troops, Marines, Border Guard, SBU, Interior Ministry, local police, National Guard or any of 46 volunteer battalions (e.g. Azov, Aydar, Dnipro-1, Donbas) and independent right-wing militias (e.g. Right Sector).

[6] To account for potential disagreement between coders, at least two coders read each training set document. Inter-coder reliability statistics, reported in the online appendix, indicate a high and statistically significant level of agreement between coders on the relevant categories, including where coders read the same documents in different languages.

In addition to violence, we collected geospatial data on several covariates common in subnational conflict research, including population density (CIESIN and Columbia University, 2005), forest cover (Loveland et al., 2000), distance to the nearest road (Defense Mapping Agency, 1992) and distance to the Russian border (Global Administrative Areas, 2012). We also include an indicator of whether a municipality was under rebel control on a given day, and the corresponding distance to the front line where rebel control ends and government control begins.[7]

Actor-specific reporting bias in Ukraine

How do the sources in Table 2 differ in their coverage of the Donbas conflict? Do all sources report the same kinds of events by the same actors, or do they focus on one group more than another? To answer these questions, we estimated the relative bias of each information provider in covering rebel versus government violence.[8] Figure 1 reports these estimates, with event reports published by the OSCE as the reference category (vertical line at zero). Positive values indicate that a source is more likely to cover rebel than government violence, and negative values indicate greater relative coverage of government violence. Where the 95% confidence intervals cover zero, relative levels of coverage were similar to reports by the OSCE.

Figure 1 reveals large systematic differences in the armed actors who receive coverage in Ukrainian, rebel, Russian and international sources. Overall, Ukrainian information providers (blue circles) devote more news coverage to rebel violence and less to government operations than any other group of sources. Four out of the five sources that systematically 'over-report' rebel attacks are Ukrainian: the military blog *Information Resistance (Sprotyv)*, and the TV channels *112*, *Espreso* and *Channel 5* – the latter owned by Ukraine's current President, Petro Poroshenko.[9]

Most sources that 'over-report' government violence are based within Russia (red circles) or

[7]We used Zhukov (2016)'s data on territorial control, which draw on three sets of sources: (1) official daily situation maps publicly released by Ukraine's National Security and Defense Council, (2) daily maps assembled by the pro-rebel bloggers 'dragon_first_1' and 'kot_ivanov,' and (3) Facebook posts on rebel checkpoint locations.

[8]The empirical model is

$$y_{jt}^{(R)} = g^{-1}(y_{j,t-1}^{(R)}\gamma + \delta_j^{(R)} + \alpha_t^{(R)} + u_{jt}^{(R)})$$
$$y_{jt}^{(G)} = g^{-1}(y_{j,t-1}^{(G)}\gamma + \delta_j^{(G)} + \alpha_t^{(G)} + u_{jt}^{(G)})$$

where $y_{jt}^{(k)}$ is the number of events of type $k \in$ (R: rebel violence; G: government violence) reported by information provider j at time t, $\delta_j^{(k)}$ is the source-specific intercept for event type k, $\alpha_t^{(k)}$ is a daily fixed effect, and $g^{-1}()$ is a quasi-Poisson inverse link function. The *relative bias* of source j is $\delta_j^{(R)} - \delta_j^{(G)}$. The δ_j term here is akin to a "house effect" in research on the pooling of public opinion data from multiple survey firms (Converse and Traugott, 1986, Jackman, 2005, Beck, Jackman and Rosenthal, 2006, Pickup and Johnston, 2008). We set $\delta_j = 0$ for $j =$ OSCE.

[9]The term "over-report" indicates that a source reports a higher share of rebel-to-government (or government-to-rebel) attacks than the OSCE.

Table 2: **Sources included in Ukraine violence dataset**. Maps show locations of all violent events in Donetsk and Luhansk between April 1, 2014 and May 2, 2016, as reported by each information provider.

Name	Map	Info	Name	Map	Info
Channel 112 (Ukraine)	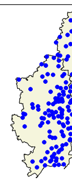 N = 3518	TV Rus/Ukr-language	Lenta.ru (Russia)	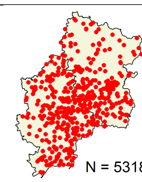 N = 5318	Online Rus-language
Channel 5 (Ukraine)	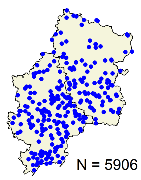 N = 5906	TV Ukr-language	NewsFront (rebel)	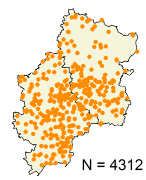 N = 4312	Online Rus-language
BFM (Russia)	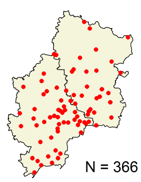 N = 366	Online Rus-language	OSCE (international)	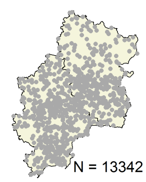 N = 13342	Online English-language
DAN (rebel)	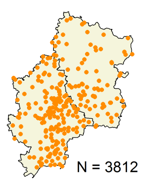 N = 3812	News agency Rus-language	RusVesna (rebel)	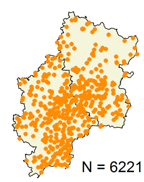 N = 6221	Online Rus-language
Dozhd/Rain (Russia)	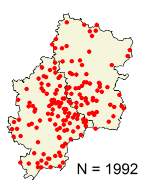 N = 1992	TV Rus-language	Sprotyv/InfoResistance (Ukraine)	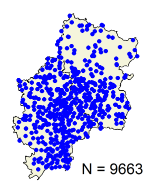 N = 9663	Online Rus-language
Espreso (Ukraine)	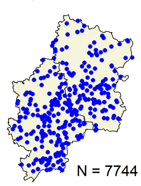 N = 7744	TV Ukr-language	Ukrinform (Ukraine)	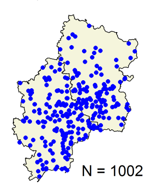 N = 1002	News agency Rus/Ukr-language
Gazeta.ru (Russia)	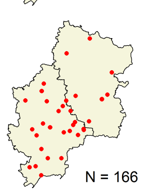 N = 166	Online Rus-language	Vesti/Rossiya-24 (Russia)	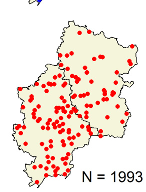 N = 1993	TV Rus-language
Interfax.ru (Russia)	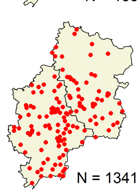 N = 1341	News agency Rus-language	Wikipedia (international)	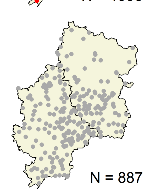 N = 887	Online Rus-language
Interfax.ua (Ukraine)	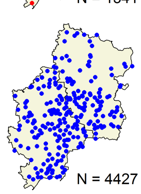 N = 4427	News agency Rus/Ukr-language			

Figure 1: **Ukrainian sources report on rebel violence, pro-Russian sources report government violence**. Dots are relative bias in reporting on rebel versus government violence. Lines are 95% confidence intervals.

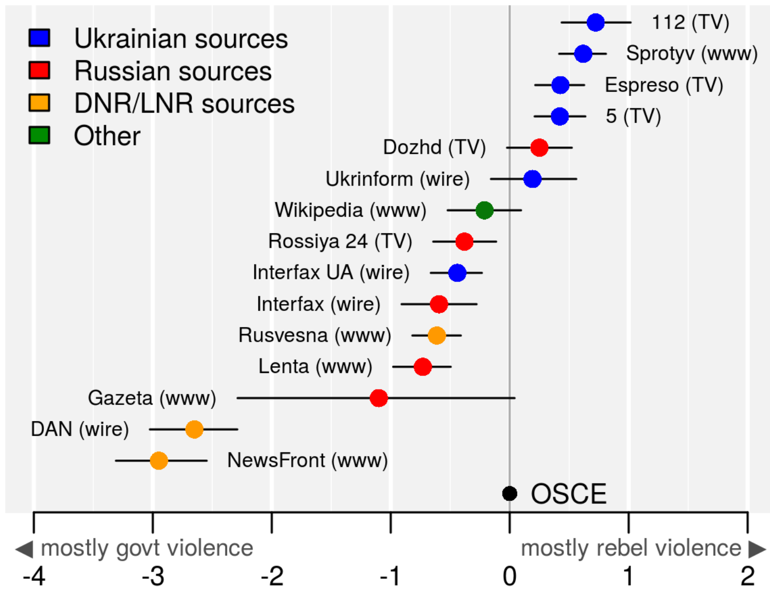

the self-proclaimed Peoples' Republics of Donetsk and Luhansk (DNR, LNR) (orange circles). DNR-based media outlets *NewsFront* and *Donetsk News Agency (DAN)*, in particular, have the most acute actor-specific bias in the data, reporting almost exclusively on violence by the government.

Russian sources have the same general direction of bias as rebel sources, but with somewhat lower magnitude. With a single exception – the independent, opposition-oriented *Dozhd TV* channel, which occupies a space on this spectrum closer to the median Ukrainian source – Russian media report disproportionately on government violence. The only Ukrainian outlet with a comparable bias in the opposite direction is Interfax-Ukraine – a Russian-owned wire service. Between rebel and Ukrainian media, there is a much clearer separation – the 'left-most' Ukrainian outlet is still to the right of the 'right-most' rebel outlet.

A very different picture appears in third-party sources, like OSCE reports and Wikipedia. These data are more 'neutral,' in the sense that they are unlikely to attribute violence to any armed group at all. The language in these reports tends to be more passive and non-specific (e.g. 'shelling was reported near village X') than language in local media. For the OSCE, this finding is consistent with anecdotal reports that – because it must maintain working relations with all sides – the monitoring organization is cautious about attributing violence to specific initiators. For Wikipedia (green circle), this pattern may reflect the crowd-sourced nature of the data: users flag entries as

Figure 2: **Which actor is "more indiscriminate"?** Lines are predicted proabilities of indiscriminate tactics appearing in reports of (a) rebel and (b) government violence. Rectangles are 95% confidence intervals.

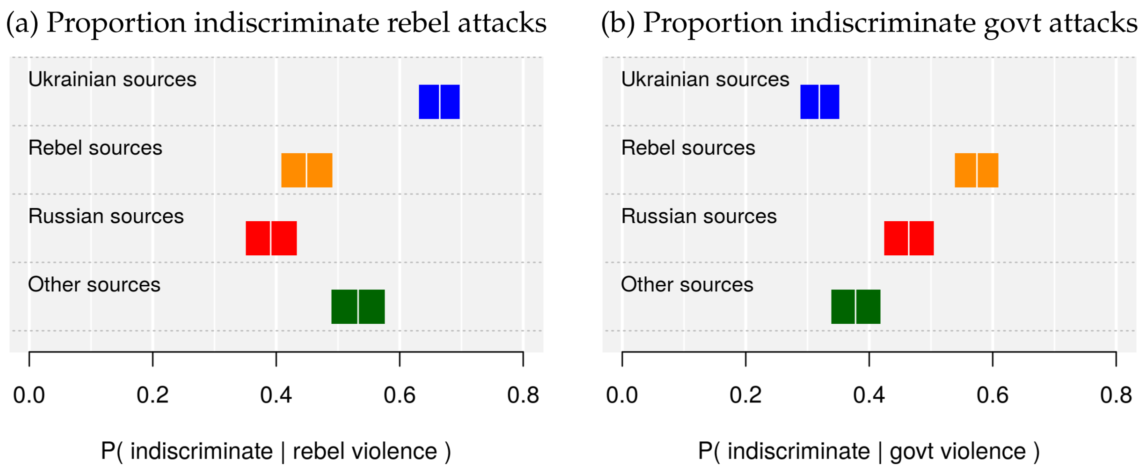

biased, remove offending information, and eventually reach a 'neutral' compromise.

Figure 2 shows that Ukrainian and rebel media not only tend to report disproportionately on violence by the 'other side,' but they disproportionately focus on *indiscriminate* violence by the 'other side' – events like artillery and rocket shelling, and the use of heavy armor. The quantities in Figure 2 have the following interpretation: how likely is an average information provider to describe an incident of rebel (or government) violence as indiscriminate?[10] Ukrainian news coverage of rebel violence cites indiscriminate tactics 66 percent of the time (95% CI: .63, .70), compared to 45 percent in rebel media (95% CI: .41, .49). Coverage of government violence is a near-mirror image: 32 percent of the government violence reported by Ukrainian sources is indiscriminate (95% CI: .29, .35), compared to 57 percent for rebel sources (95% CI: .54, .61). Russian and international sources, as before, lie somewhere in between.

[10]These quantities are predicted probabilities from a generalized additive model

$$y_{jit}^{(R-ind)} = g^{-1}(X_i\beta + \delta_j^{(R-ind)} + \alpha_t^{(R-ind)} + s(long_i, lat_i)^{(R-ind)} + u_{ijt}^{(R-ind)} | y_{jit}^{(R)} = 1)$$
$$y_{jit}^{(G-ind)} = g^{-1}(X_i\beta + \delta_j^{(G-ind)} + \alpha_t^{(G-ind)} + s(long_i, lat_i)^{(G-ind)} + u_{ijt}^{(G-ind)} | y_{jit}^{(G)} = 1)$$

where $y_{jit}^{(k)}$ is 1 if a source from set $j \in$ (Ukraine, DNR/LNR, Russia, other) reports an event of type $k \in$ (R-ind: indiscriminate rebel violence; G-ind: indiscriminate government violence) as occurring in location i at time t, conditional on $y_{jit}^{(k_0)} = 1$ (i.e. that j reports at least one event of type $k_0 \in$ (R: rebel violence; G: government violence) as occurring in i, t). Also on the right hand side are an inverse logit link function $g^{-1}()$, a vector of covariates X_{it} (population density, distance to nearest road, open terrain, distance to the front line, territorial control), daily fixed effects α_t to account for coverage fatigue, and spatial spline $s(long_i, lat_i)$ to account for location-specific biases. To identify the model, we pooled individual sources by country in our estimation of the δ_j parameters.

An information provider's county or group affiliation is, of course, not the only determinant of actor attribution. Some types of selective reporting are common to all sources. For example, in line with existing evidence on the 'supply-side' causes of selective reporting (Weidmann, 2016), we find that – for all sources – there is significantly more attribution in places with more witnesses (high population density), visibility (open terrain), and accessibility (near a major road).

This initial glance at the data reveals systematic differences in the actors whose violence individual sources cover. But how do these opposing narratives affect predictions of how the Ukrainian conflict will unfold? Do different sources yield different conclusions about what sort of equilibrium, or steady state, may emerge in the absence of outside intervention? How might reporting bias shape our expectations about the strategic interaction between government and rebels, and about which side is more likely to cooperate or escalate?

Table 3 reports the stationary distribution of violence in Ukraine, according to each set of sources.[11] The quantities in the table have the following interpretation: if the conflict were to continue to develop as reported in the press until it reaches a stable equilibrium, what proportion of time will an average location experience: (1) no violence, (2) one-sided rebel violence, (3) one-sided government violence, and (4) two-sided violence?

As the table indicates, each group of source offers its own perspective on how fighting in Ukraine is likely to unfold, and what sort of equilibrium is likely to emerge. According to Russian and outside sources (i.e. OSCE, Wikipedia), this equilibrium will be largely peaceful, with rebel-government interactions becoming non-violent about nine-tenths of the time. Local sources paint a more ominous picture. If the conflict continues to play out as reported in Ukrainian media, the two sides will be at peace just 69 percent of the time, and will experience one- or two-sided violence during the remaining 31 percent. Rebel sources are even more pessimistic, with the system staying non-violent 62 percent of the time.

Which actor is most likely to break the peace, according to each set of sources? As one might expect, the greatest disparity here is between Ukrainian and rebel sources. Ukrainian sources predict that rebels are more than twice as likely to unilaterally escalate than government troops –

[11] The online appendix provides a full derivation of the stationary distribution, which we empirically estimate here with predicted probabilities of a bivariate probit model

$$\begin{matrix} y_{R,it} \\ y_{G,it} \end{matrix} = \begin{matrix} g^{-1}(y_{G,it-1}\zeta_R + y_{R,it-1}\alpha_R + y_{G,it-1}y_{R,it-1}\gamma_R + \mathbf{x}_{R,it}\beta_R + \boldsymbol{W}y_{R,it-1}\rho_R + \epsilon_{R,it} + \eta_{it}) \\ g^{-1}(y_{R,it-1}\zeta_G + y_{G,it-1}\alpha_G + y_{R,it-1}y_{G,it-1}\gamma_G + \mathbf{x}_{G,it}\beta_G + \boldsymbol{W}y_{G,it-1}\rho_G + \epsilon_{G,it} + \eta_{it}) \end{matrix} \tag{1}$$

where $y_{k,it-1}$ is a time-lagged dependent variable for actor k, $\boldsymbol{x}_{k,it}$ is a vector of covariates (population density, distance to nearest road, open terrain, distance to the front line, rebel territorial control), α_k, β_k, ζ_k and γ_k are regression coefficients, ϵ_k is an error component unique to each actor, and η is an error component shared by the two actors. To account for spillovers of violence from neighboring locations, we include spatio-temporal lags of the dependent variable, where $\boldsymbol{W}$ is a row-normalized spatial weights matrix, and ρ_k is the autoregressive parameter.

with the probability of one-sided violence at .13 for rebels versus .05 for the government. Rebel sources predict an even stronger pattern in the opposite direction, with government troops almost ten times more likely to unilaterally escalate than the rebels (.27 versus .03).

These four sets of predictions have divergent implications for conflict resolution. In the case of outside sources like the OSCE, a news consumer or policymaker may conclude that intervention is not necessary to reduce violence. Here, violence diminishes organically over time, and neither side appears likely to unilaterally escalate in equilibrium – a situation in which a negotiated settlement may become self-sustaining. Local sources yield very different lessons: here, transgressions by one or both actors appear to be more common, and a negotiated settlement less likely to hold. For violence to decline, it follows, enforcement efforts should target whichever side is more prone to unilaterally escalate. According to Ukrainian sources, this intervention should seek to restrain rebels; according to rebel sources, it should target the government.

In sum, the direction of actor-specific reporting bias in Ukraine aligns with what one might expect if information providers were actively seeking to discredit the opposing side and mobilize public opinion against it. Internally, information consumers may doubt that an actor inclined to use unilateral violence can stick to the terms of a negotiated agreement. Externally, the relative propensity for escalation can shape perceptions over how intractable a conflict is likely to be, whether third-party enforcement is necessary to stop it, and whether that response should be impartial or directed at one side. Whether these biases can actually produce such effects is an open empirical question, which we address in the next section.

Table 3: **Which side is more likely to unilaterally escalate?** Quantities represent the stationary distribution of violence, according to data from each set of sources. 95% confidence intervals in parentheses.

Sources	Pr(no violence)	Pr(one-sided govt violence)	Pr(one-sided rebel violence)	Pr(two-sided violence)
Ukraine	0.691 (0.668,0.714)	0.049 (0.044,0.054)	0.130 (0.122,0.138)	0.130 (0.120,0.140)
DNR/LNR	0.621 (0.598,0.645)	0.267 (0.257,0.276)	0.028 (0.024,0.032)	0.084 (0.074,0.094)
Russia	0.909 (0.889,0.927)	0.034 (0.028,0.041)	0.024 (0.019,0.029)	0.033 (0.026,0.041)
Other	0.899 (0.882,0.913)	0.033 (0.028,0.038)	0.046 (0.039,0.053)	0.023 (0.019,0.027)

Implications for public opinion

The previous analysis demonstrated that actor-specific reporting bias has a substantive impact on statistical inferences, and particularly predictive models of how conflict is likely to unfold and how sustainable a negotiated settlement may be. Ukrainian sources predict an equilibrium in which unilateral escalation by rebels is much more common than unilateral escalation by government troops. Rebel sources predict an opposite equilibrium, where unilateral escalation by government troops is pervasive. Russian and external, third-party sources predict equilibria in which violence is generally less likely, and unilateral transgressions are more rare.

To more directly explore the impact of actor-specific bias on policy preferences – and thereby assess the implications of reporting bias on conflict resolution in general, and on global support for intervention in particular – we ran a survey experiment. In the experiment, we exposed subjects to news stories about a generic civil conflict, and asked them about their attitudes toward third-party intervention in that conflict. Our subject pool included 1,596 respondents from the United States. We conducted the survey in April-May 2014 (before the Ukrainian conflict had fully escalated), using Amazon Mechanical Turk to recruit participants.[12]

We employed a 2-by-2 factorial design survey to test the effect of both actor attribution (i.e. rebels or government) and war-fighting tactics (i.e. selective and indiscriminate) – the same types of biases we see in our event data. The survey for each group consisted of a short news story about violence in a generic, non-specified civil conflict, followed by a battery of questions designed to elucidate the effect of the stories on public opinion. For each version of the treatment, we kept the bulk of the text intact, modifying only the actors committing the violence, and their reported tactics. We randomly assigned participants to one of the four treatment regimes, or to a control group, which received the same text without any information about actor or targets. To help respondents comprehend the relative destructiveness of selective and indiscriminate tactics, we accompanied each text with a photo – showing either an individual arrest (selective), a destroyed house (indiscriminate), or an armed group riding on an armored personnel carrier (control). These photos were the same for rebel- and government-initiated events. Table 4 summarizes the treatments, and the corresponding versions of the text passage we asked participants to read.

In the subsequent battery of questions, we asked participants about the type of response – by the international community and their own country – they felt was appropriate, with options including 'No response,' 'Economic sanctions,' 'Military aid,' and 'Military intervention' (see appendix for full survey). In addition, we asked whether the response should be impartial, or directed against one of the two sides. To account for other potential drivers of policy preferences, we concluded the survey with a battery of general questions about participant demographics, political ideology, military background and political knowledge. We also included several 'attention

[12]Previous research has shown that respondents recruited via Mechanical Turk tend to be less representative of the U.S. population than Internet-based panels or national probability samples, but more representative than in-person convenience samples (e.g. student lab experiments). Replications of experimental results using MTurk samples are substantively similar to those found with more nationally representative samples (Berinsky, Huber and Lenz, 2012).

Table 4: **Survey instrument**. Bold text denotes differences in reported actors and tactics for each treatment.

	Treatment	Photo	Text
T1	rebel violence, selective		'Hundreds are missing in the war-torn country this week, as **rebel** forces escalated their operations against suspected **government** supporters. International monitors reported that 50 individuals died and at least 200 are missing after a **series of assassinations and detentions** by **pro-rebel** forces in provincial towns.'
T2	rebel violence, indiscriminate		'Hundreds are missing in the war-torn country this week, as **rebel** forces escalated their operations against suspected **government** supporters. International monitors reported that 50 individuals died and at least 200 are missing after a **barrage of heavy artillery shelling** by **pro-rebel** forces in provincial towns.'
T3	govt violence, selective		'Hundreds are missing in the war-torn country this week, as **government** forces escalated their operations against suspected **rebel** supporters. International monitors reported that 50 individuals died and at least 200 are missing after a **series of assassinations and detentions** by **pro-regime** forces in provincial towns.'
T4	govt violence, indiscriminate		'Hundreds are missing in the war-torn country this week, as **government** forces escalated their operations against suspected **rebel** supporters. International monitors reported that 50 individuals died and at least 200 are missing after a **barrage of heavy artillery shelling** by **pro-regime** forces in provincial towns.'
C	no info on actors, tactics		'Hundreds are missing in the war-torn country this week, as **rebel and government** forces **escalated their campaigns**. International monitors reported that 50 individuals died and at least 200 are missing after the **two sides clashed** in provincial towns.'

filter' questions, to exclude respondents who either did not read the news story or were clicking through the questionnaire at random.

The survey yields three main findings. First, news stories that disclose more information about actors and tactics increase support for intervention and reduce support for impartiality. Second, actor-specific bias has a greater effect on policy preferences than tactic-specific bias. Third, there is a greater effect on support for more militarized policy responses (e.g. military aid, direct intervention) than less militarized ones (e.g. economic sanctions).

Figure 3 reports model-based estimates of the average treatment effect of any actor/tactic specific coverage (i.e. any of the four treatment groups in Table 4) on support for intervention, and support for taking sides.[13] Relatively few respondents favored a completely 'hands-off' response to reports of violence, with more than 90 percent favoring some diplomatic or military response. Yet among participants who read a story with information on the actors involved and tactics used, average support for intervention was 2.5 percent higher, rising from 95.5 to 98 percent. The effect on respondent impartiality was stronger. Although only a quarter of respondents who read non-specific coverage favored taking sides in the conflict, one third of those in the four treatment groups did the same – an increase of 8 percent relative to the control group.

Figure 3: **Actor-specific reports increase support for intervention and taking sides**. Points are average treatment effects. Lines are 95% confidence intervals.

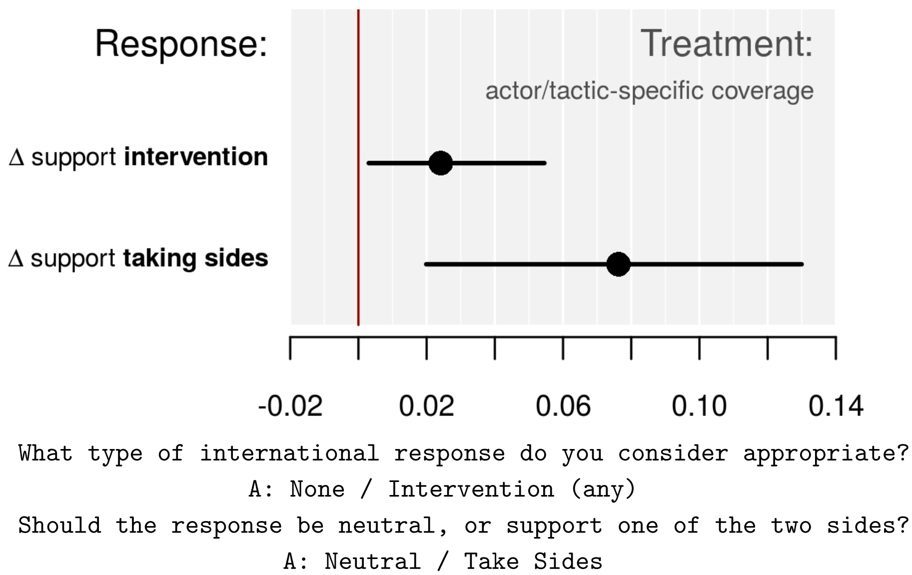

Our second finding is that actor attribution is more important for public opinion than information about tactics. Figure 4 reports average treatment effects of actor- and tactic-specific coverage on predicted levels of support for anti-rebel and anti-government interventions.[14] As the figure shows, respondents tend to favor intervention against whichever actor is reported to have ini-

[13]Formally, $ATE = E[\hat{y}_{1i} - \hat{y}_{0i}]$, where the predicted probabilities $\hat{y}$ are based on the parameters of a logit model

$$y_i = g^{-1}(\beta_0 + \beta_1 T_i + \beta_2 X_i + u_i)$$

where y_i is the individual response, T_i is treatment status (1 if $T_i \in \{T1, T2, T3, T4\}$ and 0 if $T_i = C$), and X_i is a vector of participant demographic and political ideological characteristics, including age, gender, education, social conservatism, past military service and political knowledge (the latter based on the accuracy of subjects' answers to questions about the composition of the UN Security Council, NATO and names of world leaders).

[14]Model specification is identical to that in Figure 3, but with T_i disaggregated into the four treatment groups in Table 4.

tiated the violence. Whether that violence was selective or indiscriminate matters less than the group responsible for the action. Support for intervention – against either actor – changed by about the same amount across subjects who read reports of selective and indiscriminate violence.

Figure 4: **Actor attribution matters more than tactical descriptions**. Points are average treatment effects. Lines are 95% confidence intervals.

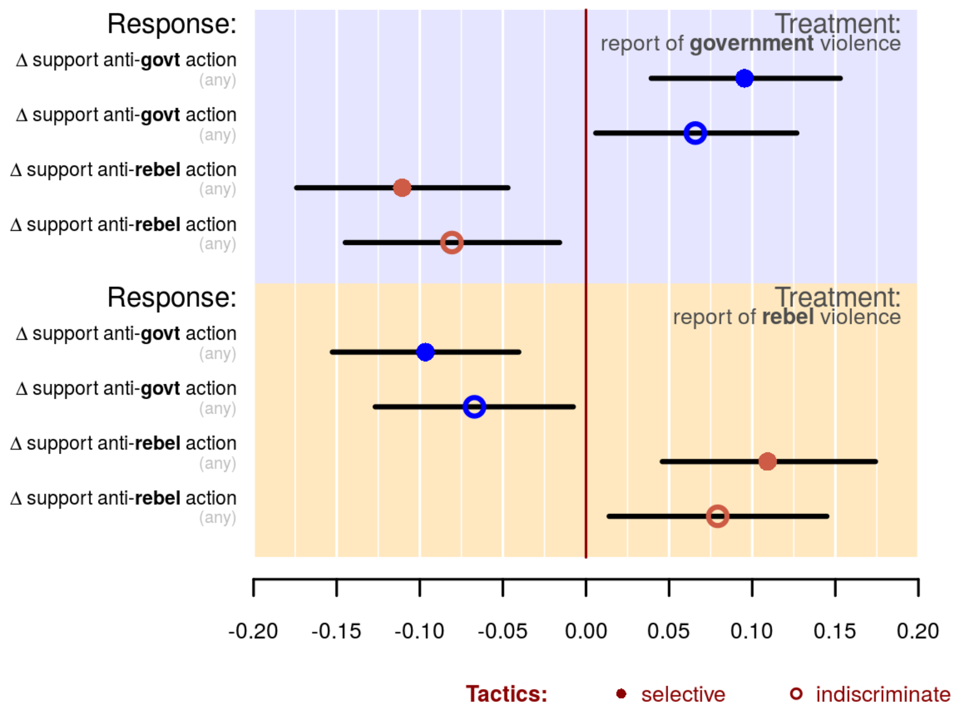

Q: Against whom should the international response be directed?
A: Government / Rebels

The relatively weak effect of tactic-specific reporting bias is surprising, since the photos we displayed in our survey instrument highlighted specifically differences between selective and indiscriminate tactics, rather than between rebel and government violence (Table 4). In theory, this emphasis should have inflated rather than attenuated the 'tactic effect.' The fact that indiscriminate violence did not have a stronger impact on respondents' policy choices suggests that the differential reporting of opponents' tactics we saw in Ukrainian and rebel sources (Figure 2) may not be the most compelling way to affect public attitudes.

The survey's third finding is that the increase in support for intervention is greater for military policy options than non-military ones. Figure 5 reports the average treatment effects of actor-specific coverage on support for three types of intervention against each side (economic sanctions, aid to opponents, and direct military intervention).[15] Respondents who read stories of rebel violence were 5.4 percent more likely (95% CI: .02, .09) to support anti-rebel international military intervention, but neither more nor less likely to support economic sanctions 0.01 (95% CI: -0.01, 0.02). Participants who read about government violence also showed more support for international military intervention than sanctions against the offending actor – yielding increases of 0.04 percent (95% CI: 0.01, 0.07) versus 0 percent (95% CI: -0.01, 0.02).

Figure 5: **Effect is greater at higher scales of intervention**. Points are average treatment effects. Lines are 95% confidence intervals.

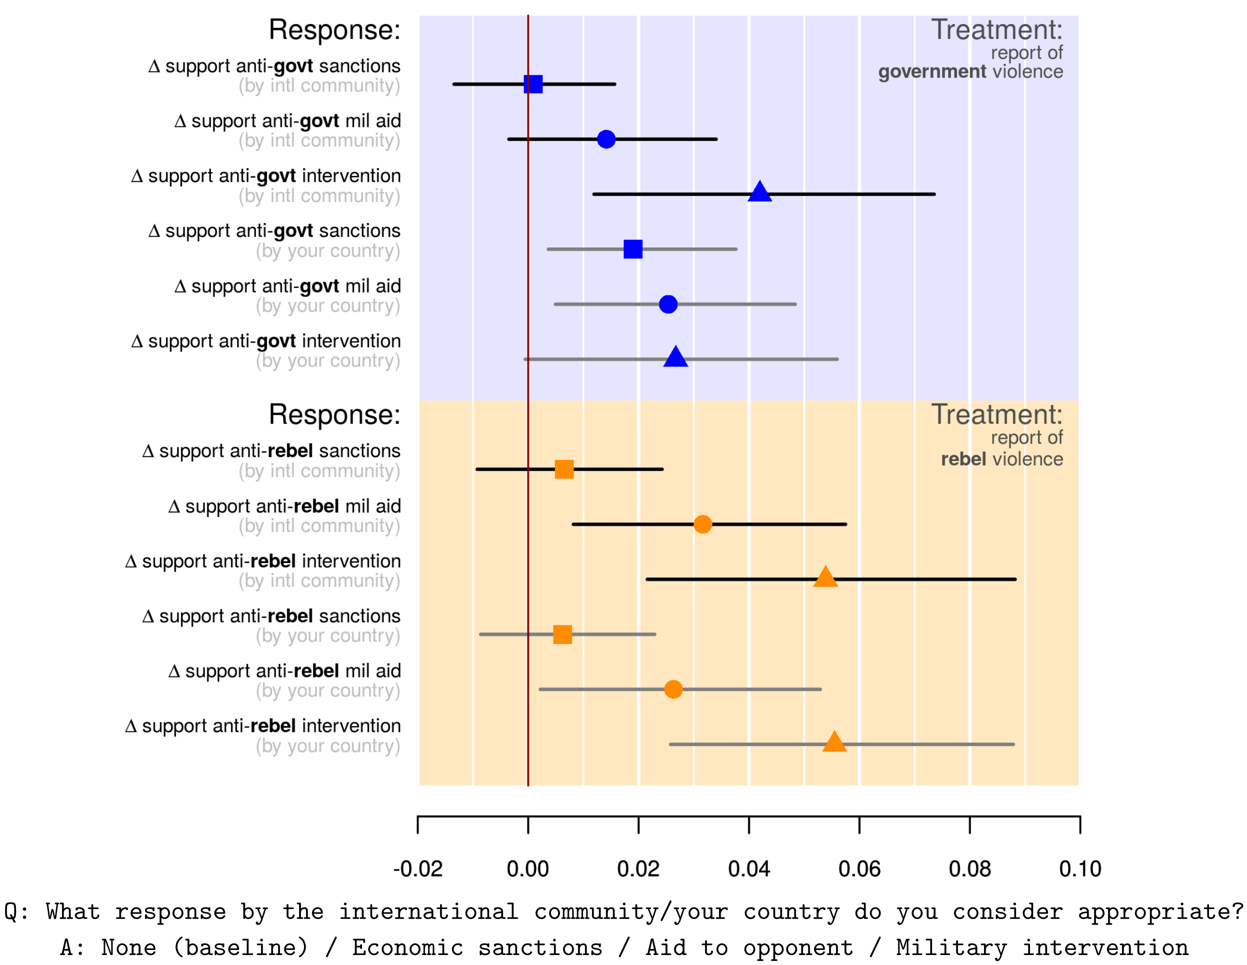

[15]Model specification is identical to that in Figure 4, but with T_i pooled by actor ({T1, T2}, {T3, T4} from Table 4).

One area where survey responses diverged was on the question of *who should intervene* – the international community or 'your country' (i.e. the United States). When considering anti-government action, respondents were more hesitant to support military intervention by their own country than by the international community, instead favoring less extreme actions, like sanctions and military aid to the government's opponent. Responses to rebel violence were more consistent, with subjects about equally likely to support military action by the international community and their own country. This pattern suggests that our subjects may have considered anti-government intervention to be more costly than anti-rebel intervention, and – even if they supported such action in principle – preferred that their country not shoulder the burden on its own.

What can these survey results tell us about the impact of reporting bias in Ukraine? If more detailed news reports generate stronger support for intervention, and lead news consumers to pick sides, then we should expect Ukrainian and rebel sources to have a greater impact on public opinion than Russian or external, third-party sources. The relatively low-information content of OSCE reports and other information providers from outside the conflict zone should dampen pro-intervention preferences. We should also expect Ukrainian and rebel sources to generate opposing preferences with respect to the direction and partiality of any intervention. The heavy focus on rebel violence in Ukrainian media is likely to generate support for anti-rebel intervention; rebel media emphasis on government violence should generate support for anti-government intervention. The scope of this intervention may be limited (e.g. sanctions) or extensive (e.g. ground invasion), but the direction of the increased support will be the same.

Our experimental results are not evidence that Ukrainian and rebel media consciously manipulate news coverage – either to attract external support or to undermine local confidence in the opponent's credibility as a negotiating partner. Yet if they indeed had such an intent, we would expect information providers to adopt the exact types of reporting biases that we observe here.

Conclusion

This study sought to advance the nascent research program on reporting bias in civil conflict, by taking a more direct look at the consequences of selective news coverage for scholarly and public knowledge. In so doing, we focused on the empirically common, but relatively understudied *actor-specific reporting bias* – the tendency of information providers to report violence by some actors more than violence by others. Unlike prior research, we treated reporting bias not merely as an inside-the-church problem for social scientists seeking to understand conflicts after the fact, but also as a source of influence on consumers of biased news reports: citizens and governments.

To that end, we conducted a survey experiment to explore how the sorts of biases revealed in our event data analysis might influence public opinion for, or against, intervening in civil conflicts.

Our results show that – by casting one actor as 'more violent' than the other – actor-specific reporting bias can have a profound impact on both statistical inference and public opinion. Data from one set of sources may predict a relatively peaceful equilibrium, where neither side is likely to unilaterally escalate the level of violence. Another source may predict a more violent equilibrium, in which violations are common, and one side is disproportionately more likely to attack than the other. These opposing perspectives on the conflict carry different implications for policy, particularly as regards the utility of outside intervention, and the actors' relative abilities to honor a negotiated agreement. These findings demonstrate that selective coverage can have an effect akin to that of propaganda, shaping public knowledge and preferences in a partisan way.

In our analysis of event data compiled from multiple information providers, we found that Ukranian news sources disproportionately emphasize violence by rebels, while rebel sources emphasize the opposite: violence by the Ukrainian government forces. Both Ukranian and rebel sources, in turn, frame their coverage of the other side's violence as overwhelmingly indiscriminate – using heavy weapons and indirect fire methods that carry a high risk of non-combatant civilians. For sources outside the conflict zone – like media outlets in Russia, and international organizations like the OSCE – we found a subtler form of bias: a tendency not to attribute responsibility for violence to either side, and frame both sides' violence as about equally indiscriminate.

Our experimental findings show that respondents were more supportive of intervening in a conflict against the side characterized as perpetrating the violence. They were less supportive of intervening when confronted with neutral coverage of the conflict. Interestingly, characterizing the violence as indiscriminate or directed against military combatants did not significantly influence participants' support for intervention. The absence of a tactic effect surprised us. The implication is that attributions of responsibility for violence loom larger in public attitudes than do the details concerning the nature of the violence which one or the other side commits.

Of course, the news that international audiences see – like that featured in our experiment – differs in content and scope from the local press reports we examined for Ukraine. News consumers are most interested in events culturally and physically proximate to themselves, and news organizations lack the incentives and resources to cover everything all the time. Yet international news organizations are still disproportionately dependent on local reporting in their coverage of war zones, particularly where – as in rebel-held Donetsk – their own reporters have limited access.

The relative direction and magnitude of actor-specific reporting biases in Ukraine represent the exact opposite of what would be needed to quickly resolve the conflict. The fact that neutral coverage – of the type which prevails in OSCE reporting of events in Ukraine – suppresses support for intervention raises troubling questions concerning the capacity of world leaders to generate the necessary public support for peacekeeping missions. When information providers 'play it down the middle' – whether due to the journalistic value of balanced, neutral coverage, or due to more cynical 'false equivalency' – consumers have more difficulty ascribing responsibility, and are less likely to support outside intervention. This dulling effect of neutral coverage on external public opinion stands in stark contrast to the polarizing effect of 'one-sided' local news coverage. The net effect of these reporting biases is that domestic audiences may become less interested in striking a bargain with the opposing side, while outside audiences become less interested in intervention.

Reversing these two sets of biases is, of course, easier said than done. In the absence of attributions of responsibility for violence, leaders and activists interested in conflict resolution will need to better inform journalists about the details of specific incidents. Where attribution exists, governments and NGOs will need to expand audiences' access to multiple sources of information.

Our study suggests that reporting bias can have a potentially significant impact on public attitudes toward conflict resolution, one that scholars and practitioners to date have largely failed to examine. Future research should thus extend this analysis to additional civil conflicts, actors, and media outlets to determine whether these findings generalize beyond Ukraine, and in doing so to further refine our estimates of the nature, extent, and consequences of selective reporting.

References

Allan, Stuart and Barbie Zelizer. 2004. *Reporting war: Journalism in wartime*. Routledge.

Arceneaux, Kevin, Martin Johnson and Chad Murphy. 2012. "Polarized political communication, oppositional media hostility, and selective exposure." *The Journal of Politics* 74(01):174–186.

Ball, Patrick, Jana Asher, David Sulmont and Daniel Manrique. 2003. *How many Peruvians have died*. Washington, DC: American Association for the Advancement of Science.

Baum, Matthew A and Samuel Kernell. 1999. "Has cable ended the golden age of presidential television?" *American Political Science Review* 93(01):99–114.

Baum, Matthew A and Tim Groeling. 2010. *War stories: The causes and consequences of citizen views of War*. Princeton: Princeton University Press.

Baum, Matthew A and Tim J Groeling. 2009. *War stories: The causes and consequences of public views of war*. Princeton University Press.

Baum, Matthew A and Yuri M Zhukov. 2015. "Filtering revolution Reporting bias in international newspaper coverage of the Libyan civil war." *Journal of Peace Research* 52(3):384–400.

Beck, Nathaniel, Simon Jackman and Howard Rosenthal. 2006. "Presidential approval: the case of George W. Bush." Presented at Annual Meeting of the Society for Political Methodology, Davis, CA.

Berinsky, Adam J, Gregory A Huber and Gabriel S Lenz. 2012. "Evaluating online labor markets for experimental research: Amazon. com's Mechanical Turk." *Political Analysis* 20(3):351–368.

Berman, Eli, Jacob N Shapiro and Joseph H Felter. 2011. "Can hearts and minds be bought? The economics of counterinsurgency in Iraq." *Journal of Political Economy* 119(4):766–819.

Broadcasting Board of Governors. 2014. Contemporary Media Use in Ukraine. Technical report Gallup.
URL: *http://www.bbg.gov/wp-content/media/2014/06/Ukraine-research-brief.pdf*

CIESIN and Columbia University. 2005. "Gridded Population of the World, Version 3 (GPWv3) Data Collection." Center for International Earth Science Information Network (CIESIN), Centro Internacional de Agricultura Tropical (CIAT).

Converse, Philip E and Michael W Traugott. 1986. "Assessing the accuracy of polls and surveys." *Science* 234(4780):1094–1098.

Dafoe, Allan and Jason Lyall. 2015. "From cell phones to conflict? Reflections on the emerging ICT–political conflict research agenda." *Journal of Peace Research* 52(3):401–413.

Danzger, M Herbert. 1975. "Validating conflict data." *American Sociological Review* 40(5):570–584.

Darden, Keith. 2014. "The War on Truth in Ukraine." *The New York Times* .

Davenport, Christian. 2009. *Media bias, perspective, and state repression: The Black Panther Party*. Cambridge University Press.

Davenport, Christian and Allan Stam. 2006. "Rashomon goes to Rwanda: Alternative accounts of political violence and their implications for policy and analysis." Working paper.

Davenport, Christian and Patrick Ball. 2002. "Views to a kill exploring the implications of source selection in the case of Guatemalan state terror, 1977-1995." *Journal of Conflict Resolution* 46(3):427–450.

Defense Mapping Agency. 1992. "Development of the Digital Chart of the World." U.S. Government Printing Office.

DellaVigna, Stefano, Ruben Enikolopov, Vera Mironova, Maria Petrova and Ekaterina Zhuravskaya. 2011. Cross-border media and nationalism: Evidence from Serbian radio in Croatia. Technical report National Bureau of Economic Research.

Djankov, Simeon, Caralee McLiesh, Tatiana Nenova and Andrei Shleifer. 2001. Who owns the media? Technical report National Bureau of Economic Research.

Drakos, Konstantinos and Andreas Gofas. 2006. "The Devil You Know but Are Afraid to Face Underreporting Bias and its Distorting Effects on the Study of Terrorism." *Journal of Conflict Resolution* 50(5):714–735.

Earl, Jennifer, Andrew Martin, John D McCarthy and Sarah A Soule. 2004. "The use of newspaper data in the study of collective action." *Annual Review of Sociology* 30(1):65–80.

Eck, Kristine. 2012. "In data we trust? A comparison of UCDP GED and ACLED conflict events datasets." *Cooperation and Conflict* 47(1):124–141.

Freedom House. 2015. *Freedom in the World 2015: The Annual Survey of Political Rights and Civil Liberties*. Rowman & Littlefield.

Galtung, Johan and Mari Holmboe Ruge. 1965. "The structure of foreign news the presentation of the Congo, Cuba and Cyprus Crises in four Norwegian newspapers." *Journal of peace research* 2(1):64–90.

Gans, Herbert J. 1980. *Deciding what's news. A study of CBS News, NBC Nightly News, Newsweek and Time*. New York: Vintage Books.

Gehlbach, Scott and Konstantin Sonin. 2014. "Government control of the media." *Journal of Public Economics* 118:163–171.

Global Administrative Areas. 2012. "GADM Database of Global Administrative Areas.".

Gohdes, Anita and Megan Price. 2012. "First Things First Assessing Data Quality before Model Quality." *Journal of Conflict Resolution* 57(6):1090–1108.

Gorodnichenko, Yuriy and Tymofiy Mylovanov. 2015. "War and Freedom of Speech in Ukraine." *VoxUkraine* .

Hallin, Daniel C. 1989. *The uncensored war: The media and Vietnam*. Univ of California Press.

Hammond, Jesse and Nils B Weidmann. 2014. "Using machine-coded event data for the micro-level study of political violence." *Research & Politics* 1(2):1–8.

Hendrix, Cullen S and Idean Salehyan. 2015. "No News Is Good News: Mark and Recapture for Event Data When Reporting Probabilities Are Less Than One." *International Interactions* 41(2):392–406.

Iyengar, Shanto and Kyu S Hahn. 2009. "Red media, blue media: Evidence of ideological selectivity in media use." *Journal of Communication* 59(1):19–39.

Jackman, Simon. 2005. "Pooling the polls over an election campaign." *Australian Journal of Political Science* 40(4):499–517.

Kumar, Deepa. 2006. "Media, war, and propaganda: Strategies of information management during the 2003 Iraq war." *Communication and Critical/Cultural Studies* 3(1):48–69.

Leetaru, Kalev and Philip A Schrodt. 2013. "Gdelt: Global data on events, location, and tone, 1979–2012." Presented at International Studies Association Annual Convention.

Levendusky, Matthew S. 2013. "Why do partisan media polarize viewers?" *American Journal of Political Science* 57(3):611–623.

Libicki, Martin. 1995. *What is Information Warfare?* Washington, DC: National Defense University Press.

Loveland, TR, BC Reed, JF Brown, DO Ohlen, Z Zhu, LWMJ Yang and JW Merchant. 2000. "Development of a global land cover characteristics database and IGBP DISCover from 1 km AVHRR data." *International Journal of Remote Sensing* 21(6-7):1303–1330.

Lyall, Jason. 2010. "Are coethnics more effective counterinsurgents? Evidence from the second Chechen war." *American Political Science Review* 104(01):1–20.

Mayer, Frederick W. 2012. Stories of climate change: Competing narratives, the media, and US public opinion 2001–2010. Technical report Joan Shorenstein Center on the Press, Politics and Public Policy.

McCarthy, John D, Clark McPhail and Jackie Smith. 1996. "Images of protest: Dimensions of selection bias in media coverage of Washington demonstrations, 1982 and 1991." *American Sociological Review* 61(3):478–499.

Moeller, Susan D. 1999. *Compassion fatigue: How the media sell disease, famine, war and death*. Psychology Press.

Morton, Linda P and John Warren. 1992. "Proximity: Localization vs. distance in PR news releases." *Journalism & Mass Communication Quarterly* 69(4):1023–1028.

Nichols, James D. 1992. "Capture-recapture models." *BioScience* 42(2):94–102.

Norris, Pippa, Montague Kern and Marion R Just. 2003. *Framing terrorism: The news media, the government, and the public*. Psychology Press.

Pariser, Eli. 2012. *Filter Bubble: Wie wir im Internet entmündigt werden*. Hanser.

Pickup, Mark and Richard Johnston. 2008. "Campaign trial heats as election forecasts: Measurement error and bias in 2004 presidential campaign polls." *International Journal of Forecasting* 24(2):272–284.

Pomerantsev, Peter. 2014. "Can Ukraine Win Its Information War With Russia?" *The Atlantic* .

Popkin, Samuel L. 1994. *The reasoning voter: Communication and persuasion in presidential campaigns*. University of Chicago Press.

Prior, Markus. 2007. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press.

Raleigh, Clionadh, Andrew Linke, Håvard Hegre and Joakim Karlsen. 2010. "Introducing acled: An armed conflict location and event dataset special data feature." *Journal of Peace Research* 47(5):651–660.

Rampton, Sheldon and John Clyde Stauber. 2003. *Weapons of mass deception: The uses of propaganda in Bush's war on Iraq*. Penguin.

Rosengren, Karl Erik. 1974. "International news: Methods, data and theory." *Journal of Peace Research* 11(2):145–156.

Shoemaker, Pamela J and Akiba A Cohen. 2012. *News around the world: Content, practitioners, and the public*. Routledge.

Snegovaya, Maria. 2015. Putin's Information Warfare in Ukraine. Technical report Institute for the Study of War.

Snyder, David and William R Kelly. 1977. "Conflict intensity, media sensitivity and the validity of newspaper data." *American Sociological Review* 42(1):105–123.

Stroud, Natalie Jomini. 2011. *Niche news: The politics of news choice*. Oxford University Press on Demand.

Sweeney, Michael S. 2001. *Secrets of Victory: The Office of Censorship and the American Press and Radio in World War II*. Univ of North Carolina Press.

Taylor, Philip M. 1992. *War and the media: Propaganda and persuasion in the Gulf War*. Manchester University Press.

The Economist. 2015. "Battle of the memes." *The Economist* .

Vikhrov, Maksim. 2014. "Zachem Ukraine 'ministerstvo pravdy'? [Why does Ukraine need a 'ministry of truth'?]." *Slon.ru* .
URL: *http://slon.ru/world/zachem_ukraine_ministerstvo_pravdy-1193446.xhtml*

Weidmann, Nils B. 2014. "On the accuracy of media-based conflict event data." *Journal of Conflict Resolution* 59(6):1129–1149.

Weidmann, Nils B. 2016. "A closer look at reporting bias in conflict event data." *American Journal of Political Science* 60(1):206–215.

Whitten-Woodring, Jenifer and Patrick James. 2012. "Fourth estate or mouthpiece? A formal model of media, protest, and government repression." *Political Communication* 29(2):113–136.

Woolley, John T. 2000. "Using media-based data in studies of politics." *American Journal of Political Science* 44(1):156–173.

Wu, Haoming Denis. 1998. "Investigating the determinants of international news flow a meta-analysis." *International Communication Gazette* 60(6):493–512.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge University Press.

Zhukov, Yuri M. 2016. "Trading hard hats for combat helmets: The economics of rebellion in eastern Ukraine." *Journal of Comparative Economics* 44(1):1–15.