1 Introduction

Our information environment is changing. Shifts in technology are facilitating fundamental transformations in the kinds of information to which we regularly have access. There are more readily available pet pictures, baby pictures, and celebratory bread-making pictures now than at any time in recent history. There is simultaneously more sensationalistic, negative, polarizing, and intentionally misleading news content. The former is at least in part a product of the latter – the Internet has created an infinite amount of space waiting to be filled with content, after all. This is a central theme of the pages that follow.[1]

This Element is partly a response to the environment in which it was started, several months into the 2020 COVID-19 pandemic, in the midst of protests following the police killing of George Floyd, and in the lead-up to the 2020 US presidential election. But the idea developed here was in evidence long before the appearance of the novel coronavirus and heightened news attentiveness to the Black Lives Matter (BLM) movement. There has for some time been concern that mass media are so inclined toward negative and/or sensationalistic and/or conflict-laden content that we will soon be entirely consumed by it. Given the past few years (or even the past few months) of news coverage, this prediction is well worth considering. There is, after all, a considerable body of work suggesting that media competition drives up negativity and sensationalism (e.g., Dunaway 2008, 2013), and an online news environment in which competition for audiences is considerable. There also is a burgeoning literature on the increasing attention given to "fake" news – news that typically capitalizes on human biases toward negative, sensationalistic, and conflict-laden information. Political polarization appears to be on the rise (e.g., Iyengar et al. 2019), and misinformation is prevalent (although thus far limited in impact; see, e.g., Guess et al. 2019). Even in the midst of a viral pandemic, wearing a face mask has turned into a contested political statement (Druckman et al. 2021).

[1] Much of our thinking on this topic has been affected by collaborative projects over the past decade. Research on outlyingness as a driver of attention to information was developed with P. J. Lamberson. Work on differences in information across media platforms was done in collaboration with Mark Daku, Lauren Guggenheim, Dan Hiaeshutter-Rice, Patrick Kraft, Kerri Milita, Josh Pasek, and John Barry Ryan. Psychophysiological studies of negativity biases were conducted with Patrick Fournier, Lilach Nir, Johanna Dunaway, Kevin Arceneaux, and Bert Bakker. All of this work features prominently in what follows. Our argument here has also benefited from conversations and advice from Sarah Bachleda Fioroni, Dan Hiaeshutter-Rice, Ariel Hasell, Patrick Fournier, Michael Wagner and Brian Weeks and from presentations in the Department of Communication and Media at the University of Michigan, the Department of Communication at the University of California, Los Angeles, the Department of Government at the University of Texas at Austin, and the Hot Politics Lab at the Amsterdam School of Communication Research, University of Amsterdam.

Is it our destiny to be surrounded by overwhelmingly negative news content? Our answer in the pages that follow is *No*. Human interests vary over time and across individuals. Technology increasingly facilitates information consumption that reflects these interests, sometimes with perverse consequences, but often with advantages as well. One such advantage is the increasing number of at-home singing and dancing videos and we're-in-this-together tweets.

We make our case in the following stages. We begin with a consideration of the valence of newspaper coverage in roughly sixty-two thousand news stories over the first eight months of 2020, with a focus on content related to both COVID-19 and BLM. We then consider the valence of more than three hundred thousand television news stories over the past thirty years. Trends in both newspaper and television coverage highlight the tendency for the valence of news to be "self-correcting" – to display very negative valence at some points, to be sure, but also to return to a more mildly negative equilibrium shortly thereafter. Why does this matter? Why does it happen? We consider the first of these questions by revisiting past work on the normative and empirical implications that the "valence" of information has for engagement, participation, preferences, and information processing. We argue that we should be interested in "valence-based asymmetries" in news coverage (i.e., the tendency for content to be predominantly positive or negative) because we know that valence matters for behavior and well-being, political and otherwise. The bulk of our monograph then focuses on the second question, and moreover on the possibility that the volume and availability of positively valenced news is likely to increase rather than decrease.

That possibility is, we believe, driven by three factors:

(1) *Valence-based asymmetries vary over time*. There is variation in valence-based asymmetries across time, driven by outlyingness, novelty, and/or adaptive processing.
(2) *Valence-based asymmetries vary across individuals*. Different people process information differently. We all have long-standing systematic biases in information selection, and in the ways in which we respond to that information.
(3) *Technology facilitates diverse news platforms catering to diverse preferences.* Technology is facilitating the ready availability of information that varies in many different ways – by topic, by political ideology, and by valence as well. The proliferation of platforms increasingly allows media consumers to move between sources in order to achieve an "ideal" balance of positive and news content. Media may well adapt by producing more

positive content; this may be seen at the level of individual media outlets, but it may be especially evident looking at "media" as a whole.

This three-part argument is somewhat more optimistic about the future of media content than is typical, at least at the present time. It also is partly a response to our own prior work, which has focused on the prevalence of negative campaigns and news content (e.g., Krupnikov 2011; Krupnikov and Piston 2015), and the related tendency for humans to be more attentive to and affected by negative rather than positive information (e.g., Soroka 2014). That work is most often focused on the central tendency – the average – across humans and over time. The argument that follows concentrates more on the variation around that average. Acknowledging this variation is an important step toward understanding why our information environment is in fact not inevitably overwhelmingly negative, and why there are at the present time reasons to expect an increasing rather than decreasing volume of good news.

Note that our argument should not be confused with one about future news events. We do not expect miraculous cures for deadly diseases or an end to injustice. Rather, we believe that even in a context marked by negativity and sadness, news content that focuses on positive events – even if these events are minor – will increasingly emerge. Put differently, our argument is not that good news will emerge because the course of human history will turn in a positive direction (this may turn out to be so, but we simply do not know), but rather that the news environment will increasingly search out and make room for more positively valenced content.

2 The Valence of News Coverage

The "valence" of news coverage refers to the degree to which the information is affectively positive or negative. In some instances, valence is easily established. News about sickness and death is almost always perceived as negative by everyone, for instance. Increasing unemployment is almost always perceived as negative news, while news about decreasing unemployment is almost always perceived as good news (Soroka 2006). But there are always individuals with different viewpoints. A wealthy, well-employed minority may see rising unemployment as advantageous insofar as it is linked to decreasing inflation, for instance. Partisan groups will often see policy outcomes through very different lenses. The valence of information is rarely purely objective. It can be more or less subjective, but it is almost always at least partly a matter of perspective.

Distinguishing between primarily *objective* versus primarily *subjective* valence is possible, and often critical in research that seeks to examine

responses to negative and positive information. Acknowledging that a large amount of mediated information may be positive or negative to different people is important. The distribution of valence in mediated information is different for everyone, at least in part because our perceptions of the valence of information vary. There nevertheless are limits to this subjectivity. It is very unlikely that people look at news coverage in the midst of a pandemic, with rising unemployment and the current second/third wave of a virus, and think that news content is mostly positively valenced.[2] Indeed, our shared assessment of the valence of mediated information is revealed in the considerable body of work that links media coverage to public attitudes. Regardless of the direction of the causal effect (see, e.g., Wlezien and Soroka 2018), either positive and negative media coverage has an impact on public attitudes, or positive and negative public attitudes are regularly reflected in the language of news. In order for either of these to be the case – and plenty of evidence suggests that they are (e.g., Dalton et al. 1998; Hopmann et al. 2010; Shaw 1999; Soroka et al. 2009, 2015) – there *has* to be some shared sense of what information is negative and what information is positive.

There are in addition good reasons to consider the valence of news above and beyond its substantive content. Living in complex, information-rich environments requires that we develop simple, fast mechanisms for categorizing and paying attention to (or not) information. The human brain is thus finely tuned to quickly recognize and respond to the valence of information (e.g., Zajonc 1980), and in some instances affective responses may matter more to our decision-making than "cognitive" assessments of incoming information (e.g., Loewenstein et al. 2001). The quick recognition of valence is central to what Kahneman (2013) famously describes as System 1 (fast, automatic, affective, unconscious) thinking versus System 2 (slow, deliberative, conscious) thinking. And his argument builds on a considerable body of prior work, including Epstein's (1994) discussion of "experiential" versus "analytical" thinking, Damasio's research on "somatic markers" (2005), and Slovic et al.'s (2002) work on the "affect heuristic." A growing body of neurological work further highlights the relative significance of and speed with which the human brain responds to the valence of information (e.g., Bayer et al. 2010; Feng et al. 2014; Gianotti et al. 2008).

This is not to say that we do not have substantive interests – some people are interested in news about football and others are not. But our decision to look for, select, and/or respond to information will often be conditioned in the first

[2] Petersen et al. (2020) suggest that some citizens purposefully circulate fake news out of a "need for chaos." They may consequently derive some enjoyment from circulating negative information. But even these readers likely recognize that the information is negatively valenced.

instance by the valence of that information. Indeed, in many cases we may prioritize the valence of news over the actual substance; or, to push this line of thinking even further, sometimes the valence is the substance. To paraphrase Marshall McLuhan: *the mood is the message*. For this reason, a consideration of (just) the valence of news is of some importance.

What, then, is the valence of current-affairs news coverage? Note that current-affairs news coverage is a small proportion of media content, particularly when "media content" includes information circulating on social media. In the closing section we consider what our argument might look like in a multiplatform media environment that is predominantly not about current affairs. For the time being, however, we focus on news, primarily in "traditional" or "legacy" newspaper and broadcast outlets. And on this front there are vast literatures detailing the tendency for this content to be sensationalistic (typically in a negative way; see, e.g., Davie and Lee 1995; Harmon 1989; Hofstetter and Dozier 1986; Ryu 1982), problem-oriented (Altheide 1997), and negative more broadly (e.g., Benoit et al. 2005; Diamond 1978; Fallows 1997; Farnsworth and Lichter 2007; Just et al. 1996; Kerbel 1995; Lichter and Noyes 1996; Niven 2000; Patterson 1994; Sabato 1991).

Has this negativity been increasing over time? Let us consider two different exhibits based on news coverage over the past forty years. The first is relatively recent. Figure 1 shows the frequency and sentiment of newspaper coverage of COVID-19 and the BLM movement over the first eight months of 2020. We focus on this coverage for several reasons. First, COVID-19 and BLM form the two uniquely defining news stories of 2020. Second, these two stories are founded on negative events. Coverage of COVID-19 begins with a global pandemic that costs millions of people their lives; coverage of BLM begins with police brutality and racial injustice. These stories, then, form notable examples – journalists would need to work deliberately to include positivity in this coverage.

The data in Figure 1 are drawn from the Lexis-Nexis full-text news archive and include all stories in the front sections of the *New York Times*, *Washington Post*, and *LA Times*, alongside all coverage in the *Atlanta Journal-Constitution*, *Minneapolis Star Tribune*, *Philadelphia Inquirer*, and *USA Today*. (We focus on front sections for the first three papers due to the size of those papers; the latter three have, in contrast, much less content to download and analyze.) From that database, we identify every news story about either COVID-19 or the BLM movement using a simple keyword search.[3] Any article that does not fall into one of these categories is placed in the "other" category.

[3] The keyword search is very straightforward. Any article with more than one mention of "COVID," "coronavirus," or "pandemic" is identified as a COVID-19 article. Any article with more than one mention of "Black Lives Matter" or "BLM" is identified as a BLM article. Note

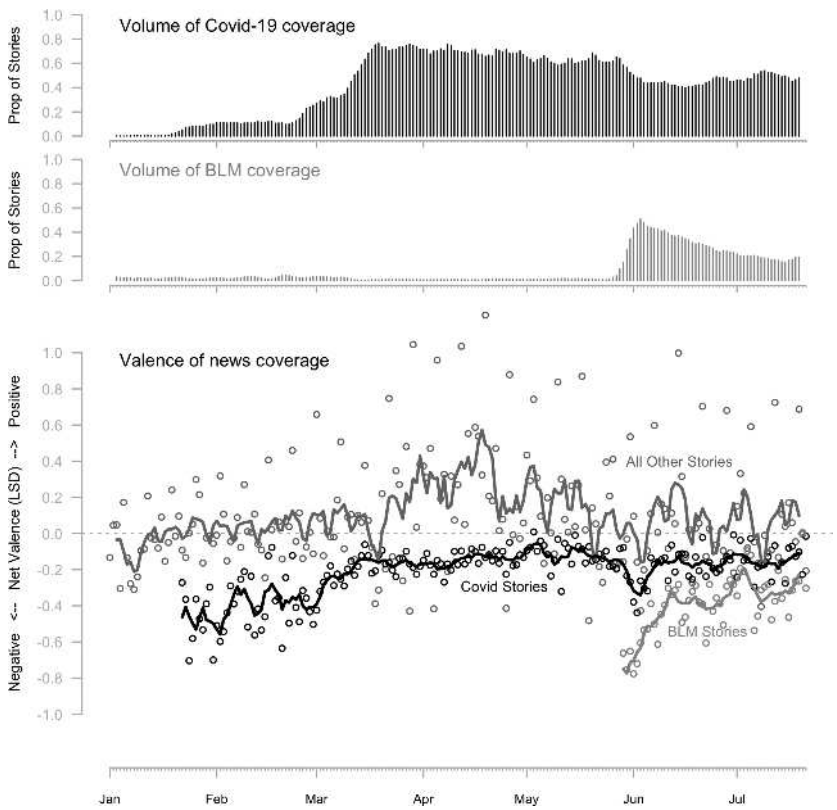

Figure 1 The sentiment of newspaper coverage, 2020

Based on COVID and BLM stories in the *Atlanta Journal-Constitution*, *Los Angeles Times*, *New York Times*, *Minneapolis Star-Tribune*, *Philadelphia Inquirer*, *USA Today*, and *Washington Post.* The line shows a rolling center-weighted five-day average. Sentiment is based on the LSD, as implemented in quanteda in R. Net sentiment is calculated as follows: log [(pos. counts + 0.5) / (neg. counts + 0.5)].

Data are analyzed at the story level, and the valence of stories is estimated using the Lexicoder Sentiment Dictionary (LSD) (Young and Soroka, 2012). The LSD is composed of roughly five thousand words, half positive and half negative. We estimate counts of positive and negative words using the implementation of the LSD in the quanteda package in R (Benoit et al. 2018). We then produce an estimate of "net sentiment" using a measure suggested in Lowe et al. (2011) and used in Proksch et al. (2019), as follows: log [(pos. counts + 0.5) /

that we require *more than one* mention of either of these words in order to categorize an article as COVID-19 or BLM. This is because a large number of articles during this period mentioned one or the other topic very briefly in passing. (This was especially true for COVID-19.) Our intention here was thus to separate out the articles that actually discuss one or the other topic.

(neg. counts + 0.5)]. This is an empirical logit, slightly smoothed toward zero, and very highly correlated with another standard approach: (pos. counts − neg. counts) / total words. The resulting net sentiment is greater than zero when positive words outnumber negative words, and less than zero when negative words outnumber positive ones.[4]

The top panels of Figure 1 do not focus on the valence but rather on the volume of news coverage of both COVID-19 and BLM. In mid-March, nearly 80 percent of our sample of news coverage mentions COVID-19 multiple times, and even at the end of July nearly 50 percent of news coverage is focused on COVID-19. This is to be expected. Although cases of the coronavirus had been documented in the United States as early as January, cases spiked in March, leading states to implement numerous mitigation policies. Attention to BLM peaks in early June, when roughly 50 percent of coverage includes multiple mentions of BLM. This timing is also reflective of news events: George Floyd was killed on May 25, 2020, which led to protests in late May. (Note that the COVID-19 and BLM categories are not mutually exclusive – articles can be and regularly are about both.)

What is the valence of these articles, and how does it compare with the valence of articles *not* about either COVID-19 or BLM? The bottom panel of Figure 1 shows the average valence of our three different categories of news articles daily. Circles show daily averages; lines show a rolling center-weighted five-day average for each category. There are in our view three trends especially worth highlighting.

First, although the valence of COVID-19 coverage is markedly negative to start, it becomes more positive through March and then remains at that moderately negative level for the subsequent four months. In fact, comparing the top and bottom panels of Figure 1, the valence of COVID-19 coverage improves at roughly the same time as it becomes highly salient. The less negative valence of COVID-19 coverage is also accompanied by roughly sixty days (from mid-March to mid-May) of decidedly positive *not* COVID-19 coverage. The end result is that the complexion of news coverage at the height of the pandemic is

[4] Past work suggests that the LSD is as reliable if not more reliable for news content than other general-purpose sentiment dictionaries (Young and Soroka 2012). Recent work highlights some advantages of sentiment tools that rely on a combination of human coding and machine learning (Van Atteveldt et al. n.d.). Note that corpus-specific machine learning-based tools should in most cases be more accurate than a general-purpose dictionary. The dictionary has typically been tested on another corpus, after all, while the machine learning typically relies on the specific corpus under investigation. To be clear: the central difference may be less about dictionaries versus machines, and more about general-purpose versus corpus-specific approaches. Regardless, at the very high level of aggregation used here the differences between various dictionaries and machine-learning approaches tend to be minimized.

not as negative as we might anticipate, and on balance it is slightly more positive than before the pandemic.

A similar dynamic is evident for coverage of BLM. Coverage of BLM is initially very negative, but that negativity is cut in half after the first two weeks of protests. There is no point at which BLM coverage is positive on average, and it is more negative than COVID-19 coverage throughout June and July. But like coverage of COVID-19, the initial negativity of BLM coverage dissipates relatively quickly.

What accounts for the seemingly fast moderation of the negativity in news coverage of two ongoing, long-standing, and mostly not improving phenomena? There are surely multiple drivers. Widespread misunderstandings about the magnitude and severity of the COVID-19 pandemic were likely one source of decreased negativity in COVID-19 coverage, for instance, just as systematic racism likely muted prolonged attentiveness to the problems highlighted by the BLM movement. But we also suspect that the positive shift in the valence of 2020 news coverage is a common phenomenon – evident in media reactions to a wide range (if not all) of the major problems over the past few decades.

Consider, for instance, Figure 2, which shows the trend in the valence of television news coverage over the past thirty years. The analysis in this instance is based on the full transcript of every evening newscast from 1990 onward on ABC, CBS, and NBC, extracted from the Lexis-Nexis database. As in Figure 1, circles in Figure 2 show the average sentiment across all stories on all three

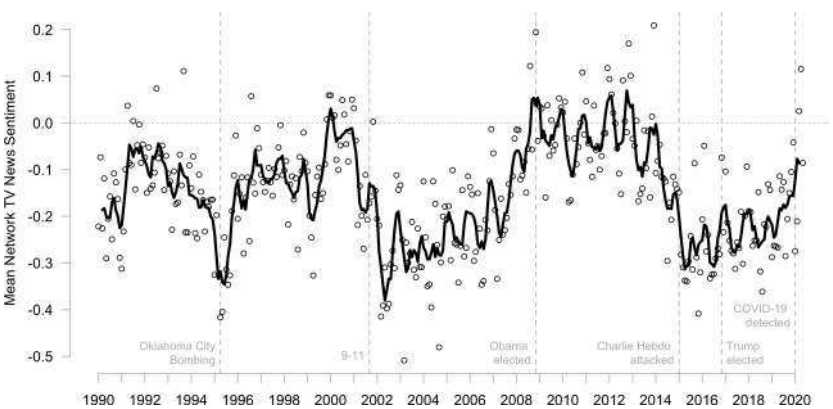

Figure 2 The sentiment of network news, 1990–2018

Circles show the monthly averages for all stories on national nightly news programs on ABC, CBS, and NBC. The line shows a rolling center-weighted five-month average. Sentiment is based on the LSD, as implemented in quanteda in R. Net sentiment is calculated as follows: log [(pos. counts + 0.5) / (neg. counts + 0.5)].

networks monthly. The line shows a rolling five-month center-weighted average of these monthly values.

Note that the LSD does not distinguish between different types of negativity (see, e.g., Lengauer et al. 2012). Sentiment expressed by interviewees, commentary and evaluations by journalists, and the valence of the events themselves (i.e., deaths versus heroics) are all lumped together in this very general measure of the valence of the language of newscasts. The measure also does not distinguish between topics or themes – the economy, health care, the environment, and celebrity breakups are all pooled in the measure of sentiment shown here. This is purposeful. There most certainly are important differences between the information conveyed in reporting versus commentary, or political versus entertainment news. But all of this information has a valence – positive, negative, or neutral – and there are good reasons to believe that both news production and news consumption are powerfully affected by this critical aspect of content.

With this in mind, then, there are three features of Figure 2 worth noting. The first is prosaic but methodologically important: the peaks and troughs in Figure 2 correspond with events that we should expect to have a marked impact on media sentiment. Some of the most important moments are shown in Figure 2 using vertical lines. The Oklahoma City bombing, 9/11, and much of 2015 (beginning with the *Charlie Hebdo* shooting, followed by school shootings) stand out as particularly negative moments in TV network news. The month Obama was elected is the second most positive month in thirty years of TV reporting (surpassed only by the month in which Nelson Mandela's passing produces a large number of positive remembrances). George W. Bush begins his term in office with relatively positive media coverage, although this collapses quickly and never recovers. Obama's terms in office see comparatively positive media coverage, although coverage sours in 2015 due to events abroad and at home. All of these trends offer simple but important confirmations of the concurrent validity of the sentiment measure.

The second feature of Figure 2 worth noting is that the vast majority of months have an average sentiment that is well below zero. (Zero is shown in Figure 2 by a horizontal dotted line.) Of the 366 months for which we have data, just 42 (11 percent) have an average sentiment above zero. The tendency for negative words to outnumber positive words in media coverage is readily apparent in these data.

Is this a bad thing or a good thing? We really do not know. There are certainly reasons for concern about biases in the valence of news coverage. Consider first the normative concerns about systematically negative news coverage and political engagement. There are long-standing worries that negative news coverage

contributes to declining trust in political institutions and disengagement in politics (Farnsworth and Lichter 2007; Moy and Pfau 2000; Patterson 1994). Similar themes echo in work on "media malaise" in which the sentiment of news coverage decreases trust and political efficacy (e.g., Robinson 1976), and in work on the "spiral of cynicism" in which conflict-oriented coverage contributes to public cynicism about politics (e.g., Cappella and Hall-Jamieson 1997; Valentino et al. 2001).

That said, there also are reasons to expect negative news to increase attention to and engagement with politics. Indeed, work evaluating both the media malaise and the spiral of cynicism hypotheses has regularly finds evidence of the opposite possibility (e.g., Newton 1999; Strömbäck and Shehata 2010; de Vreese 2005). This makes sense given that people seem more interested in negative news content (e.g., Trussler and Soroka 2014). It may also be that the impact of affectively negative content has heterogeneous effects, depending in part on the specific negative emotion that content produces (e.g., Brader 2006; Nabi 1999; Valentino et al. 2009). There is of course a vast literature chronicling the heterogeneous and context- and candidate-dependent effects of negative political advertising (e.g., Fridkin and Kenney 2011; Krupnikov 2011, 2014; Krupnikov and Piston 2015). It is similarly unclear whether negative news content will on average increase or decrease our engagement with politics. We suspect that there is likely an "ideal point" where negativity in media coverage is concerned: there should be enough negativity to keep us interested, but not so much over time that we withdraw from news consumption and politics (Soroka 2014). Of course, we do not know what that "ideal point" is.

It also is unclear whether a bias toward negatively valenced content reflects journalism run amok or journalists doing their job. One account of negativity in news content emphasizes a post-Watergate increase in journalistic skepticism that produces systematically biased and cynical reporting (e.g., Patterson 1994). Another is that we expect media to act as a "Fourth Estate," monitoring and highlighting errors in representative government. Doing this necessarily produces media content focused on negative outcomes. Indeed, negative content may reflect greater levels of information than positive content insofar as that negativity reflects a critique or contrasting of political positions (e.g., Geer 2008). Moreover, there are times when media content quite clearly should be negative because current events are negative. It would be hard to sustain the argument that during a global pandemic media content is too negative. Sometimes, perhaps often, negatively valenced news accurately reflects the state of the world.

In short, it needn't be the case that negatively valenced media content will drive all citizens to withdraw from politics, nor is it necessarily a signal that