UNIVERSITÀ COMMERCIALE "LUIGI BOCCONI"
PHD SCHOOL

PhD program in Economics and Finance

Cycle: 33º

Disciplinary Field (code): SECS-P/01

Essays in Media and Political Economics

Advisor: Guido Tabellini

Co-Advisor: Eliana La Ferrara

PhD Thesis by

Jaime Marques Pereira

ID number: 3049318

Year 2023

Contents

Chapter 1

Do TV News Chase Twitter?

A High-Frequency Analysis

Abstract

In order to understand if cable news outlets cover information disclosed on social media, I analyze U.S. cable news Trump-related coverage in narrow time windows centered around President Trump's tweets. Using high-frequency data on cable television news to describe how were @realDonaldTrump tweets covered by cable news. I find evidence in favor of cable new outlets having covered President Trump's tweets in real time. Using an exhaustive dataset on online news, I then study if this latter coverage was related or not to past news. I find that on average, Fox News and MSNBC's coverage of President Trump's tweets was unrelated to recent news, an indication that President Trump was able to temporarily shift these outlets' agenda simply by tweeting. Lastly, I take advantage of recent advancements in natural language processing techniques to classify President Trump's tweets into a set of interpretable topics. I then allow for cable outlets to react differently to different types of Trump tweets. Here I find that President Trump's power over cable news attention was not specific to any particular topic. This result suggests that cable news stations tended to differentiate themselves not by covering different Trump statements but instead by slanting differently a relatively similar distribution of stories.

1.1. Introduction

U.S. politicians increasingly rely on social media to issue public statements (Pew, 2021b). This communication strategy allows politicians, in theory, to directly reach a significant share of U.S. adults without intermediaries. In fact, as much as half of U.S. adults use social media as one of their news sources (Pew, 2021d). However, this half is mainly composed by young adults, a generation which is still a minority in U.S. presidential elections (Pew, 2021a). Indeed, U.S. media consumption today is segmented across ages - young adults rely heavily on social media for news, old adults instead use television (Pew, 2021c).

In a context such as this one - in which U.S. television plays a major role as a primary news source for a majority of U.S. voters - it seems inefficient for politicians to use social media as their main means of communication. Nonetheless, U.S. television outlets recurrently cover politicians' statements on social media. This paper focuses exactly on this type of coverage. A coverage that amplifies a selected set of political statements to an audience that would otherwise not know about them. Most importantly, a coverage that does so with an audience that is more likely to vote and so, is expected to be more sensitive to this type of political messaging.

In this paper, I currently describe how this coverage is cast. In future work, I intend to quantify by how much does this coverage affect politicians' behavior on social media and general social media discussions. To be more specific, in this paper I intend to focus on three questions. How are politicians' social media statements covered by television news outlets? Are these statements shaped by this same coverage and if so, how? What effect does this type of coverage have on social media forums?

To answer these questions I study a relevant case study - U.S. cable news outlets' coverage of President Trump's tweets. I turn to this case for a variety of reasons.

To start, due to its likely high relevance from an international perspective. It focuses on a singular political figure, one that has most likely served as reference to other foreign rising populist politicians. Be it through his use of Twitter as a main communication tool - in fact, throughout his presidency, Donald J. Trump recurrently used Twitter to, among other acts, disclose cabinet nominations (e.g., Trump, 2019a) and issue first-reactions to statements from other figures (e.g., Trump, 2019b) - or through his active engagement with television news outlets - indeed, President Trump openly used Twitter to react live to given cable news shows (e.g., Gertz, 2018).

Second, given how sizeable of an effect this type of coverage is likely to have had on an array of relevant outcomes. Cable news coverage by itself has been shown to be able to politically persuade individuals to espouse conservative views and, more recently, to determine health behaviors (DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2017; Ash et al., 2021; Allcott et al., 2020; Bursztyn et al., 2020). Trump tweets on another hand have been associated to increases in racial hate speech, surges in political violence, decreases in trust on electoral institutions (respectively, Müller and Schwarz, 2018; Brown and Sanderson, 2020; Clayton et al., 2021).

To answer these questions causally, I take advantage of an exhaustive dataset of timestamped transcripts on cable news television. I use this dataset to construct two Trump-related high-frequency coverage measures - (a) amount of time cable outlets covered Trump-related stories and (b) similarity in content between cable news coverage and Donald J. Trump's tweets. Then, I study both measures in narrow time intervals centered around Trump tweets. I am able to pin-down how much of a causal impact did President Trump's tweets had on cable news coverage decisions (assuming that any change in cable coverage minutes after a tweet can be only due to that tweet). I find that cable news outlets tended, on average, to turn their attention towards those issues tweeted by President Trump in a matter of minutes.

Afterwards, I leverage on an exhaustive dataset of news posted on Facebook to understand whether these shifts in coverage could simply be explained by President Trump being systematically faster at reacting to last minute events (e.g., a news event happens at period t-1, President Trump reacts at period t and cable outlets react at period t+1, this late reaction by television news ultimately implying a convergence in content between cable news and Trump tweets in period t+1). I compare each shift in coverage with past online news. I find that for certain outlets (Fox News and MSNBC), these shifts in media attention are unrelated to recent breaking events. These results seem to suggest that President Trump was able to temporarily set Fox News and MSNBC's agenda through his tweets.

In a third analysis, I take advantage of recent advancements in Natural Language Processing (NLP) techniques to classify President Trump's tweets into an array of interpretable topics. Then, I allow for cable outlets to react differently to different types of Trump tweets. I find that President Trump tended to shift cable news attention irrespective of which topic he tweeted about (with some minor exceptions). This result seems to suggest that television stations tended to differentiate themselves not by covering different Trump statements but, instead, by slanting differently a relatively similar distribution of Trump stories.

In future iterations of this chapter, I will use an extensive dataset on tweets posted in reply to President Trump to study whether cable outlets' reactions were a function of how was a tweet received on Twitter minutes after being posted. This analysis is intended to shed light on which factors explain outlets' coverage decisions (à la Cagé et al., 2020; Zhuravskaya et al., 2021). These become important in a context in which President Trump acts strategically, aligning his tweets with these outlets' editorial criteria.

In addition, I will take advantage of previous analyses to identify within each outlet those shows that actively followed President Trump's tweets. Then, I intend to use this information to study not only if Donald J. Trump actively timed his tweets to moments in which this type of show was being aired but, also, whether he changed which topics he addressed within a day according to which content each show was more likely to cover (conditional on its past history). This exercise is aimed at understanding not only if but also how did cable news coverage shape President Trump's tweeting behaviors (connected to a literature focused on identifying strategic behaviors from political actors; see Gratton et al., 2018; Djourelova and Durante, 2019).

Concerning this chapter, it is organized as follows. Section 1.2 describes the data sources used until now and the variables constructed. Section 1.3 describes the empirical setting and specifications estimated. Section 1.4 presents the results and robustness exercises implemented until now. Section 1.5 concludes and discusses ongoing work.

1.2. Data

1.2.1 Sources

U.S. cable news transcripts. Timestamped transcripts for the three main cable news stations in the U.S. - CNN, Fox News and MSNBC. This dataset covers close to the universe of shows broadcasted from January 2017 to January 2021. It was kindly provided by the TV News Archive (Link).

Tweets by @realDonaldTrump. Timestamped tweets posted by Donald J. Trump's personal Twitter account, @realDonaldTrump (Link)[1]. This dataset covers the universe of tweets posted by Donald J. Trump from January 2017 to January 2021. It was collected and later on made publicly available by the Trump Twitter Archive (Link).

Facebook news posts by U.S. outlets. Text, timestamp and other post-specific metrics for the universe of Facebook posts released by a comprehensive subset of U.S. national news outlets. This dataset has been collected by CrowdTangle (Link). It covers every post released from January 2017 to January 2021.

[1] President Trump's personal Twitter account. This account was created in March 2009. It issued a first set of tweets in May 2009. It was permanently suspended by Twitter on January 8 2021.

1.2.2 Variables

@realDonaldTrump tweets

I construct a count variable defined as:

$$D_t = |\{TrumpTweet: \ Timestamp(TrumpTweet) \in t\}| \tag{1.1}$$

where $TrumpTweet$ stands for a Trump tweet; $Timestamp(TrumpTweet)$ stands for when in time was $TrumpTweet$ posted; t stands for a 15-minute absolute time period (e.g., first 15 minutes of 1pm).

D_t is a quarter-hourly count of Donald J. Trump tweets. It counts tweets belonging to a selected sample of Trump statements: it does not count retweets as it is intended to point towards original Trump tweets; it excludes short tweets, to filter out statements with little information and so, of little general interest.[2]

Event windows

I will tend to focus on event windows - i.e., time intervals centered around Trump tweets.

In formal terms, as illustrated above, let a treatment period be a period in which Donald J. Trump tweeted at least once (i.e., a period in which $D_t > 0$). This treatment period defines an event window w composed of 3 pre-treatment and 2 post-treatment periods (i.e., a window of 1 hour and 30 minutes).

[2]See Appendix 1.6.1 for different descriptive statistics and more technical details on D_t.

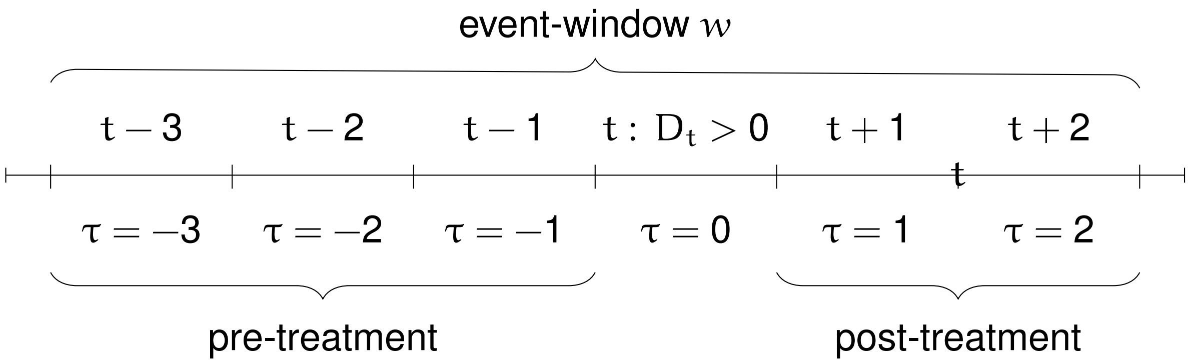

Figure 1.1: **Event window.** The figure illustrates a generic event window with 6 periods, 1 treatment ($\tau = 0$) and 5 relative-to-treatment periods ($\tau \neq 0$); t stands for absolute time; τ stands for relative-to-treatment time (within event window w).

Television coverage

I construct a cable news coverage measure defined as:

$$C_{n,w,\tau} = \frac{\sum_{i \in I_{n,w,t}} \left(\mathbb{1}\left[\text{“Trump”} \in \mathrm{Text}_i\right] \times \mathrm{Duration}_i\right)}{\left(\sum_{i \in I_{n,w,t}} \mathrm{Duration}_i\right) = 900} \tag{1.2}$$

where i stands for a set of uninterrupted sentences spoken by one person; $I_{n,w,t}$ stands for the set of i's spoken in network n during relative time period τ of window w; Text_i stands for the text of i; $\mathbb{1}\left[\text{“Trump”} \in \mathrm{Text}_i\right]$ stands for an indicator variable equal to one when “Trump” was mentioned in i; $\mathrm{Duration}_i$ stands for the duration in seconds of i.

$C_{n,w,\tau}$ is the share of time network n devoted to Trump-related issues during relative time period τ of event window w. It should be interpreted as a lower bound for Trump-related coverage given that it does not take into account segments related to President Trump where his surname was not explicitly mentioned.[3]

[3]See Appendix 1.6.1 for different descriptive statistics and more technical details on $C_{n,w,\tau}$.

Similarity between television and tweets

I construct a textual similarity measure defined as follows:

$$S_{n,w,\tau} = sim\Big(n3\left(Transcripts_{n,w,\tau}\right), n3\left(Tweets_{w}\right)\Big) \tag{1.3}$$

where sim stands for a Jaccard similarity, $n3\left(Transcripts_{n,w,\tau}\right)$ stands for the 3-word phrases used on network n, during relative time period τ of window w and $n3\left(Tweets_{w}\right)$ stands for the 3-word phrases used in those tweets posted during the treatment period of event window w.

$S_{n,w,\tau}$ is the number of 3-word expressions used on both (i) network n during relative time period τ of window w and (ii) Trump tweets posted during window w.

It is intended as a conservative indicator for moments in which Trump tweets, or topics mentioned on Trump tweets, are being discussed by network n. It is conservative in nature as it does not take into account instances in which a tweet is being implicitly discussed through other words.[4]

Similarity between television and online news

I relate cable television with online news as follows:

$$\mathcal{S}_{n,w,\tau} = sim\Big(n3\left(Transcripts_{n,w,\tau}\right), \bigcup_{k=1}^{4}\left\{n3\left(OnlineNews_{w,\tau-k}\right)\right\}\Big) \tag{1.4}$$

where sim stands for a Jaccard similarity, $n3\left(Transcripts_{n,w,\tau}\right)$ stands for the 3-word phrases used on network n during relative time period τ of window w and term $\bigcup_{k=1}^{4}\left\{n3\left(OnlineNews_{w,\tau-k}\right)\right\}$ stands for the 3-word phrases featuring in news posted on Facebook in the hour preceding relative time period τ of window w.

[4]See Appendix 1.6.1 for different descriptive statistics and more technical details on $S_{n,w,\tau}$.

$\mathcal{S}_{n,w,\tau}$ is the number of 3-word expressions used on (i) network n during relative time period τ of window w and (ii) recent news posted on Facebook by national news outlets.

This measure should be considered as an indicator for moments in which cable news are discussing a recent news event that has already been newscasted online (here, on Facebook). It should be interpreted with caution - it does speak to instances in which an event is discussed differently on TV, relative to how it was addressed online.[5]

Topics addressed on @realDonaldTrump tweets

I fit a Biterm Topic Model (an unsupervised topic model designed for short texts; see Yan et al., 2013) on a selected corpus of @realDonaldTrump tweets (those used to construct D_t, defined in Equation 1.1) in order to cluster Donald J. Trump's tweets into 10 different topics (each of these topics being associated to a unique set of words).[6]

Then, I build a set of indicator variables to distinguish between different event windows:

$$\mathcal{T}_w^{\nu} = \mathbb{1}\left[\mathrm{Mo}\left(\bigcup_{t \in Tweets_w} \mathrm{Topic}(t)\right) = \nu\right] \tag{1.5}$$

where t stands for tweet, $Tweets_w$ stand for tweets posted at relative time period 0 of window w, $\mathrm{Topic}(t)$ stands for topic of tweet t and $\mathrm{Mo}\left(\bigcup_{t \in Tweets_w} \mathrm{Topic}(t)\right)$ stands for mode of topics addressed during window w (given that within a window multiple tweets of multiple topics can be posted).

$\mathcal{T}_w^{\nu}$ is an indicator variable that points towards those windows in which topic ν was most addressed by President Trump. It does not provide any information on how many

[5]See Appendix 1.6.1 for different descriptive statistics and more technical details on $\mathcal{S}_{n,t}$.

[6]I chose to cluster President Trump's tweets in 10 topics. This was an ad hoc decision, taken for tractability purposes - to work with a fewer number of topics, more general and so, easier to be interpreted. In future work, I intend to inform this modelling choice with different information criteria purposed by NLP scholars. These criteria are all built to act as measures for how semantically coherent a model's topics (i.e., clusters of documents associated to unique sets of words) are.

tweets of topic ν were posted during window w.[7]

1.3. Empirical Setting

1.3.1 Identification

Non-overlapping event window (over time). Since President Trump tweeted on average at a shorter frequency than that of a quarter-hourly frequency (see Figure 1.4), there are several instances in which event windows partially overlap over time.

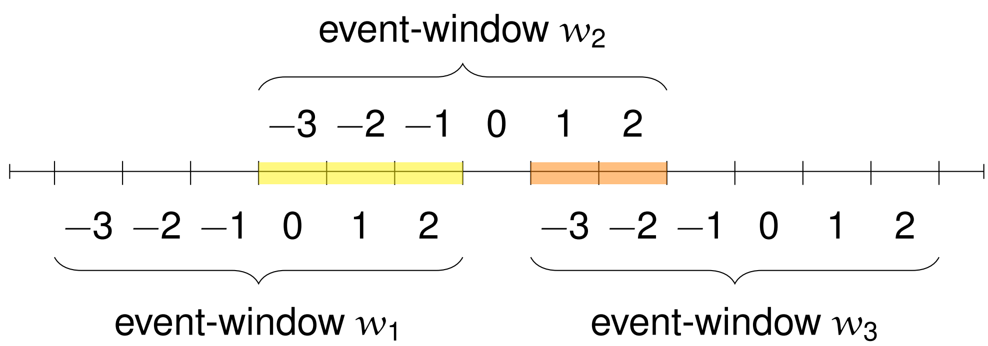

Figure 1.2: **Overlapping event windows (over time)**. The figure illustrates an hypothetical scenario in which 3 event windows partially overlap over time (shaded areas stand for an overlap). Each number stands for a within window relative time period.

As illustrated above, windows that overlap over time share different relative time periods. This is problematic from an identification standpoint for outcomes that are based on time, such as Trump-related coverage (defined in Section 1.2.2). These types of outcomes do not vary across overlapping time periods, making it impossible for one to distinguish between pre and post treatment intervals (see Baker et al., 2021).

This is a concern in the current setting - it does not allow for an unbiased assessment of coverage dynamics before and after a tweet. As such, in what follows, I restrict my

[7]See Appendix 1.6.1 for different descriptive statistics and more technical details on $\mathcal{T}_w^\nu$.

sample of study, when focusing on Trump-related coverage, to event windows that do not overlap over time. I will discuss in more detail the implications of this identification restriction in Section 1.3.2.[8]

Non-overlapping event window (over content). The identification problem described above is not of necessary concern for outcomes that are built with event window specific information such as similarity between television and tweets (in Section 1.2.2).

A partial overlap of two event windows over time is of concern for this last class of outcomes if and only if that information that is specific to each event window overlaps in any way. In this case, a component of these outcomes will again not change across overlapping periods and pre and post treatment times will not be distinguishable.

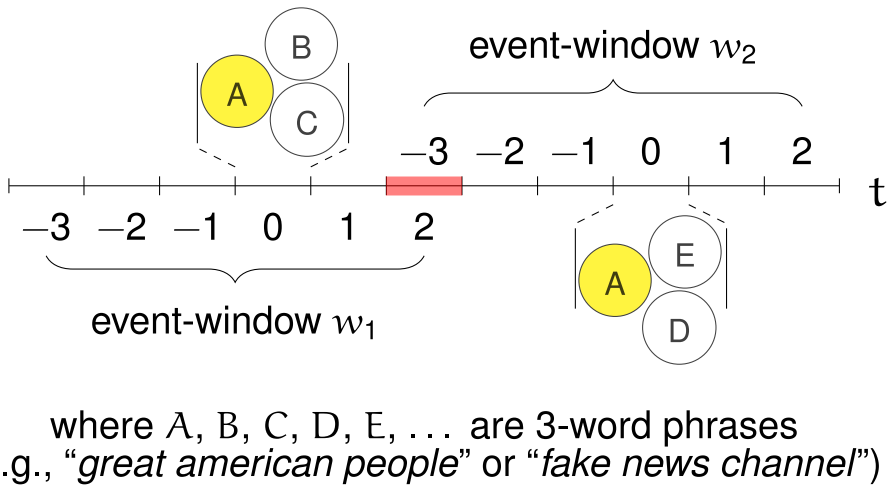

Figure 1.3: **Overlapping event windows (over content)**. The figure shows a scenario in which 2 windows partially overlap over content. Shaded areas stand for overlaps. Numbers stand for relative time periods. Letters stand for 3-word phrases posted in tweets.

This type of overlap (illustrated in Figure 1.3) is a concern when I study how cable news content compared to Donald Trump's tweets. In this context, those windows that overlap over time and have tweets that partially overlap over content (here, 3-word

[8] I present different descriptive statistics on this class of event windows in Section 1.6.1.

expressions) will provide a biased snapshot of how similarity evolved minutes before and after a tweet.

To circumvent these concerns, throughout this paper I will restrict my sample of study to two classes of event windows when studying how cable news content converged (or not) towards Trump tweets - I will focus on (1) event windows that do not overlap over time and (2) event windows with tweets that do not overlap over content. I discuss the implications of these identification restrictions in Section 1.3.2.[9]

1.3.2 Specification

To investigate how did cable news outlets react to President Trump's tweets, I estimate a standard event-study specification (see Schmidheiny and Siegloch, 2019 and Baker et al., 2021 for an extensive methodological review):

$$Y_{n,w,\tau} = (\alpha_n \times \delta_w) + \sum_{\eta \in \{C,F,M\}} \left(\mathbb{1}\left[n = \eta\right] \left[\sum_{k=-3,k\neq-1}^{3} \beta_k^\eta \left(\mathbb{1}\left[\tau = k\right] \times D_{w,0}\right) \right] \right) + \varepsilon_{n,w,\tau} \tag{1.6}$$

where $Y_{n,w,\tau}$ stands for an outcome variable specific to network n and relative time period τ of event window w; α_n stands for a network fixed effect; δ_w stands for an event window fixed effect; $\mathbb{1}(n = \eta)$ stands for an indicator variable equal to one if network n is network η (where η can be CNN (C), Fox News (F) or MSNBC (M)); $\mathbb{1}(t = \tau)$ stands for an indicator variable equal to one if relative time period τ is equal to τ; $D_{w,0}$ stands for a treatment variable indicating how many tweets President Trump posted during the treatment period of event window w; $\varepsilon_{n,w,\tau}$ is an idiosyncratic term specific to network n, event window w and relative time period τ.

$\{\beta_k^\eta\}_{k\in\{0,\ldots,3\}}$ are my coefficients of interest. These ought to be interpreted differently depending on which outcome is under investigation. If $Y_{n,w,\tau}$ stands for how much

---

[9] I present different descriptive statistics on these classes of event windows in Section 1.6.1.

time network n allocated to Trump-related issues during relative time period τ of event window w (as defined in Section 1.2.2), β_k^η should be interpreted as, by how much, on average, did Trump-related coverage by network η varied, relative to its pre-tweet value, k minutes after one Trump tweet was posted. If $Y_{n,w,\tau}$ instead stands for how similar network n's content, during relative time period τ of event window w, was to those tweets that were posted during event window w (as defined in Section 1.2.2), β_k^η should be interpreted as by how much, on average, did network η's content converge to a recently posted tweet, k minutes after that tweet was posted.

Turning to these coefficients causal interpretation, I restrict myself to studying cable news coverage in narrow time intervals exactly to argue in favor of an identification assumption that is necessary for these estimates to be interpreted as causal - that of variations in $Y_{n,w,\tau}$ within high-frequency periods neighboring Trump tweets being only attributable to Trump tweets (other relevant network-specific and macro factors kept constant). This assumption is not testable. In addition, if $\{\beta_k^\eta\}_{k\in\{0,...,3\}}$ are to be interpreted as causal effects, then, $Y_{n,w,\tau}$ should not exhibit any abnormal behavior prior to each tweet (so-called parallel trends assumption). This is an identification assumption that I test for by estimating a set of pre-treatment coefficients: $\{\beta_k^\eta\}_{k\in\{-3,-2\}}$. In what follows, I will always plot these coefficients together with $\{\beta_k^\eta\}_{k\in\{0,...,3\}}$, to argue that, indeed, television coverage did not seem to behave abnormally moments prior to @realDonaldTrump tweets.

Lastly, I impose different identification restrictions on Equation 1.6 depending on which outcome I focus on. As discussed in Section 1.3.1, these restrictions are related to how frequently President Trump tended to tweet during his mandate, creating a setting in which event windows tended to overlap over different dimensions:

- When studying how Trump-related coverage evolved moments prior and after a Trump tweet, I estimate Equation 1.6 by only using those event windows that did not overlap over time. This allows me to unbiasedly estimate pre and post-treatment coverage dynamics. Still, it has relevant implication on how $\{\beta_k^\eta\}_{k\in\{0,...,3\}}$

are to be interpreted. It constrains $\{\beta_k^\eta\}_{k\in\{0,\dots,3\}}$ to be at most interpretable as local average treatment effects, focused on those tweets that tended to generate event windows that did not overlap over time. This class of tweets accounts for approximately 20% of Donald Trump's presidential tweets (see Figure 1.7) and is more likely to be posted during the evening (see Figure 1.8).

- When studying how did cable news content related to Trump tweets text moments prior and after a Trump tweet, I estimate Equation 1.6 by using separately (1) those event windows that did not overlap over time and (2) those windows that did not overlap over content. As before still, both identification restrictions constrain $\{\beta_k^\eta\}_{k\in\{0,\dots,3\}}$ to be at most interpretable as local average treatment effects, focused on those tweets that tended to generate (1) event windows that did not overlap over time and (2) that did not overlap over content. On this last class of tweets, contrary to before, this accounts for approximately 80% of Donald Trump's presidential tweets (see Figure 1.25) and follows closely President Trump's posting patterns throughout his mandate (see Figure 1.12).

One last remark concerning the fixed effect specification laid out in Equation 1.6 - coefficients are estimated while controlling for network-specific event window fixed effects ($\alpha_n \times \delta_w$). This specification choice implicitly assumes that each cable news outlet reacted differently to common macro factors other than President Trump's tweets. This is an assumption that is aligned with previous empirical findings - Groseclose and Milyo (2005) and more recently Martin and Yurukoglu (2017) have documented significant differences in these outlets' editorial criteria.[10]

[10] Note that, in Appendix 1.6.2 I present estimates for specifications with coarser time fixed effects. Still, this specification choice - that of controlling for network-specific time factors - is kept constant, exactly due to its empirical grounding.

1.4. Results and Robustnesses

1.4.1 Results

Reaction through Coverage

I start by studying how did Trump-related coverage behave minutes before and after a Trump tweet. This exercise can speak to at least two different scenarios:

- On one hand, President Trump could have been more likely to tweet moments after a specific outlet spent an abnormally high (or low) amount of time covering him. These tweets could either be posted as a reaction to said coverage or, as an attempt to generate Trump-related coverage (e.g., President Trump could have been more likely to tweet right after abnormally low levels of coverage by Fox News). In this case, pre-treatment coefficients ought to be different from zero.
- On the other hand, cable outlets could have reacted in real time to @realDonaldTrump tweets. In this scenario, pre-treatment coefficients are expected to be still while post-treatment coefficients are expected to be significantly different from zero (but not necessarily positive; e.g., Fox News could have diverted coverage to non-Trump topics right after @realDonaldTrump's tweet - given how sensitive these statements were on average (e.g., see Müller and Schwarz, 2018)).

Figure 1.26 shows the estimated coefficients from Equation 1.6 for when $Y_{n,w,\tau}$ stands for how much time network n devoted to Trump-related issues.[11] $\{\beta_k^n\}_{k\in\{-3,\ldots,3\}}$ are all statistically insignificant. This result holds across outlets. It is also robust to (1) controlling for different outlet-specific macro factors and (2) to clustering standard errors at different levels (see Figure 1.27 and 1.28). Taking these estimates by their face value,

[11] In addition, the coefficients reported in 1.26 have been estimated using only event windows that did not overlap over time (hence, these estimates should be interpreted as how cable coverage evolved minutes before and after a specific class of Trump tweets; see Section 1.6.1 for descriptives on this particular class of statements).

these seem to indicate that, on average, throughout Donald J. Trump's mandate, (1) the President did not seem to tweet in reaction to cable news coverage and (2) Trump-related coverage did not change in any significant way, minutes after a Trump tweet.

Nonetheless, these results deserve significantly more study. In fact, given how Equation 1.6 is specified, null coefficients such as those plotted in Figure 1.26 can be a product of unobserved heterogeneity. Taking scenario (2) for illustration purposes - assume that an outlet's short run reaction to a Trump tweet varied according to when a tweet was posted throughout his presidency - e.g., cable outlets could have learnt over time which tweets were more interesting for their audiences, varying their coverage decisions accordingly. In this setting, each coefficient in Equation 1.6 should be interpreted as a weighted average of different types of outlet-specific reactions. For a plot such as that in Figure 1.26 to materialize, it would only be required that these reactions varied in such a way over time such that their averages would equal zero.

Reaction through Content

In a second analysis, I turn to how did cable news content relate to President Trump's tweets moments before and after a tweet. As before, this analysis can be motivated with different scenarios:

- On one hand, President Trump could have reacted to cable news coverage in real time. In this scenario, pre-treatment coefficients are expected to be significantly different from zero and positive (as the President would tweet at period $\tau = 0$ about those issues being discussed on cable outlets during periods $\tau < 0$).
- In a second scenario, cable outlets could have been covering in real time @realDonaldTrump tweets. As before, post-treatment coefficients would be expected to be positive and significant while pre-treatment coefficients should, on a contrary, not differ from zero.

Figure 1.29 shows the estimated coefficients from Equation 1.6 for when $Y_{n,w,\tau}$ stands for how similar was network n's content to that of @realDonaldTrump's most recent tweet(s).[12] Contrary to before, $\{\beta^{\eta}_{k}\}_{k\in\{-3,...,3\}}$ follow a pattern consistent with scenario (2). Pre-treatment coefficients do not differ from zero, hence, the parallel trends assumption seems to hold on average, meaning that post-treatment coefficients can a priori be interpreted as average causal effects. Post-treatment coefficients are significantly positive This result holds across outlets. It is robust to (1) controlling for different outlet-specific macro factors and (2) to clustering standard errors at different levels (see Figure 1.30 and 1.31).[13]

These estimates are suggestive of President Trump having been able to, on average, temporarily shift cable news outlets' attention through his tweets. Still, this "suggestion" should be taken with caution. Previous concerns regarding unobserved heterogeneities hold. In addition, an additional caveat in this case relates to an external factor that is currently not being taken into account in the current regressions. I discuss and test for this second caveat in Section 1.4.2.

1.4.2 Robustnesses

Content Reaction due to Breaking News

As discussed in Section 1.4.1, previous results concerning how television coverage converged in content towards Trump tweets minutes after a tweet, can be attributed not only to unobserved heterogeneities but also to a second uncontrolled factor - pressing events that happened within (or close to) each event window. In fact, a possible explanation for those estimates presented in Section 1.4.1 is that President Trump could have systematically reacted to recent breaking news faster than cable outlets (e.g., a news event happens at period t-1, President Trump reacts at period t and cable outlets

[12]The coefficients reported in 1.29 have been estimated using only event windows that did not overlap over time (hence, as before, these estimates should be interpreted as how cable coverage evolved minutes before and after a Trump tweet a).

[13]These results also hold qualitatively when using event windows that do not overlap over content, instead of time (still, effects decrease by 4 orders of magnitude; see Figure 1.32, 1.33 and 1.34).

react at period t+1, this late reaction by television news ultimately implying a convergence in content between cable news and Trump tweets in period t+1).

In this section, I take advantage of an exhaustive corpus of online news (described in Section 1.2.1) to test for this hypothesis. In doing so, I assume that in a framework with 3 players, (a) online news outlets, (b) cable news outlets and (c) @realDonaldTrump, online news outlets should be those fastest at covering a breaking news. Under this assumption, I compare cable news coverage with recent online news content within event windows in which television news converged towards President Trump's tweets, minutes after a tweet. If President Trump tended to react faster to a breaking news than television, then, I ought to see a systematic convergence in content of television content with past online news minutes after a tweet was posted.

To be more specific, I estimate Equation 1.6 with $Y_{n,w,\tau}$ being a similarity measure between current cable news transcripts and past online news ("*past*" $\equiv$ last hour; as defined in Section 1.2.2). To understand if previous results were driven by President Trump having been systematically faster at reacting to breaking news, I estimate Equation 1.6 using only observations belonging to event windows in which cable news content converged to @realDonaldTrump tweets minutes after a tweet.[14] I turn to $\{\beta_k^n\}_{k\in\{0,\dots,3\}}$ (i.e., post-tweet coefficients) to understand if indeed TV news converged in content to past online news moments after a Trump tweet or not (in other words, I estimate $\{\beta_k^n\}_{k\in\{0,\dots,3\}}$ to understand if these are significantly different from zero and positive).

Figures 1.35 and 1.37 show the estimated coefficients from Equation 1.6 for when $Y_{n,w,\tau}$ stands for how similar was network n's content to recent news posted on Facebook. Both figures point to a similar result - CNN tended to converge towards Trump tweets that were related to past online news; Fox News and MSNBC on the contrary

[14]In particular, I proceed as follows: (1) demean similarity measure between television and tweets at a network $\times$ event window level; (2) sum within event window demeaned outcome by pre and post tweet periods; (3) take difference between post and pre tweet aggregated demeaned outcome and last (4) estimate coefficients using observations belonging to windows where difference (3) is positive.

tended to shift their attention towards topics tweeted by President Trump that were unrelated to recent online news.[15] These results seem to suggest that President Trump was able to temporarily set Fox News and MSNBC's agenda through his tweets. This agenda setting power is of interest as it could have been used by President Trump to shift these outlets' attention towards or away from specific topics at particular times.

1.4.3 Results (2)

Reaction per Topic

Previous results are derived within a framework in which an outlet's reaction to a Trump tweet is assumed to be homogeneous across different dimensions. This is a theoretically far-fetched assumption. Indeed, a factor that is likely to play a role on how a news outlet tends to react to a Trump tweet is which topic President Trump addressed on that same tweet. This can be due to supply (i.e., an outlet's idiosyncratic editorial biases), demand (i.e, an outlet's audience and its demand for news on certain issues), or both.

To give an example, a conservative outlet such as Fox News (see Martin and Yurukoglu, 2017) can be thought of as being more likely to cover a topic that is a priori friendlier towards President Trump. This can be rationalized by a supply argument - e.g., Fox News' editors having an idiosyncratic objective of improving President Trump's ratings through their coverage - or a demand motif - e.g., this outlet's audience, while conservative, drawing utility from consuming news that flatter key conservative figures.

In this section, I empirically test for whether or not certain topics caused a relatively larger shift in content for specific outlets. This is an exercise that is important from at least two different perspectives. On one hand, understanding whether specific outlets were consistently more drawn to particular topics can potentially help us understand

---

[15] Results are robust to different FE specifications (see Figure 1.36 and 1.38).

President Trump's tweeting patterns, possibly helping us shedding light on which objective function dictated his statements on Twitter. On another hand, describing how different outlets tended to cover a specific issue can provide us with more information about which dimension of coverage is normally used by outlets as a differentiation factor, relative to their competitors. This second angle speaks to an open question in media bias studies - which dimension of coverage differentiates cable news stations and is thus determinant at explaining cable news persuasive effects on voting (see DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2017).

To test whether or not certain topics caused a relatively larger shift in content for specific outlets, I extend Equation 1.6. In particular, I allow for each station's reaction to vary across an array of topics that President Trump consistently addressed on Twitter, throughout his mandate:

$$Y_{n,w,\tau} = \ldots + \sum_{\nu \in \mathcal{N}} \mathcal{T}_w^{\nu} \left[\sum_{\eta \in \{C,F,M\}} \left(\mathbb{1}\,[n = \eta] \left[\sum_{k=-3,k\neq-1}^{3} \beta_k^{\eta,\nu} \left(\mathbb{1}\,[\tau = k] \times D_{w,0} \right) \right] \right) \right] \tag{1.7}$$

where $Y_{n,w,\tau}$ stands for an outcome variable specific to network n and relative time period τ of event window w; "..." stands for a network $\times$ window fixed effect and an idiosyncratic term; ν stands for topic; $\mathcal{N}$ stands for set of topics addressed by President Trump throughout his mandate; $\mathcal{T}_w^{\nu}$ stands for an indicator variable equal to one if topic ν was that most discussed by President Trump during window w (as defined in Section 1.2.2); $\mathbb{1}(n = \eta)$, $\mathbb{1}(t = \tau)$ and $D_{w,0}$ are defined as in Equation 1.6.

$\{\beta_k^{\eta,\nu}\}_{k \in \{0,\ldots,3\}}$ stand for how much $Y_{n,w,\tau}$ varied after a Trump tweet of topic ν was posted. As before, these ought to be interpreted as causal estimates conditional on two different assumptions: (1) any variation in $Y_{n,w,\tau}$ minutes after a Trump tweet is solely attributable to that same tweet; (2) $Y_{n,w,\tau}$ does not exhibit any abnormal pattern minutes before a Trump tweet. Assumption (1) is not testable. Assumption (2) is testable - I test for (2) by estimating different pre-treatment coefficients $\{\beta_k^{\eta,\nu}\}_{k \in \{-3,-2\}}$.

Figures 1.39, 1.40 and 1.41 plot how Trump-related coverage evolved minutes before and after tweets from different topics were posted. Overall, previous results seem to hold across topics. In other words, irrespective of which topic one focuses on, Trump tweets did not cause a significant increase in cable outlets' Trump-related coverage. This piece of evidence casts aside previous concerns on how much was this absence of a response from Trump-related coverage a product of unobserbed heterogeneity. It is suggestive of a world in which cable networks tended to follow President Trump's focus on Twitter without explicitly naming him.

In similar fashion, Figures 1.42, 1.43 and 1.44 plot how cable news content compared to President Trump's tweets minutes before and after a Trump tweet was posted. Here, I arrive at different conclusions: (1) certain tweets caused all outlets to shift their coverage minutes after a tweet (these were tweets related to domestic events, foreign policy, collusion charges and the White House); (2) a second set of topics was mostly associated to shifts in Fox News and MSNBC's coverage (tweets related to news media, immigration and economic topics); (3) a specific set of tweets did not cause significant shifts in cable news coverage (tweets related to trade policy, shootings and disasters)[16]

Put together, these results seem to suggest that, on average, cable networks tended to cover a relatively similar set of Trump tweets. This seems to suggest that, in this context, television outlets tended to differentiate themselves not by covering different Trump statements but, instead, by slanting differently a relatively similar distribution of news stories.

1.5. Conclusion

I have taken advantage of an exhaustive dataset on cable news transcripts and President Trump's tweets to study cable news coverage in short time intervals centered

[16]These results keep mainly constant both if I focus on tweets that generate event windows that do not overlap over time (1.42, 1.43 and 1.44) or over content (Figures 1.45, 1.46 and 1.47).

around Trump's tweets. In doing so, I have concluded that cable news outlets tended on average to shift their focus to topics related to President Trump's tweets minutes after these tweets were posted.

In a second exercise, I turned to an exhaustive corpus of news posted on Facebook by U.S. national news outlets to understand if cable outlets tended to purposefully follow President Trump's tweets in content or, instead, if this convergence was simply due to President Trump tending to react faster to recent breaking news (relative to television). Here, I found that Fox News and MSNBC seemed to purposefully follow President Trump's tweets, independent of recent online news.

This latter result is interesting from different perspectives. On one hand, it suggests that different outlets had differing coverage criteria when focusing on President Trump's tweets. Understanding how did these criteria differed across outlets is important as it can help us understand how is social media covered by mass media and how does bias plays a role in that coverage.

Motivated by this result, I went on and studied how did each outlet cover different types of Trump tweets (where by types I mean topics; e.g., tweets related to domestic events, tweets on immigration, tweets on foreign policy). Here, I find that President Trump tended to shift cable news attention irrespective of which topic he tweeted about (with some minor exceptions). This result seems to suggest that television stations tended to differentiate themselves not by covering different Trump statements but, instead, by slanting differently a relatively similar distribution of Trump stories.

On another hand, it suggests that President Trump was able to temporarily set Fox News and MSNBC's agenda through his tweets. This agenda setting power is of interest as it could have been used by President Trump to shift mass media and, thus, public attention, towards or away from specific topics at specific times. This hypothesis is of relevance in nowadays context as it can possibly shed light on how populist politicians strategically interacted with mass media through social media.

Bibliography

Allcott, H., Boxell, L., Conway, J., Thaler, M., and Gentzkow, M. (2020). Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic. *Journal of Public Economics*, 191:104254. [URL].

Ash, E., Galletta, S., Pinna, M., and Warshaw, C. (2021). The Effect of Fox News Channel on U.S. Elections: 2000-2020. *SSRN Research Paper Series*. SSRN 3837457 [URL].

Baker, A., Larcker, D. F., and Wang, C. C. Y. (2021). How Much Should We Trust Staggered Difference-In-Differences Estimates? *SSRN Research Paper Series*. SSRN 3794018 [URL].

Brown, M. and Sanderson, Z. (2020). How Trump Impacts Harmful Twitter Speech: A Case Study in Three Tweets. Brookings Institution, Washington D.C. [URL].

Bursztyn, L., Rao, A., Roth, C., and Yanagizawa-Drott, D. (2020). Misinformation during a Pandemic. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-44). [URL].

Cagé, J., Hervé, N., and Mazoyer, B. (2020). Social media and newsroom production decisions. *SSRN Research Paper Series*. SSRN 3663899 [URL].

Clayton, K., Davis, N. T., Nyhan, B., Porter, E., Ryan, T. J., and Wood, T. J. (2021). Elite Rhetoric Can Undermine Democratic Norms. *Proceedings of the National Academy of Sciences*, 118(23).

CrowdTangle (Link). https://www.crowdtangle.com/ [URL].

DellaVigna, S. and Kaplan, E. (2007). The Fox News Effect: Media Bias and Voting.

The Quarterly Journal of Economics, 122(3):1187–1234. [URL].

Djourelova, M. and Durante, R. (2019). Media Attention and Strategic Timing in Politics: Evidence from U.S. Presidential Executive Orders. *SSRN Research Paper Series*. SSRN 3464486 [URL].

Gertz, M. (2018). "*The president just threatened a nuclear strike while live-tweeting a Fox News segment. Left, Fox, 7:37 pm. Right, Trump, 7:49 pm.*" [Tweet by @MattGertz]. Retrieved from Twitter [URL].

Gratton, G., Holden, R., and Kolotilin, A. (2018). When to Drop a Bombshell. *The Review of Economic Studies*, 85(4):2139—2172. [URL].

Groseclose, T. and Milyo, J. (2005). A Measure of Media Bias. *The Quarterly Journal of Economics*, 120(4):1191–1237. [URL].

Martin, G. J. and Yurukoglu, A. (2017). Bias in Cable News: Persuasion and Polarization. *American Economic Review*, 107(9):2565–99. [URL].

Müller, K. and Schwarz, C. (2018). From Hashtag to Hate Crime: Twitter and Anti-minority Sentiment. *SSRN Research Paper Series*. SSRN 3149103 [URL].

Pew (2021a). Behind Biden's 2020 Victory. Written by Ruth Igielnik, Scott Keeter and Hannah Hartig, Washington D.C. [URL].

Pew (2021b). Charting Congress on Social Media in the 2016 and 2020 Elections. Written by Sono Shah, Andrew Grant and others, Pew Research Center, Washington D.C. [URL].

Pew (2021c). More than Eight-in-Ten Americans get News from Digital Devices. Written by Elisa Shearer, Pew Research Center, Washington D.C. [URL].

Pew (2021d). News Consumption Across Social Media in 2021. Written by Mason Walker and Katerina Eva Matsa, Pew Research Center, Washington D.C. [URL].

@realDonaldTrump (Link). https://twitter.com/realdonaldtrump/ [URL].

Schmidheiny, K. and Siegloch, S. (2019). On Event Study Designs and Distributed-lag Models: Equivalence, Generalization and Practical Implications. *CEPR Discussion Paper Series*. No. DP13477 [URL].

Trump, D. J. (2019a). "*Acting Secretary of Defense Patrick Shanahan, who has done a wonderful job, has decided not to go forward with his confirmation process so that he can devote more time to his family........I thank Pat for his outstanding service and will be naming Secretary of the Army, Mark Esper, to be the new Acting Secretary of Defense. I know Mark, and have no doubt he will do a fantastic job!*" [Tweet thread by @realDonaldTrump]. Retrieved from Twitter [URL1 and URL2].

Trump, D. J. (2019b). "*So interesting to see "Progressive" Democrat Congresswomen, who originally came from countries whose governments are a complete and total catastrophe, the worst, most corrupt and inept anywhere in the world (if they even have a functioning government at all), now loudly..........and viciously telling the people of the United States, the greatest and most powerful Nation on earth, how our government is to be run. Why don't they go back and help fix the totally broken and crime infested places from which they came. Then come back and show us how........it is done. These places need your help badly, you can't leave fast enough. I'm sure that Nancy Pelosi would be very happy to quickly work out free travel arrangements!*" [Tweet thread by @realDonaldTrump]. Retrieved from Twitter [URL1, URL2 and UR3].

Trump Twitter Archive (Link). https://www.trumptwitterarchive.com/archive/ [URL].

TV News Archive (Link). https://archive.org/details/tv/ [URL].

Yan, X., Guo, J., Lan, Y., and Cheng, X. (2013). A Biterm Topic Model for Short Texts. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 1445–1456. [URL].

Zhuravskaya, E., Hatte, S., and Madinier, E. (2021). Reading Twitter in the Newsroom: How Social Media Affects Traditional-Media Reporting of Conflicts. *Mimeo*. [URL].

1.6. Appendix

1.6.1 Data

Tweets posted by @realDonaldTrump

Table 1.1: **Descriptive statistics**

	'17-'20	'17	'18	'19	'20
Total tweets	16,632	2,292	3,104	4,946	6,290
Total 15-minutes	11,914	1,925	2,553	3,489	3,947
Tweets per 15-minutes	'17-'20	'17	'18	'19	'20
Min.	1	1	1	1	1
p25	1	1	1	1	1
Median	1	1	1	1	1
p75	2	1	1	2	2
Max.	18	8	5	13	18

Notes: The table shows descriptive statistics on a sample of @realDonaldTrump tweets that does not include retweets nor short tweets. (tweets have been winsorized according to their dimension, measured by the number of characters; bottom ten percentiles dropped).

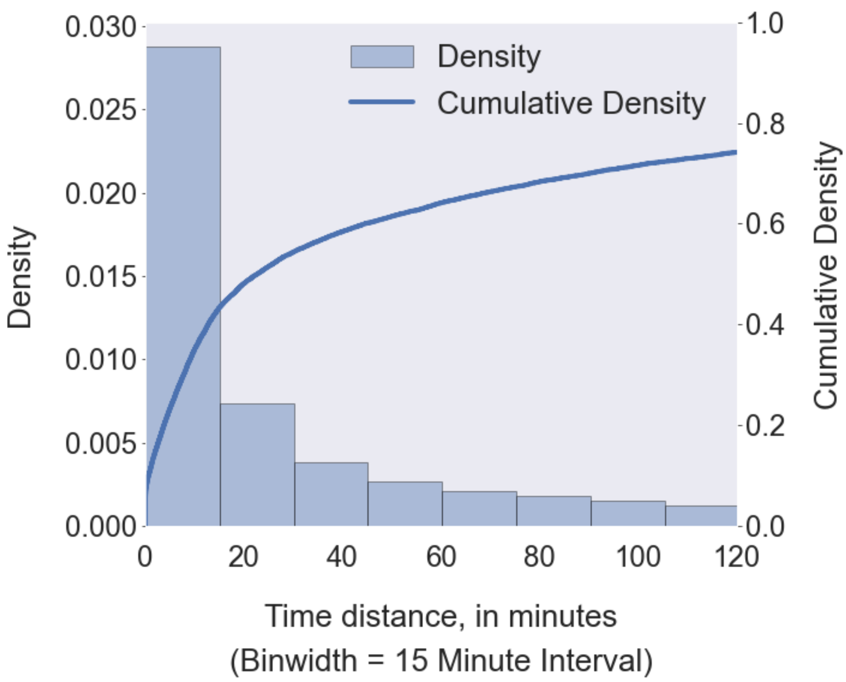

Figure 1.4: **Time distance between @realDonaldTrump tweets.** The figure shows the probability and cumulative density functions for the time distance measured in a specific class of Trump tweets (as in Table 1.1; no retweets, no short tweets, from 2017 to 2020).

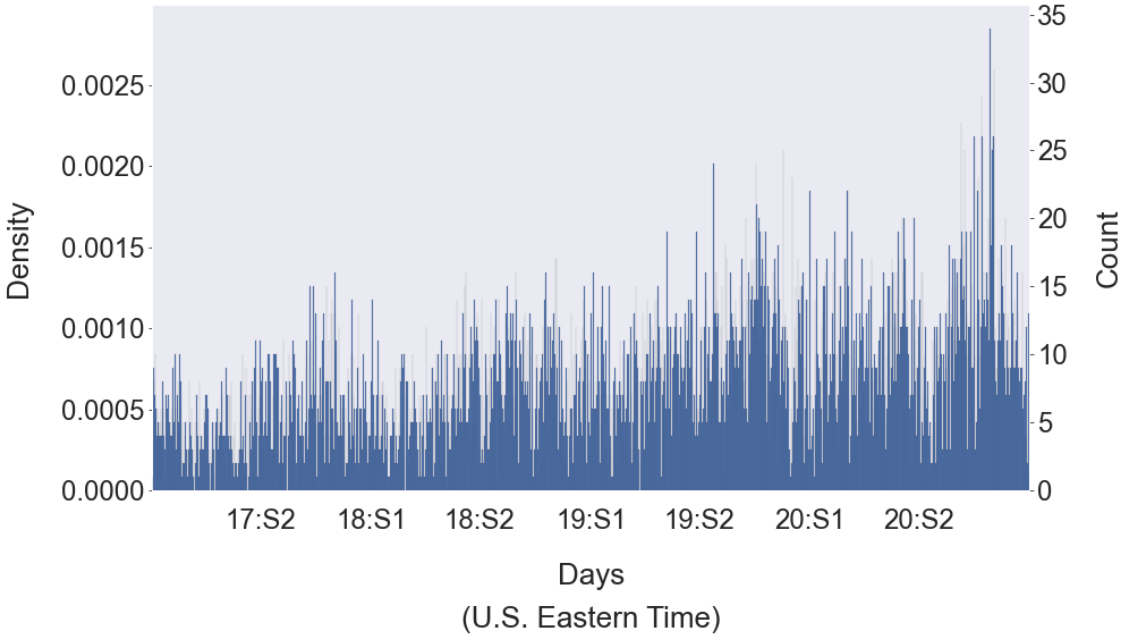

Figure 1.5: **@realDonaldTrump tweets per day**. The figure plots the number of Trump tweets posted per day, covering a period that goes from 2017 to 2020. It focuses on a specific class of Trump tweets (as in Table 1.1; no retweets, no short tweets).

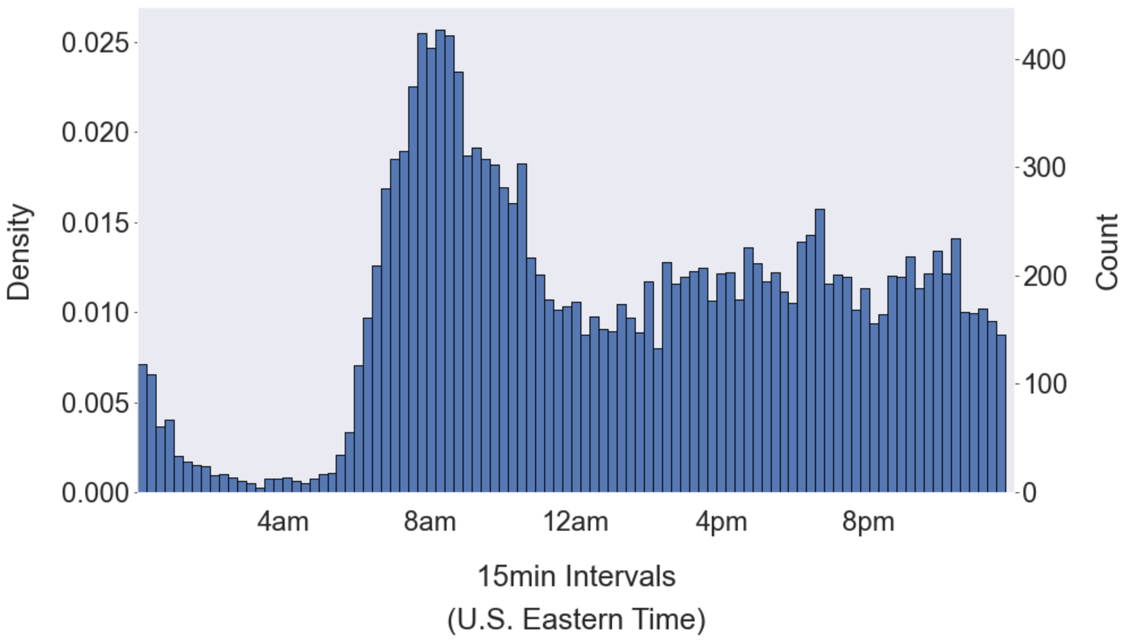

Figure 1.6: **@realDonaldTrump tweets within a day**. The figure plots the number of Trump tweets posted within a day, per 15 minute interval. It focuses on a specific class of Trump tweets (as in Table 1.1; no retweets, no short tweets, from 2017 to 2020).

Non-overlapping @realDonaldTrump tweets

1.1.2.1. Over time

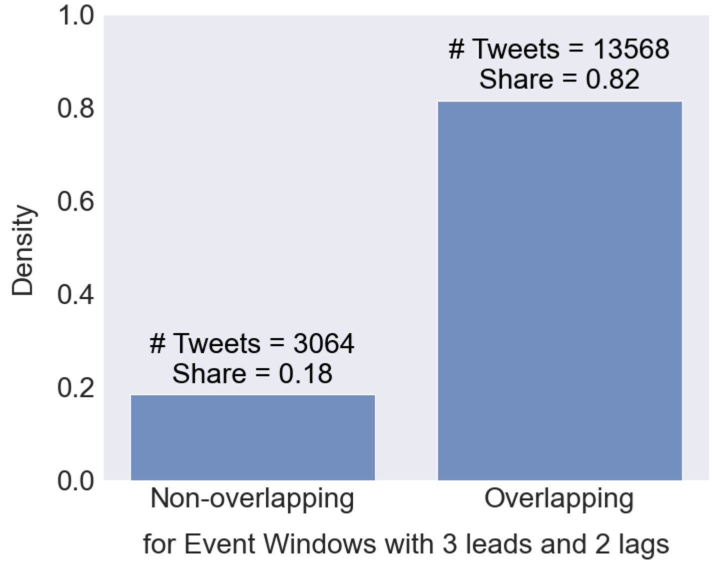

Figure 1.7: **Number/share of overlapping and non-overlapping tweets.** The figure shows how many of Trump's tweets overlapped and non-overlapped over time (from 2017 until 2020; without counting retweets nor short tweets).

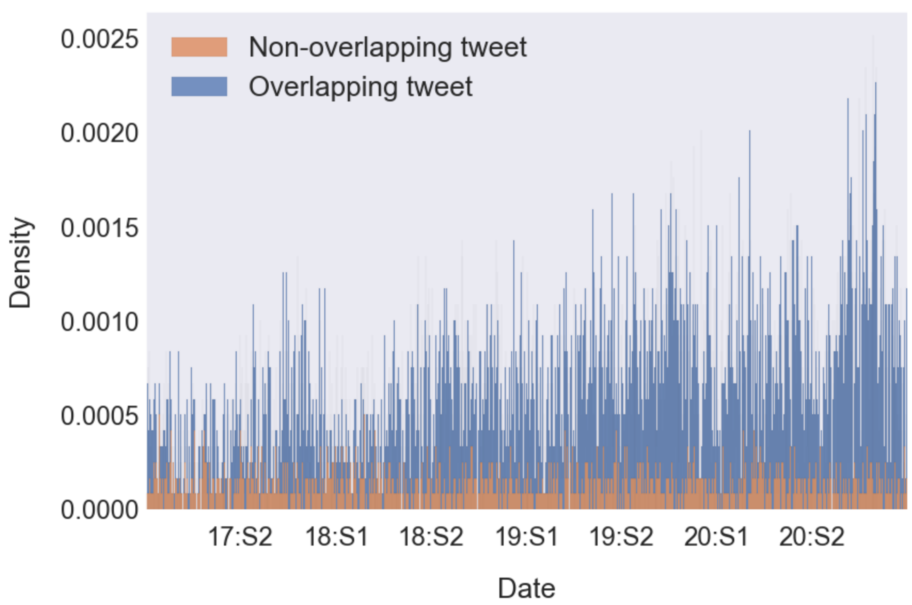

Figure 1.8: **Overlapping and non-overlapping tweets per day.** The figure plots the number of tweets that generated overlapping and non-overlapping event windows (over time) per day (from 2017 until 2020; without counting retweets nor short tweets).

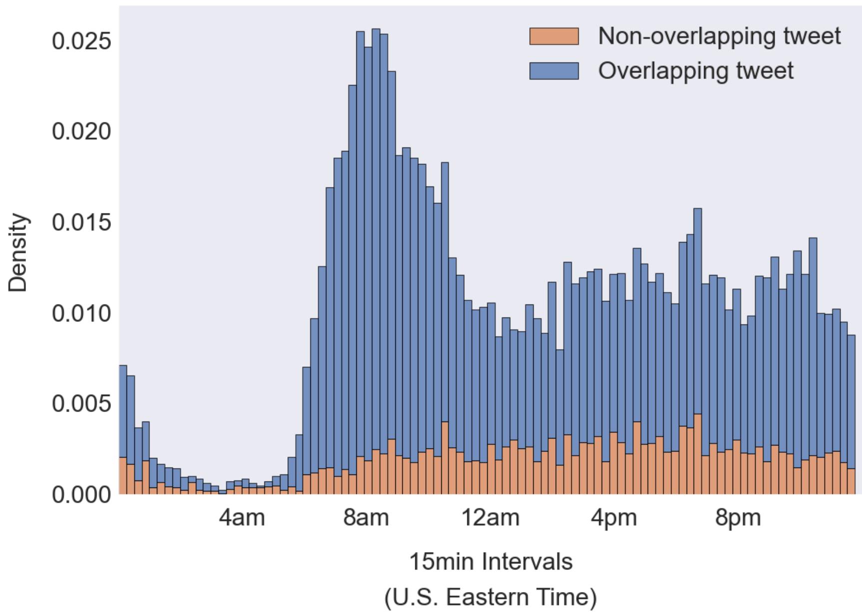

Figure 1.9: **Overlapping and non-overlapping tweets within day.** The figure plots the number of tweets that generated overlapping and non-overlapping event windows (over time) within a day, per 15 minute interval (from 2017 until 2020; without counting retweets nor short tweets).

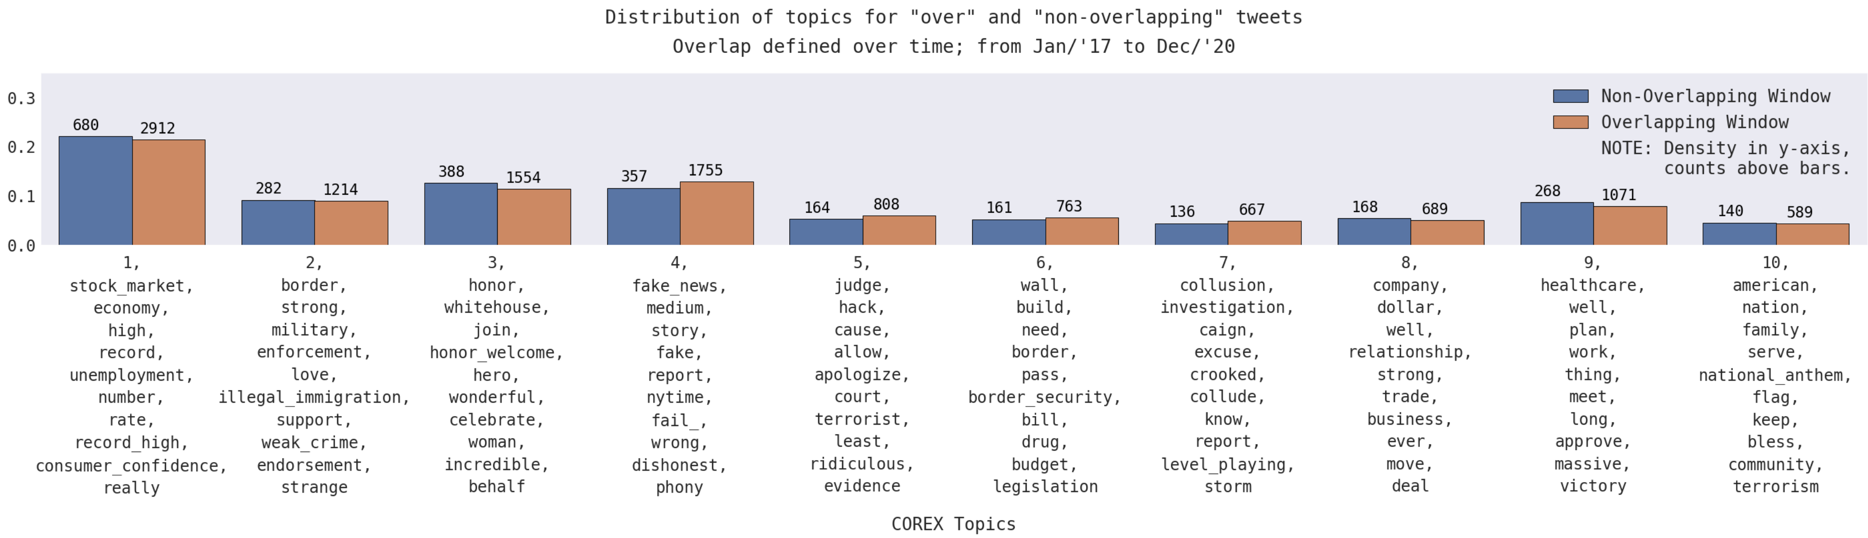

Figure 1.10: **Topic distribution for overlapping and non-overlapping tweets**. The figure plots the distribution of topics (inferred from a Biterm Topic Model, explained in Section 1.2.2 and described in Section 1.6.1) for overlapping and non-overlapping tweets (across time; from 2017 until 2020, without counting retweets nor short tweets).

1.1.2.2. Over content

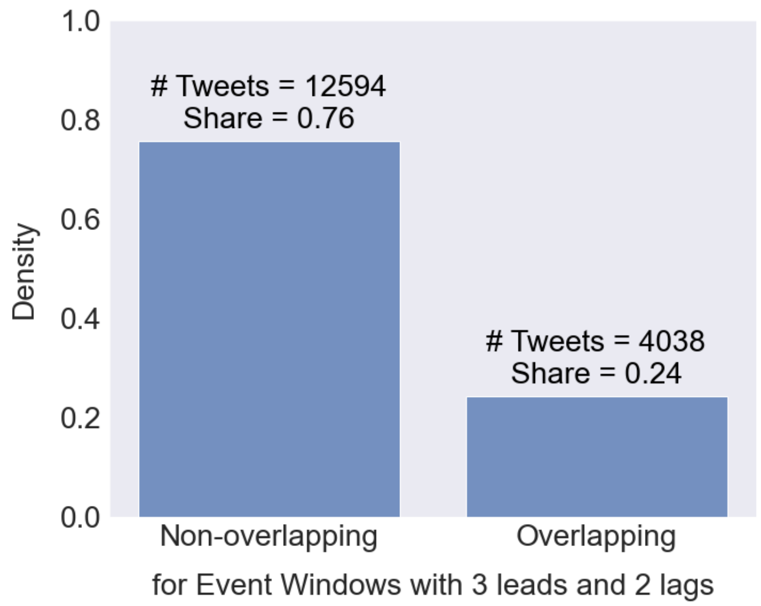

Figure 1.11: **Number/share of overlapping and non-overlapping tweets**. The figure plots different descriptive statistics for those tweets included in event windows that overlap and do not overlap over content (number of tweets and share relative to total).

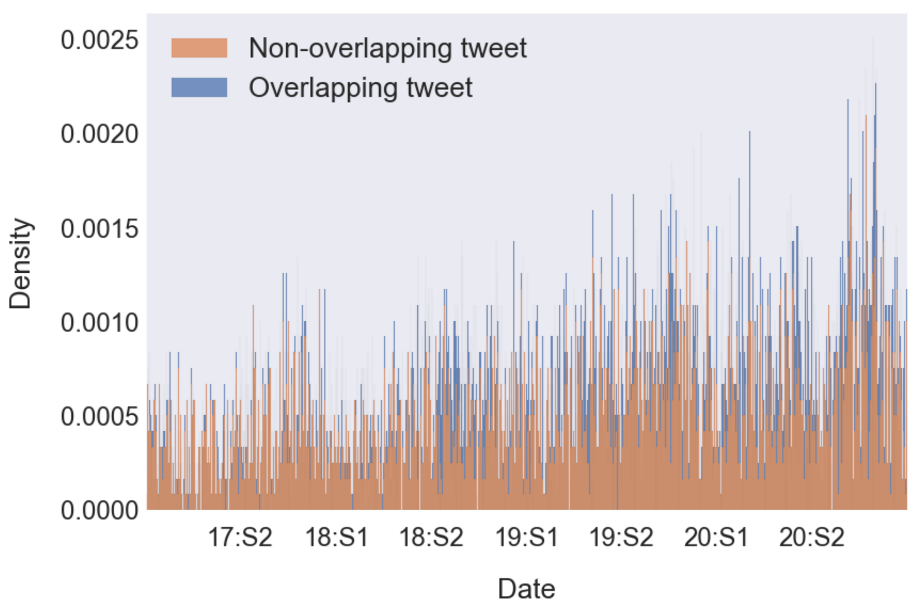

Figure 1.12: **Overlapping and non-overlapping tweets per day**. The figure plots the number of tweets that generated overlapping and non-overlapping event windows (over content) per day (from 2017 until 2020; without counting retweets nor short tweets).

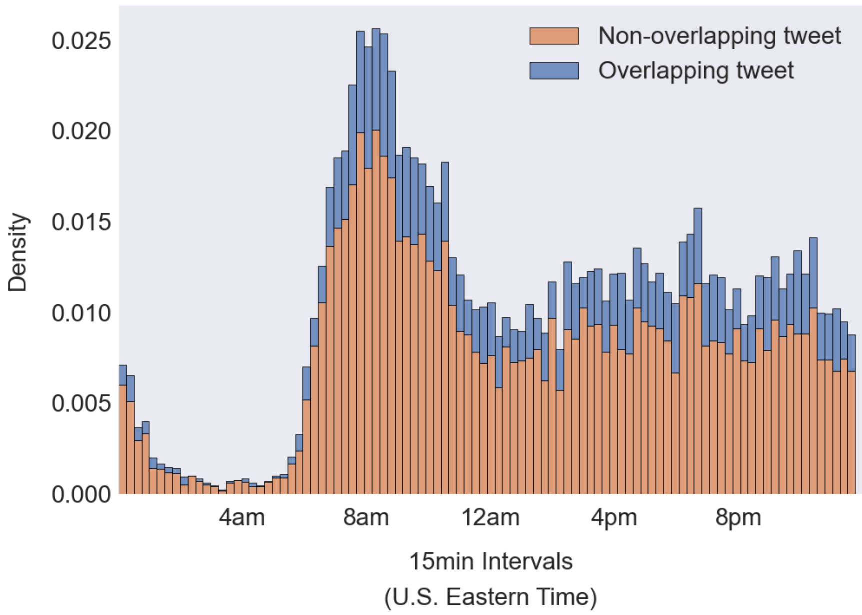

Figure 1.13: **Overlapping and non-overlapping tweets within day**. The figure plots the number of tweets that generated overlapping and non-overlapping event windows (over content) within a day, per 15 minute interval (from 2017 until 2020; without counting retweets nor short tweets).

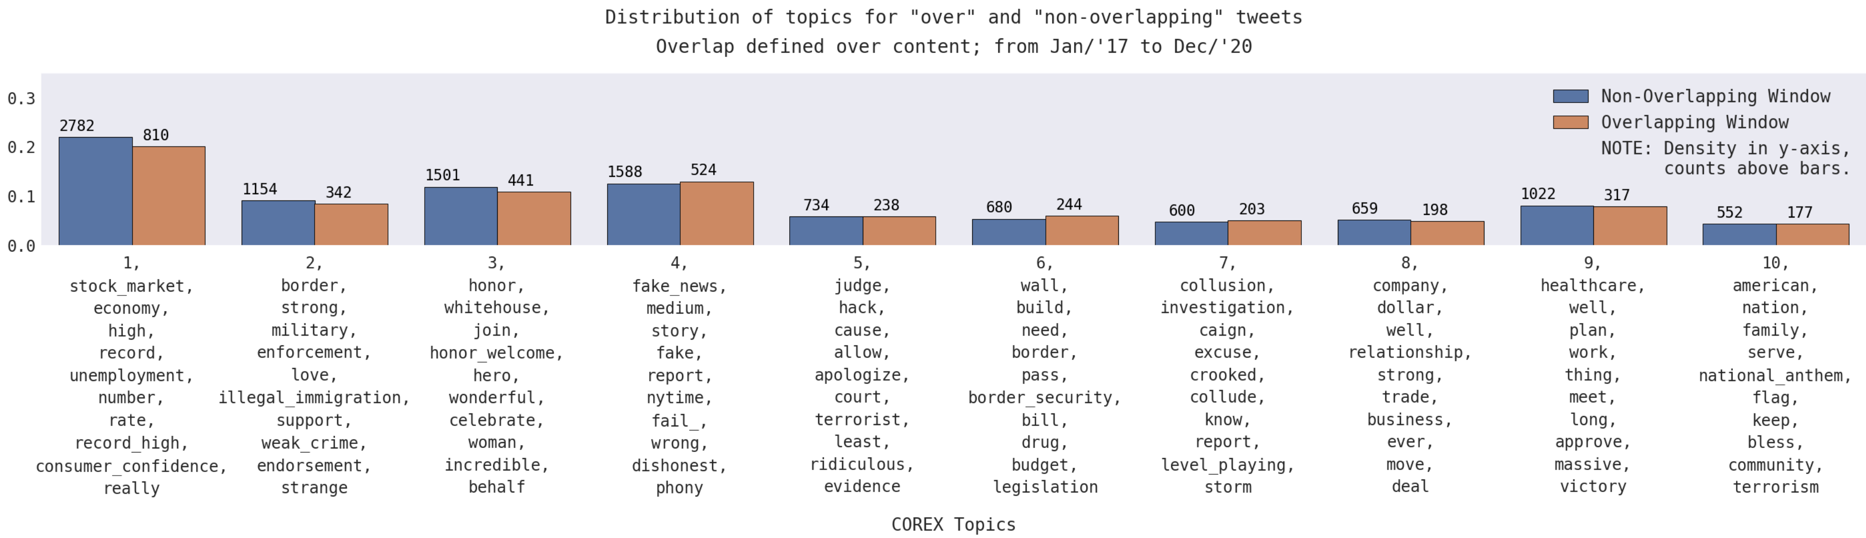

Figure 1.14: **Topic distribution for overlapping and non-overlapping tweets**. The figure plots the distribution of topics (inferred from a Biterm Topic Model, explained in Section 1.2.2 and described in Section 1.6.1) for overlapping and non-overlapping tweets (across content; from 2017 until 2020, without counting retweets nor short tweets).

Television news coverage of Trump-related issues

Table 1.2: **Descriptive statistics**

	Mean	Std. Dev.	Min.	p25	Median	p75	Max.	Observations
CNN	0.21	0.18	0.00	0.04	0.17	0.33	0.88	13,794
FNC	0.21	0.19	0.00	0.05	0.16	0.33	0.88	13,722
MSN	0.27	0.21	0.00	0.09	0.25	0.42	0.88	13,680

Notes: The table shows statistics built using only observations belonging to event windows that did not partially overlap over time. These amount to a total of 2,346 windows, some for which it was not possible to construct a coverage measure for specific outlets - TV News Archive (Link) transcripts cover an average of 98% of all content broadcasted by CNN, Fox News and MSNBC. As expressed in column “Observations”, these event windows cover a total of 10,299 hours of cable content (≈ 429 days).

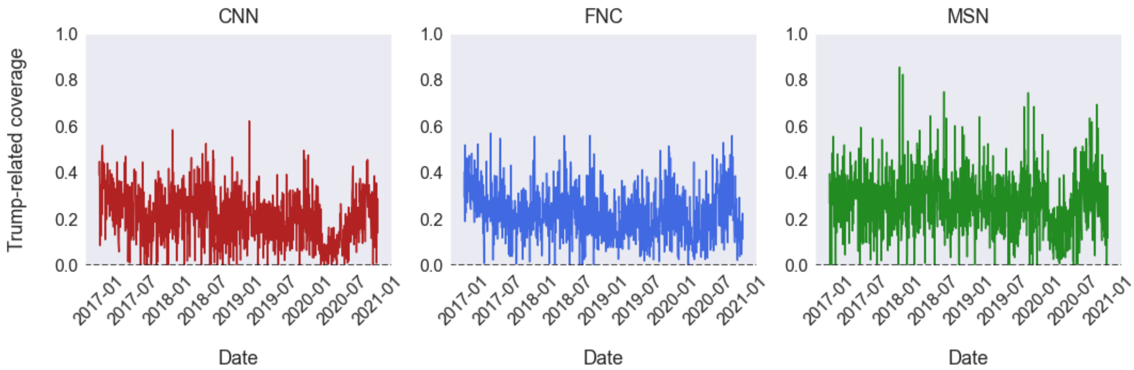

Figure 1.15: **Measure across sample**. The figure plots the average share of time devoted to Trump-related issues, per day, by network, during non-overlapping event windows (over time).

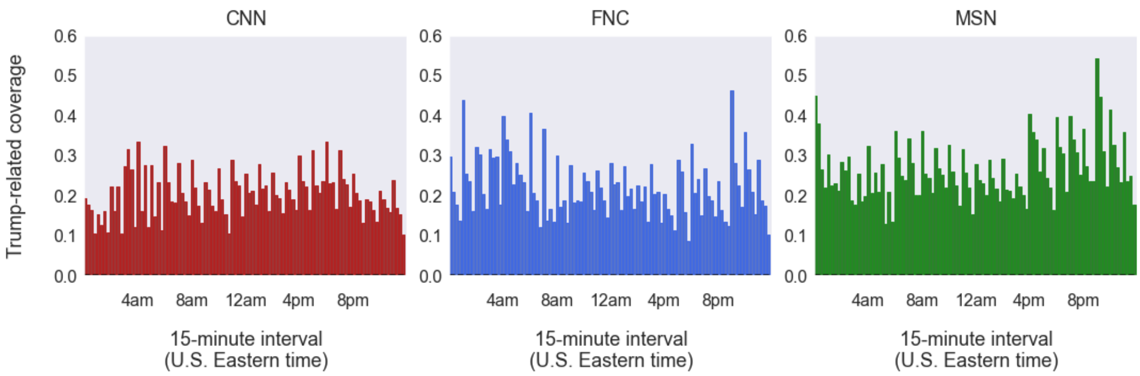

Figure 1.16: **Measure within day**. The figure plots the average share of time devoted to Trump-related issues, within a generic day, per 15 minute interval, by network, during non-overlapping event windows (over time).

Similarity between television news and Trump tweets

1.1.4.1. During non-overlapping event windows (over time)

Table 1.3: **Descriptive statistics**

	Mean	Std. Dev.	Min.	p25	Median	p75	Max.	Observations
CNN	0.07	0.36	0.00	0.00	0.00	0.00	4.06	13,794
FNC	0.08	0.36	0.00	0.00	0.00	0.00	3.74	13,722
MSN	0.06	0.33	0.00	0.00	0.00	0.00	4.13	13,680

Notes: The table shows statistics built using only observations belonging to event windows that did not partially overlap over time. These amount to a total of 2,346 windows, some for which it was not possible to construct a coverage measure for specific outlets - TV News Archive (Link) transcripts cover an average of 98% of all content broadcasted by CNN, Fox News and MSNBC. As expressed in column "Observations", these event windows cover a total of 10,299 hours of cable content (≈ 429 days).

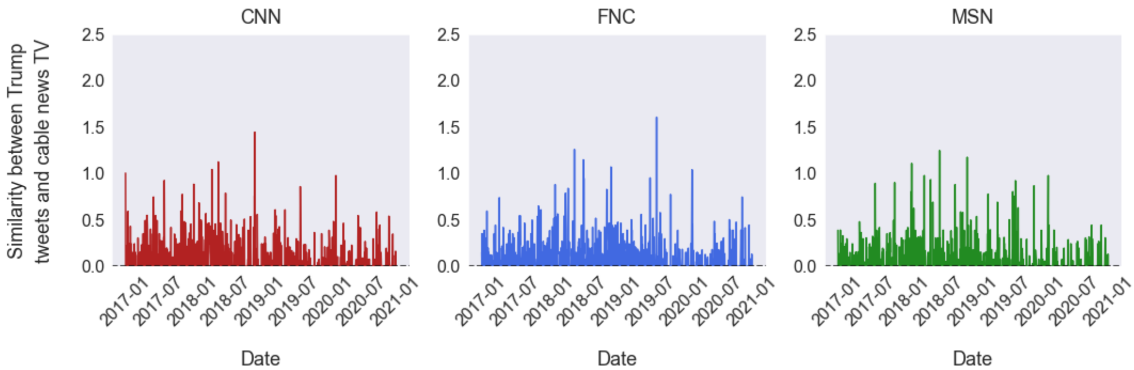

Figure 1.17: **Measure across sample**. The figure plots the average similarity between cable news and tweets, per day, by network, during non-overlapping event windows (over time).

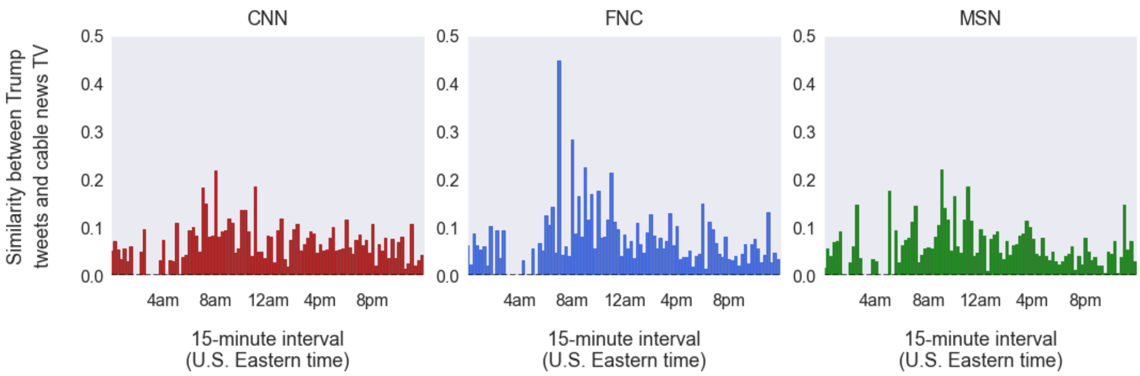

Figure 1.18: **Measure within day**. The figure plots the average similarity between cable news transcripts and Trump tweets, within day, per 15 minute interval, by network, during non-overlapping event windows (over time).

1.1.4.2. During non-overlapping event windows (over content)

Table 1.4: **Descriptive statistics**

	Mean	Std. Dev.	Min.	p25	Median	p75	Max.	Observations
CNN	0.06	0.34	0.00	0.00	0.00	0.00	4.06	53,856
FNC	0.07	0.36	0.00	0.00	0.00	0.00	4.22	53,250
MSN	0.06	0.32	0.00	0.00	0.00	0.00	4.13	53,190

Notes: The table shows statistics built using only observations belonging to event windows that did not partially overlap over content. These amount to a total of 9,141 windows, some for which it was not possible to construct a coverage measure for specific outlets - TV News Archive (Link) transcripts cover an average of 98% of all content broadcasted by CNN, Fox News and MSNBC. As expressed in column "Observations", these event windows cover a total of 40,074 hours of cable content (≈ 1,670 days).

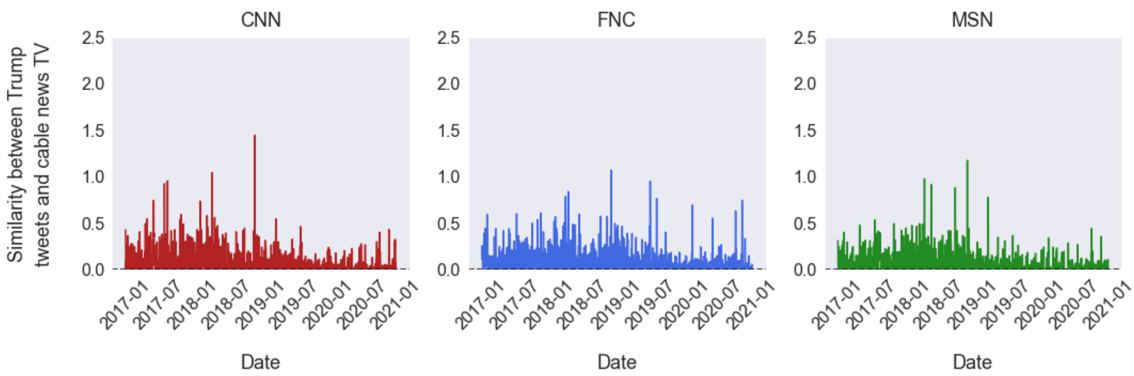

Figure 1.19: **Measure across sample**. The figure plots the average similarity between cable news transcripts and Trump tweets, per day, by network, during non-overlapping event windows (over content).

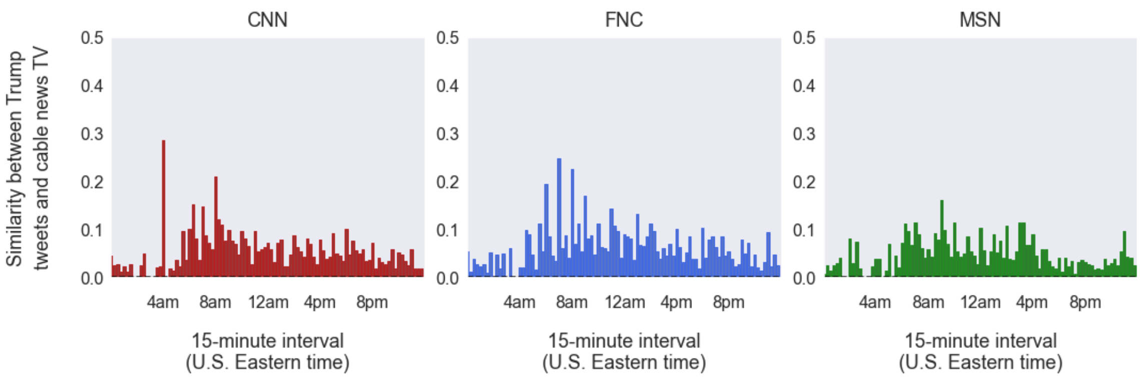

Figure 1.20: **Measure within day**. The figure plots the average similarity between cable news transcripts and Trump tweets, within day, per 15 minute interval, by network, during non-overlapping event windows (over content).

Similarity between television news and past Facebook news

1.1.5.1. During non-overlapping event windows (over time)

Table 1.5: **Descriptive statistics**

	Mean	Std. Dev.	Min.	p25	Median	p75	Max.	Observations
CNN	1.18	0.97	0.00	0.00	0.88	2.09	3.69	492
FNC	0.29	0.60	0.00	0.00	0.00	0.00	3.09	546
MSN	0.28	0.58	0.00	0.00	0.00	0.00	3.00	390

Notes: The table shows statistics built using only observations (1) belonging to event windows that did not partially overlap over time and in which (2) TV news converged in content towards Trump tweets. Considering (2), windows have been selected by (2.i) demeaning similarity between TV and Trump tweets (at a network × event window level) and (2.ii) selecting windows where demeaned outcome was higher after a tweet (relative to before). These amount to a total of 238 windows (357 hours, ≈ 15 days).

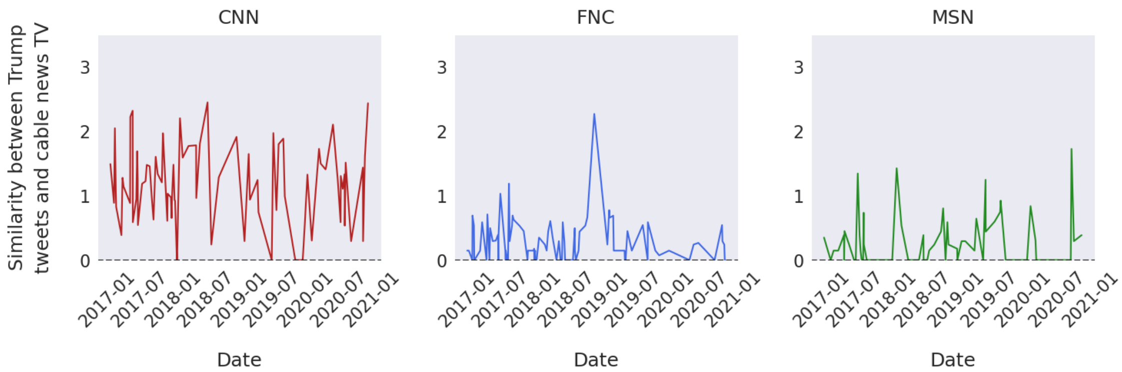

Figure 1.21: **Measure across sample**. The figure plots the average similarity between TV and past FB news, per day, by network, during windows that (1) do not overlap over time and in which (2) TV converged in content towards Trump tweets.

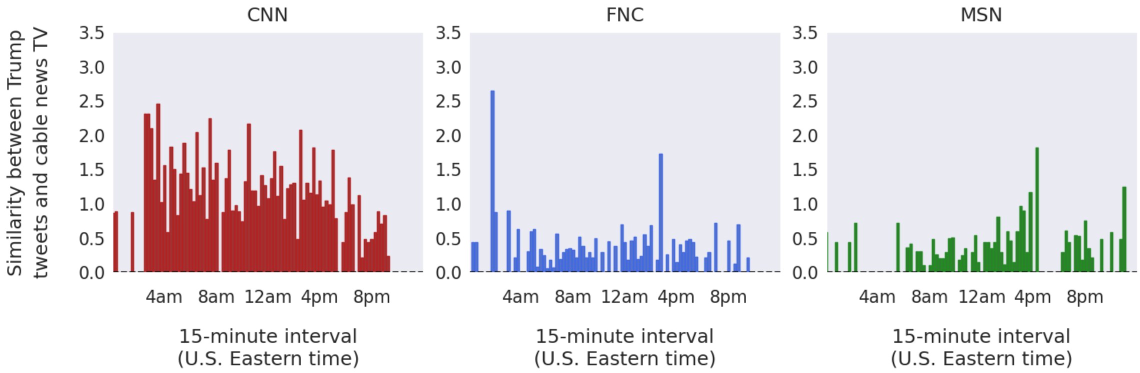

Figure 1.22: **Measure within day**. The figure plots the average similarity between TV and past FB news, within day, per 15 minute interval, by network, during windows that (1) do not overlap over time and in which (2) TV converged in content towards Trump tweets.

1.1.5.2. During non-overlapping event windows (over content)

Table 1.6: **Descriptive statistics**

	Mean	Std. Dev.	Min.	p25	Median	p75	Max.	Observations
CNN	1.13	0.99	0.00	0.00	0.88	1.82	4.06	1,872
FNC	0.33	0.65	0.00	0.00	0.00	0.00	3.58	2,322
MSN	0.29	0.61	0.00	0.00	0.00	0.00	3.64	1,686

Notes: The table shows statistics built using only observations (1) belonging to event windows that did not partially overlap over content and in which (2) TV news converged in content towards Trump tweets. Considering (2), windows have been selected by (2.i) demeaning similarity between TV and Trump tweets (at a network $\times$ event window level) and (2.ii) selecting windows where demeaned outcome was higher after tweet (relative to before). These amount to a total of 980 windows (1,470 hours, $\approx$ 61 days).

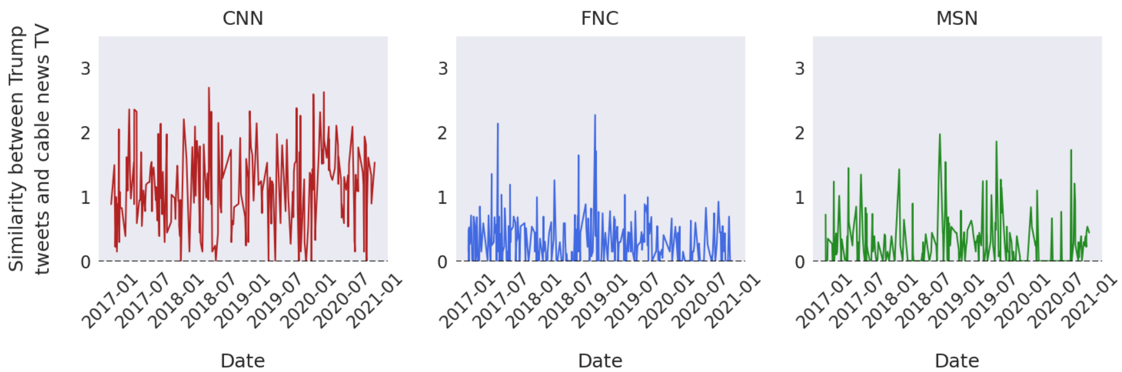

Figure 1.23: **Measure across sample**. The figure plots the average similarity between TV and past FB news, per day, by network, during windows that (1) do not overlap over time and in which (2) TV converged in content towards Trump tweets.

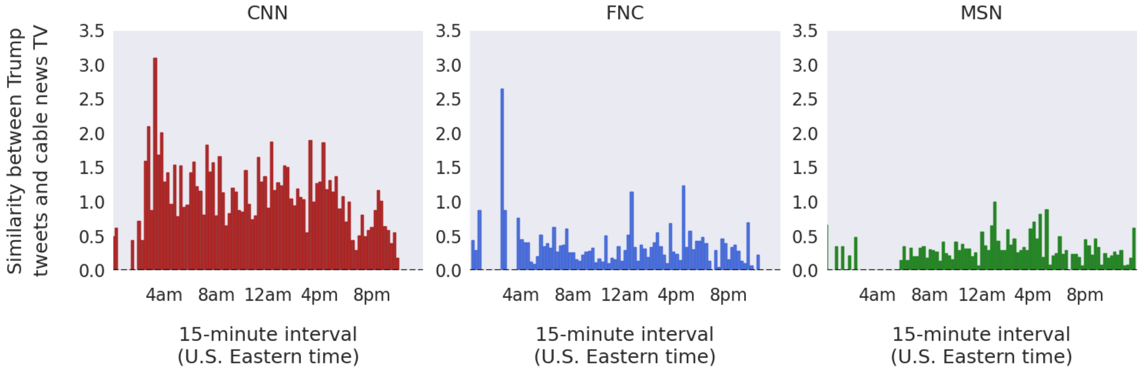

Figure 1.24: **Measure within day**. The figure plots the average similarity between TV and past FB news, within day, per 15 minute interval, by network, during windows that (1) do not overlap over time and in which (2) TV converged in content towards Trump tweets.

Topics addressed on @realDonaldTrump tweets

Table 1.7: **Topic-word distribution for @realDonaldTrump tweets**

Topic	Theme	Words
1	Domestic Events	endorsement, governor, alabama, state, border, senator, luther_strange, congressman, fantastic, strong_crime, georgia, taxis, tough_crime, love_military, weak_crime, endorse, want_raise, race, republican, need
2	Foreign Policy	north_korea, meeting, prime_minister, great_honor, meet, china, today, south_korea, honor_welcome, host, japan, peace, summit, whitehouse_today, dinner, leader, delegation, launch, white_house, turkey
3	Collusion Charges	crooked_hillary, witch_hunt, russia, caign, collusion, clinton, investigation, russian, comey, hillary_clinton, mueller, phony, dossier, report, information, server, james_comey, hoax, election, obama
4	Immigration	democrats, republican, obamacare, daca, bill, need, immigration, wall, senate, border_security, democrat, healthcare, border, pass, southern_border, repeal_replace, crime, congress, house, republican_senator
5	Veterans / Military	great_honor, today, hero, welcome, honor, whitehouse, nation, woman, service, veteran, american, life, sacrifice, celebrate, serve, america, memorial, national, brave, vietnam
6	News Media	fake_news, medium, story, fake, nytimes, dishonest, write, report, wrong, fail_, news, reporting, mainstream_medium, fact, cnn, washington_post, book, hate, know, media
7	Economy	stock_market, economy, unemployment, record, number, plant, american, high, company, prosperity, manufacturing, record_high, worker, growth, economic, america, optimism, business, regulation, create
8	Trade Policy	china, north_korea, trade, deal, tariff, farmer, dollar, trade_deal, disrespect, negotiation, iran, relationship, trade_deficit, nafta, missile, many_year, united_state, product, negotiate, money
9	White House	join, melania, crowd, white_house, tonight, evening, maga, land, beautiful, rally, incredible, south_carolina, wonderful, first_lady, chion, last_night, flotus_melania, head, paris, arrive
10	Shootings / Disasters	first_responder, thought_prayer, hurricane, family, fema, school, shooting, bless, california, teacher, heart, storm, london, pray, terrible, train, prayer, victim, tragedy, state_local

Note: Topic-word distribution (20 words most likely to feature in a document of that topic) for a 10 topic Biterm Topic Model (Yan et al., 2013) fitted on a selected pre-processed corpus of @realDonaldTrump tweets.

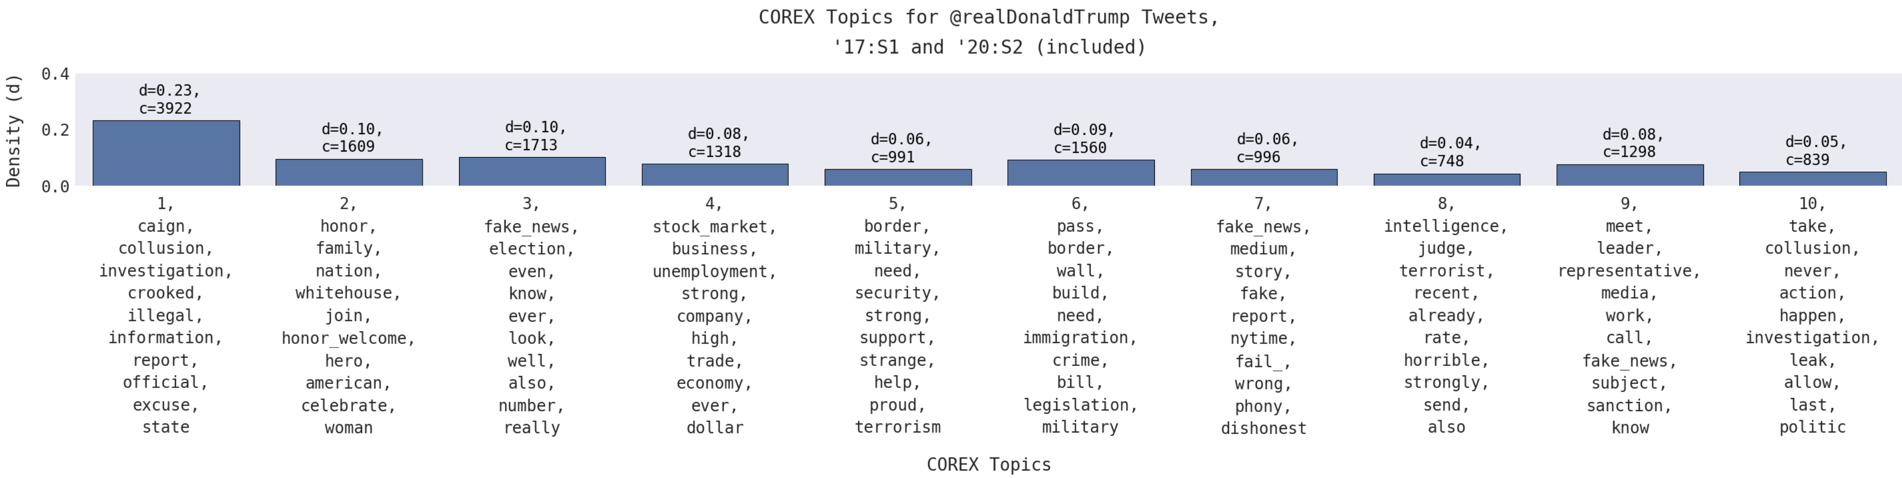

Figure 1.25: **Topic distribution for @realDonaldTrump tweets**. The figure plots the distribution of topics (inferred from a Biterm Topic Model, explained in Section 1.2.2) for a selected sample of @realDonaldTrump posted from 2017 until 2020 - without counting retweets nor short tweets.

1.6.2 Results and Robustnesses

Results. Reaction through Coverage

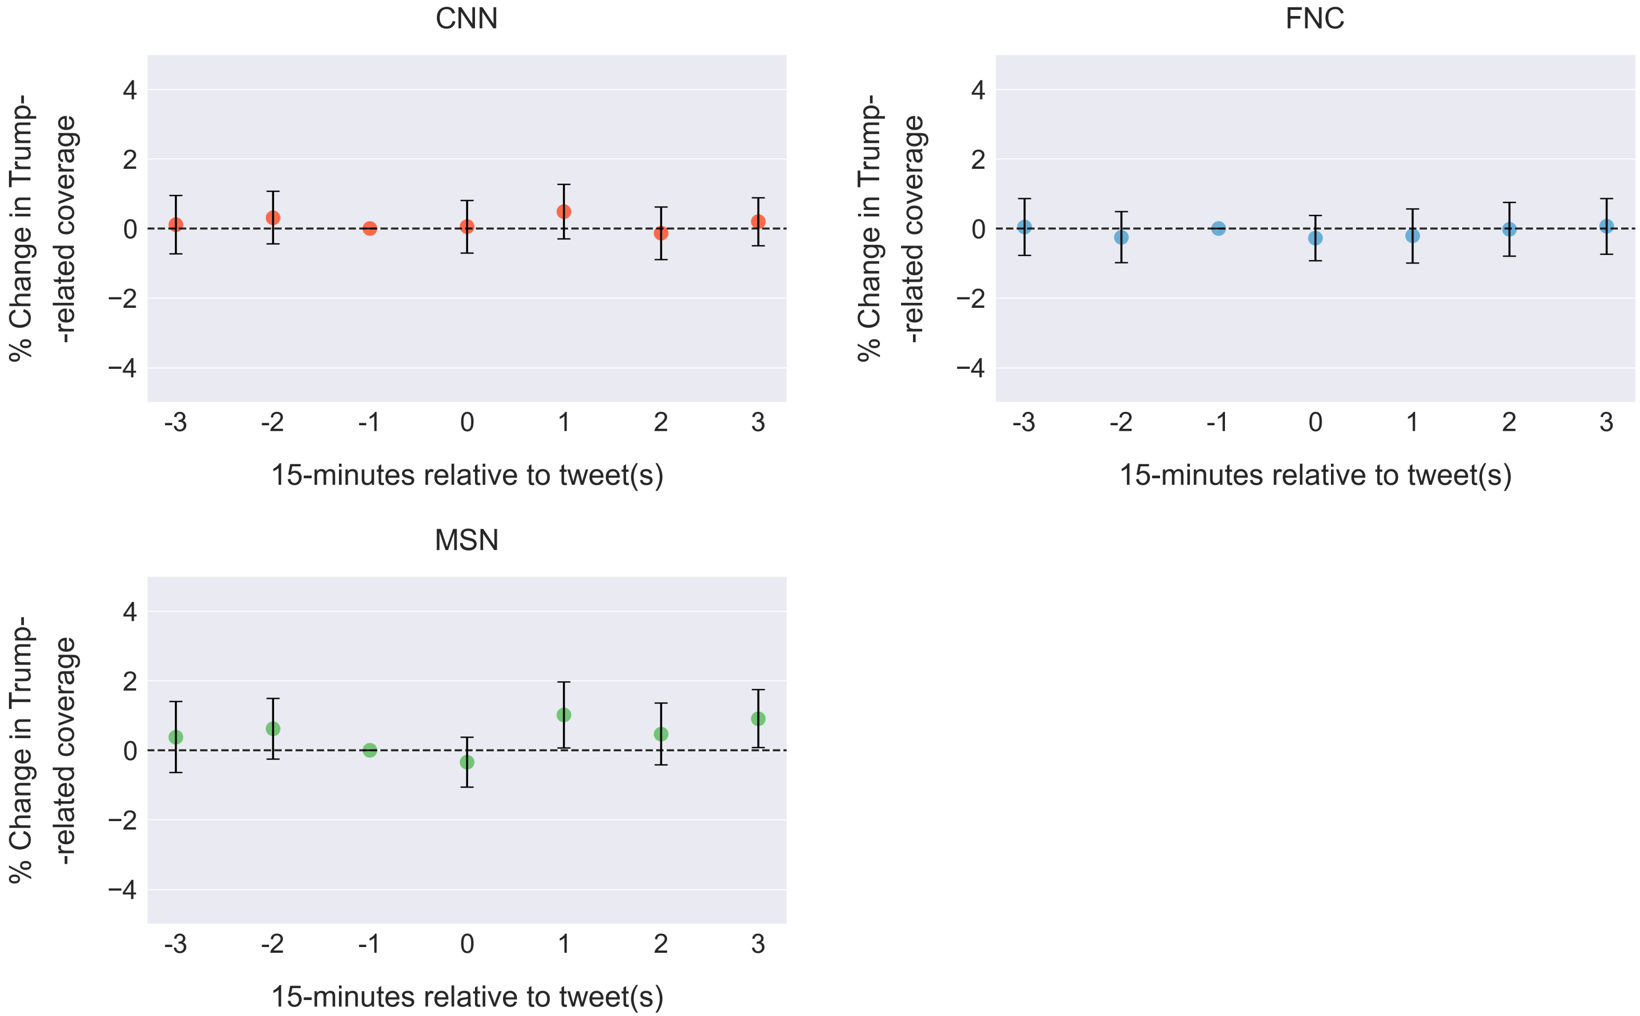

Figure 1.26: **Coverage before and after a @realDonaldTrump tweet**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level. Dependent variable has been subjected to an inverse hyperbolic sine transformation.

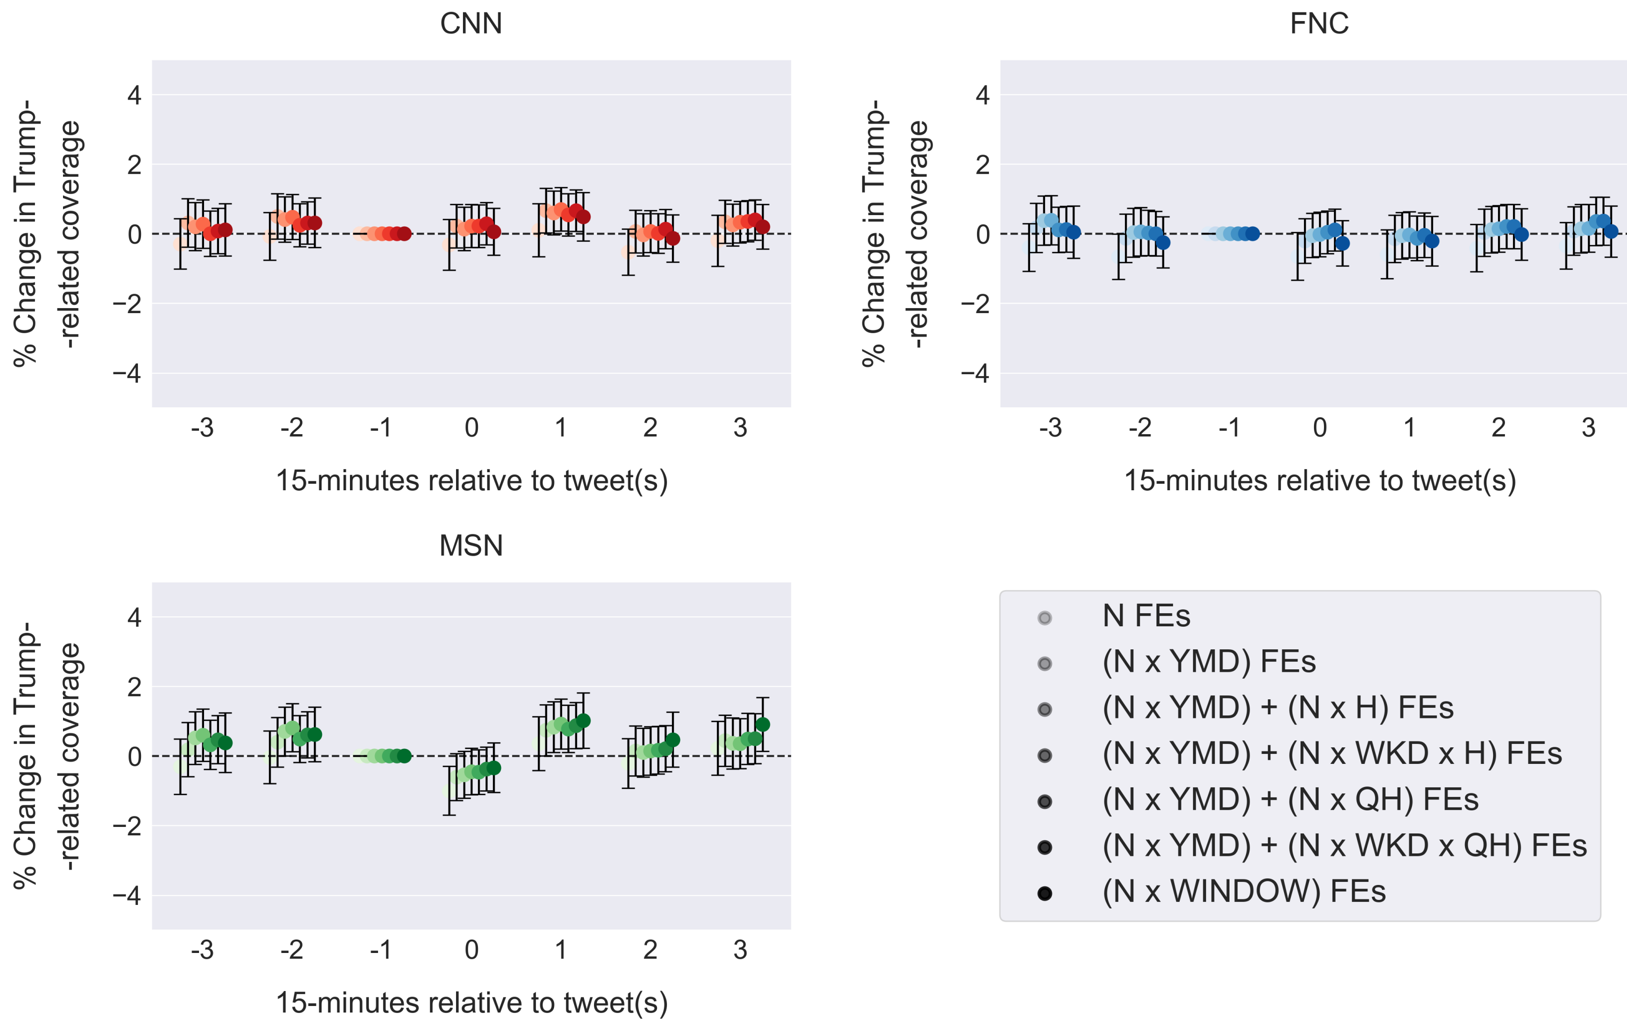

Figure 1.27: **Coverage before and after a tweet // Varying FEs specifications**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated by controlling for unit-specific macro factors computed using only those observations belonging to non-overlapping event windows (across time). From left to right: (i) network FEs; (ii) network x date FEs; (iii) network x date and network x hour-of-day FEs; (iv) network x date and network x week-day x hour-of-day FEs; (v) network x date and network x quarter-hour-of-day FEs; (vi) network x date and network x week-day x quarter-hour-of-day FEs and (vii) network x event window FEs. Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level.

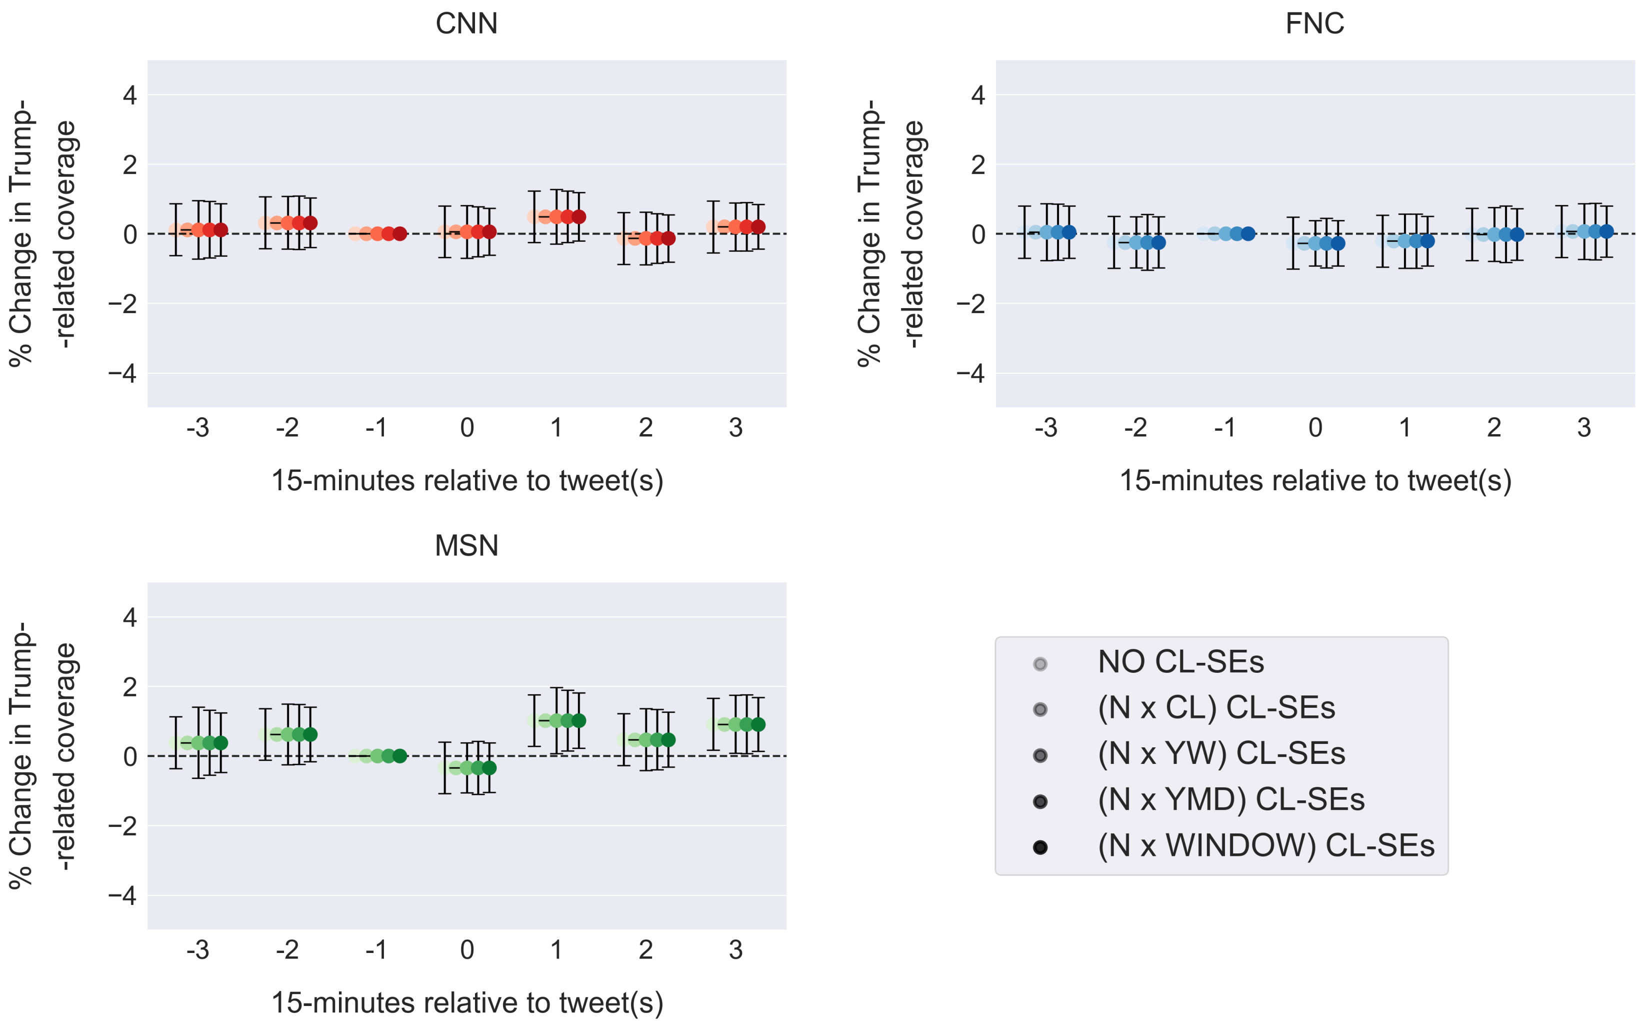

Figure 1.28: **Coverage before and after a tweet // Varying SEs clusters**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% confidence intervals computed with clustered SEs. From left to right: (i) non-clustered SEs; (ii) SEs clustered by network; (iii) SEs clustered by network x week-of-the-year; (iv) SEs clustered by network x date; (v) network x window.

Results. Reaction through Content

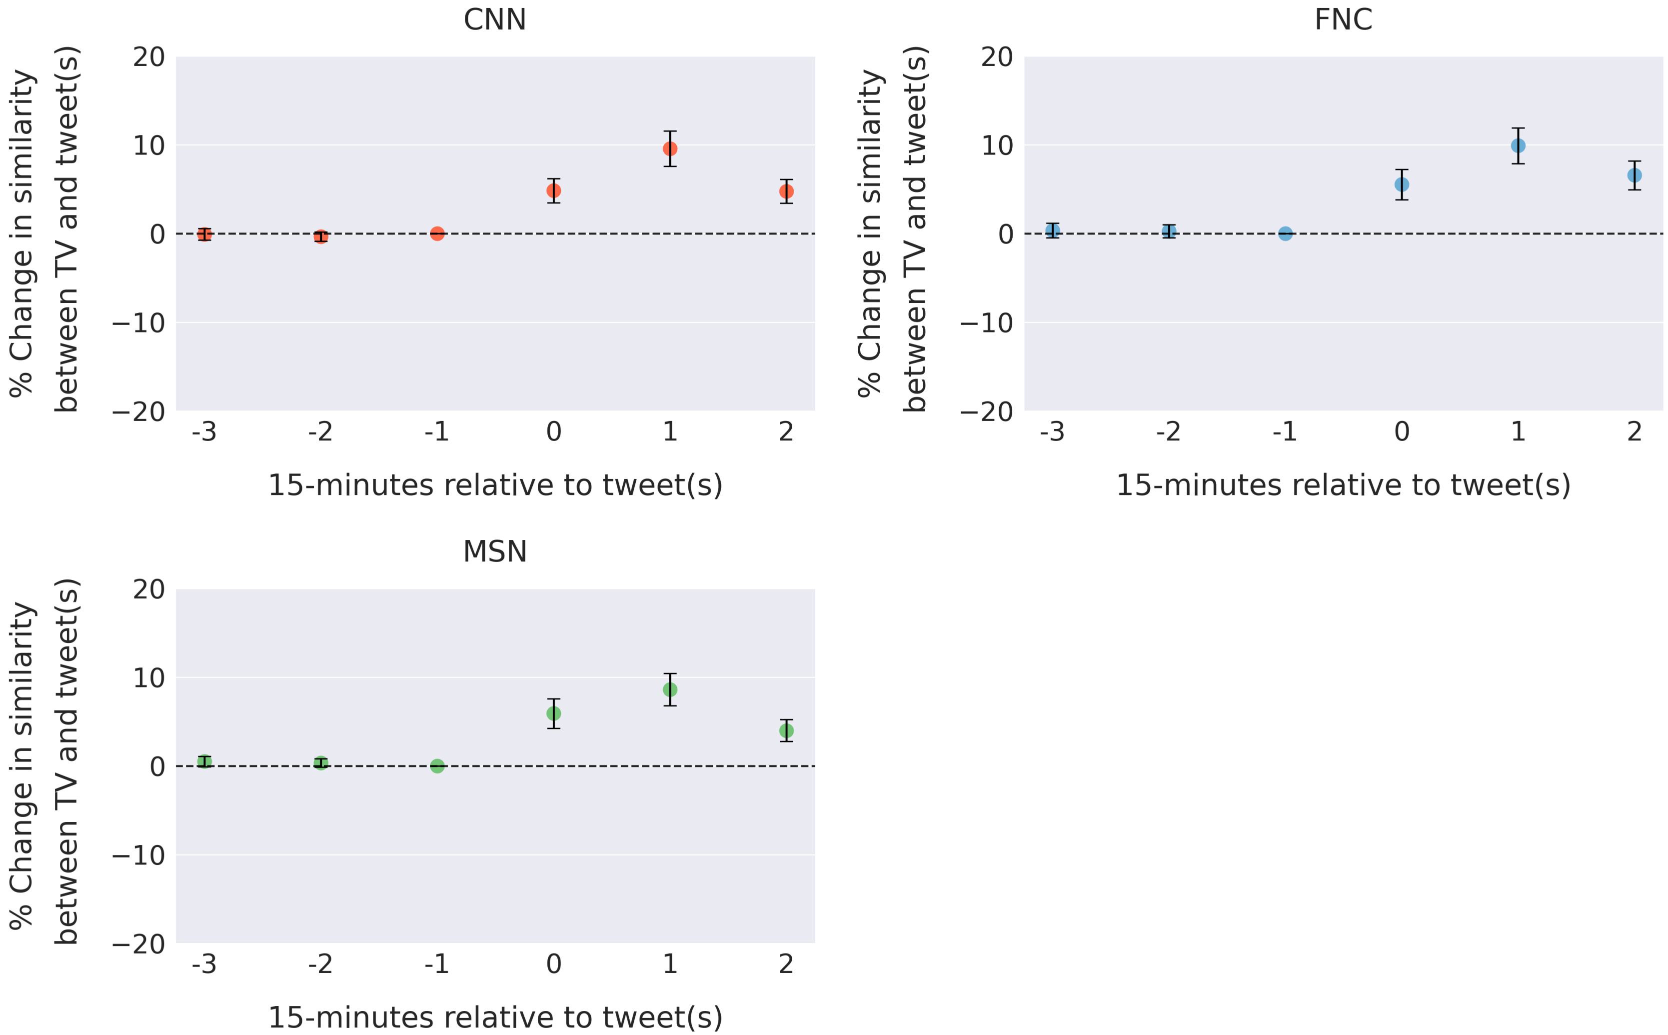

Figure 1.29: **Similarity before and after a @realDonaldTrump tweet // Non-overlapping windows across time**. Estimates refer a to selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level. Dependent variable has been subjected to an inverse hyperbolic sine transformation.

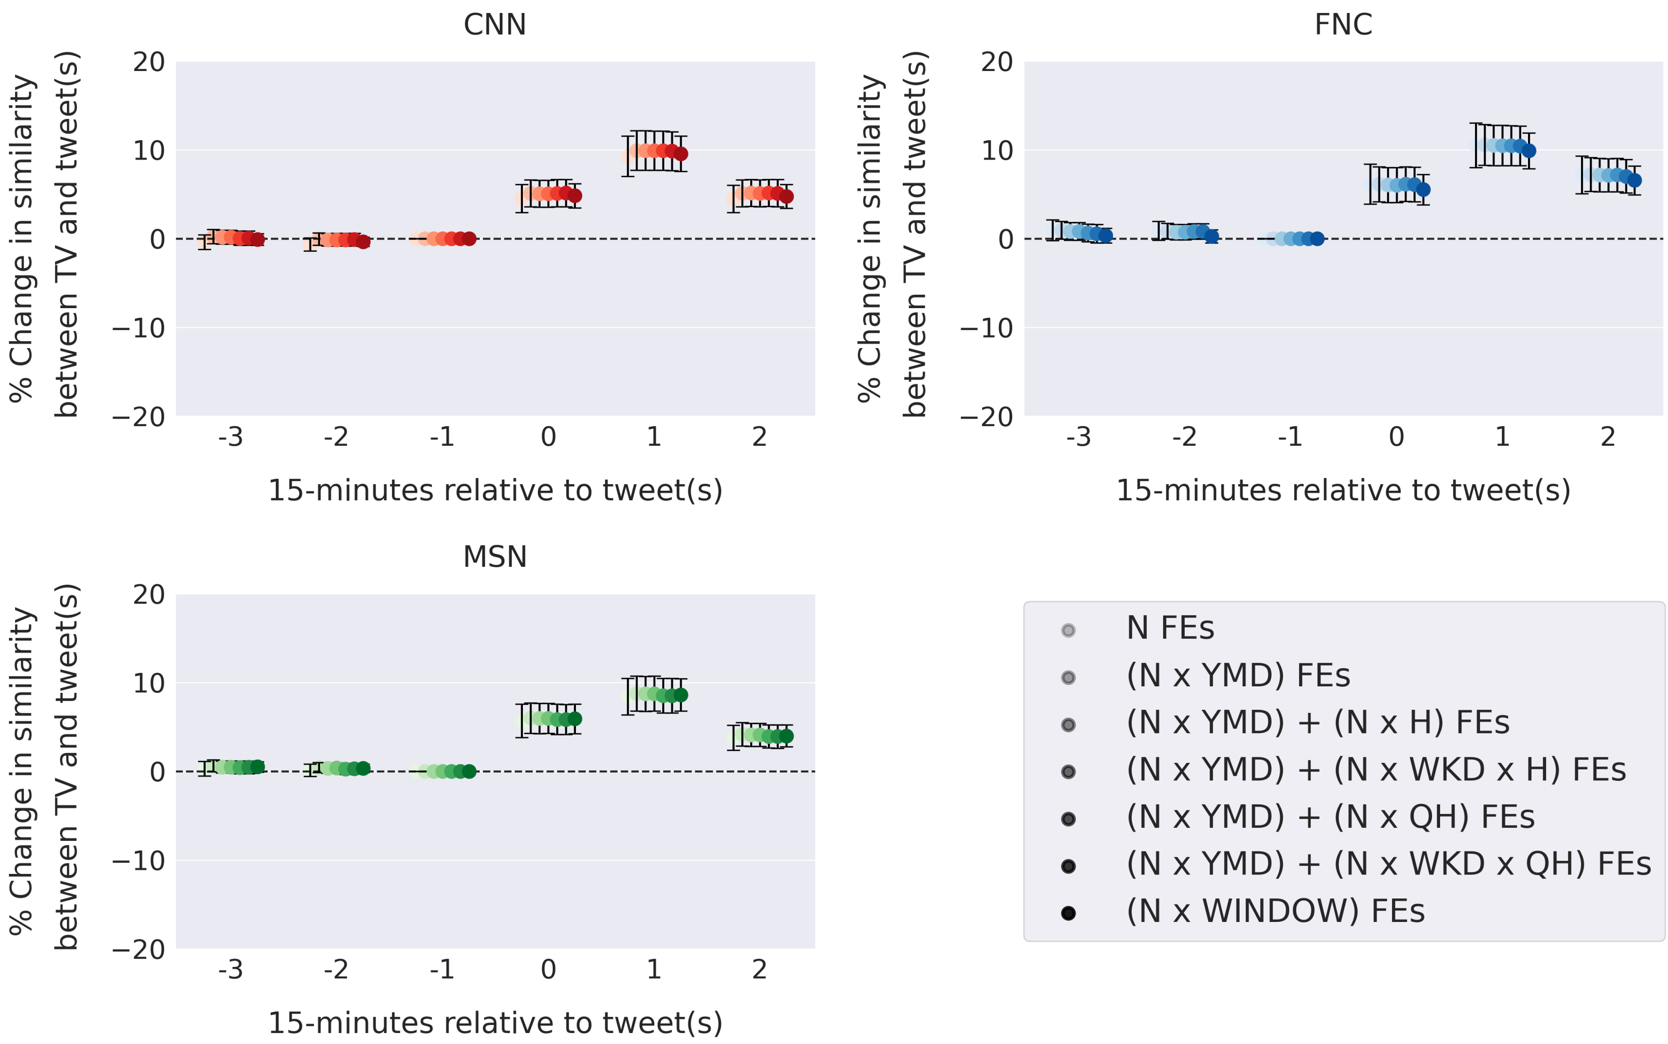

Figure 1.30: **Similarity before and after a tweet // Non-overlaps across time // Varying FEs specifications**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated by controlling for unit-specific macro factors computed using only those observations belonging to non-overlapping event windows (across time). From left to right: (i) network FEs; (ii) network x date FEs; (iii) network x date and network x hour-of-day FEs; (iv) network x date and network x week-day x hour-of-day FEs; (v) network x date and network x quarter-hour-of-day FEs; (vi) network x date and network x week-day x quarter-hour-of-day FEs and (vii) network x event window FEs. Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level.

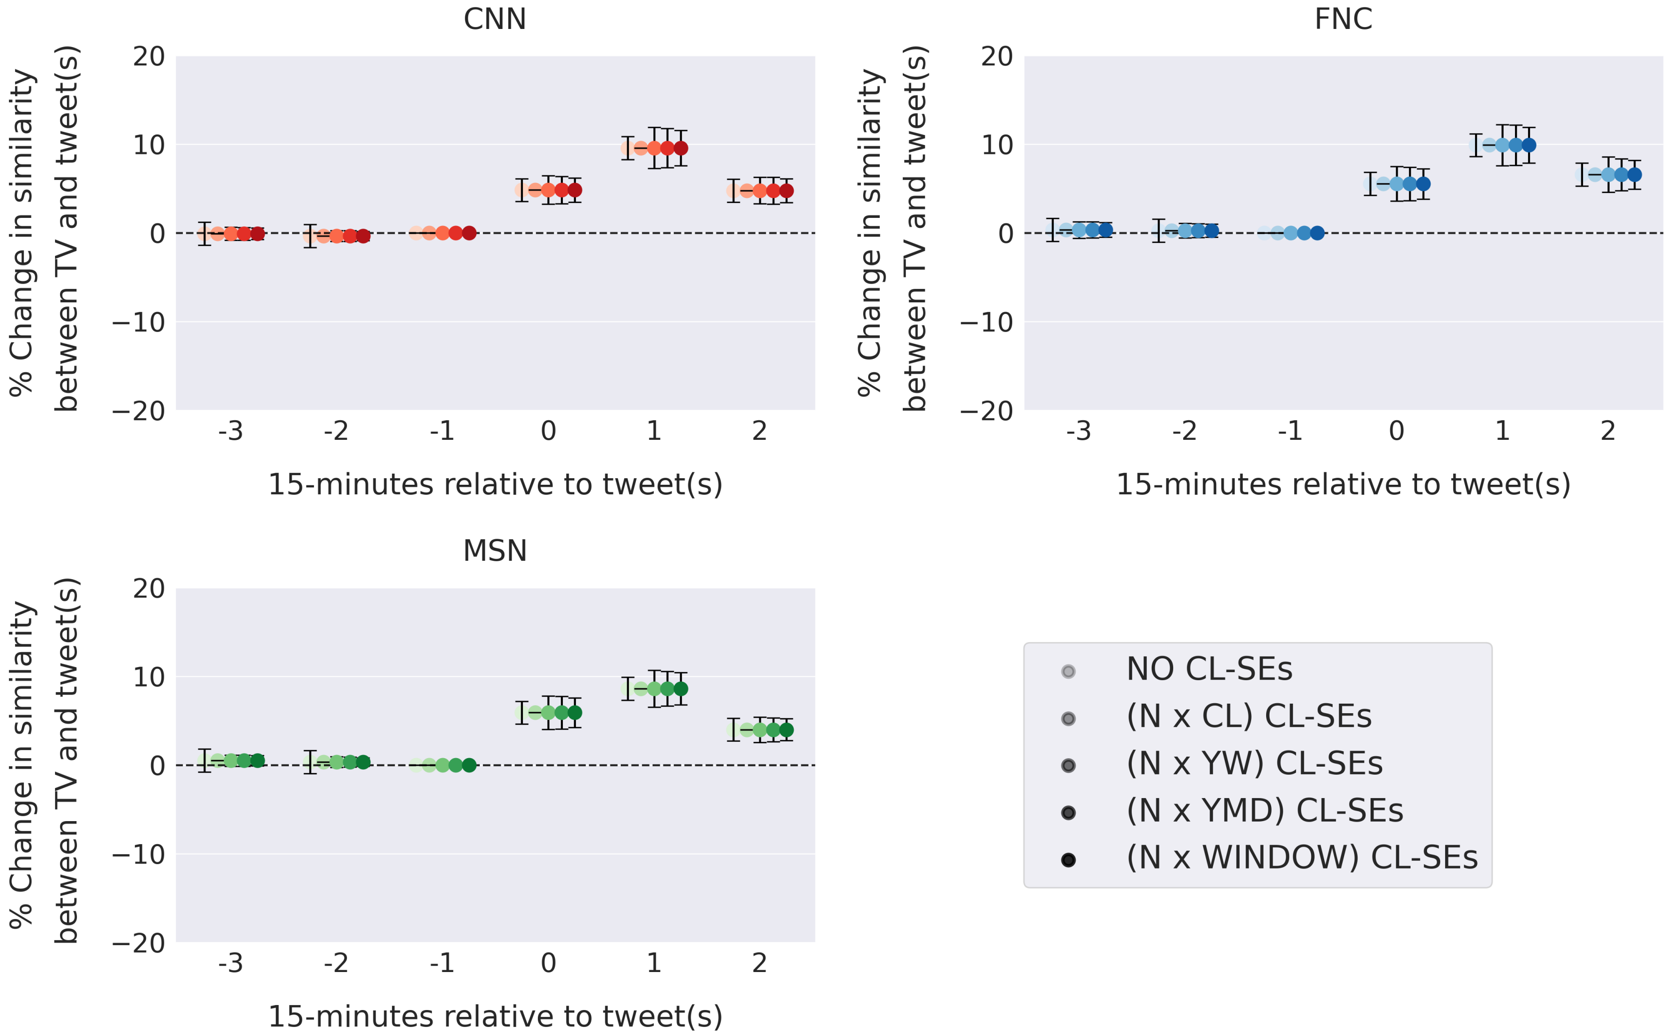

Figure 1.31: **Similarity before and after a tweet // Non-overlaps across time // Varying SEs clusters**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% confidence intervals computed with clustered SEs. From left to right: (i) non-clustered SEs; (ii) SEs clustered by network; (iii) SEs clustered by network x week-of-the-year; (iv) SEs clustered by network x date; (v) network x window.

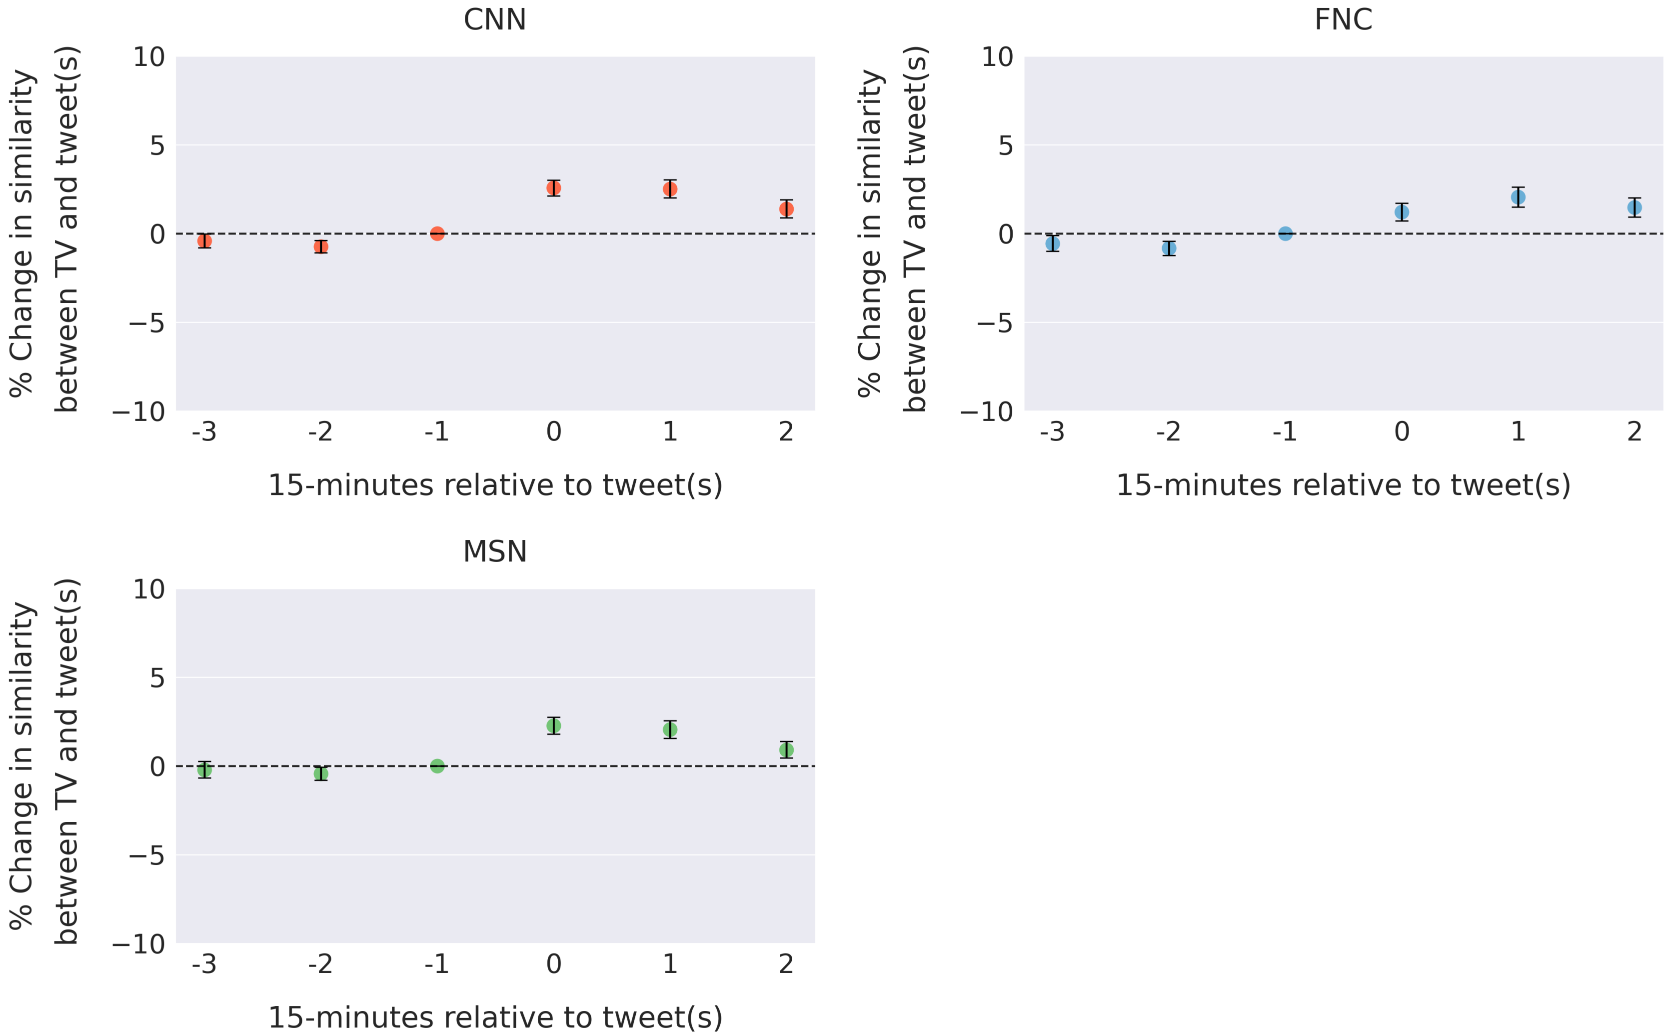

Figure 1.32: **Similarity before and after a tweet // Non-overlapping windows across content**. Estimates refer a to selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations belonging to non-overlapping event windows (across conent). Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level. Dependent variable has been subjected to an inverse hyperbolic sine transformation.

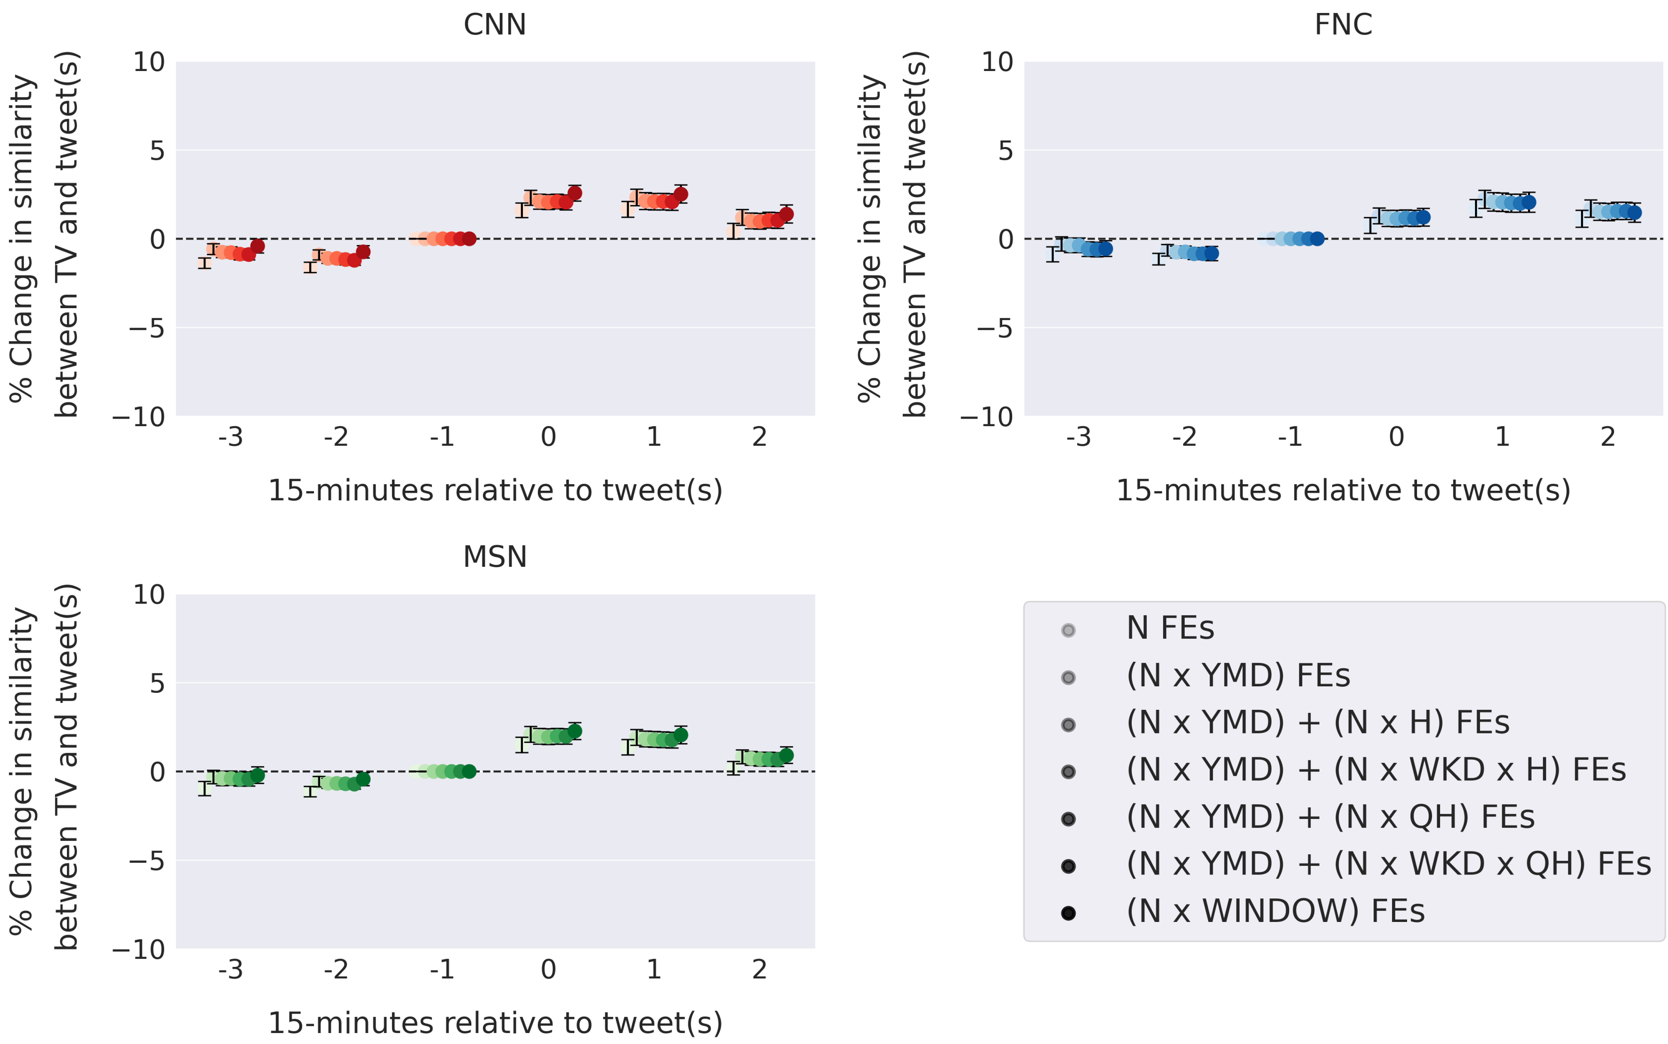

Figure 1.33: **Similarity before and after a tweet** // **Non-overlaps across content** // **Varying FEs specifications**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated by controlling for unit-specific macro factors computed using only those observations belonging to non-overlapping event windows (across content). From left to right: (i) network FEs; (ii) network x date FEs; (iii) network x date and network x hour-of-day FEs; (iv) network x date and network x week-day x hour-of-day FEs; (v) network x date and network x quarter-hour-of-day FEs; (vi) network x date and network x week-day x quarter-hour-of-day FEs and (vii) network x event window FEs. Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level.

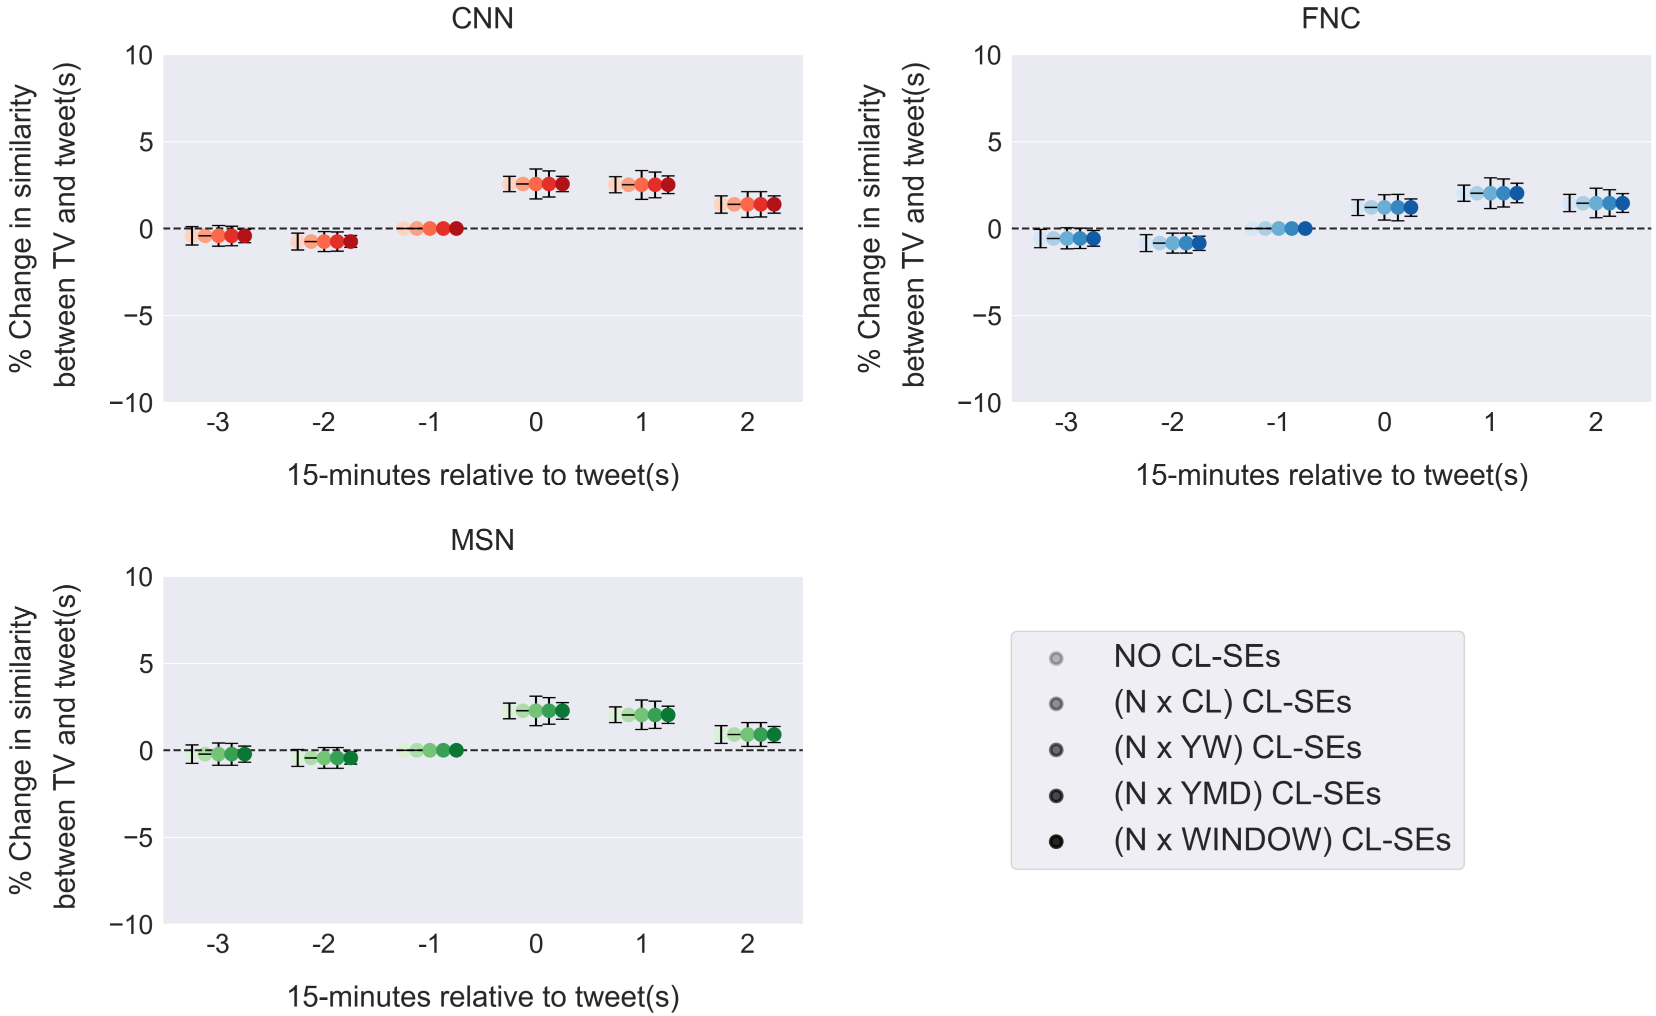

Figure 1.34: **Similarity before and after a tweet // Non-overlaps across content // Varying SEs clusters**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations belonging to non-overlapping event windows (across content). Error bars stand for 95% confidence intervals computed with clustered SEs. From left to right: (i) non-clustered SEs; (ii) SEs clustered by network; (iii) SEs clustered by network x week-of-the-year; (iv) SEs clustered by network x date; (v) network x window.

Robustness. Content Reaction due to Breaking News

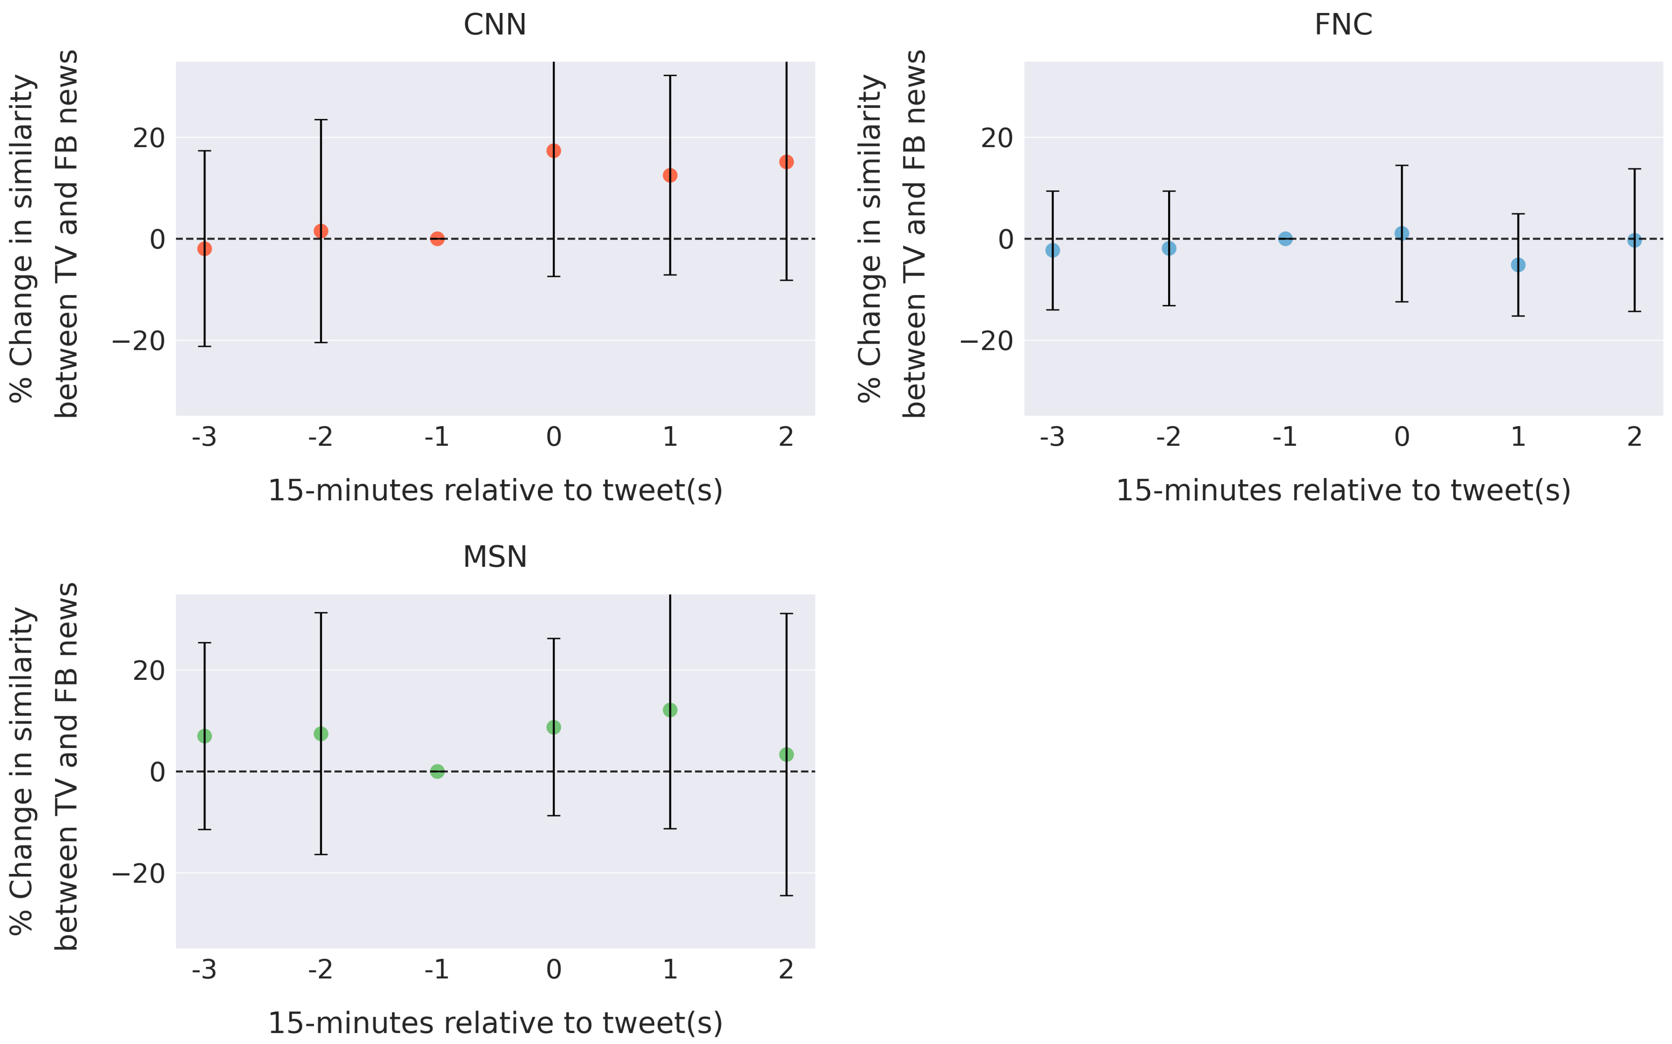

Figure 1.35: **Similarity before and after a tweet between TV and Facebook (FB) news // Non-overlapping windows across time**. Estimates refer a to selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped; (iii) tweets for which cable news does not converge in content minutes after a tweet are dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations that belong to (a) non-overlapping event windows (across time) and (b) non-overlapping event windows where cable news converged in content towards Donald J. Trump's tweets, minutes after these. Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level. Dependent variable subjected to an inverse hyperbolic sine transformation.

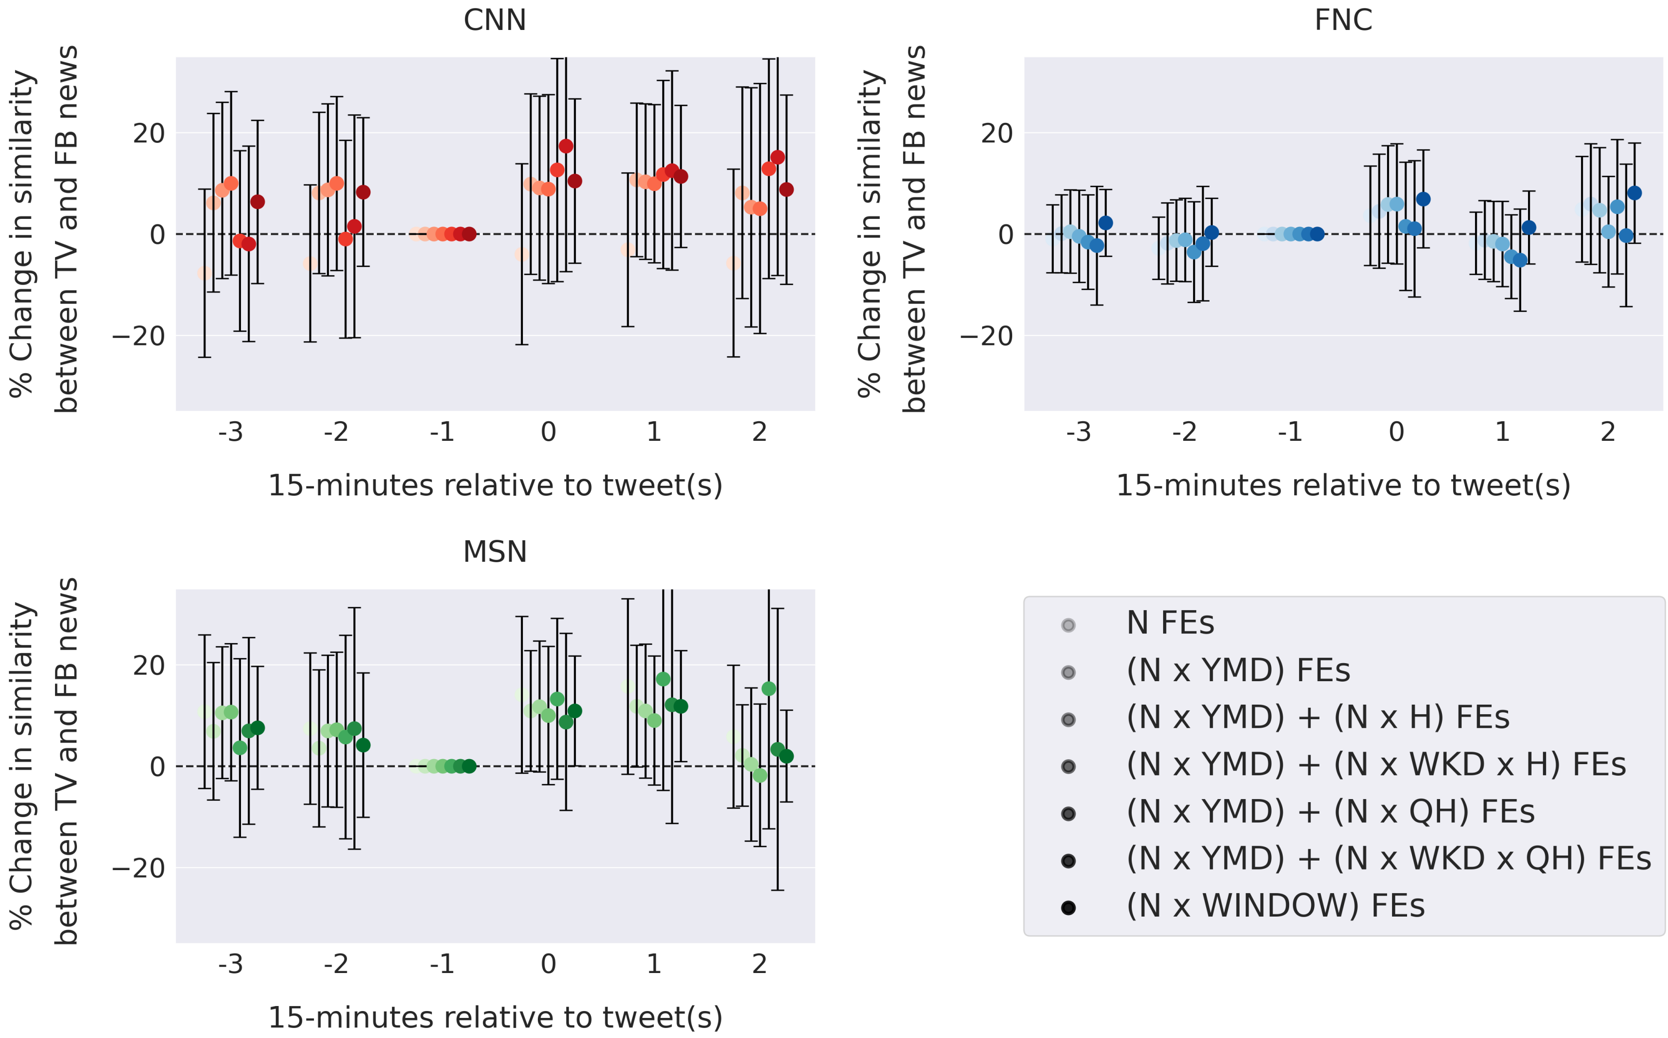

Figure 1.36: **Similarity before and after a tweet between TV and Facebook News // Non-overlaps across time // Varying FEs specifications.** Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over time have been dropped; (iii) tweets for which cable news does not converge in content minutes after a tweet are dropped. Coefficients have been estimated by controlling for unit-specific macro factors computed using only those observations belonging to non-overlapping event windows (across time). From left to right: (i) network FEs; (ii) network x date FEs; (iii) network x date and network x hour-of-day FEs; (iv) network x date and network x week-day x hour-of-day FEs; (v) network x date and network x quarter-hour-of-day FEs; (vi) network x date and network x week-day x quarter-hour-of-day FEs and (vii) network x event window FEs. Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

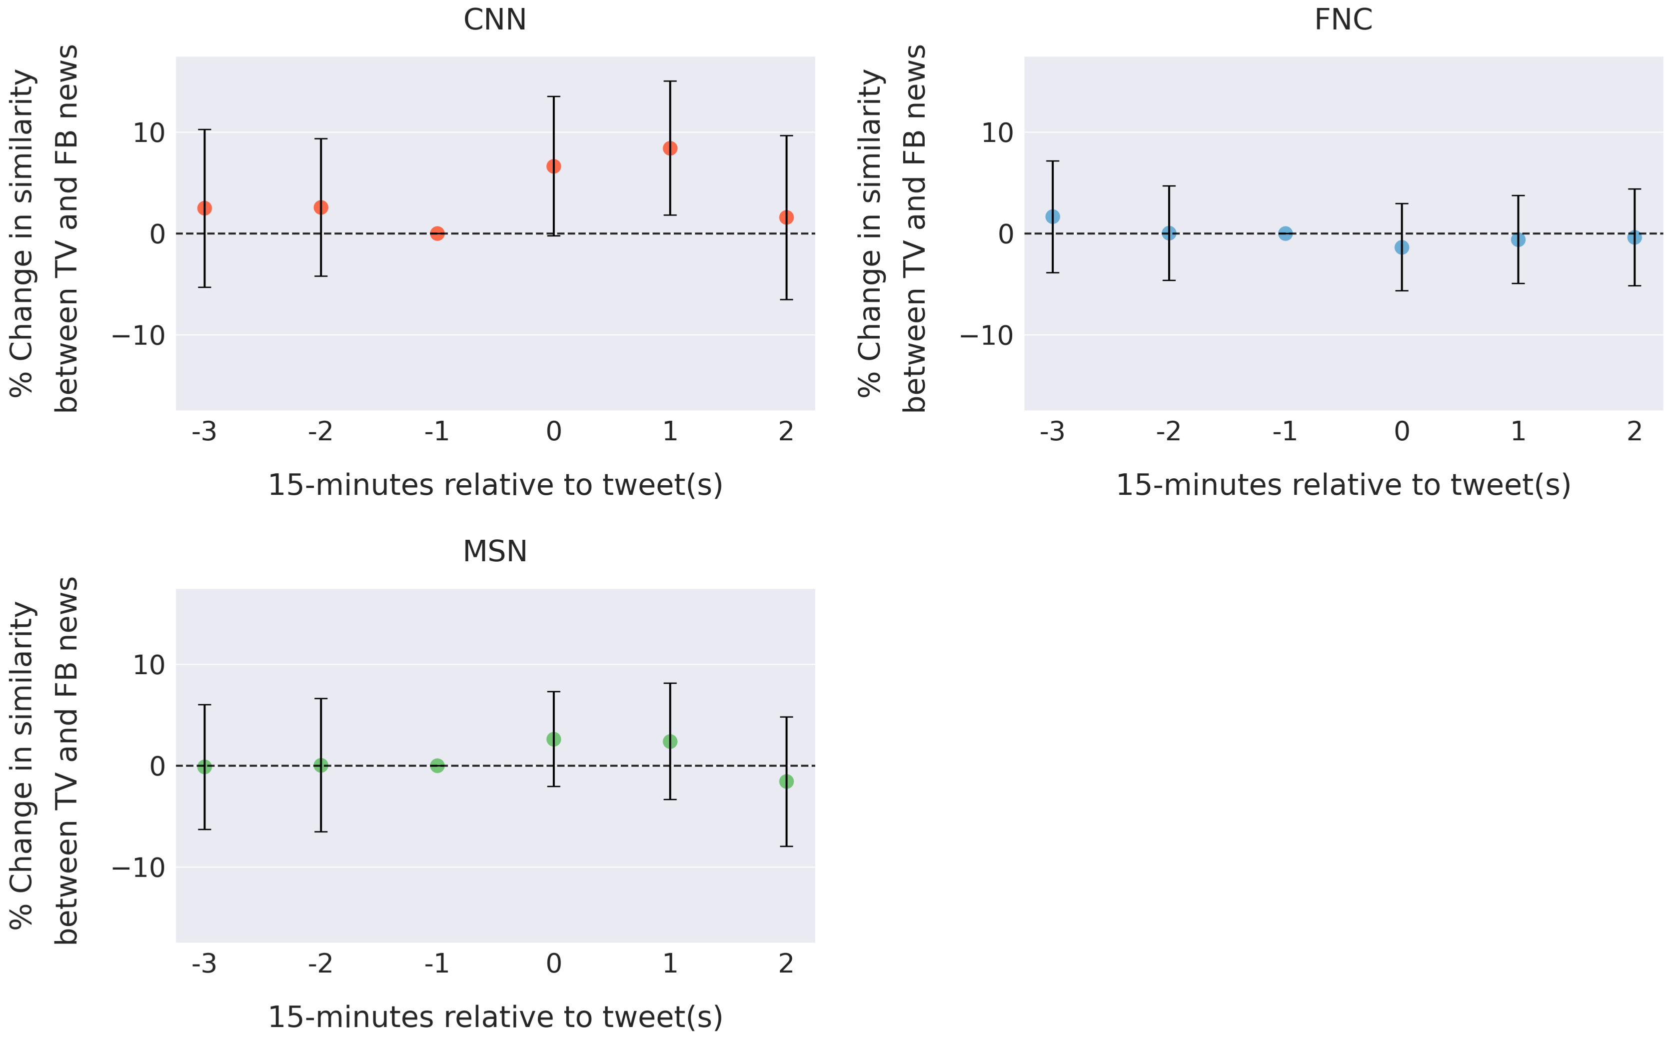

Figure 1.37: **Similarity before and after a tweet between TV and Facebook (FB) news // Non-overlapping windows across content**. Estimates refer a to selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over content have been dropped; (iii) tweets for which cable news does not converge in content minutes after a tweet are dropped. Coefficients have been estimated controlling for network x event-window FEs computed using only those observations that belong to (a) non-overlapping event windows (across content) and (b) non-overlapping event windows where cable news converged in content towards Donald J. Trump's tweets, minutes after these. Error bars stand for 95% confidence intervals computed with SEs clustered at a network-window level. Dependent variable subjected to an inverse hyperbolic sine transformation.

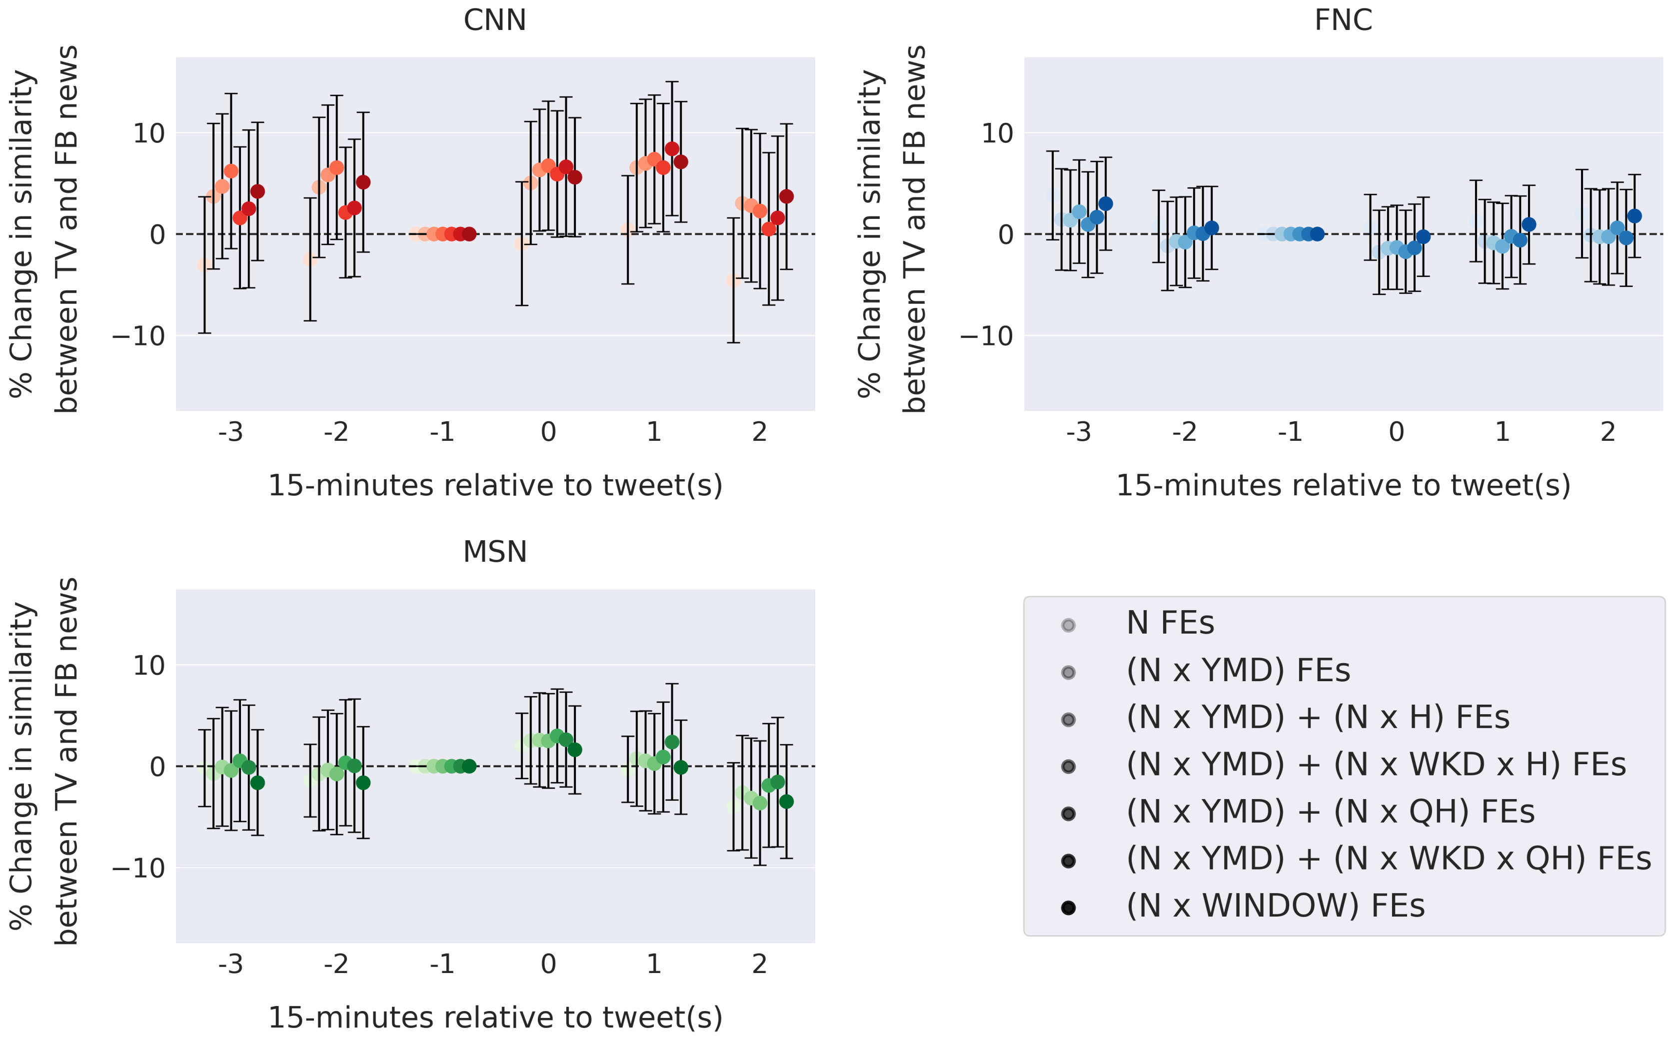

Figure 1.38: **Similarity before and after a tweet between TV and Facebook News // Non-overlaps across content // Varying FEs specifications.** Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020 - (i) short tweets have been dropped; (ii) tweets that generate partially overlapping event windows over content have been dropped; (iii) tweets for which cable news does not converge in content minutes after a tweet are dropped. Coefficients have been estimated by controlling for unit-specific macro factors computed using only those observations belonging to non-overlapping event windows (across content). From left to right: (i) network FEs; (ii) network x date FEs; (iii) network x date and network x hour-of-day FEs; (iv) network x date and network x week-day x hour-of-day FEs; (v) network x date and network x quarter-hour-of-day FEs; (vi) network x date and network x week-day x quarter-hour-of-day FEs and (vii) network x event window FEs. Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

Results. Reactions through Content per Topic

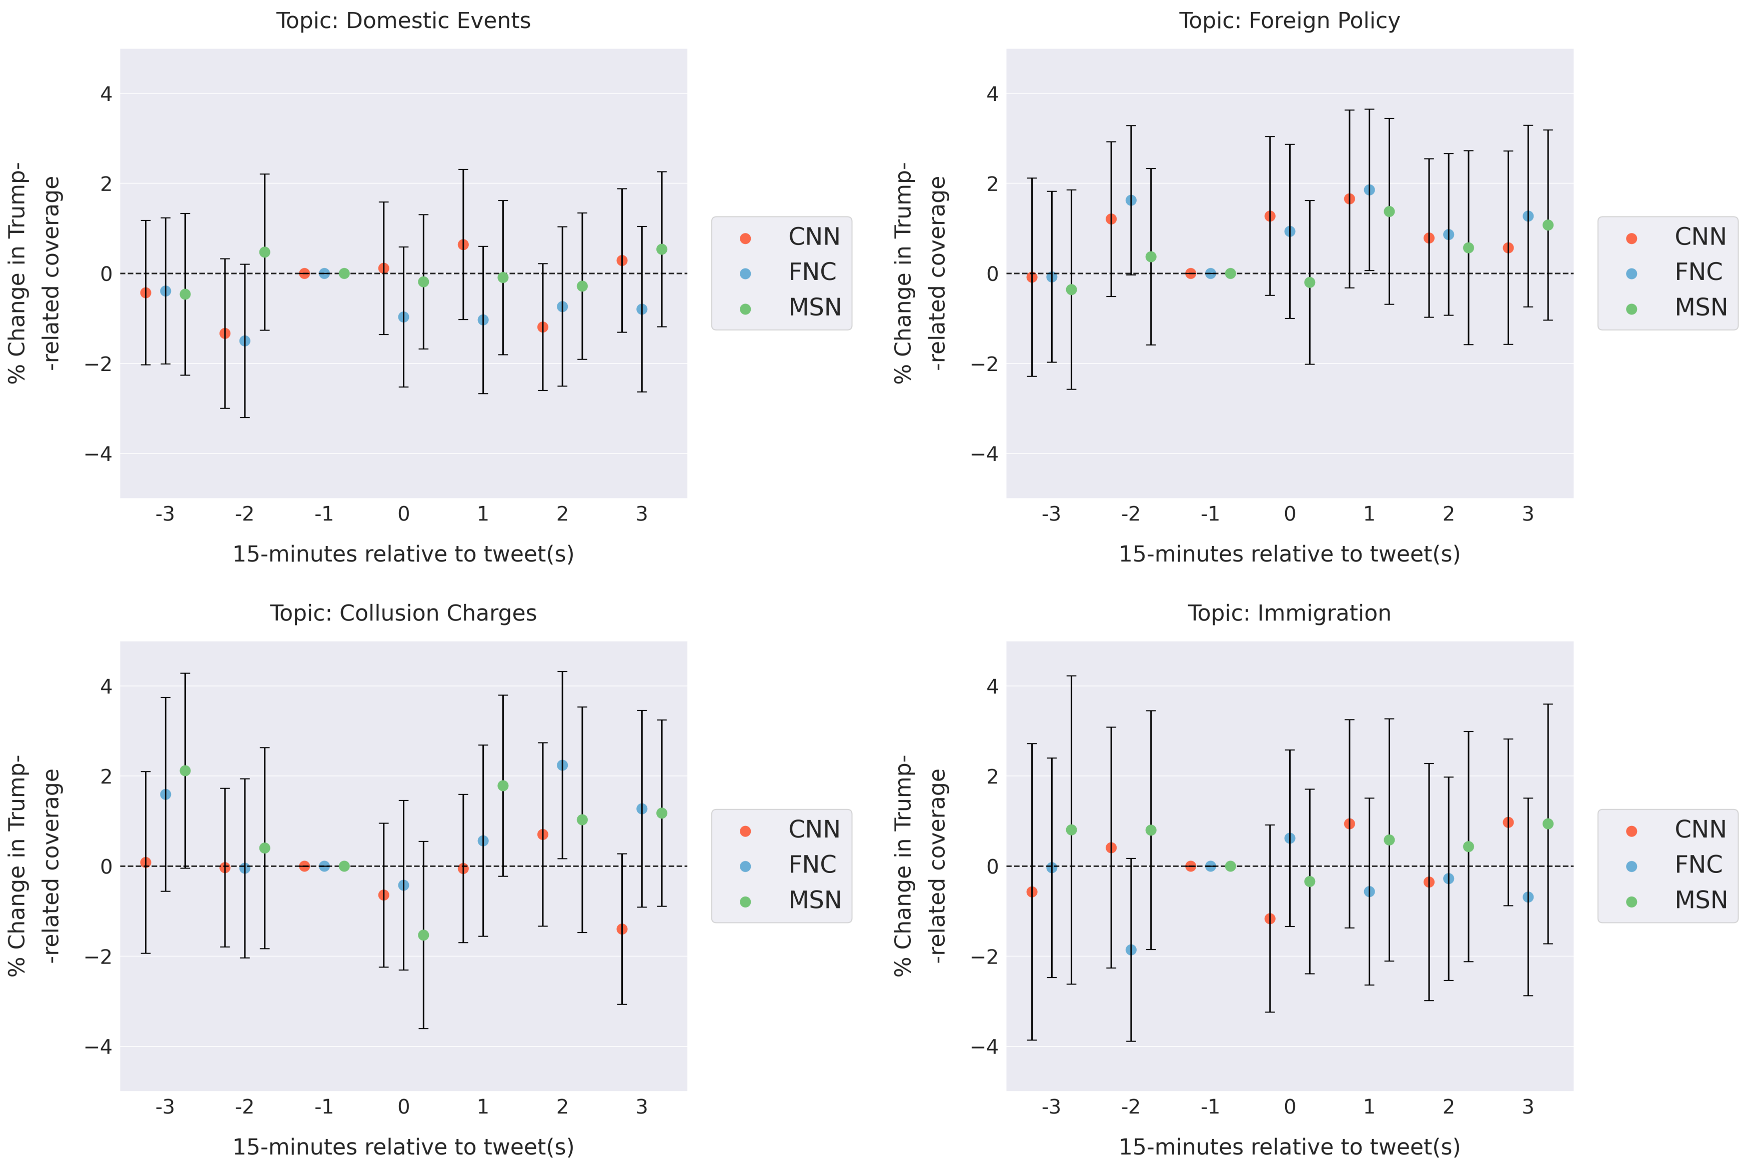

Figure 1.39: **Trump-related coverage before and after a tweet // Non-overlaps across time // Reaction per topic (1).** Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (1) Domestic Events, (2) Foreign Policy, (3) Collusion Charges and (4) Immigration. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

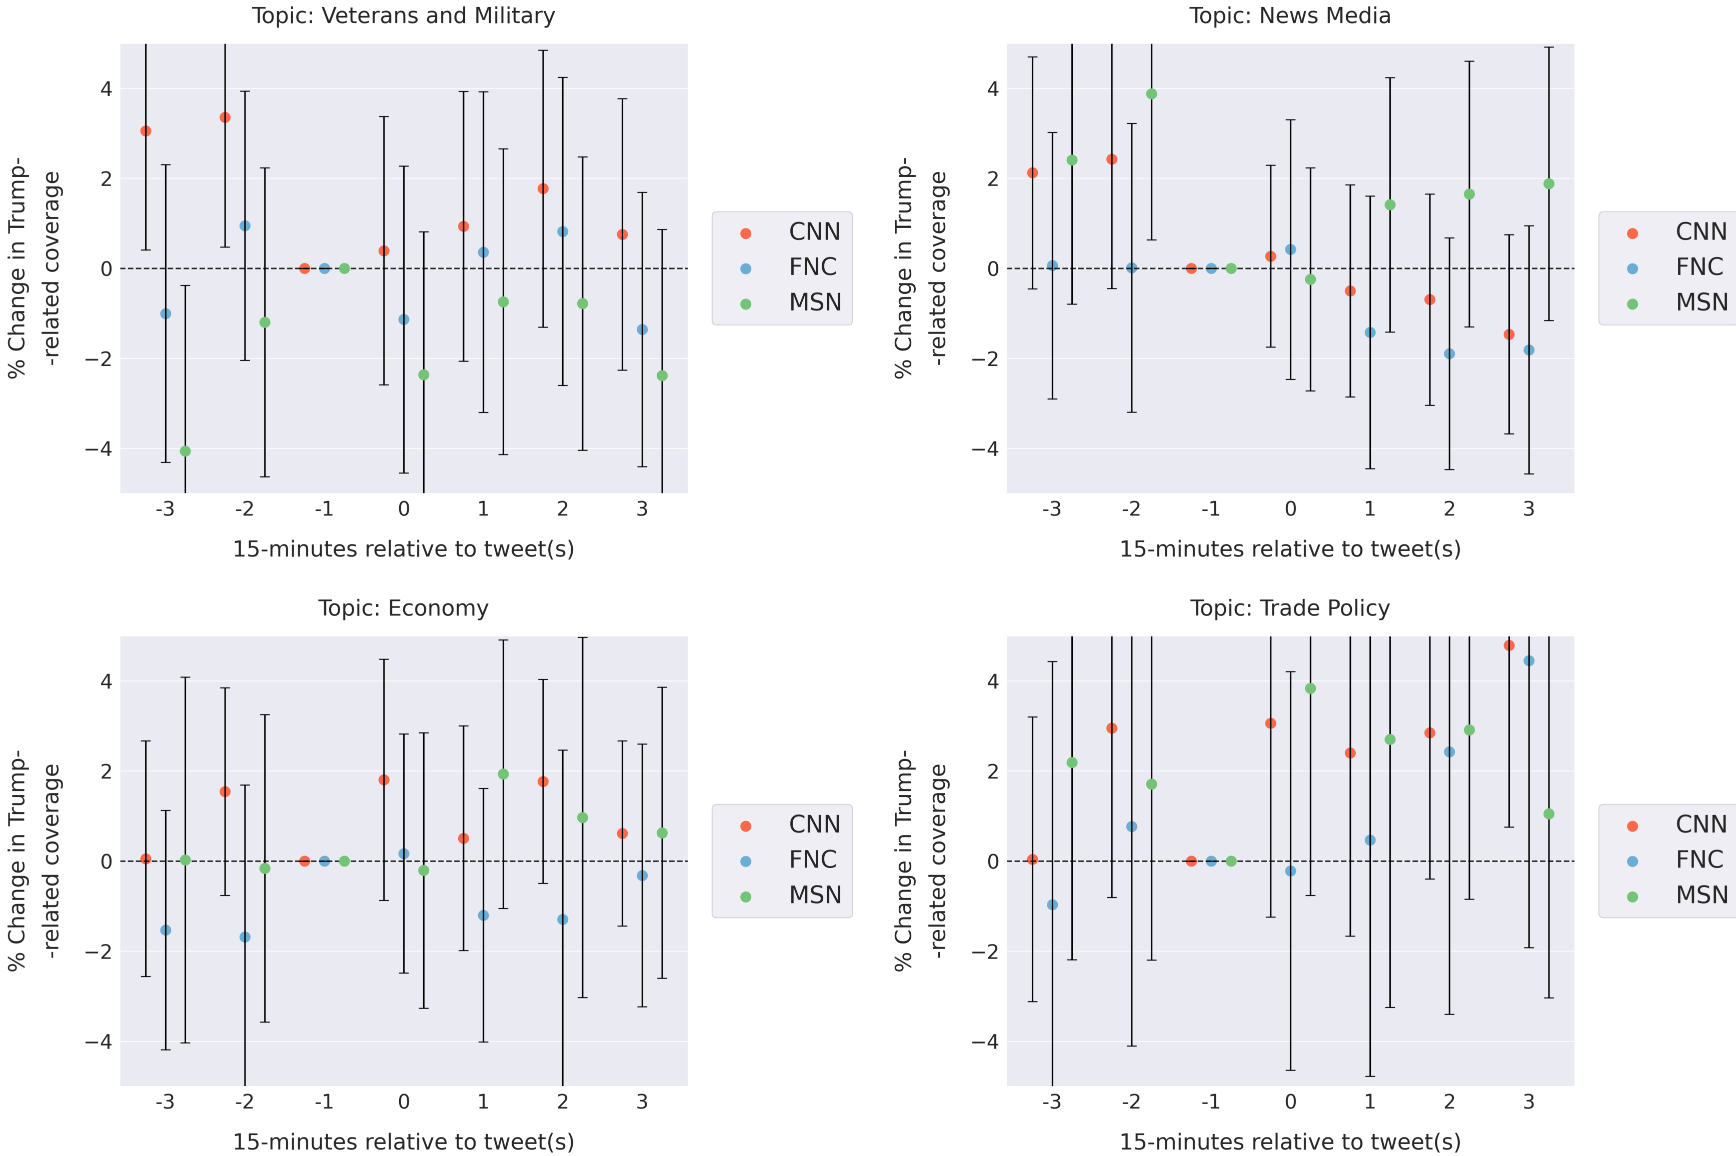

Figure 1.40: **Trump-related coverage before and after a tweet // Non-overlaps across time // Reaction per topic (2)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (5) Veterans and Military, (6) News Media, (7) Economy and (8) Trade Policy. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

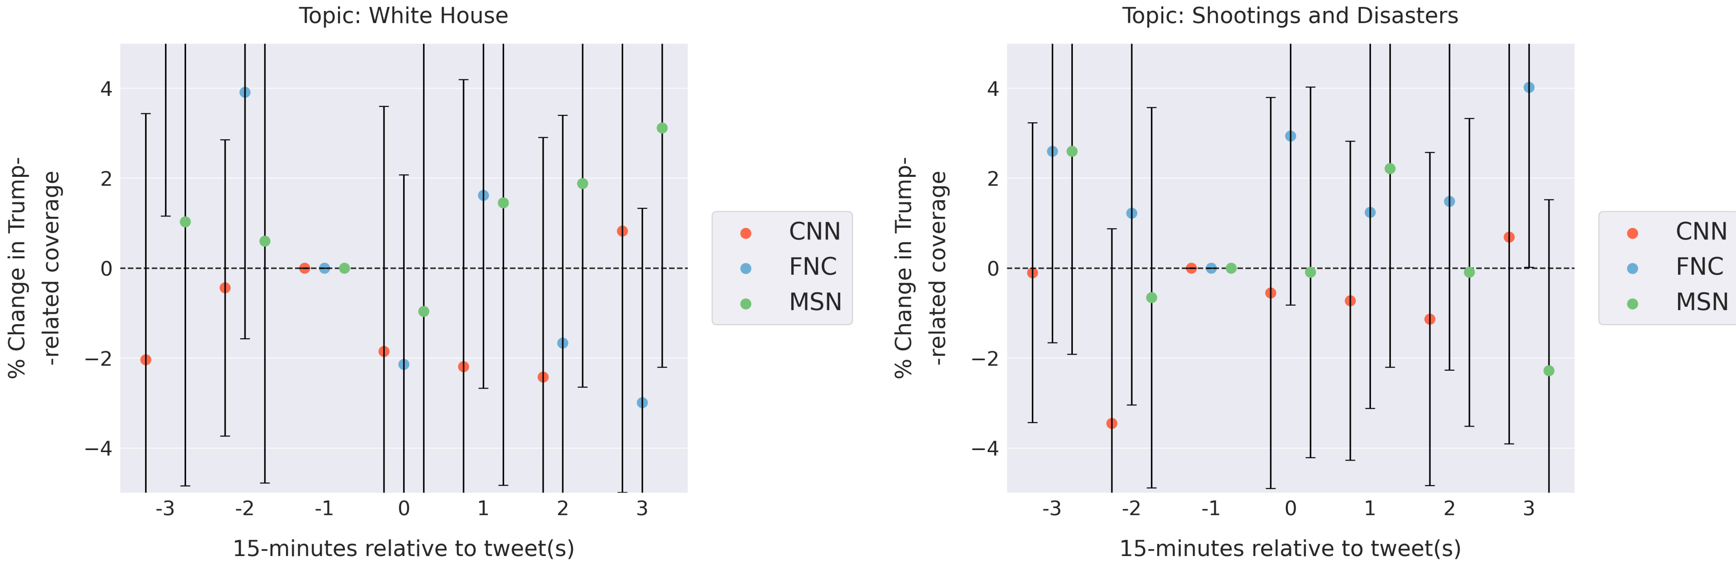

Figure 1.41: **Trump-related coverage before and after a tweet** // **Non-overlaps across time** // **Reaction per topic (3)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (9) White House, (10) Shootings and Disasters. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

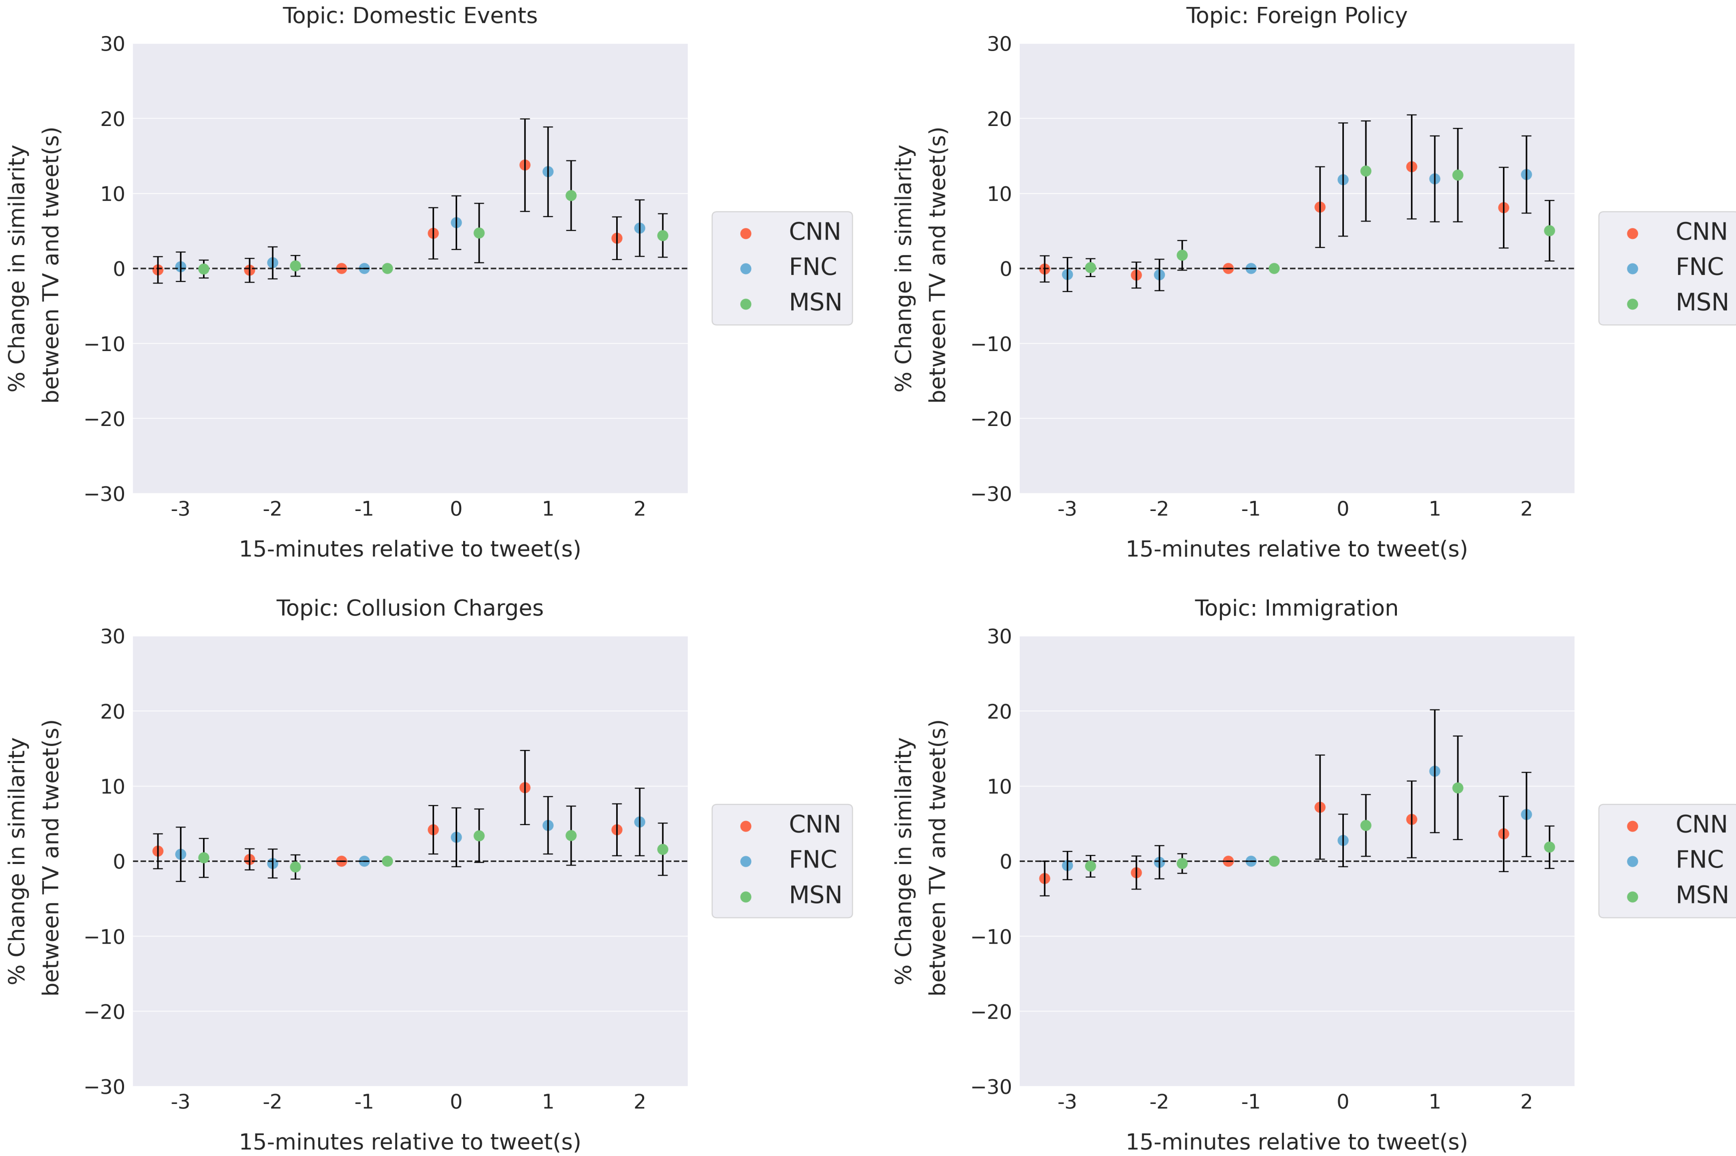

Figure 1.42: **Similarity before and after a tweet // Non-overlaps across time // Reaction per topic (1)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (1) Domestic Events, (2) Foreign Policy, (3) Collusion Charges and (4) Immigration. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

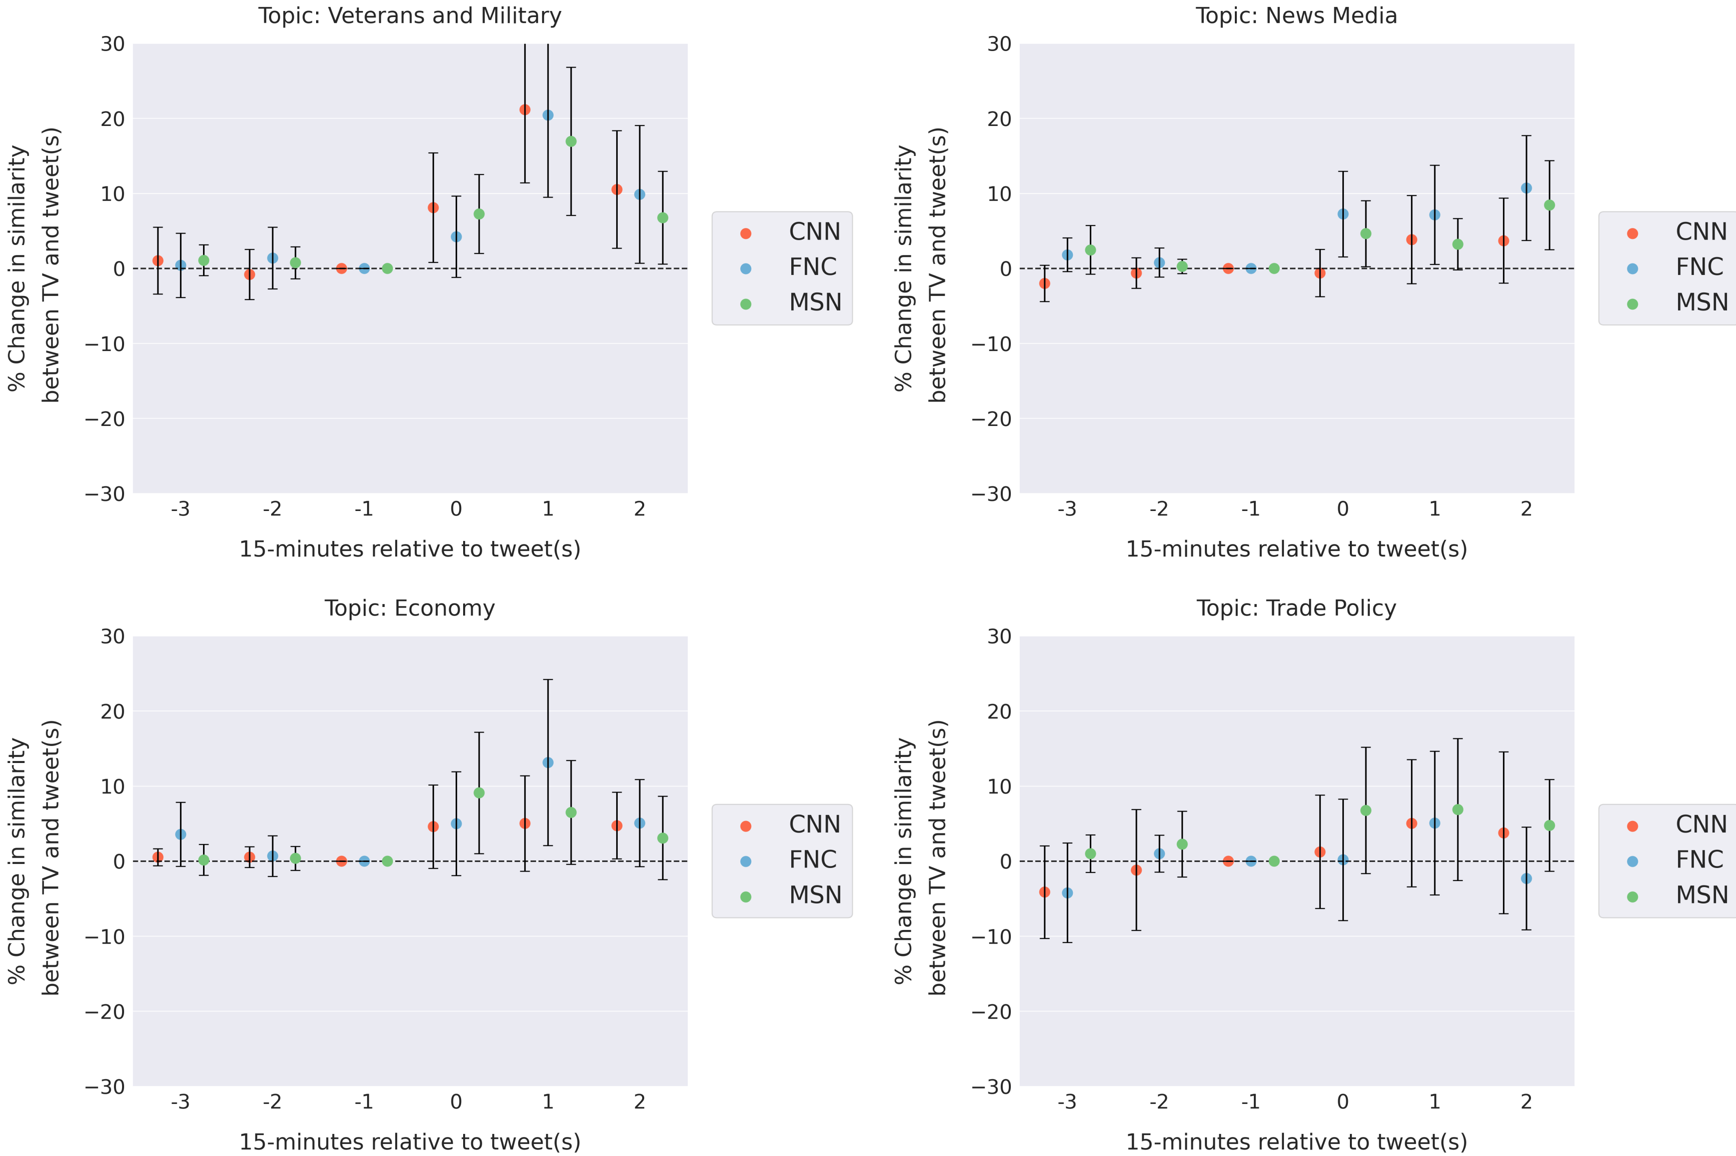

Figure 1.43: **Similarity before and after a tweet // Non-overlaps across time // Reaction per topic (2)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (5) Veterans and Military, (6) News Media, (7) Economy and (8) Trade Policy. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

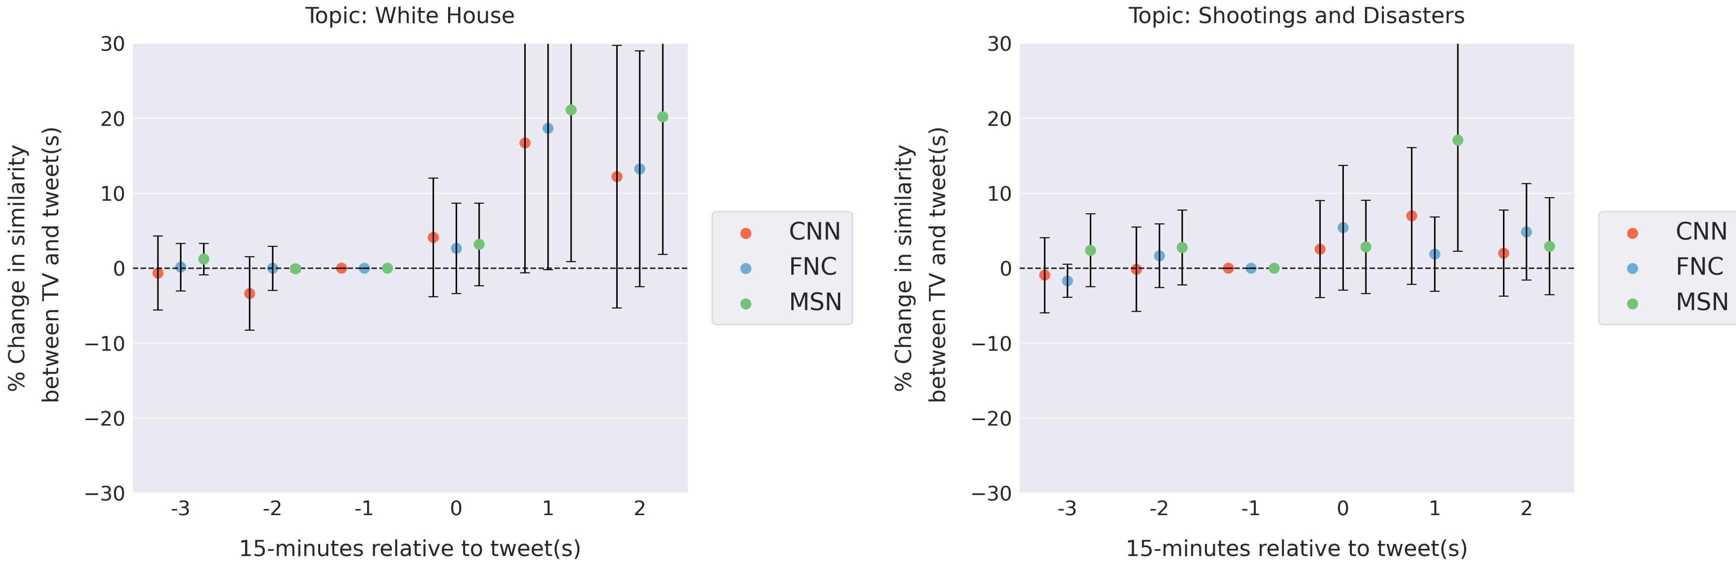

Figure 1.44: **Similarity before and after a tweet** // **Non-overlaps across time** // **Reaction per topic (3)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (9) White House, (10) Shootings and Disasters. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across time). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

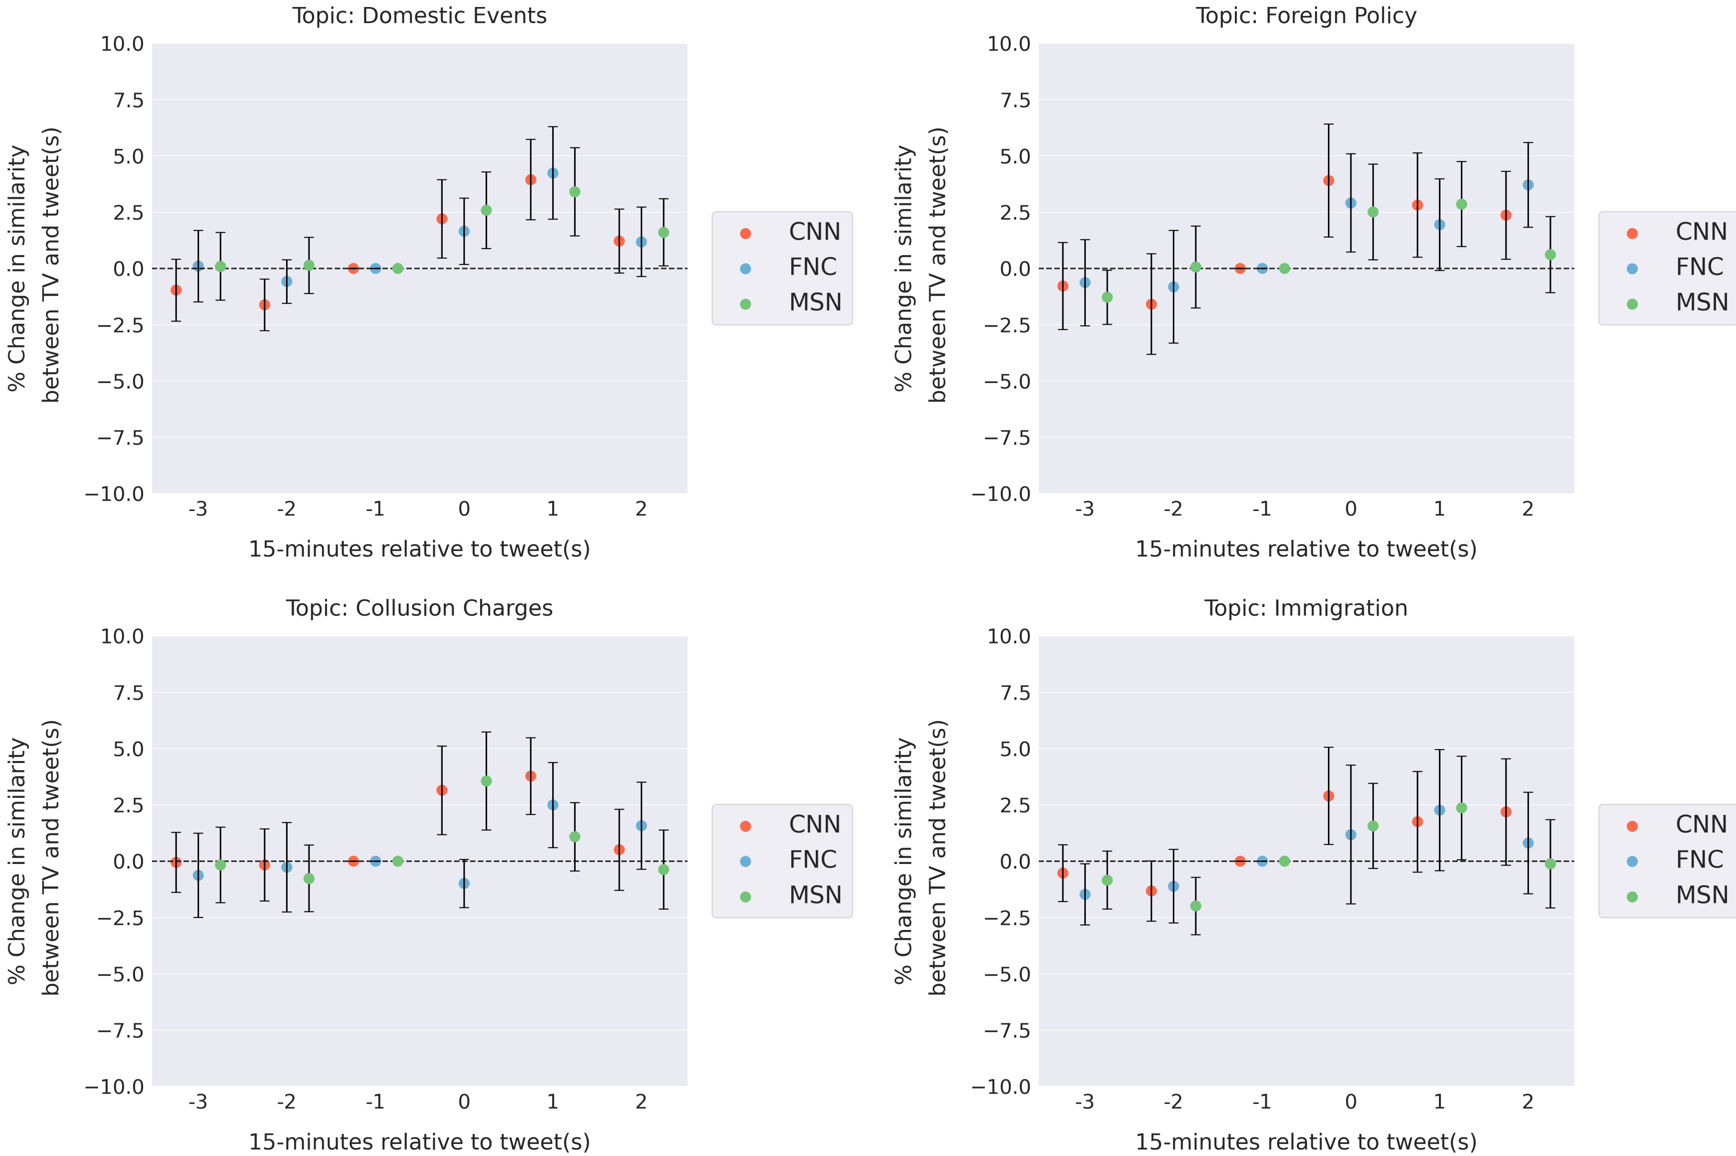

Figure 1.45: **Similarity before and after a tweet // Non-overlaps across content // Reaction per topic (1)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (1) Domestic Events, (2) Foreign Policy, (3) Collusion Charges and (4) Immigration. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across content). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

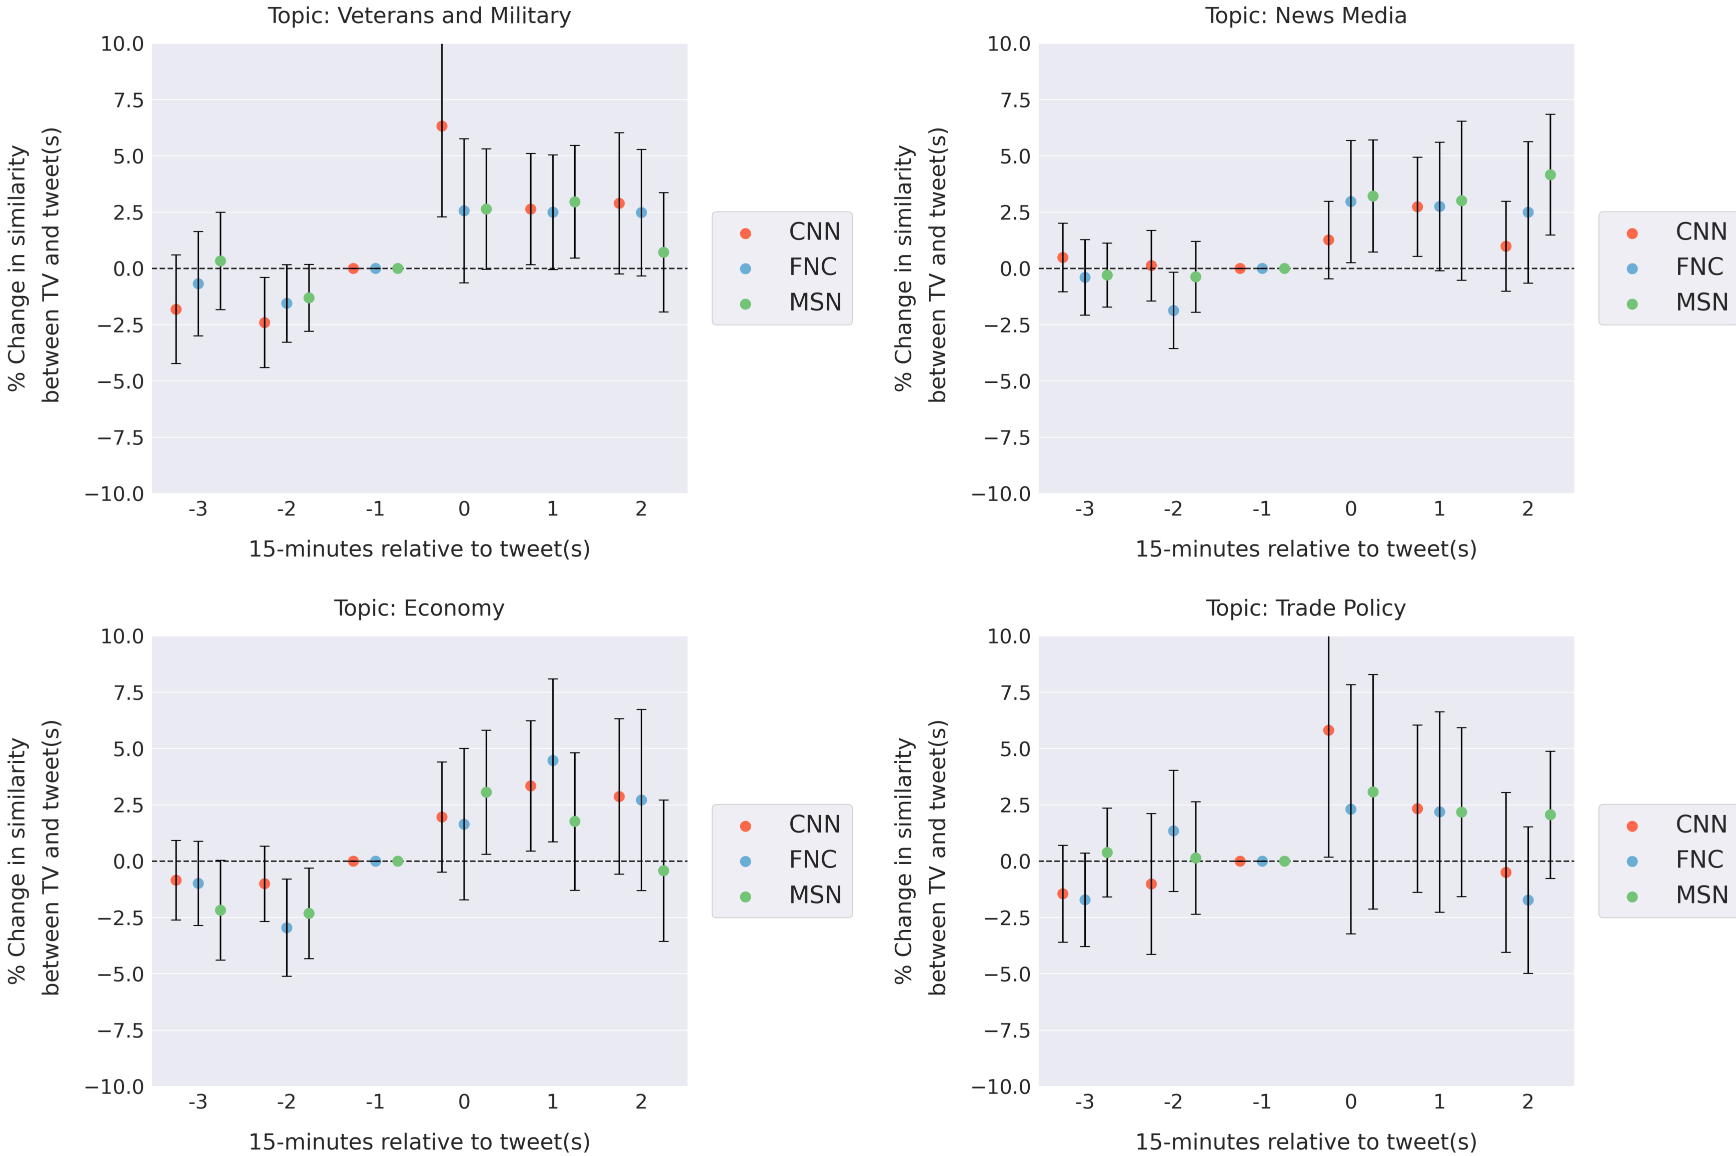

Figure 1.46: **Similarity before and after a tweet // Non-overlaps across content // Reaction per topic (2)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (5) Veterans and Military, (6) News Media, (7) Economy and (8) Trade Policy. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across content). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

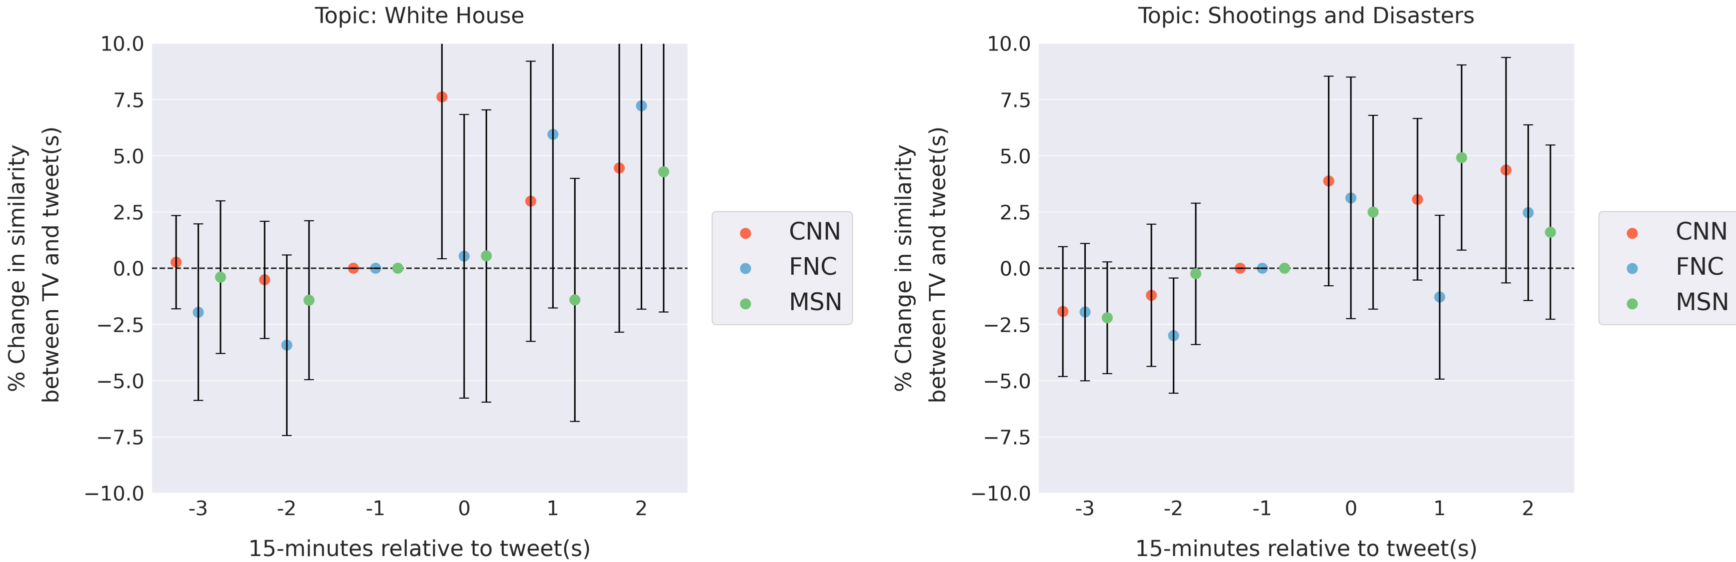

Figure 1.47: **Similarity before and after a tweet** // **Non-overlaps across content** // **Reaction per topic (3)**. Estimates refer to a selection of @realDonaldTrump tweets posted between 2017 and 2020, as in previous figures. In this figure I plot those estimates referent to 4 out of 10 topics consistently addressed by President Trump on Twitter - (9) White House, (10) Shootings and Disasters. Coefficients have been estimated by controlling for network $\times$ window FEs, computed using only those observations belonging to non-overlapping event windows (across content). Error bars stand for 95% conf. intervals computed with SEs clustered at a network-window level.

Chapter 2

Is Congress Watching TV News?

Evidence from Congressional Tweets

Abstract

Evidence in support of an agenda setting power from mass media on the agenda of politicians is based on correlations. Using novel intra-day content measures for television news and congress-members' statements, I provide a causal assessment of traditional mass media's agenda-setting power. I start by constructing an original dataset with precise timestamps and transcripts for each instance that a breaking news was covered on one-year of cable television news. Timestamps are retrieved from video data, images referring to a breaking news being identified through an image retrieval algorithm. Transcripts are retrieved from cable networks' closed captions. Then, I construct a dataset covering the universe of tweets posted by U.S. congress-members' in that same year. Lastly, I study if and how congress-members' react on Twitter to television breaking news in real-time, using narrow time windows.

2.1. Introduction

Are politicians reacting to an agenda set by traditional mass media? Existing evidence in support of this hypothesis is based on correlations and focuses only on social media (Barberá et al., 2019).

In this paper, using intra-day content measures for U.S. legislators' reactions and cable news television, I intend to causally identify an agenda-setting relationship between a traditional mass medium, television, and politicians' public agenda. In doing so, I will contribute to strands of literature that measure the effects of television coverage on policy and political outcomes, such as: foreign aid provision (Eisensee and Strömberg, 2007), political accountability (Snyder Jr and Strömberg, 2010), political polarization (Campante and Hojman, 2013), among others.

In addition, if politicians are reacting to an agenda set by cable news television, are they reacting to alternative outlets depending on their political affiliation?

I will leverage on having content measures for media outlets recurrently classified as liberal (CNN and MSNBC) and conservative (FOX NEWS) to study this hypothetical relationship (Groseclose and Milyo, 2005; Puglisi and Snyder Jr., 2015; Martin and Yurukoglu, 2017). Here I will be contributing to literature relating partisan media with a variety of outcomes: voting behavior (DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2017), political discourse (Ash and Labzina, 2019), judicial outcomes (Ash and Poyker, 2019) and, more recently, health behavior (Allcott et al., 2020; Bursztyn et al., 2020).

I construct a novel data of intra-day content measures for cable news television - timestamps and transcripts for breaking news released by CNN, FOX NEWS and MSNBC. I take advantage of having access to one year of data related to the universe of content broadcasted by CNN, FOX NEWS and MSNBC - video and timestamped transcripts, provided by TV News Archive. Breaking news are identified from video data through computer vision techniques. Transcripts are retrieved from closed captions - subtitles made available by networks to individuals with hearing disabilities.

Then, I retrieve each elected congress-member's tweets from Tweets of Congress. These are used as real-time reactions from congress-members to ongoing newsworthy events, constituting at the same time a standardized representation of their expressed issue agenda in other communication channels such as Facebook and press releases (Casas and Morar, 2015).

To understand if congress-members react in real-time to breaking news put forward by cable news outlets, I study their tweets in narrow time windows centered in each released breaking news. To control both for macro factors and for time-invariant member-specific characteristics I use within-congress-member variation in the sorts of topics addressed on Twitter: I classify breaking news in topics through an unsupervised classification method; then, I retrieve an array of keywords for each topic; I classify congress-members' tweets in news topics through the generated keywords; at last, I study if and how legislators react on Twitter to breaking news related to topics of their own interest (according to their posting history).

While using computer vision techniques to time instances where breaking news are broadcasted, I am contributing to a recent literature in political science and economics that introduces computer vision methods to retrieve data from images (see Joo and Steinert-Threlkeld, 2018 for a survey) and video (Dietrich, 2020). In addition, by classifying both tweets and breaking news in common topics, I contribute to a strand of economics literature that uses text as data (Gentzkow and Shapiro, 2010; Gentzkow et al., 2019).

The remainder of the paper is organized as follows. Section 2.2 describes the data sources; Section 2.3 describes the empirical setting, defines treatment and outcome variables and presents the empirical specification to be estimated; Section 2.4 outlines future steps.

2.2. Data

I am constructing a high-frequency panel dataset linking in real-time representatives' Twitter activity with cable news television content. To do so, I am leveraging on three different data sources:

Congressional timestamped tweets. Text and timestamp for the universe of tweets posted by U.S. Congress members present on Twitter - made available by Tweets of Congress. Information on tweets is from Twitter's API, encompassing not only text and timestamps but also if a tweet was sent from a computer, smartphone or social management app.

Additional information on the sources for congressional tweets data, together with preliminary descriptive statistics, is provided in Appendix A.

Cable news timestamped transcripts. Text and timestamps for the universe of dialogues broadcasted by U.S. cable news networks - courtesy of TV News Archive. Retrieved from closed captions, i.e., transcriptions of dialogue taken place in short windows of time (no more than 10 seconds), made available for individuals with hearing disabilities.

Cable news breaking news timestamps. Timestamps for breaking news broadcasted by U.S. cable news networks. Retrieved from cable news broadcast videos - courtesy of TV News Archive. Stories labelled as "*breaking*" or "*alert*" are identified within videos through computer vision techniques, as described in Appendix C and Appendix D.

Examples of closed captions and images are provided in Appendix B. These are provided for illustrative, and thus non-consumptively, purposes.

Data provided by *TV News Archive* covers the universe of content broadcasted by CNN, FOX News and MSNBC. I have been given access to a one-year sample - from July 2018 to June 2019 (included).

2.3. Empirical Strategy

I intend to study if breaking news being broadcasted on cable news television (*treatment*) have a causal impact on posts and content posted on Twitter (*outcome*) from elected U.S. congress-members (*units*).

In this section I formally describe my empirical setting (2.3.1); I define my treatment variable (2.3.2); I define my outcome variables (2.3.3) and, I outline my empirical specification (2.3.4).

2.3.1 Setting

Television

Take a cable station s; denote a piece of news as n.

Each station s will put forward C_s number of breaking stories throughout a period of time, thus being mapped to a vector of breaking news N_s,

$$N_s = (n_s^1, n_s^2, ..., n_s^{C_s - 1}, n_s^{C_s}). \tag{2.1}$$

Each breaking news n will be released at period e_n, for a particular duration of time d_n, with a specific choice of wording w_n and on a topic k_n. Define then a breaking news as a quadruplet of features,

$$n \equiv (e_n, d_n, w_n, k_n). \tag{2.2}$$

Each feature is retrieved from television data, through alternative methods - starting time t^n and duration d^n are collected from stations' video footages (described in Appendix C and D); wording w^n is recovered, using t^n and d^n, from stations' closed captions; topic k^n is recovered by fitting an unsupervised topic classification model on a corpus containing each identified "*breaking news*".

Twitter

Then, take a congress-member i; denote a tweet as r.

Each congress-member will post C_i number of tweets throughout a period of time, thus being mapped to a vector of tweets R_i,

$$R_i = (r_i^1, r_i^2, ..., r_i^{C_i-1}, r_i^{C_i}). \tag{2.3}$$

Each tweet r, irrespective of the congress-member, will be posted at period e_r, on a topic k_r and with a particular choice of wording w_r. Hence, define a tweet as a triplet of features,

$$r \equiv (e_r, k_r, w_r). \tag{2.4}$$

These features are available only for elected congress-members: starting time t^r and wording w^r are retrieved from Twitter's API; topic k^r is assessed by filtering tweets according to breaking news topics keywords (these keywords are retrieved by fitting an unsupervised topic classification model on "*breaking news*", as described previously).

Television and Twitter

Topics from breaking news and tweets belong to a common set of topics $\mathbb{K}$,

$$k_n \text{ and } k_r \in \mathbb{K}, \forall\, n \in \mathbb{N} \text{ and } \forall\, r \in \mathbb{R}, \tag{2.5}$$

where $\mathbb{N}$ stands for all breaking news broadcasted on cable news television; where $\mathbb{R}$ stands for all tweets posted by congress-members on Twitter.

2.3.2 Treatment

Treatment can be defined differently, according to which assumptions are made regarding both breaking news and tweets topics. If one assumes:

Homogeneous breaking news and tweets. Homogeneous breaking news imply that treatment should be understood as any piece of breaking news being put forward by a television station. Homogeneous tweets imply that, if we abstract from congress-member-specific characteristics, treatment is assumed to have an identical intensity across members.

In formal terms, treatment is defined via a treatment indicator b_t^j equal to 1 in period t if and only if any breaking news is broadcasted in period $t+j$,

$$b_t^j \equiv \mathbb{1}[t+j = e_n, \forall n \in \mathbb{N}], \tag{2.6}$$

Heterogeneous breaking news and homogeneous tweets. Heterogeneous breaking news imply that treatment differs depending on a news topic. At the same time, if we abstract again from member-specific characteristics, homogeneous tweets imply an identical intensity of each treatment across representatives.

Treatment is to be defined as a treatment indicator b_t^{jk} equal to 1 in period t if and only if a breaking news of topic k is broadcasted in period $t+j$,

$$b_t^{jk} \equiv \mathbb{1}[t+j = e_n \text{ and } k = k_n, \forall n \in \mathbb{N}]. \tag{2.7}$$

Heterogeneous breaking news and tweets. Treatment differs in terms of news topics. Additionally, if congress-members cover through their tweets different distributions of topics, treatment is expected to have a varying intensity across congress-members - members are expected to react more promptly to topics of their own taste.

Treatment is then formally defined as a treatment indicator b_{it}^{jk} equal to 1 for congress-member i in period t if and only if a breaking news of a topic k, a topic recurrently tweeted by congress-member i, is broadcasted in period $t+j$,

$$b_{it}^{jk} \equiv \mathbb{1}[t+j = e_n, \;\; k = k_n, \;\; \text{and} \;\; k \in Q_4^i, \forall n \in \mathbb{N}], \tag{2.8}$$

where Q_4^i stands for a fourth quartile of a distribution of topics tweeted by congress-member i, topics being ordered according to tweets frequencies.

In this paper I take breaking news as heterogeneous, between and within networks (to be assessed as soon as breaking news are identified). At the same time, I assume tweets are heterogenous between congress-members (in line with findings from Hemphill et al. 2019).

Hence, from here onwards, treatment is defined as in Equation 2.8.

2.3.3 Outcome

To answer the question - "*are politicians reacting to an agenda set by traditional mass media?*" - I assess how similar are congress-members' tweets to recent breaking news.

In order to do so, I define outcome y_{it} as the textual similarity between tweets of congress-member i posted in period t with the breaking news closest in time to period t.

In formal terms,

$$y_{it} \equiv sim(w_r, w_n), \; \forall r \in R_{it} \;\; \text{and} \;\; n := argmin_n \{|e_n - t|\}, \tag{2.9}$$

where $sim(w_r, w_n)$ is a textual similarity metric (e.g., Jaccard's similarity metric, defined as the intersection over the union of w_r and w_n).

2.3.4 Design

General

To study if breaking news have a causal impact on congress-members' Twitter agenda, I will implement an event-study design (for a methodological review, see Schmidheiny and Siegloch, 2019).

In particular, I will estimate a standard event-study specification:

$$y_{it} = \sum_{k \in \mathbb{K}} \left(\sum_{\underline{j}}^{\bar{j}} \beta_j^k \cdot b_{it}^{kj} \right) + \mu_i + \theta_t + \varepsilon_{it}, \tag{2.10}$$

where i stands for congress-member i; t stands for time period t; y_{it} is an outcome variable (as defined in Equation 2.9); k stands for a news/tweet topic; $\mathbb{K}$ stands for the set of news/tweets topics; $\underline{j}$ and $\bar{j}$ are bins for each breaking news window; b_{it}^{kj} stands for a treatment indicator (as defined in Equation 2.8); μ_i and θ_t are congress-member and time fixed effects (FEs), respectively.

With regards to identification, member FEs, time FEs and, consequently, treatment effects, will be identified by exploiting *topic* and *time* variation in (a) breaking news and (b) congressional tweets. As made explicit in Equation 2.10 above, both types of variation will be circumscribed to event windows centered around time periods when a breaking news is released. I illustrate an event window in Figure 2.1.

Figure 2.1: ***Illustration of an Event Window.***

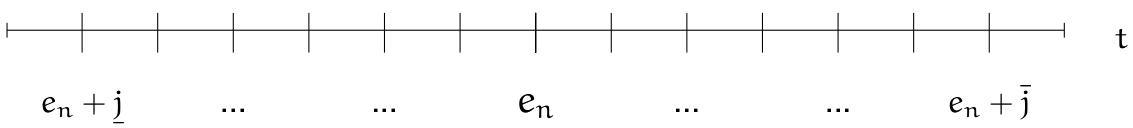

`Note:` e_n stands for the time period in which a breaking news n has been put forward by a cable news television station; $e_n + \underline{j}$ and $e_n + \bar{j}$ stand for the first and last observation of the event window, respectively; t stands for time.

To interpret β_j^k as a real-time causal effect of breaking news on congress-members' Twitter agenda, I have to argue for exogeneity - i.e., I must argue that any identified treatment effect is only caused by event windows' breaking news. To do that, I will study members' Twitter activity in narrow event windows ($\pm 15 min$). In particular, time periods t will correspond to blocks of between 1 and 2.5 minutes.

In addition, I cast aside any possibility that event windows partially overlap in time - if I do not, I would not be able to distinguish between pre and post breaking news periods (a crucial distinction, necessary to test for pre-trends and simultaneously assess if causal effects are, in some way, dynamic).

Figure 2.2: ***Illustration of Partially Overlapping Windows.***

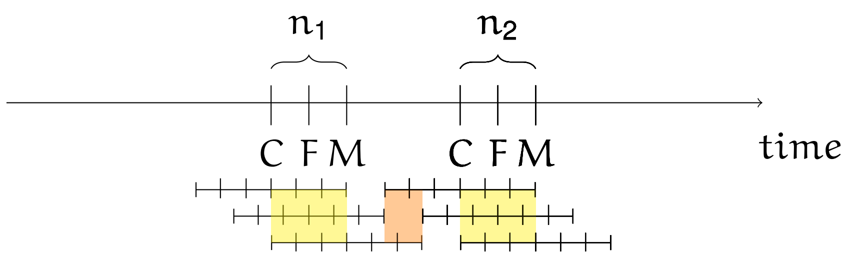

Note: C, F and M stand for CNN, FNC and MSN; n_1 and n_2 stand for different breaking news; partial overlaps of event-windows are signaled in yellow and orange; t stands for time.

As illustrated in Figure 2.2, partial overlaps are likely to happen in this setting: breaking news are expected to be released by different networks in a staggered but close in time fashion. To rule these out, I will: (1) identify clusters of breaking news, in time and content; (2) define treatment with first breaking news within each cluster.

Last but not least, β_j^k coefficients will refer to treatments that will *not only* affect a subset of congress-members *but also* affect the same congress-member repeatedly across time. Hence, β_j^k should be interpreted as an average treatment effect on treated (ATT) across treated units *and* within treated units (across time). With that in mind, a causal interpretation of β_j^k is possible after testing for two alternative hypothesis:

(1) treatment effects are heterogeneous across units; tested by estimating ATTs conditional on dimensionalities that ex-ante would be expected to determine different reactions to breaking news;

(2) treatment effects are unstable within units, across time; tested by estimating average treatment on treated conditional on different time periods.

Partisan

To understand if congress-members from different political affiliations react differently on Twitter to breaking news, I will estimate different event-study specifications.

(1) To study if congress-members are more or less reactive to particular news topics, depending on their political color, I will estimate two specifications identical to that in Equation 2.10.

In essence, I will estimate Equation 2.10 with the following different selected samples:

1.1. treated units composed only of *Republican* legislators; control units identical to those used to estimate Equation 2.10 in Subsection 2.3.4;

1.2. treated units composed only of *Democrat* legislators; control units identical to those used to estimate Equation 2.10 in Subsection 2.3.4.

I estimate specifications identical to Equation 2.10 with selected treated samples so to have treatment effects that are comparable across studies.

(2) To study if congress-members, conditional on their political family, are reacting more or less promptly to breaking news depending on which station broadcasted that news, I will estimate two different specifications.

In line with classifications from past literature:

2.1. To test if Republicans are abnormally reactive to FOX NEWS:

$$y_{it} = \sum_{k \in \mathbb{K}} \left(\sum_{\underline{j}}^{\bar{j}} \beta_j^k \cdot b_{it}^{kj} + \sum_{\underline{j}}^{\bar{j}} \beta_j^{kf} \cdot b_{it}^{kj} \cdot F_t \right) + \mu_i + \theta_t + \varepsilon_{it}, \tag{2.11}$$

where F_t stands for an indicator variable equal to 1 if and only if breaking news closest in time to period t has been released by FOX.

In terms of interpretation, Equation 2.11 allows me to distinguish between an average treatment effect related to breaking news released by cable news television stations in general and a FOX-specific treatment effect encompassed in β_j^{kf}.

2.2. To test if Democrats are significantly sensitive to CNN/MSNBC:

$$y_{it} = \sum_{k \in \mathbb{K}} \left(\sum_{\underline{j}}^{\bar{j}} \beta_j^k \cdot b_{it}^{kj} + \sum_{\underline{j}}^{\bar{j}} \beta_j^{knf} \cdot b_{it}^{kj} \cdot (1 - F_t) \right) + \mu_i + \theta_t + \varepsilon_{it}, \tag{2.12}$$

where F_t stands for an indicator variable equal to 1 if and only if breaking news closest in time to period t has been released by FOX.

As above, Equation 2.12 allows me to distinguish between an average treatment effect related to breaking news released by generic cable news television stations and a CNN/MSNBC-specific treatment effect encompassed in β_j^{knf}.

To estimate treatment effects comparable across analyses, I will estimate Equations 2.11 and 2.12 using selected panels. In particular, I will estimate Equations 2.11 and 2.12 with treated units being composed exclusively of Republican and Democrat representatives, respectively. In both specifications I will also use control units identical to those used in Equation 2.10.

2.4. Conclusion

I have collected and pre-processed data referent to 6 months of video footage (17TB of video footage downloaded and later pre-processed into 2.3TB of one-second frames).

I have written different image retrieval algorithms, capable of distinguishing between news and breaking news frames (early validation exercises point for accuracies above 95%, when retrieving news frames).

In future iterations of this chapter, outside of the scope of this thesis, I intend to arrive at a first set of results, using data referent only to 2019:S1 only. To arrive at these, I will:

Retrieve breaking news timestamps and transcripts: retrieve exact timestamps for breaking news broadcasted by CNN, FNC and MSN throughout January and June, 2019; after timing instances where breaking news have been broadcasted, I will retrieve these periods' respective transcripts using networks' closed captions;

Classify breaking news and tweets into topics: with breaking news transcripts, classify news into topics by employing an unsupervised classification model (a LDA topic model) on the corpus of transcripts; after, filter tweets according to keywords for each news topics - to classify tweets in news topics;

Construct outcome and treatment variables and estimate event-study specification: as soon as I have access to news and tweets texts, topics and timestamps, construct outcome and treatment variables as defined in Section 2.3; then, construct a panel dataset with outcome and treatment variables for U.S. congress-members and estimate different empirical specifications as described in Equations 2.10, 2.11 and 2.12.

Validate image retrieval algorithms: perform an extensive validation exercise for both image retrieval algorithms. In particular: extract a comprehensive random sample of shows from data in storage; (2) skim manually through these to classify frames as news/ads and breaking/standard news; (3) map both classifications, human and computer-based, to infer on how accurate both image retrieval algorithms are.

Bibliography

Allcott, H., Boxell, L., Conway, J., Thaler, M., and Gentzkow, M. (2020). Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic. *Journal of Public Economics*, 191:104254. [URL].

Ash, E. and Labzina, E. (2019). Does politicized media distort political discourse? evidence from us cable news channels. *Center for Law & Economics Working Paper Series*, 5. [URL].

Ash, E. and Poyker, M. (2019). Conservative News Media and Criminal Justice: Evidence from Exposure to Fox News Channel. *Columbia Business School Research Paper*. [URL].

Barberá, P., Casas, A., Nagler, J., Egan, P., Bonneau, R., Jost, J., and Tucker, J. (2019). Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public using Social Media Data. *American Political Science Review*, 113(4):883–901. [URL].

Bursztyn, L., Rao, A., Roth, C., and Yanagizawa-Drott, D. (2020). Misinformation during a Pandemic. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-44). [URL].

Campante, F. R. and Hojman, D. A. (2013). Media and Polarization: Evidence from the Introduction of Broadcast TV in the United States. *Journal of Public Economics*, 100:79–92. [URL].

Casas, A. and Morar, D. (2015). Different Channel, Same Strategy? Filling Empirical Gaps in Congress Literature. In *Proceedings of the Annual Meeting of the American Political Science Association*. Eds. Layna Mosley and Alvin Bernard Tillery. San Francisco: American Political Science Association. 1–21. [URL].

DellaVigna, S. and Kaplan, E. (2007). The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics*, 122(3):1187–1234. [URL].

Dietrich, B. J. (2020). Using Motion Detection to Measure Social Polarization in the U.S. House of Representatives. *Political Analysis*. Forthcoming. [URL].

Eisensee, T. and Strömberg, D. (2007). News Droughts, News Floods, and US Disaster Relief. *The Quarterly Journal of Economics*, 122(2):693–728. [URL].

Gentzkow, M. and Shapiro, J. M. (2010). What Drives Media Slant? Evidence from US Daily Newspapers. *Econometrica*, 78(1):35–71. [URL].

Gentzkow, M., Shapiro, J. M., and Taddy, M. (2019). Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4):1307–1340. [URL].

Groseclose, T. and Milyo, J. (2005). A Measure of Media Bias. *The Quarterly Journal of Economics*, 120(4):1191–1237. [URL].

Hemphill, L., Russell, A., and Schopke, A. (2019). The Rhetorical Agenda: What Twitter Tells Us About Congressional Attention. In *Proceedings of the Midwest Political Science Association Annual Meeting*. Chicago, IL, USA. [URL].

Joo, J. and Steinert-Threlkeld, Z. C. (2018). Image as Data: Automated Visual Content Analysis for Political Science. *arXiv preprint arXiv:1810.01544*. [URL].

Lewis, J. B., Poole, K., Rosenthal, H., Boche, A., Rudkin, A., and Sonnet, L. (2020). Voteview: Congressional Roll-call Votes Database. *https://voteview. com*. [URL].

Martin, G. J. and Yurukoglu, A. (2017). Bias in Cable News: Persuasion and Polarization. *American Economic Review*, 107(9):2565–99. [URL].

Poole, K. T. and Rosenthal, H. (1985). A Spatial Model for Legislative Roll Call Analysis. *American Journal of Political Science*, pages 357–384. [URL].

Puglisi, R. and Snyder Jr., J. M. (2015). Empirical Studies of Media Bias. In *Handbook of Media Economics*, volume 1, pages 647–667. [URL].

Schmidheiny, K. and Siegloch, S. (2019). On Event Study Designs and Distributed-lag Models: Equivalence, Generalization and Practical Implications. *CEPR Discussion Paper Series*. No. DP13477 [URL].

Snyder Jr, J. M. and Strömberg, D. (2010). Press Coverage and Political Accountability. *Journal of Political Economy*, 118(2):355–408. [URL].

Appendix.

Appendix A. Congressional Tweets - Sources and Description

Appendix B. Cable News Television - Sources and Description

Appendix C. Image Retrieval Algorithm - News

Appendix D. Image Retrieval Algorithm - Breaking News

Appendix A. Congressional Tweets

(back to Appendix)

`Sources:` Universe of tweets posted by U.S. congress-members made available by Tweets of Congress. Metadata for each member (e.g., party) are provided by @UnitedStates Project.

Table A.1: ***Descriptive Statistics for Cong. Tweets, Total and Across Parties***. (*sample period* - July 2018 to June 2019, included).

	Total	Party Democrats	Republicans
# of members	632	305	324
# of tweets	674,109	411,058	256,032
# of tweets per date	2.9	3.7	2.2
# of chars. per tweet	211	220	196

`Note:` *# of members* stands for number of congress-members; *# of tweets* stands for number of tweets; *# of tweets per date* stands for number of tweets per date, i.e., average number of tweets a generic congress-member posts by date; *# of chars. per tweet* stands for number of characters per tweet, i.e., average number of characters of a generic tweet; URL links are not counted.

Figure A.1: ***Number of Cong. Tweets, Within Day and Across Parties***. (*sample period* - July 2018 to June 2019, included).

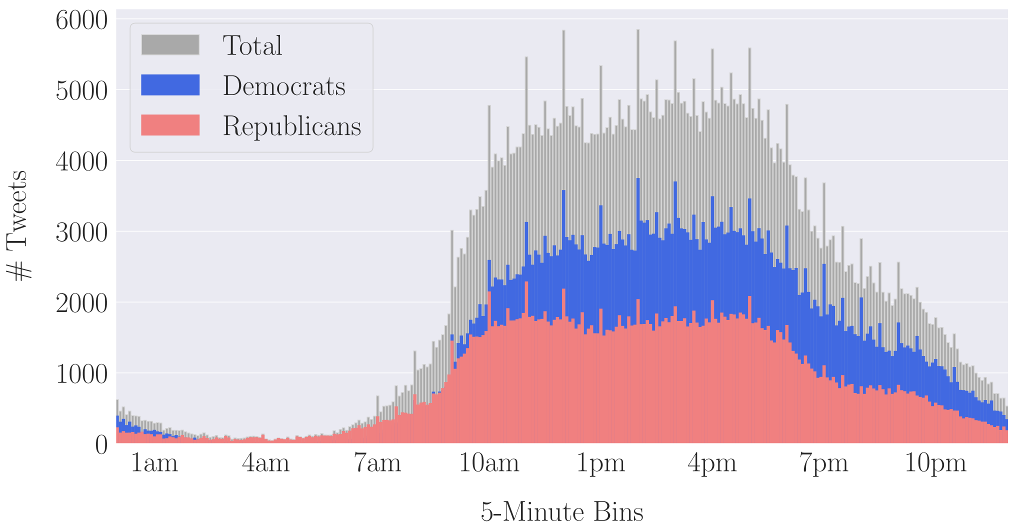

`Note:` *# Tweets* stands for number of tweets; *5-Minute Bins* stands for quarter-thirds of an hour; total number of tweets posted in each 5-minute intervals of a generic in-sample date are plotted; *Democrats* points for tweets by Democratic congress-members; *Republicans* points for tweets by Republican congress-members; *Total* stands for tweets posted by Democrats and Republicans.

Table A.1: ***Top 20 Partisan Phrases, Bigrams and Trigrams***
(*sample period* - July 2018 to June 2019, included).

2-word democratic phrases	2-word republican phrases	3-word democratic phrases	3-word republican phrases
trump claim	watch gt	children still separ	women serv countri
black brown	dont like	woman elect congress	across finish line
fight health	foreign aid	appoint suprem court	honor men women
famili senior	religi liberti	increas minimum wage	today introduc legisl
accur count	go toward	roll back protect	work around clock
need invest	sanctiti life	common sens gun	tax cut perman
rich power	trump deliv	power back peopl	speaker nanci pelosi
fight women	current law	coequal branch govern	secur crisi southern
show trump	govern spend	zero toler polici	border patrol agent
sens gun	state emerg	care right privileg	sovereignti golan height
ensur equal	across border	say presid trump	introduc bill would
barr must	pass first	smart border secur	secur border amp
labor movement	ronald reagan	health care amp	lowest unemploy rate
latino commun	big govern	ensur everi american	spoke senat floor
johnson sign	first quarter	gun violenc epidem	border secur fund
presid hous	evid collus	massiv tax cut	unemploy rate fell
still separ	great job	american peopl amp	uphold rule law
social justic	beat expect	trump administr want	speech hous floor
women die	feder spend	cut medicar medicaid	protect american peopl
propos cut	protect life	civil right act	need secur border

Note: Table showcases phrases used on Twitter that better predict a representative's ideological score. In steps: (1) tweets are pre-processed- i.e., text is lower-cased; retweets, URLs, hashtags, emojis and stopwords are eliminated; words are tokenized and stemmed to their morphological roots; (2) universe of 2 and 3 word phrases are identified and "common" phrases are selected - i.e., phrases that are present in at least 100 different tweets only; (3) an elastic net regression is estimated, to predict each congress-member ideology with hers or his respective (relative) use of "common" 2-word/3-word phrases; (4) coefficients are ranked in terms of their magnitude, top 20 coefficients are shown above. Democratic phrases stand for positive coefficients with largest magnitude; Republican phrases are negative coefficients with largest absolute magnitude. DW-NOMINATE used as congress-members' ideological score, from Lewis et al. (2020) and introduced by Poole and Rosenthal (1985); procedure to identify partisan phrases drawn from Martin and Yurukoglu (2017) and Gentzkow and Shapiro (2010).

Appendix B. Cable News Television

(back to Appendix)

Sources: Video, audio and closed captions for CNN, FOX NEWS and MSNBC are provided by TV News Archive for a one-year sample, from July 2018 to June 2019 (included).

Figure A.2: ***Examples of Video Frames for CNN, FOX NEWS and MSNBC***. (*sample period* - 1-week pilot dataset, November 18-24, 2018).

(a) CNN

(b) FOX NEWS

(c) MSNBC

Figure A.3: ***Examples of Closed Captions for CNN, FOX NEWS and MSNBC***. (*sample period* - 1-week pilot dataset, November 18-24, 2018).

```
[000:28:28;573] >>> PRESIDENT TRUMP AND FIRST
[000:28:29;207] LADY MELANIA WILL CAP OFF THEIR
[000:28:31;276] STATE VISIT IN JAPAN BY VISITING
[000:28:34;045] TROOPS.
[000:28:34;412] BACK HERE AT HOME VICE PRESIDENT
[000:28:38;917] PENCE LAID A WREATH AT ARLINGTON
[000:28:41;986] NATIONAL CEMETERY.
[000:28:43;188] THAT'S WHERE WE FIND BARBARA
```

(a) CNN

```
[001:07:52;551] >> IF YOU LOOK AT THE DEAL
[001:07:56;522] THAT BIDEN AND PRESIDENT
[001:07:57;857] OBAMA SIGNED, THEY WOULD
[001:07:59;191] HAVE ACCESS, FREE ACCESS TO
[001:08:02;661] NUCLEAR WEAPONS WHERE THEY
[001:08:03;529] WOULDN'T EVEN BE IN
[001:08:04;597] VIOLATION IN JUST A VERY
[001:08:06;132] SHORT PERIOD OF TIME.
```

(b) FOX NEWS

```
[000:26:40;231] >>> IN WASHINGTON, ON THIS
[000:26:40;999] MEMORIAL DAY WEEKEND, THE ANNUAL
[000:26:42;334] ROLLING THUNDER GATHERING OF
[000:26:43;735] BIKERS, MANY OF THEM VETERANS,
[000:26:45;670] HONORING POWs AND MISSING
[000:26:47;372] SERVICE MEMBERS.
[000:26:47;939] BECAUSE OF RISING COSTS AND
[000:26:50;976] LOGISTICAL ISSUES, IT LOOKED
```

(c) MSNBC

Appendix C. Image Retrieval Algorithm - News

(back to Appendix)

In this subsection I describe each step that I take to implement an image retrieval algorithm capable of classifying cable TV news frames as news or ads-related frames. When describing each step, I point out which software tools I use for those interested in replicating these steps with different data. In addition, I provide different examples to better illustrate each step taken.

In theory. Instances where news are broadcasted on cable news television are identifiable in time if (1) each news piece is always broadcasted with a station-specific logo and if (2) each station exhibits that same logo in a constant set of frame coordinates. If both assumptions hold:

- **(a)** take a frame where a news piece has been broadcasted (call it a seed frame); crop the station's logo from that seed frame, which will be displayed in a fixed position; then,
- **(b)** take a random frame; crop, from that frame, the exact area where the station's logo is displayed in your seed frame; it follows that,
- **(c)** crops from frames where a piece of news has been covered are expected to display a high level of similarity (measured through a similarity metric of choice) with respect to its respective station's logo, cropped from your seed frame.

In practice. I start by collecting, selecting and processing video data for each TV station:

- **(1)** I download videos that are broadcasted in periods where congress-members are significantly active on Twitter. In particular, I focus on videos linked to shows not broadcasted through dawn (Eastern Time) - i.e., videos displayed between 1am and 7am ET are left out;
- **(2)** I split each video into 1-second frames (using `FFMPEG`, a C tool accessible through the command line and used to automatize video and audio editing tasks);
- **(3)** I skim through each set of frames, to identify videos where time is constantly displayed either through a bottom-left or a bottom-right digital clock. I focus on these videos to retrieve accurate timestamps for each frame. Videos that do not display any time are thus left out;

Then, I confirm that news content is broadcasted with a constant set of image identifiers:

(4) I extract a random sample of frames (using `random`, a Python module used to randomly select files from a list) and then skim manually through these frames to confirm that **(4.1)** cable news TV stations display their own logos while broadcasting news and that **(4.2)** these logos are shown in constant or close to constant sets of frame coordinates.

After confirming that news content is *always* followed by a station-specific logo displayed in a constant image position, I proceed to gather information about each station's logo. To do that, I use the random sample of frames extracted in step (4) as follows:

(5) I manually skim through these frames to retrieve for each station a set of seed frames, where news and consequently logos are being displayed in a clear manner; then,

(6) I examine each seed frame and retrieve for each logo a set of constant frame coordinates (using `opencv`, a C-based Python module used to solve computer vision problems).

Regarding logos, CNN and FNC have 2 different logos; MSN has 1 unique logo. As for frame coordinates, CNN and FNC alternate between 2 positions to display their respective logos; MSN shows its logo in 1 constant position.

In what follows I leave examples of (1) seed frames and their respective (2) station-specific logo:

Figure A.4: ***Examples of Seed Frames for CNN, FNC and MSN***. (*sample period* - 1-week pilot dataset, November 18-24, 2018).

Seed (CNN-news_wrb)

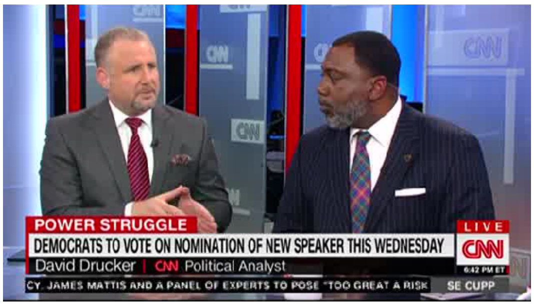

Cropped logo from seed (CNN-news_wrb)

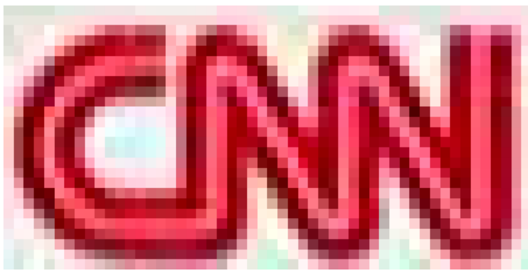

(a) CNN

Seed (FNC-news_bwa)

Cropped logo from seed (FNC-news_bwa)

(b) FNC

Seed (MSN-news)

Cropped logo from seed (MSN-news)

(c) MSN

After retrieving a set of seed frames for each station:

(7) I identify a similarity threshold - i.e., a metric of choice threshold below which an image is significantly similar to its respective station logo - for each cable news station (using `opencv`; a detailed step-by-step guide for this particular step is provided below).

To identify each similarity threshold I have to choose a metric that allows me to compare pairs of images and evaluate their similarity: for this I use the L1-distance as it is computationally inexpensive to compute, an important feature given that I will compare a large volume of frames.

In formal terms: take two images of identical dimension - image A and B of width w and height h; convert both images into grayscale images, i.e., images that assume only tones of gray; both images can be represented as matrices of identical dimension - columns w and rows h.

Each element represents a pixel in an image and assumes a value between 0 (black) and 255 (white); the L1-distance between image A and B is given by,

$$d_1(A, B) = \sum_{j=1}^{w} \left(\sum_{i=1}^{h} |A_{ij} - B_{ij}| \right),$$

where d_1 stands for L1-distance, i stands for the row of the matrix and j stands for the column of the matrix; d_1 is bounded from below at zero and from above at $w \cdot h \cdot 255$; higher values of d_1 point for larger differences between images, or , in other words, a lower degree of similarity.

Then, using L1-distance as a similarity metric, I identify similarity thresholds as follows:

(7.1) I resize both standard and seed frames into images of 640 by 360 pixels (using `opencv`);

(7.2) I pre-process frames to minimize high-frequency image noise (through `opencv`);

(7.3) I convert frames into matrices (through `opencv`), each matrix element representing a pixel;

(7.4) I crop from each frame their station's logo (with `numpy`, an array-processing Python library); [`remark`: note that I crop each frame using logo coordinates found in step **(6)**];

(7.5) I compute L1-distances between each frame and their station's seed (through `numpy`);
[`remark`: I compute L1-distances using crops from seed frames retrieved in step **(5)**];

(7.6) I skim manually through each video to identify frames where news were shown; then, I compare computed distances of frames with and without a logo to infer on station-specific similarity thresholds, upon which frames are displaying a station's logo.

After collecting for each station their respective (i) seed frames, (ii) logo frame coordinates and (iii) L1-distance similarity threshold, I proceed with classifying each frame as news or ads:

(8.1) I re-run steps (8.1) to (8.5) for each station's frames;

(8.2) I label frames as news if at least 1 of their L1-distances (relative to their station's seeds) is below their station's similarity threshold [`remark`: these latter are identified in step **(7)**].

In what follows I leave examples of unclassified and classified L1-distances:

Figure A.5: ***Examples of L1-Distances for CNN, FNC and MSN.***
(*sample period* - 1-week pilot dataset, November 18-24, 2018).

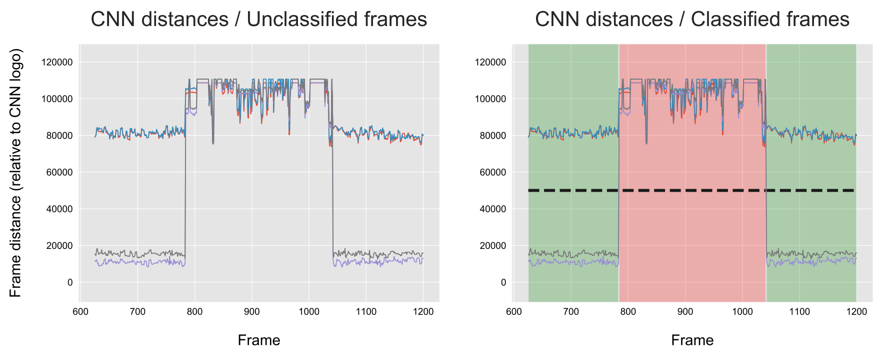

(a) CNN

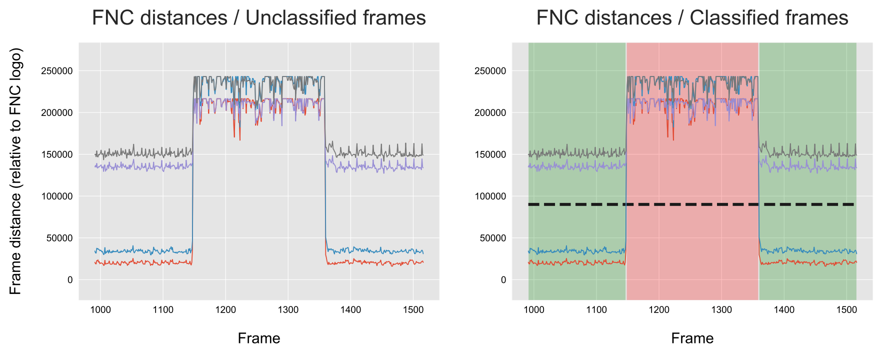

(b) FNC

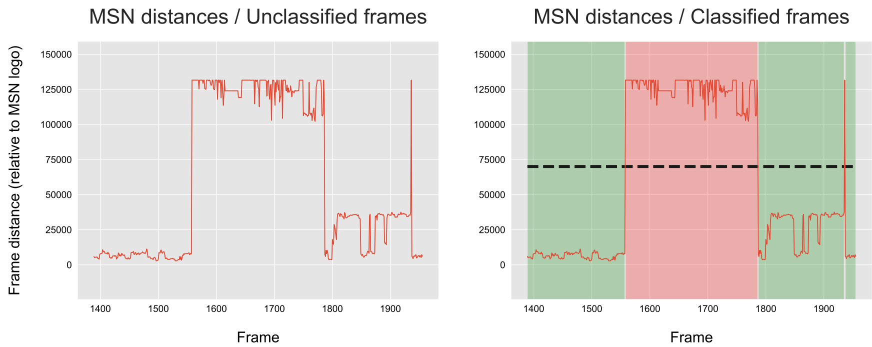

(c) MSN

`Note`: Left sub-figures report distances between (1) a set of cropped frames and (2) seed crops containing station logos. Right sub-figures highlight which frames are classified as news (in green) and ads (in red). Similarity thresholds are represented in right sub-figures, as dashed lines. Sub-figures for CNN and FNC report multiple distances as both stations display different logos in alternate image positions: (1) CNN/FNC show 2 different logos while broadcasting news; (2) these logos are shown in 2 alternate image positions; (3) I crop from each CNN/FNC frame both image positions and compute for each crop its respective distance to 2 cropped seeds; (4) as seen in both CNN/FNC sub-figures, I compute 4 distances for each CNN/FNC frame; (5) frames for which at least 1 distance is below their station's similarity threshold is labelled as news (in green).

Appendix D. Image Retrieval Algorithm - Breaking News

(back to Appendix)

In this subsection I enumerate each step taken to implement an image retrieval algorithm capable of classifying cable TV news frames as breaking news or standard news plus ads. In addition, I provide different examples to better illustrate each step.

In theory. As described in Appendix C, instances where breaking news are broadcasted on cable news television are identifiable in time if (1) each breaking news piece is always shown with a station-specific warning and if (2) each station shows that same warning in a constant position:

(a) take a seed frame where a breaking news piece has been broadcasted; crop from that seed frame its station's warning, which will be displayed in a constant position; then,

(b) take a random frame; crop from that frame the exact area where in your seed frame you have displayed that station's breaking news warning; it follows that,

(c) crops from frames where a piece of breaking news has been covered are expected to display a high level of similarity (measured through a similarity metric of choice) with respect to its respective station's warning, cropped from your seed frame.

In practice. I proceed as follows:

(1) I collect, select and process video data for each cable news TV station;

Regarding selecting video data, I proceed as described in Appendix C:

(1.1) I focus on shows that are broadcasted in periods where congress-members are significantly active on Twitter - i.e., shows not broadcasted through dawn (Eastern Time);

(1.2) in addition, to retrieve accurate timestamps for each frame, I focus on shows that constantly display time through either a bottom-left or a bottom-right digital clock;

(2) I run an image retrieval algorithm, as described in detail in Appendix C, to identify instances (frames) where general news were being broadcasted on cable news television;

(3) I extract a random sample of *news-related* frames and then skim manually through these to confirm that **(3.1)** cable news TV stations display their own breaking warning while broadcasting breaking news and that **(3.2)** these warnings are shown in constant or close to constant sets of frame coordinates;

(4) I manually skim through these random news frames to retrieve for each station a set of seed frames, where breaking news and consequently breaking warnings are displayed;

(6) I examine each seed frame and retrieve for each warning a set of constant frame coordinates.

Concerning warnings, CNN, FNC and MSN have 1 unique breaking warning each. In terms of frame coordinates, CNN and FNC alternate between 2 and 4 positions, respectively, to display their breaking warnings; MSN broadcasts breaking warnings in a constant unique image position.

Next, I leave examples of seed frames and station-specific warnings:

Figure A.6: ***Examples of Seed Frames for CNN, FNC and MSN***.
(*sample period* - 1-week pilot dataset, November 18-24, 2018).

Seed (CNN-brk_above)

Cropped logo from seed (CNN-brk_above)

(a) CNN

Seed (FNC-brk_above_left)

Cropped logo from seed (FNC-brk_above_left)

(b) FNC

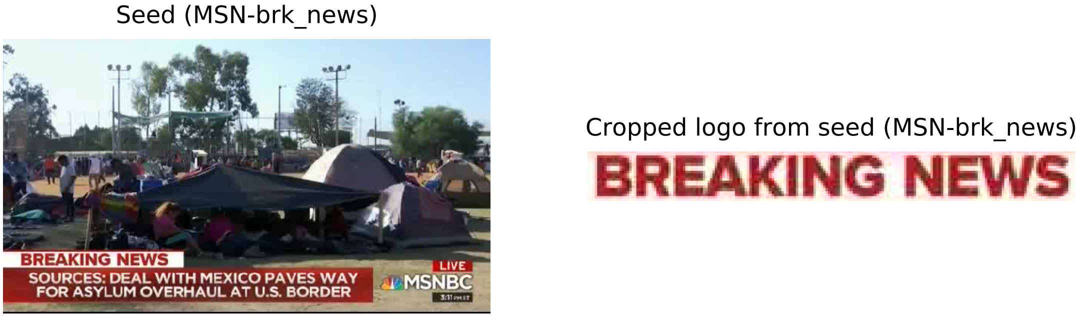

(c) MSN

After collecting for each station a set of seed frames and their warnings' frame coordinates:

(7) I identify similarity thresholds for each news station, as described in extent in Appendix C;

(8) I classify frames as breaking news or not, **(8.1)** re-running each image processing step included in step (7) and then **(8.2)** labelling frames as breaking news if at least 1 of their L1-distances (relative to their station's seeds) is below their station's similarity threshold.

Below I leave examples of unclassified and classified distances:

Figure A.7: ***Examples of L1-Distances for CNN, FNC and MSN.***
(*sample period* - 1-week pilot dataset, November 18-24, 2018).

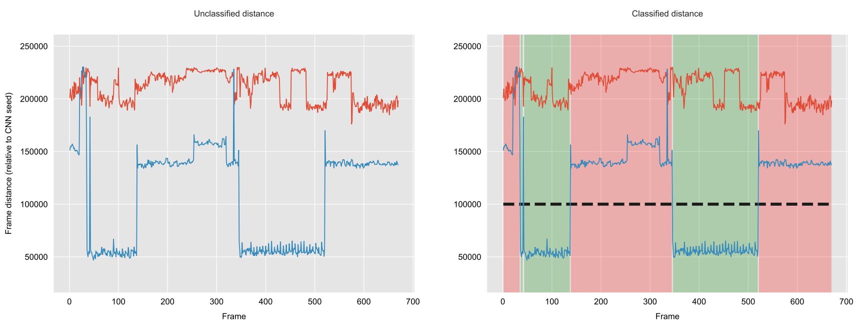

(a) CNN

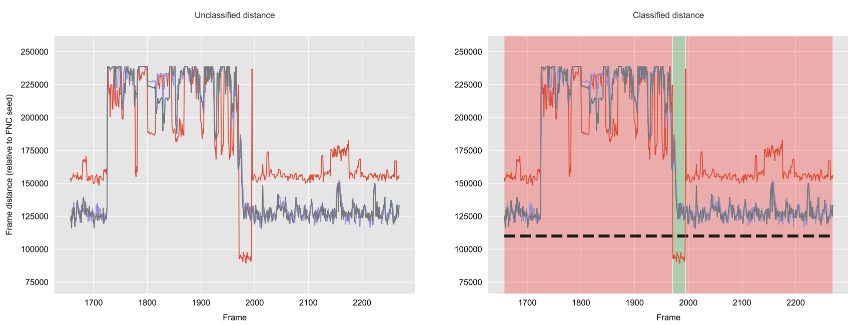

(b) FNC

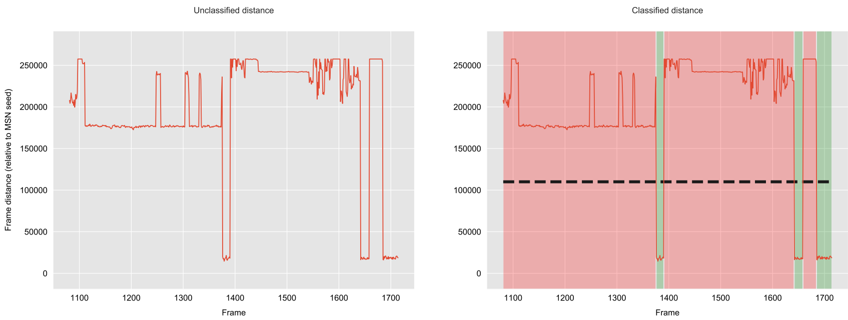

(c) MSN

`Note`: Left sub-figures report distances between (1) a set of cropped frames and (2) seed crops containing breaking news logos. Right sub-figures highlight which frames are classified as breaking news (in green) and standard news or ads (in red). Similarity thresholds are represented in right sub-figures, as dashed lines. Sub-figures for CNN and FNC report multiple distances as both stations display their respective breaking news warning in alternate image positions: (1) CNN/FNC show breaking news warnings in 2/4 alternate image positions; (3) I crop from each CNN/FNC frame each image position and compute for each crop its respective distance to their respective seeds; (4) frames for which at least 1 distance is below their station's similarity threshold are labelled as breaking news.

Chapter 3

Are TV News Slanting Audio?

Analyzing Prime-Time News

Abstract

Literature in media bias has measured bias on television news only through verbal-based measures. Non-verbal measures, while crucial to rate slant on television, are nonexistent. I take advantage of computer audition techniques to expand existing measures with audio-based outcomes. In doing so, I improve our understanding of how are news topics slanted by television outlets. I rate how news anchors emotionally address liberal and conservative topics. Word expressions are labelled as liberal and conservative according to how often these are used on social media by Democrat and Republican congress-members, respectively. The emotional tone of a news anchor is measured through that commentator's vocal pitch.Those liberal or conservative expressions that are most associated with abnormal variations in an anchor's vocal pitch are identified through a regression analysis.

3.1. Introduction

Media outlets are increasingly biased in their news coverage (Martin and Yurukoglu, 2017; Kim et al., 2022). There is different evidence that this type of coverage persuades individuals into voting for specific political parties, thus, affecting economic policy (DellaVigna and Kaplan, 2007; Hopkins and Ladd, 2014; Martin and Yurukoglu, 2017; Ash et al., 2021b).

To date, studies quantify media bias on television based on text transcripts. Nonetheless, other dimensions of television coverage (e.g vocal tone) are fundamental for slanting messages (Wang et al., 2021) and thus, shaping audiences' opinions and behaviors. In this paper, I provide a novel measure of bias for television news outlets, focused entirely on an audio-based outcome.

More specifically, I rate how news anchors behave emotionally when mentioning a liberal or conservative expression. To classify an expression as liberal or conservative, I take advantage of an exhaustive dataset of tweets posted by U.S. congress-members between 2019 and 2020, inclusively. More specifically, I label a word expression as liberal or conservative according to how often that expression was used on Twitter by a Democrat or a Republican congress-member (see Martin and Yurukoglu, 2017).

To measure how news anchors behave emotionally, I match an exhaustive dataset of audio clips for prime-time cable news shows to text transcripts for each of these shows. I then rate the emotional behavior of a news anchor by studying that individual's vocal pitch across time[1]. In particular, I measure how much emotional intensity a news commentator allocates to a given word expression by measuring how that commentator's vocal pitch differs from their standard vocal tone when mentioning a specific word expression (see Dietrich et al., 2018, 2019a,b).

This paper contributes to two main strands of economic literature.

[1]This is motivated by literature in psychology; Puts et al., 2006; Goudbeek and Scherer, 2010.

First, the current study contributes to a body of work that has focused on measuring bias in media coverage. Past works tended to classify media outlets as liberal or conservative according to different text-based criteria (Groseclose and Milyo, 2005; Gentzkow and Shapiro, 2010; Greenstein and Zhu, 2012; Martin and Yurukoglu, 2017).[2] A recent strand of this literature has started to expand existing measures of bias to non-verbal outcomes (see e.g., Boxell, 2020; Caprini, 2021; Ash et al., 2021a; Kim et al., 2022; Ash and Caliskan, 2022). However, these studies have focused exclusively on images. While audio is an important vehicle to slant messages and thus persuade individuals (Wang et al., 2021), it has not yet been studied. This study provides a first audio-based measure of bias.

Second, after measuring how cable news actors address different issues vocally, I plan to contribute to a strand of literature that has studied the effects of cable television over different social outcomes. In particular, I intend to apply this measure to study which dimension of TV coverage is determinant to setting policy agendas (Eisensee and Strömberg, 2007; Snyder Jr and Strömberg, 2010), persuading individuals into voting for particular political parties (DellaVigna and Kaplan, 2007; Gentzkow and Shapiro, 2010; Ash et al., 2021b), or politicizing social behaviors (see e.g., Allcott et al., 2020; Bursztyn et al., 2020, for health-related behaviors).

The rest of the paper is organized as follows: in Section 3.2, I outline the data sources and method that I use; in Section 3.3, I describe the empirical strategy followed; in Section 3.4, I conclude by outlining potential applications.

[2] Note that these references are not exhaustive and refer only to studies that measure implicit forms of bias. See Puglisi and Snyder Jr. (2015) for an exhaustive survey of empirical studies of media bias focused exclusively on text-based outcomes.

3.2. Data

To understand how prime-time TV news anchors vocally address liberal / conservative expressions, I take advantage of three datasets:

Cable television news timestamped transcripts. Text, timestamps and news/ads identifier for universe of dialogues broadcasted by cable news networks - from TV News Archive.

Cable television audio and image. Data on audio and images broadcasted by CNN, FOX News and MSNBC, from January 2020 to December 2020 (included). Both datasets have been constructed using raw video data on cable news broadcasts, kindly made available by TV News Archive.

Congressional tweets and ideological scores. Data on tweets and political ideology, for each congress-member, from January 2020 to December 2020 (included). Data on congressional tweets from Tweets of Congress; data on each representative's political ideology from Voteview.

3.2.1 Identifying Political Expressions

I construct a novel dictionary on liberal and conservative expressions, using tweets from U.S. congress-members. Then, I use this same dictionary to construct one main set of indicator variables, intended to point to instances in time where a markedly liberal and/or conservative term was mentioned.

Dictionary for Political Expressions

To position word expressions in a left-and-right political spectrum, I analyze every tweet posted by U.S. representatives from January 2019 to December 2020 (included).

In particular, I rate word expressions in terms of their usage by Democrat and Republican congress-members, respectively (see Gentzkow and Shapiro, 2010; Martin and Yurukoglu, 2017; Gentzkow et al., 2019, for identical analyses with congressional speeches).

In particular, I rank those expressions tweeted by U.S. congress-members according to how these are used (in terms of frequency of use) by accounts from Democrat and Republican congress-members, respectively. Expressions that are commonly tweeted by Democrat congress-members are labelled as "*liberal*". Instead, phrases that are normally used by Republicans are classified as "*conservative*". I provide additional details on this procedure in Appendix 3.5.1.

Table 3.1 shows those set of trigrams tweeted by U.S. congress-member that were most associated with Democrat and Republic representatives, respectively. Tables 3.5 and 3.6 provide a more exhaustive list of markedly liberal and conservative n-grams for 2019 and 2020, respectively.

In total, merging those expressions that are most associated with Democrat and Republican tweets during 2019 and 2020, I identify 2,633 unique expressions that I label either as liberal or conservative. Table 3.2 provides a full description of how many 1, 2 and 3-word phrases are labelled as liberal and conservative.

Indicators of Political Expressions

After assembling a dictionary for political expressions, I construct one main set of indicator variables. These are intended to point to instances in time where liberal or conservative expressions were mentioned on cable TV news.

More specifically, I construct two anchor-specific indicator variables, each focused on liberal and conservative expressions, respectively:

(A) ***Top 5 Trigrams during 2019***

Democrat	Republican
foreign govern interfer	need secur border
bill hous pass	bill would allow
deserv equal pay	thank presid trump
skyrocket cost prescript	last week discuss
mexico would pay	american peopl see

(B) ***Top 5 Trigrams during 2020***

Democrat	Republican
trump administr attempt	new york post
fight racial justic	peac middl east
ask suprem court	need held account
mitch mcconnel block	border wall system
fight justic equal	get back work

Table 3.1: ***Top 5 Partisan Trigrams***. Top 5 trigrams (i.e., 3-word expressions) used on Twitter by U.S. congress-members, that best predict a representative's ideological score (DW-NOMINATE; Lewis et al., 2020 and Poole and Rosenthal, 1985) during 2019 (Panel A) and 2020 (Panel B).

	1-Word	2-Words	3-Words
Liberal	256	948	467
Conservative	187	528	247

Table 3.2: ***Dimension of Political Dictionaries***. Number of 1, 2 and 3-word phrases that are labelled either as liberal (i.e., associated with tweets by Democrat congress-members) or conservative (symmetric for Republican politicians).

$$X^C_{i,t} = \begin{cases} 1 & \textit{if } (\#\,\mathtt{conservative})_{i,t} > 0 \\ 0 & \textit{if } (\#\,\mathtt{conservative})_{i,t} = 0 \end{cases} \quad \text{and} \quad X^L_{i,t} = \begin{cases} 1 & \textit{if } (\#\,\mathtt{liberal})_{i,t} > 0 \\ 0 & \textit{if } (\#\,\mathtt{liberal})_{i,t} = 0 \end{cases} \tag{3.1}$$

where $(\#\,\mathtt{conservative})_{i,t}$ stands for the number of times a conservative term was mentioned by anchor i during time interval t (and vice-versa for liberal phrases).

To construct these variables, I take advantage of a pre-trained speech recognition algorithm, to align the transcripts of cable news outlets with their respective audio tracks (see `Gentle`). This allows me to pin-point each transcribed expression in time, at a secondly frequency.

An important remark regarding $X^C_{i,t}$ and $X^L_{i,t}$: these variables refer only to a subset of cable news dialogues in which a news anchor is speaking. I provide additional details on how these anchor-specific news segments are identified in Section 3.2.2.

3.2.2 Measuring Emotional Intensity

I draw on psychology, phonetics and political science studies to build an audio-based measure for emotional intensity. In particular, I take reference from psychology to use variations in vocal pitch as a measure for emotional intensity.

To be more specific, an individual's voice pitch is determined by a person's fundamental frequency (F_0), defined below as in Titze and Martin (1998):

$$F_0 = \frac{1}{2L}\sqrt{\frac{\sigma}{\rho}} \tag{3.2}$$

where L is the vocal fold length, σ is the longitudinal stress on the vocal folds, and ρ is the vocal fold tissue density. Variations in vocal fold length (L) and density ρ are determined by genetics (see Debruyne et al., 2002); changes in longitudinal stress are instead related to idiosyncratic factors, such as emotions.[3]

Bearing these results in mind, I take advantage of both literature and software in phonetics to measure the fundamental frequencies of prime-time news anchors' (Jadoul et al., 2018; Boersma and Weenink, 2021).[4][5] In particular, I compute for each prime-time news anchor:

$$Y_{i,t} = \frac{F_0^{i,t} - \mu^{i,T}}{\sigma^{i,T}} \tag{3.3}$$

where $F_0^{i,t}$ stands for the estimated fundamental frequency of anchor i's in centisecond t. $\mu^{i,T}$ stands for the average fundamental frequency of anchor i's, during time interval T (e.g., a week or a day). $\sigma^{i,T}$ stands for a standard deviation of anchor i's fundamental frequency (during time interval T).

$Y_{i,t}$ measures by how much is a news anchor's vocal tone deviating from that anchor's baseline pitch at a given moment in time (these deviations being scaled according to a dynamic standard deviation, to account for voice changes that are time-specific – e.g., changes in a person's pitch driven by health reasons).

In addition, note that $Y_{i,t}$ is computed for a non-exhaustive set of anchor-specific interventions which are identified through a careful analysis of the closed captioned transcripts of cable news outlets.

In particular, closed captions (CCs) are endowed with two particularities:

[3]More specifically, Puts et al. (2006) have shown that emotional activation triggers abnormal increases in F_0, through σ. Goudbeek and Scherer (2010) provided additional evidence in this direction.

[4]These frequencies are measured in adjustable time steps, following Boersma et al. (1993)'s algorithm on how to compute a sound segment's fundamental frequency. I follow past political science literature, that applied these same methods to other settings, and compute F_0 at a centisecond frequency.

[5]In addition, I take into account that an individual's pitch is related to idiosyncratic factors such as e.g, gender and health status - I standardize variations in an individual's voice pitch according to that same individual's vocal pitch standard deviation and average (see Dietrich et al., 2019a,b).

(i) CCs from news and advertisements are showcased in different formats;

(ii) CCs segment dialogues into speaker-specific pieces.

I first use information (i) to identify those dialogues that are strictly related to news content. I then take advantage of information (ii) to identify so-called "*long*" interventions (i.e. in total amount of time) made by a single individual.[6] I then compute $Y_{i,t}$ using these *long* news segments only, assuming that longer news-related dialogues are more likely to be attributable to a show's news anchor.[7]

	CNN	FNC	MSN
8PM	42	30	47
9PM	33	26	29
10PM	36	37	35

(A) Number of Dialogues

	CNN	FNC	MSN
8PM	227	158	262
9PM	203	171	378
10PM	180	160	206

(B) Minutes of Dialogues

Table 3.3: ***Prime-time News-Anchor Dialogues***. Panel A refers to number of "*long*" dialogues that took place in between January 2020 and June 2020 (included) during prime-time shows. Panel B shows how much time these same dialogues took, in terms of minutes. "*FNC*" stands for Fox News while "*MSN*" stands for MSNBC.

In total, by applying this selection criteria over those prime-time news shows that ran from January to June 2020 (included), I identify a total of 315 "*long*" dialogues that amount for approximately 32 hours. Panels (A) and (B) of Figure 3.3, above, provide additional information about these dialogues.[8]

[6]More specifically, I rank each network's single-person interventions according to their duration in time. Then, I label an intervention as "*long*" if this same intervention has a duration above a 99th network-specific percentile.

[7]To validate this assumption, in future work I will cross-check if each segment is indeed attributable to these shows' anchors. To do so, I will run a facial recognition algorithm, pre-trained to recognize faces from celebrities, on each segment's video frames (see Boxell, 2020).

[8]Regarding Table 3.3, each statistic refers to a particular news anchor. More specifically: (CNN, 8PM) Anderson Cooper in "*Anderson Cooper 360*"; (CNN, 9PM) Andrew Cuomo in "*Cuomo Prime Time*"; (CNN, 10PM) Don lemon in "*CNN Tonight With Don Lemon*"; (FNC, 8PM) Tucker Carlson in "*Tucker*

Figure 3.1 plots average vocal pitches for each outlet's prime-time news anchors (measured by employing a pitch detection algorithm onto audio tracks referent to "*long*" prime-time dialogues – briefly described in Table 3.3). Interestingly, within a prime-time slot, news-anchors do not seem to differ substantially in terms of their vocal tone (a slight exception can be found for Fox News at 10pm).

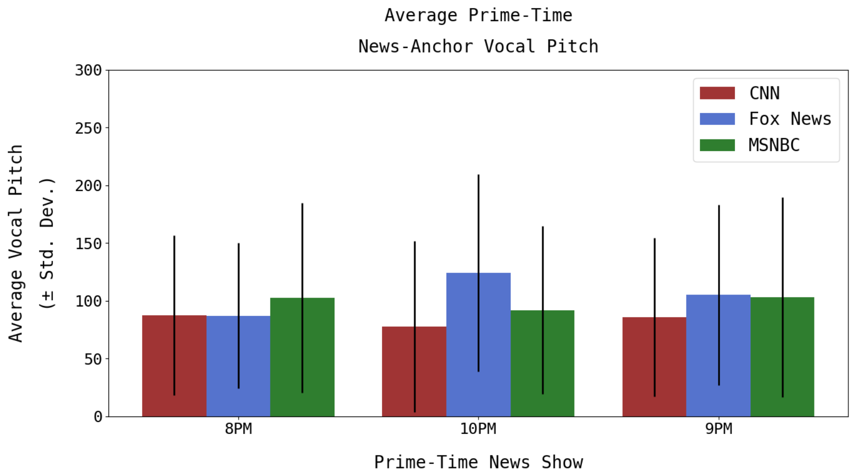

Figure 3.1: ***Prime-Time News-Anchors Average Pitch***. Each bar refers to a network-show average vocal pitch (inferred from a set of "*long*" dialogues, aired between January and June 2020 (included), during each network-show). Error bars refer to each anchor's standard deviation (i.e., in terms of their vocal pitch).

After inferring for each news anchor that commentator's baseline pitch, it is possible to measure how is an anchor's vocal tone deviating from that person's standard voice over time (i.e., one can go on and compute $Y_{i,t}$).

Carlson Tonight"; (FNC, 9PM) Sean Hannity in "*Hannity*"; (FNC, 10PM) Laura Ingraham in "*The Ingraham Angle*"; (MSN, 8PM) Chris Hayes in "*All In With Chris Hayes*"; (MSN, 9PM) Rachel Maddow in "*The Rachel Maddow Show*"; (MSN, 10PM) Lawrence O'Donnel in "*The Last Word With Lawrence O'Donnel*".

3.3. Empirical Strategy

I intend to rate how news-anchors, within TV outlets, emotionally address different types of political expressions. To do so, I will estimate a standard OLS regression:

$$Y_t^{i,T} = \delta_T + \beta_C X_t^C + \beta_C X_t^L + \varepsilon_t^{i,T}, \tag{3.4}$$

where $Y_t^{i,T}$ stands for an emotional intensity outcome measuring how news anchor i is behaving vocally during time period t (as defined in Equation 3.3). δ_T stands for a time fixed effect, intended to control for macro factors that can affect how all prime-time news-anchors behave emotionally (e.g., a school shooting that becomes news during a common prime-time slot, causing all anchors to become more emotional in their comments). X_t^C (X_t^L) is an indicator variable equal to 1 in period t if a conservative (liberal) expression is mentioned during period t (as defined in Equation 3.1).

The coefficients of interest, β_j, estimates by how much does an anchor's vocal pitch deviates from that commentator's baseline tone whenever a conservative ($j = C$) or liberal ($j = L$) expression is mentioned by that same anchor (note: this coefficient provides an *average* estimate of how a news anchor reacts to a partisan expression, relative to a so-called *neutral* word phrase).

Ideally, these expressions ought to be compared with phrases that are similarly used during prime-time (at least in terms of frequency, ideally also considering which words or topics are normally addressed close to each expression). With this in mind, in future analyses, I will estimate Equation 3.4 by restricting this equation's estimation sample to moments in time in which similar word expressions are being mentioned by a news anchor.

3.4. Conclusion

I intend to expand existing measurements of political bias in television news with audio-based outcomes. By measuring emotional intensity through vocal pitch, I will study if television news anchors allocate different emotional charges to liberal or conservative features. In doing so, I contribute to a recent literature employing computer vision and computer audition techniques to study television cable news: from slant in images (Ash et al., 2021a) to differences in speech according to gender (Hong et al., 2020).

This study can be extended to a number of interesting applications. Recent efforts to measure bias in different perspectives (text, image and, here as a proposal, audio) can provide researchers with tools to understand which dimensions of slant (on television) are more influential at dictating public opinion and thus, (i) influencing policy agendas (Eisensee and Strömberg, 2007; Snyder Jr and Strömberg, 2010), (ii) persuading individuals to vote for specific political platforms (DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2017) and (iii) affecting political polarization (Campante and Hojman, 2013).

Bibliography

Allcott, H., Boxell, L., Conway, J., Thaler, M., and Gentzkow, M. (2020). Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic. *Journal of Public Economics*, 191:104254. [URL].

Ash, E. and Caliskan, C. (2022). Fight or Flight on Fox?: Partisan Fear Responses on U.S. Cable News Shows. *SSRN Research Paper Series*. SSRN 4204973 [URL].

Ash, E., Durante, R., Grebenshchikova, M., and Schwarz, C. (2021a). Visual Representation and Stereotypes in News Media. *SSRN Research Paper Series*. SSRN 3934858 [URL].

Ash, E., Galletta, S., Pinna, M., and Warshaw, C. (2021b). The Effect of Fox News Channel on U.S. Elections: 2000-2020. *SSRN Research Paper Series*. SSRN 3837457 [URL].

Boersma, P. et al. (1993). Accurate Short-Term Analysis of the Fundamental Frequency and the Harmonics-To-Noise Ratio of a Sampled Sound. In *Proceedings of the Institute of Phonetic Sciences*, volume 17, pages 97–110. Citeseer.

Boersma, P. and Weenink, D. (2021). Praat: Doing Phonetics by Computer [Computer Program].

Boxell, L. (2020). Slanted Images: Measuring Nonverbal Media Bias. *Mimeo*. [URL].

Bursztyn, L., Rao, A., Roth, C., and Yanagizawa-Drott, D. (2020). Misinformation during a Pandemic. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-44). [URL].

Campante, F. R. and Hojman, D. A. (2013). Media and Polarization: Evidence from the Introduction of Broadcast TV in the United States. *Journal of Public Economics*, 100:79–92. [URL].

Caprini, G. (2021). Visual Bias. *Mimeo (Job Market Paper)*. [URL].

Debruyne, F., Decoster, W., Van Gijsel, A., and Vercammen, J. (2002). Speaking Fundamental Frequency in Monozygotic and Dizygotic Twins. *Journal of Voice*, 16(4):466–471. [URL].

DellaVigna, S. and Kaplan, E. (2007). The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics*, 122(3):1187–1234. [URL].

Dietrich, B. J., Enos, R. D., and Sen, M. (2019a). Emotional Arousal Predicts Voting on the US Supreme Court. *Political Analysis*, 27(2):237–243. [URL].

Dietrich, B. J., Hayes, M., and O'Brien, D. Z. (2019b). Pitch Perfect: Vocal Pitch and the Emotional Intensity of Congressional Speech. *American Political Science Review*, 113(4):941–962. [URL].

Dietrich, B. J., Schultz, D., and Jaquith, T. (2018). This Floor Speech Will Be Televised: Understanding the Factors that Influence When Floor Speeches Appear on Cable Television. *Mimeo*. [URL].

Eisensee, T. and Strömberg, D. (2007). News Droughts, News Floods, and US Disaster Relief. *The Quarterly Journal of Economics*, 122(2):693–728. [URL].

Gentzkow, M. and Shapiro, J. M. (2010). What Drives Media Slant? Evidence from US Daily Newspapers. *Econometrica*, 78(1):35–71. [URL].

Gentzkow, M., Shapiro, J. M., and Taddy, M. (2019). Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4):1307–1340. [URL].

Goudbeek, M. and Scherer, K. (2010). Beyond arousal: Valence and potency/control cues in the vocal expression of emotion. *The Journal of the Acoustical Society of America*, 128(3):1322–1336. [URL].

Greenstein, S. and Zhu, F. (2012). Is Wikipedia Biased? *American Economic Review*, 102(3):343–48. [URL].

Groseclose, T. and Milyo, J. (2005). A Measure of Media Bias. *The Quarterly Journal of Economics*, 120(4):1191–1237. [URL].

Hong, J., Crichton, W., Zhang, H., Fu, D. Y., Ritchie, J., Barenholtz, J., Hannel, B., Yao, X., Murray, M., Moriba, G., et al. (2020). Analyzing Who and What Appears in a Decade of US Cable TV News. *arXiv preprint arXiv:2008.06007*. [URL].

Hopkins, D. J. and Ladd, J. M. L. (2014). The Consequences of Broader Media Choice: Evidence from the Expansion of Fox News. *Quarterly Journal of Political Science*, 9(1):115–135. [URL].

Jadoul, Y., Thompson, B., and de Boer, B. (2018). Introducing Parselmouth: A Python interface to Praat. *Journal of Phonetics*, 71:1–15. [URL].

Kim, E., Lelkes, Y., and McCrain, J. (2022). Measuring Dynamic Media Bias. *Proceedings of the National Academy of Sciences*, 119(32):e2202197119. [URL].

Lewis, J. B., Poole, K., Rosenthal, H., Boche, A., Rudkin, A., and Sonnet, L. (2020). Voteview: Congressional Roll-call Votes Database. *https://voteview. com*. [URL].

Martin, G. J. and Yurukoglu, A. (2017). Bias in Cable News: Persuasion and Polarization. *American Economic Review*, 107(9):2565–99. [URL].

Poole, K. T. and Rosenthal, H. (1985). A Spatial Model for Legislative Roll Call Analysis. *American Journal of Political Science*, pages 357–384. [URL].

Puglisi, R. and Snyder Jr., J. M. (2015). Empirical Studies of Media Bias. In *Handbook of Media Economics*, volume 1, pages 647–667. [URL].

Puts, D. A., Gaulin, S. J., and Verdolini, K. (2006). Dominance and the Evolution of Sexual Dimorphism in Human Voice Pitch. *Evolution and Human Behavior*, 27(4):283–296. [URL].

Snyder Jr, J. M. and Strömberg, D. (2010). Press Coverage and Political Accountability. *Journal of Political Economy*, 118(2):355–408. [URL].

Titze, I. R. and Martin, D. W. (1998). Principles of Voice Production.

Wang, X., Lu, S., Li, X., Khamitov, M., and Bendle, N. (2021). Audio Mining: The Role of Vocal Tone in Persuasion. *Journal of Consumer Research*, 48(2):189–211. [URL].

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320.

3.5. Appendix

3.5.1 Congressional Tweets

Sources: Universe of tweets posted by U.S. congress-members during 2019 and 2020 - collected and made available by Tweets of Congress. Metadata for each congress-member (e.g., party, age, gender, ideological score) - data on party, age and gender comes from @unitedstates; data on each congress-member's DW-NOMINATE score comes from Voteview (see Lewis et al., 2020).

Table 3.4: ***Congressional Tweets – Descriptive Statistics***

		Gender		Age[1]		Party[2]	
	Total	M	F	< 61	> 61	D	R
# of Members	586	445	141	259	327	295	287
# of Tweets	1,603,888	1,115,263	488,625	777,521	826,367	998,938	588,483
# of Tweets per Date	3.7	3.4	4.7	4.1	3.5	4.6	2.8
# of Chars. per Tweet	221	221	224	220	223	228	210

Note: (0) "*# of Members*" stands for number of Twitter accounts related to a congress-member; "*# of Tweets*" stands for number of tweets posted from January 2019 until December 2020, included; "*# of Tweets per Date*" stands for average number of tweets posted during sample of analysis; "*# of Chars. per Tweet*" stands for average number of characters per tweet during sample of analysis. (1) Age labelled as below or above median age, median here being 61 years old. (2) Congress-members that are either *Independent* or *Libertarian* are not included (these account for 12,824 (0.8%) and 3,643 (0.2%) tweets, respectively).

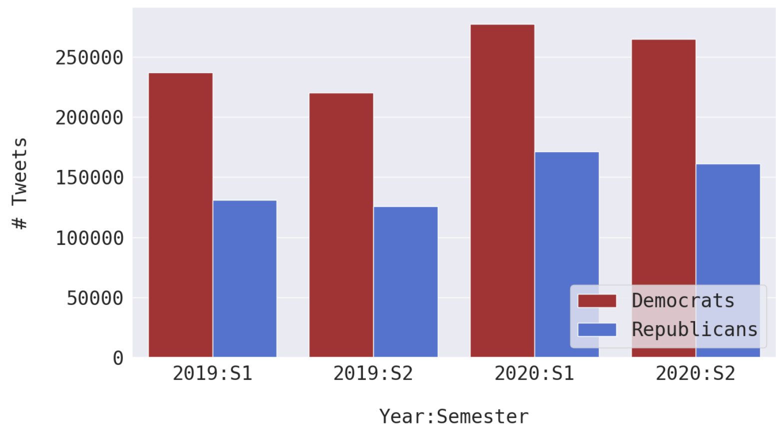

Figure 3.2: ***Congressional Tweets – By Semester***. Number of tweets posted by Democrat and Republican congress-members, by semester, during 2019 and 2020.

Table 3.5: ***Top 20 Partisan Phrases for 2019 (1-Word Expressions, Bigrams and Trigrams)***

1-Word Dem.	1-Word Rep.	2-Word Dem.	2-Word Rep.	3-Word Dem.	3-Word Rep.
pop	burdensom	labor movement	govern spend	foreign govern interfer	need secur border
reded	beto	choos pay	big govern	bill hous pass	bill would allow
decenc	hunter	trump claim	current law	deserv equal pay	thank presid trump
cynic	theft	support medicar	religi liberti	skyrocket cost prescript	last week discuss
nonhispan	lawabid	union worker	one small	mexico would pay	american peopl see
rudi	delegitim	worker go	protect life	hold administr account	million new job
skip	evil	propos cut	free market	one law even	border patrol agent
beneath	bureaucrat	latino commun	lawabid citizen	dirt polit oppon	may never forget
preval	slam	right privileg	human life	alreadi pass hous	protect american peopl
succe	amnesti	stay healthi	thank men	investig polit oppon	last two year
harder	unlimit	need invest	demand transpar	releas tax return	put america first
nixon	smuggler	call senat	prayer breakfast	proud cosponsor act	alway great see
formerli	sanctiti	coequal branch	life liberti	afford prescript drug	one thing clear
furlough	regulatori	cosponsor act	countri like	violenc prevent bill	social media platform
doral	tyranni	sure voic	deal china	proud support bill	citizenship question censu
nielsen	creator	barr must	u sen	militari construct project	vote articl impeach
urgent	liberti	need fight	free speech	access clean water	sinc took offic
pantri	bless	impact trump	year row	anoth mass shoot	human right violat
count	christian	tri cover	time bring	make sure voic	duli elect presid
africanamerican	electron	everi commun	go toward	presid hous floor	everi singl one

Note: Table showcases phrases used on Twitter that best predict a representative's ideological score (DW-NOMINATE) during 2019 (i.e., only making use of tweets posted during 2019). "*Dem.*" stands for Democrat while "*Rep.*" stands for Republican. In steps: (1) tweets are pre-processed- i.e., text is lower-cased; retweets, URLs, hashtags, emojis and stop-words are eliminated; words are tokenized and stemmed to their morphological roots; (2) universe of 2 and 3 word phrases are identified and "common" phrases are selected - i.e., phrases that are present in at least 100 different tweets only; (3) an Elastic Net regression (Zou and Hastie, 2005) is estimated to predict each congress-member ideology with hers or his respective (relative) use of "common" 1/2/3-word phrases; (4) coefficients are ranked in terms of their magnitude (only first 20 coefficients are shown above). Democratic phrases stand for positive coefficients with largest magnitude; Republican phrases are negative coefficients with largest absolute magnitude. DW-NOMINATE used as congress-members' ideological score, from Lewis et al. (2020) and introduced by Poole and Rosenthal (1985); procedure to identify partisan phrases drawn from Martin and Yurukoglu (2017) and Gentzkow and Shapiro (2010).

Table 3.6: ***Top 20 Partisan Phrases for 2020 (1-Word Expressions, Bigrams and Trigrams)***

1-Word Dem.	1-Word Rep.	2-Word Dem.	2-Word Rep.	3-Word Dem.	3-Word Rep.
grim	emir	labor movement	alway good	trump administr attempt	new york post
heartless	evil	countri demand	govern spend	fight racial justic	peac middl east
mubarak	rioter	stand shoulder	agreement israel	ask suprem court	need held account
photoop	arab	without health	tech censorship	mitch mcconnel block	border wall system
supremacist	regulatori	ask suprem	today may	fight justic equal	get back work
empathi	unborn	die coronaviru	latest episod	worker put live	america small busi
presidentelect	bureaucrat	care everi	death per	use tear ga	join discuss latest
gut	censorship	sexual orient	debat amend	matter life death	state local offici
inequ	liberti	safeti protect	free market	social secur check	unit arab emir
orient	riot	expand food	york post	join us virtual	senat democrat block
hysterectomi	growth	trump tell	legisl process	worker risk live	speaker pelosi hous
ego	greatest	account offic	offer amend	access ballot box	pack suprem court
gender	articl	fight racial	spi american	mitch mcconnel desk	legal vote count
uninsur	regul	commit peac	would love	trump admin must	join morn discuss
undercount	mob	trump go	protect freedom	continu fight justic	direct deposit inform
sabotag	anarchi	cut medicar	nation great	biden vice presidentelect	kid back school
equit	lockdown	protect pregnant	senat say	commun commun color	thank presid trump
disenfranchis	success	protect planet	patrol agent	end gun violenc	everi legal vote
background	freedom	patient protect	busi week	sign join us	thank men women
unconscion	progrowth	get cover	join talk	end polic brutal	sadden hear pass

Note: Table showcases phrases used on Twitter that best predict a representative's ideological score (DW-NOMINATE) during 2020 (i.e., only making use of tweets posted during 2020). "*Dem.*" stands for Democrat while "*Rep.*" stands for Republican. In steps: (1) tweets are pre-processed- i.e., text is lower-cased; retweets, URLs, hashtags, emojis and stop-words are eliminated; words are tokenized and stemmed to their morphological roots; (2) universe of 2 and 3 word phrases are identified and "common" phrases are selected - i.e., phrases that are present in at least 100 different tweets only; (3) an Elastic Net regression (Zou and Hastie, 2005) is estimated to predict each congress-member ideology with hers or his respective (relative) use of "common" 1/2/3-word phrases; (4) coefficients are ranked in terms of their magnitude (only first 20 coefficients are shown above). Democratic phrases stand for positive coefficients with largest magnitude; Republican phrases are negative coefficients with largest absolute magnitude. DW-NOMINATE used as congress-members' ideological score, from Lewis et al. (2020) and introduced by Poole and Rosenthal (1985); procedure to identify partisan phrases drawn from Martin and Yurukoglu (2017) and Gentzkow and Shapiro (2010).