2021

Bias, inequality, and polarization in modern digital information systems

BOSTON UNIVERSITY

QUESTROM SCHOOL OF BUSINESS

Dissertation

BIAS, INEQUALITY, AND POLARIZATION IN

MODERN DIGITAL INFORMATION SYSTEMS

by

KAI ZHU

B.S., Beijing Language and Culture University, 2011
B.A., Peking University, 2013
A.M., Indiana University, 2015

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2021

Approved by

First Reader ____

Dylan T. Walker, Ph.D.
Assistant Professor of Information Systems

Second Reader ____

Marshall Van Alstyne, Ph.D.
Questrom Professor in Management
Professor of Information Systems

Third Reader ____

David G. Rand, Ph.D.
Erwin H. Schell Professor
Associate Professor of Management Science
Associate Professor of Brain and Cognitive Sciences
Massachusetts Institute of Technology

BIAS, INEQUALITY, AND POLARIZATION IN

MODERN DIGITAL INFORMATION SYSTEMS

KAI ZHU

Boston University Questrom School of Business, 2021

Major Professor: Dylan T. Walker, Ph.D., Assistant Professor of Information Systems

ABSTRACT

Digital technology has the potential to "democratize information" – making ideas, opinions, and knowledge accessible anywhere, anytime, and to everyone. But is this potential truly realized or will it ever be realized? Do systems enabled by digital technology exhibit or even enhance information bias, skewness, and polarization? How can we overcome them? In this dissertation, I investigate these questions in two major but distinct digital information systems: open collaboration systems (i.e., Wikipedia) and mass media broadcast networks (i.e., broadcast television in the United States).

Open collaboration platforms have fundamentally changed the way knowledge is produced, disseminated, and consumed. Wikipedia is arguably one of the most successful examples of such platforms, serving millions of information seekers daily. Despite many benefits provided by the decentralization of knowledge production on Wikipedia, does the open nature and lack of broad oversight and coordination leave the question of information poverty and skewness to the mercy of the system’s natural dynamics? And if so, what can be done to address this? In Chapter 1, I examined this question using both causal inference from a natural experiment and empirically informed diffusion simulations.

Another important and pervasive information system is that of televised mass media. Whereas Wikipedia is relatively open and does not have strong information gatekeeping, televised mass media has various forms of information gatekeeping, particularly through media ownership, government regulation and journalistic practice. But how does this gatekeeping affect skewness and polarization in the real-world information that is conveyed to the public? To investigate these questions, I study televised news information systems in the United States with a massive scale unstructured text data and various state-of-the-art text mining techniques in Chapter 2 and Chapter 3 of this dissertation. The text transcripts include the complete televised content from more than 800 television channels across all 210 designated media markets in the United States over a 5-year period between 2013 and 2018.

Chapter 2 of this dissertation examines how media ownership impact political slant and information diversity in the news using massive-scale text transcripts. I found that when large owners act coherently, they can skew information to emphasize views, perspectives and framing that they advocate. This is important because previous studies have shown that broadcast media can have a dramatic impact on political and social outcomes and undeniably shapes the national dialogue surrounding important issues.

In Chapter 3 of this dissertation, I study the skewed coverage of gun violence incidents in local televised news. I found that some types of gun violence, such as suicide, accidents, domestic violence and sex crimes are systematically covered less relative to other types such as assault weapon incidents, are systematically covered more. Importantly, areas of high vs. low gun ownership received different exposure to different incident types through local news coverage, further dividing an already divided population. I conducted ting a nationally representative survey found that the general public's view on different type of gun violence is skewed in a manner that is consistent with "the warped mirror" that our media conveys.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER ONE

Content Growth and Attention Contagion in Information Networks: Addressing Information Poverty on Wikipedia

Introduction

Wikipedia is one of the most successful examples of open collaboration platforms, serving millions of information seekers daily. It is both a repository of free knowledge and the most-visited educational resource on the planet[1]. By the end of 2017, a mere sixteen years since its inception, the English language Wikipedia alone contained over 5.5 million articles and a total of over 3.1 billion words, over 60 times as many as the next largest English-language encyclopedia, *Encyclopædia Britannica*[2]. It consists of millions of articles written by a global network of volunteers and is accessible to anyone with an internet connection. Wikipedia represents a new generation of internet-based collaborative tools that strives to be open, accessible, and egalitarian.

However, Wikipedia's reliance on open and distributed collaboration as well as community governance is not without its problems. As noted by Wikipedia itself, volunteers don't always contribute to the content that people need the most[3]. A large proportion of articles are incomplete or insufficiently supported with references[4]. Because of Wikipedia's open and distributed production model, it is difficult to direct contributors' attention to articles that most need improvement. Hence, not only are these

---

[1] It is the 5th most visited website in the world, according to Alexa.
[2] https://en.wikipedia.org/wiki/Wikipedia:Size_of_Wikipedia
[3] https://wikiedu.org/changing/wikipedia/
[4] http://time.com/4180414/wikipedia-15th-anniversary/

articles incomplete, but they are likely to remain so. As a consequence, the coverage and depth of knowledge in Wikipedia articles is uneven. While well-developed articles are considerably longer than their analogues in *Encyclopædia Britannica*, many articles are still of poor quality and are on average half as long as their professionally edited analogues[5]. Importantly, coverage also appears to be uneven across both geographical areas and knowledge domains (Graham et al. 2014, Halavais and Lackaff 2008, Kittur et al. 2009). For example, Wikipedia has strong coverage of military history and political events in America, but articles on biology, law, medicine, and information on developing countries are often absent or underdeveloped[6].

Left unchecked, the societal implications of uneven coverage are deeply troubling. Despite the openness of Wikipedia, there are growing concerns that geographical areas and knowledge domains that are left out or underrepresented will remain so or become even further underrepresented relative to the growing knowledge base in a kind of poor-get-poorer phenomenon. Geographical informational skews can act to further limit our understandings of, attention to, and interactions with impoverished areas in terms of regional economic, social, political, and cultural concerns (Forman et al. 2012, Graham et al. 2014, Norris 2001, Yu 2006). Knowledge-domain information skews can compound insularity, lead to domain-based siloing, and push information seekers towards alternative, domain-specific information platforms that are less open and not free. Informational skew may reinforce or even compound existing biases in worldviews

[5] https://en.wikipedia.org/wiki/Wikipedia:Size_comparisons
[6] https://en.wikipedia.org/wiki/Criticism_of_Wikipedia

and exacerbate information poverty. Existing research has shown that information (un)availability has a surprisingly strong impact on real-world outcomes in financial markets, scientific advancement, and the tourist industry (Hinnosaar et al. 2017, Thompson and Hanley 2017, Xiaoquan and Lihong 2015, Xu and Zhang 2013). These studies further emphasize the salience of the skewed coverage problem in Wikipedia. Importantly, while we focus on Wikipedia, concerns of uneven coverage exist in a variety of platforms that facilitate collaborative content production, including open-source software (e.g. GitHub), knowledge markets (e.g. Stack Overflow or Quora), and product reviews (e.g. Amazon or Steam).

It is unclear whether Wikipedia's uneven coverage is driven by selection effects on the part of Wikipedia editors due to their intrinsic interests (Kuznetsov 2006, Nov 2007), natural emerging trends and exogenous factors (Kämpf et al. 2012, 2015, Keegan et al. 2013) or a systematic tendency for well-developed articles to continue to receive more attention via the "rich-get-richer" dynamic (Aaltonen and Seiler 2016, Barabási and Albert 1999). Most existing work on knowledge contribution behavior on Wikipedia has focused primarily on the motivation of its *editors* (Gallus 2016, Harhoff et al. 2003, Lampe et al. 2012, Nov 2007, Zhang and Zhu 2011, Zhu et al. 2013). However, it is critical that we understand the factors that govern the evolution and lifecycle of *articles,* which are central to the dynamics of Wikipedia as a system. Such factors are also likely important determinants of uneven coverage. Unfortunately, our understanding of how open collaboration platforms evolve and attract attention is still very limited.

There are three streams of research in the literature that are relevant to our study.

The first stream of research emphasizes the dynamic co-evolution of knowledge consumption and knowledge production. The open collaboration model allows consumers of knowledge to react to existing content and potentially also become contributors. But, how does production and consumption of knowledge interact in this complex dynamic system (Kämpf et al. 2012, Wilkinson and Huberman 2007)? Aaltonen and Seiler (2016) find that longer Wikipedia articles tend to receive more editing in the future. Kummer (2019) studied how attention shocks arising from natural disasters affect contributions. Kane and Ransbotham (2016) investigate the feedback loop between consumption and contribution of articles in WikiProject Medicine and find that the state of content moderates this feedback loop. It is noteworthy that they argue that this feedback loop in open collaboration platforms has been under-researched and that a deeper understanding is warranted.

The second stream of research emphasizes the network perspective by recognizing that, similar to the web as a whole, Wikipedia is an information network of hyperlinked articles. This has important implications: at least some of the traffic (attention) arriving at a particular article flows outward along links to other downstream articles. The importance of this network perspective derives from a long tradition of relating a node's relative importance to its network properties -- an assumption that is implicit to the well-known PageRank algorithm. The overall exposure of an article in Wikipedia is determined by the various ways that an information seeker can arrive at it via both external (e.g., search engines) and internal sources (upstream Wikipedia articles). Previous research has shown that the network position of an article is correlated with its content consumption and production (Kane 2009, Kummer et al. 2016, Ransbotham et al. 2012). Moreover, the structural embeddedness of an article in the content-contributor network is positively related to its viewership and information quality (Kane and Ransbotham 2016, Ransbotham et al. 2012). Beyond information networks, Lin et al. (2017) examined a product recommendation network and found that both network diversity and stability are significantly associated with product demand. These findings suggest that articles that are disadvantaged in terms of network position may receive less attention, further limiting their future evolution.

The third stream of research focuses on attention flow or spillover in information networks and policies to optimally leverage spillover. West and Leskovec (2012) used an experimental game to study the dynamics of attention flow in Wikipedia through the lens of goal-oriented search. Kummer (2014) studied spillovers from articles that are featured on the home page of German Wikipedia. Wu and Huberman (2007) study the dynamics of attention to articles on the news aggregator Digg.com and show how attention to articles decays with their novelty. Several works have focused on how content, and particularly perception of its importance, can drive attention. Salganik et al. (2006) conducted a series of randomized online experiments to determine the impact of music track ranking on consumption. Muchnik et al. (2013) demonstrated that perceived popularity of comments not only attract attention and additional votes but can lead to herding phenomena where "likes" beget additional "likes." Carmi et al. (2017) carried this idea further and studied how demand shocks generate not only attention but attention spillover in the product recommendation networks of Amazon.com, yielding substantial benefits to downstream recommended products. Finally, Aral et al. (2013) studied seeding strategies for policies that leverage spillover in the context of social networks. These studies suggest that attention spillover has a significant impact on real-world outcomes and policies that leverage spillover can be beneficial.

While all three streams of research have enriched our understanding of knowledge production and consumption in information networks, much of the work on open collaboration platforms like Wikipedia relies on endogenous observational data, making it difficult to draw valid causal conclusions. In addition, existing work has focused only on the local direct effect of attention spillover. It has not addressed how heterogeneous characteristics of articles moderate spillover. Nor has it considered the systemic effect of spillover and its broader policy implications.

Yet, a rigorous understanding of the dynamics at play in the Wikipedia network and collaborative information systems in general is indispensable for understanding how information evolves in these systems. Such an understanding is vital to the mission of global empowerment through open knowledge production and dissemination. Moreover, it is an important precursor to the development of sound policies, such as incentivizing contributions to achieve more robust coverage[7]. Randomized controlled experiments are the gold standard for causal inference but are difficult to conduct on platforms like Wikipedia. Apart from the technical challenges and ethical concerns associated with experiments in this context, the continued survival and operations of these platforms depend completely upon the community of contributors, who are highly sensitive to

[7] https://meta.wikimedia.org/wiki/Research:Increasing_article_coverage

sudden and unvetted policy changes. On the other hand, natural experiments that create exogenous variation in otherwise endogenous relationships can also permit valid causal inference.

In this study, we leverage a natural experiment to examine how exogenous content contributions to a Wikipedia article affect future activities surrounding the article in terms of both pageview dynamics and editing behavior. More interestingly, we examine how the attention an article attracts can spill over to other articles it links to and hence further propagate through the network. Furthermore, we consider the broader policy implications of spillover. We conduct policy simulations to understand how spillovers concentrated in the clusters of the network, which we term *attention contagion*, could impact the evolution of Wikipedia as a system and how it could be harnessed and incorporated into policies to address impoverished regions in information networks.

The goal of the policy simulation is to integrate our findings into an empirically-calibrated attention diffusion model and to guide policy decisions through the analysis of counterfactuals. While the platform can answer some policy questions through analysis of observational data and through experimentation, many relevant counterfactuals for policy recommendation are not directly recoverable from direct estimates. They may be too costly or even impossible to test. In our context, interpreting the spillover effect of individual articles on the whole system is not straightforward. In particular, the effect of spillovers might be amplified when editorial efforts are directed at a group of interconnected articles. The key idea behind the policy simulation approach is that reduced-form analysis is used to estimate parameters of a model of the system so that the model can be used to extrapolate findings to more complex or more interesting policies, at the cost of imposing additional model assumptions (Taylor and Eckles 2018).

Our study provides three major contributions. First, we confirm and obtain causal estimates of the feedback loop between contribution and attention. We find that contribution drives sustained increase in future attention (12% on average, with stronger impact for more significant contributions) and future contributions (3.6 more edits and 2 more unique editors over a 6-month period). Second, we determine the article and network characteristics that most amplify spillover or attention contagion. We find that spillovers have the most impact (as much as 22%) for less popular articles that are hyperlinked from focal articles through newly created links. Third, we provide insights from comparisons of policies to address information-impoverished regions of the network based on analytic derivation and empirically-calibrated simulations. We demonstrate that a policy designed to leverage attention contagion can yield substantial increases in attention (as much as a twofold) to impoverished regions of information networks. These results are directly relevant to concerns of societal equity and have managerial importance for collaborative information platforms.

Natural Experiment and Data

Since 2010, the Wikipedia Education Foundation has been collaborating with university course instructors to encourage students in the United States and Canada to expand and improve Wikipedia articles through course assignments. The mission of this endeavor is to cultivate students' skills such as media literacy, writing, and critical thinking, while leveraging student effort to fill content gaps on Wikipedia. Since its launch, university instructors participating in the program have guided their students to add content to approximately 46,000 course-related articles on Wikipedia. About 35,000 students have contributed more than 35 million words to Wikipedia, equivalent to 22 volumes of a printed encyclopedia. These student-edited articles have collectively received 282 million views by the end of 2017[8].

In this study, we leverage the exogenous content contributions that result from this campaign to enrich our understanding of the dynamics in open collaboration platforms. The identification derives from the assumption that the content contributions by students are exogenous to the natural evolution of the articles and would not have occurred during the same time period in the absence of the Wiki Education campaign. This is likely to hold for two reasons: first, many of the treated articles pertain to topics that do not naturally relate to current events (e.g., detailed topics in fundamental sciences, such as properties of molecules, etc.); Second, the timing of contribution is exogenous. The content addition occurs during a fixed time period that corresponds to an arbitrary class period – that is to say that the contribution would not have occurred during the same

[8] https://wikiedu.org/changing/wikipedia/

time period in the absence of the assignment. We seek to learn three things from this natural experiment: First, whether efforts that focus on developing underdeveloped pages can lead to long-term, sustained impact; Second, more generally, how contribution and attention dynamically interact and how this interaction depends upon article attributes; Third, whether and to what extent attention propagates through the information network, i.e. the phenomenon of *attention contagion*. Finally, we seek to combine insights in order to synthesize and assess policies that address information poverty and skewness.

For this study, we collected all the articles that received content contribution from students through this campaign in the year of 2016[9]. For each article, we retrieved its title, URL, the time period of the course (i.e., the shock period), and the number of characters added to the article by the assigned student from the website of Wiki Education Dashboard[10]. In our analysis, we retain only articles that existed prior to the campaign (excluding new articles created by students) and those that received substantive contributions (of at least 500 added characters during the shock period). This leaves us with 3,296 unique treated articles in the sample.

To assess the impact of the content shock, we consider the number of pageviews of an article, a widely-used measure of information consumption. In addition, we parse the complete revision history of each article to obtain the time series of edits and authorship (i.e., the number of unique editors that worked on the article over time). Both

[9] Wikimedia changed their measurement of "pageviews" in May 2015 to better filter out bot traffic and incorporate the visits from mobile devices. Looking at the articles edited in 2016 guarantee we have a consistent measure of pageviews in the 6 months before and after the content shock.
[10] https://dashboard.wikiedu.org/

the pageviews and revisions are collected through the public API developed and maintained by the Wikimedia Foundation[11].

Matching and control group

Rates of Wikipedia content creation and consumption are subject to seasonality and other temporal patterns. A simple comparison of quantities of interest (e.g., pageviews and revisions) before and after the content shock may therefore be misleading. Observed changes can be attributed to alteration of the page content, but also to naturally occurring trends. Statistical modeling techniques alone are often insufficient to fully account for seasonality and other complex temporal patterns of article activity. We address this issue by constructing a sample of treated and control articles, matched across multiple attributes. The control group is used to identify the average outcomes corresponding to the counterfactual state that would have occurred for articles in the treatment group had they not received the content contribution during the shock period.

The control group is chosen via the following procedure. First, we pick candidates for the control group by choosing a random sample of 100,000 Wikipedia articles that did not receive content contribution from students. Next, we define the hypothetical shock period for each control article by randomly sampling from the pool of shock periods of treated articles and measure the pre-shock article characteristics for control articles. Finally, we use Coarsened Exact Matching (CEM) (Iacus et al. 2012) based on each article's pre-shock characteristics of tenure, size and popularity (calculated

[11] https://www.mediawiki.org/wiki/API:Main_page

based on average historical pageviews) to obtain a matched sample by pruning articles that have no close match in the treated and control group. We opt for a k-to-k matching solution (i.e., an equal number of treated and control units), which is accomplished by pruning observations from a CEM solution within each stratum until the solution contains the same number of treated and control units in all strata. Pruning occurs within a stratum through nearest neighbor selection using a Euclidean distance function.

Matching is a frequently used technique for drawing causal conclusions from observational data based on the assumption of selection on observables (Ho et al. 2007, Rosenbaum and Rubin 1983). It emulates a randomized experiment, after the data has been collected, by constructing a balanced dataset in which samples in the control group are similar to the samples in the treated set in observed characteristics. We confirm that the constructed control group closely mirrors the treatment group in seasonality and natural time trends. This can be verified in the model-free plots of pageviews over time that we will provide later and by comparing article attributes in each group as displayed in Table 1. The average of all three covariates are very close across groups and t-tests fails to reject the null hypothesis that they have the same mean value. In addition, this between-group panel research design lends itself neatly to a standard Difference-in-Difference estimation of the effect of content contribution.

Table 1: Balanced Check for Matched Sample

	Size (characters)	Popularity (weekly pageviews)	Tenure (weeks)
Control	16,228	1,575	506
Treatment	16,255	1,574	506
t-test (p-value)	0.70	0.93	0.51

Table 1 illustrates the quality of our matching procedure. It compares pre-shock characteristics of articles in the matched groups. T-tests indicate that we cannot reject the null hypothesis that articles in treatment and control group have the same mean across all three characteristics.

The above procedure yields 2,766 pairs of matched treated and control articles. For each article, we construct a panel of weekly pageviews from 26 weeks before the shock to 26 weeks after (excluding the shock period itself). Our final sample consists of a balanced panel of 52 periods for 5,532 articles or 287,664 observations at the article-week level. Our results are robust to other matching procedure choices. For example, we evaluated an alternative matching procedure that incorporates matching on article topic and find that the direct effect results are qualitatively similar with only small changes in the magnitude of effect sizes. In addition, we also demonstrate that our results are robust to matching based on network characteristics of articles (see Appendix for further details).

Links and hyperlink articles

Because we are also interested in attention spillovers from treated articles to downstream hyperlinked articles, we parse content revisions to retrieve the outgoing hyperlinks from focal articles. Following the links, we retrieve all articles linked to by treated and control articles. There are millions of such hyperlinked articles. To avoid confounds that may arise from multiple exposures to the treatment, we retain only hyperlinked articles that are linked to from one and only one treated article (Walker and Muchnik 2014). For parity, we treat articles downstream of control articles in the same manner. This allows us to obtain a clean estimate of the spillover effect from each link. This procedure yields 131,974 hyperlinked articles that are downstream from directly treated articles. The spillover treated and spillover control articles constitute our sample for analyzing the spillover effect of the content contribution. This is illustrated in Figure 1.

Figure 1: Research Design - Direct Effect and Spillover Effect

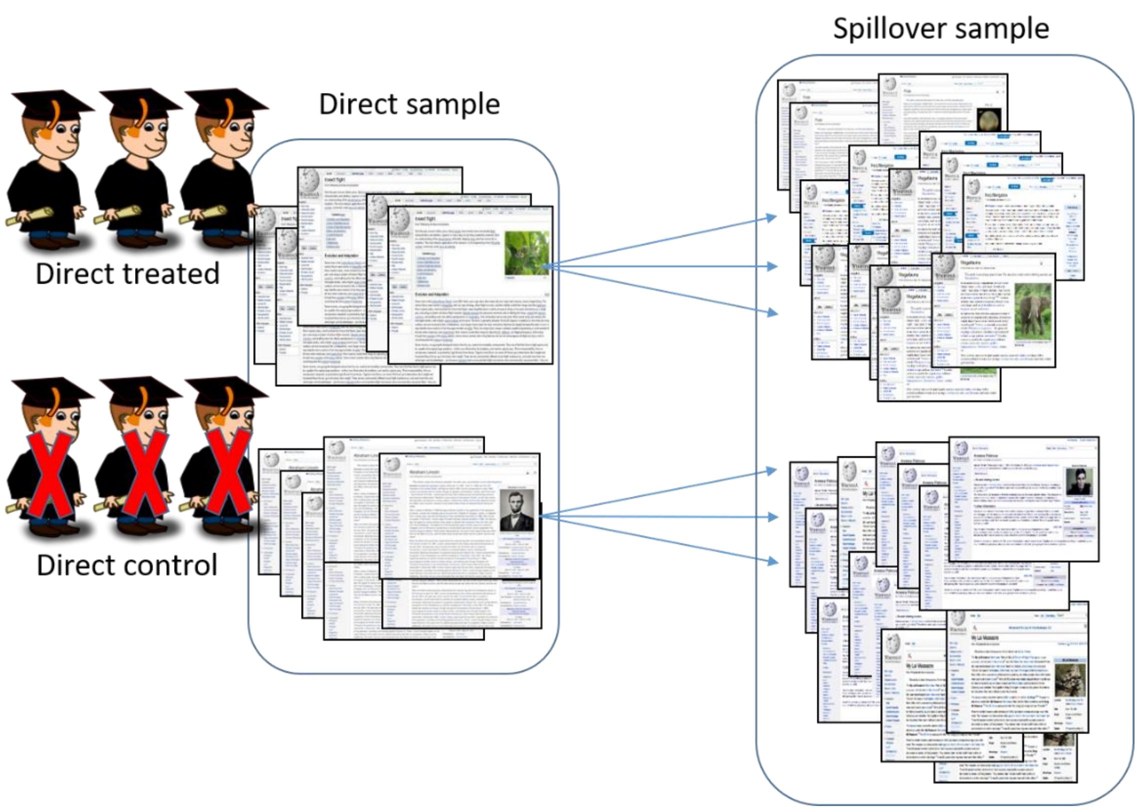

Figure 1 illustrates the direct treated and direct control articles, which constitute our matched sample for analyzing the direct effect of the treatment. Similarly, the spillover treated and spillover control articles constitute our sample for analyzing the spillover effect of the content contribution.

Model-free evidence

In this section, we present model-free evidence regarding the direct and spillover impact of the content shock, in terms of both pageview dynamics and editing behavior. A model-free examination of the evidence can reveal important effects while avoiding modeling assumptions.

Pageviews dynamic

Because articles are highly heterogeneous, they experienced a large variance in activities (such as pageviews) even prior to treatment, a phenomenon that is typical for complex social systems (Muchnik et al. 2013) To compensate for large baseline variation, we scaled pageviews for each article relative to its own pre-shock popularity, which is computed as average weekly pageviews over 26 weeks (about 6 months) prior to its shock period[12]:

$$scaledPageview_{i,t} = \frac{pageview_{i,t}}{preShockPopularity_i} \ (eq\ 1)$$

Where $preShockPopularity_i = 1/26 \sum_{\mu=1}^{26} pageview_{i,\tau-\mu}$ and τ is the week when the content shock begins for article i. Because courses in our sample begin at different weeks and have different durations, we align their start dates and exclude the duration of shock period itself from the analysis. We consider relative time before or after the shock. Figure 2 plots the mean and standard deviation of weekly scaled pageviews in the 6

[12] Note that this normalization simply scales the time series of pageviews of each article by a constant. Examination of the model-free evidence for scaled and unscaled pageviews reveals that this scaling is appropriate.

months prior to and after the shock period for treated and control articles.

Figure 2 Impact of Content Shock on Pageviews

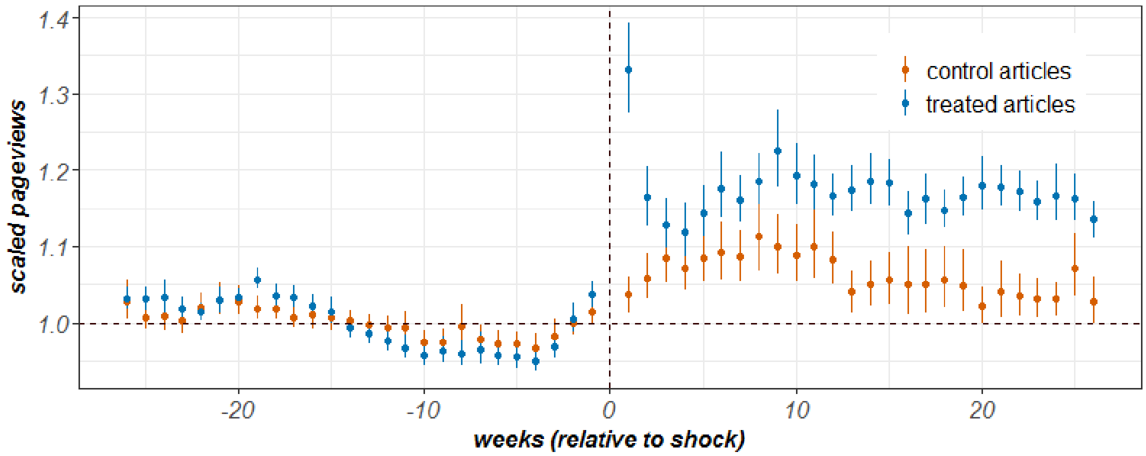

Figure 2 displays the pageviews dynamics for articles in the treatment and control group. Time is measured relative to the shock period (which is excluded), up to 26 weeks before and after. Dots and whiskers represent the mean and standard deviation of scaled pageviews in each bin, respectively.

This model-free view of the data displays a clear seasonal trend for both treatment and control group articles, indicating the need for careful construction of a control group as a counterfactual. Prior to the shock, articles in the control group mimic the time trend of those in the treatment group well, highlighting the success of our CEM procedure. We can also see the significant and relatively long-lasting impact of the treatment on post-shock pageviews. Treated articles received approximately 10% more traffic than control articles, and this effect persisted for at least 26 weeks after the contribution shock. Evidently, Wikimedia's campaign efforts to develop underdeveloped pages both worked and had a relatively long-term impact, suggesting the potential for a policy approach to fill impoverished regions in Wikipedia's information network.

Figure 3 plots the mean and standard deviation of weekly scaled pageviews in the 26 weeks prior to and after the shock period for articles in the spillover treated and spillover control groups. While pageviews of spillover treated articles seem to exceed those of spillover control articles after week 10, it is unclear from this model-free evidence alone whether the effect is significant. It should be noted that there is little doubt that spillover of attention occurs on Wikipedia — this can be seen explicitly from published clickstream data of actual traffic flowing over hyperlinks from one article to another (see *Sources of Increased Attention* in section 3 for further discussion). What is unclear is the extent and heterogeneity of treatment spillover effect and whether it can be teased out. Downstream articles, by virtue of being selectively linked to, tend to be more popular and have a larger variance in pageviews, suggesting that the effect, if it exists, may require econometric strategies to uncover. For example, it could be the case that the spillover is significant for only less popular articles, which may themselves be underdeveloped.

Figure 3 Spillover Effect on Hyperlinked Articles

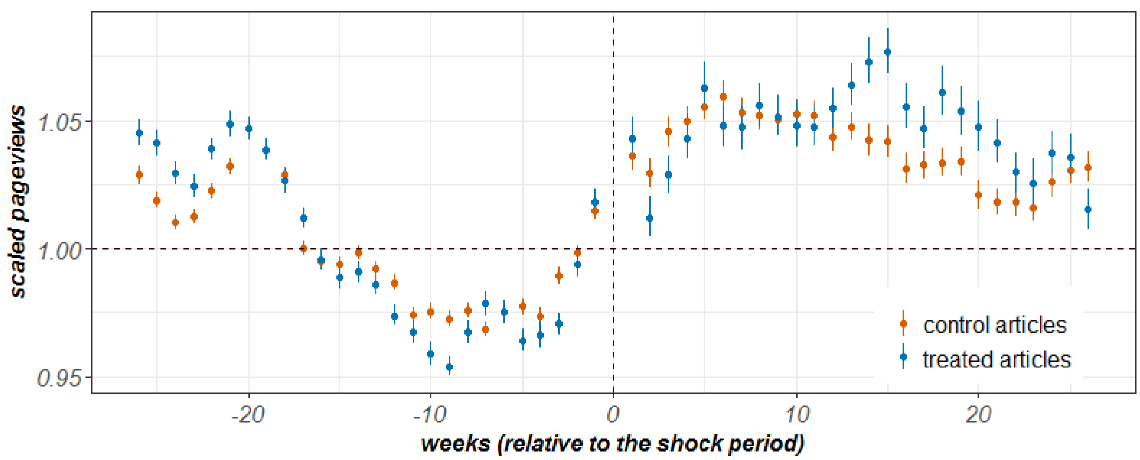

Figure 3 displays the pageviews dynamics for articles to which treatment and control group articles link. Time is measured relative to the shock period (which is excluded), up to 26 weeks before and after. Dots and whiskers represent the mean and standard deviation of scaled pageviews in each bin, respectively.

During the shock period, students also added new links to downstream pages, as part of their contribution efforts. Newly added links are interesting in terms of attention spillover, because they may function to "open the valve" of attention flow between articles. Intuitively, old links can convey only changes in attention to downstream articles. In contrast, a newly added link can convey the totality of attention to downstream articles. This is illustrated in a simple conceptual model:

$$\Delta pageviews_{i,j}^{spillover} \propto pageviews_i * newLink_{i,j} + \Delta pageviews_i^{treated} \quad (eq\ 2)$$

Where $newLink_{i,j}$ can be thought of as an indicator variable (equal to 1 for new links, and 0 for old links). This suggests that attention spillover may be more clearly visible in model-free evidence if we look only at newly-linked downstream articles (i.e., those downstream articles that were linked to from treated articles *during the shock period*). Figure 4 is similar to Figure 3 but distinguishes spillover populations by whether the link from the directly treated article was pre-existing (old link) or was added during the shock period (new link). New link articles in the spillover control group are not displayed because they did not receive sufficient new links during the shock period.

Figure 4 Spillover Effect – New Link

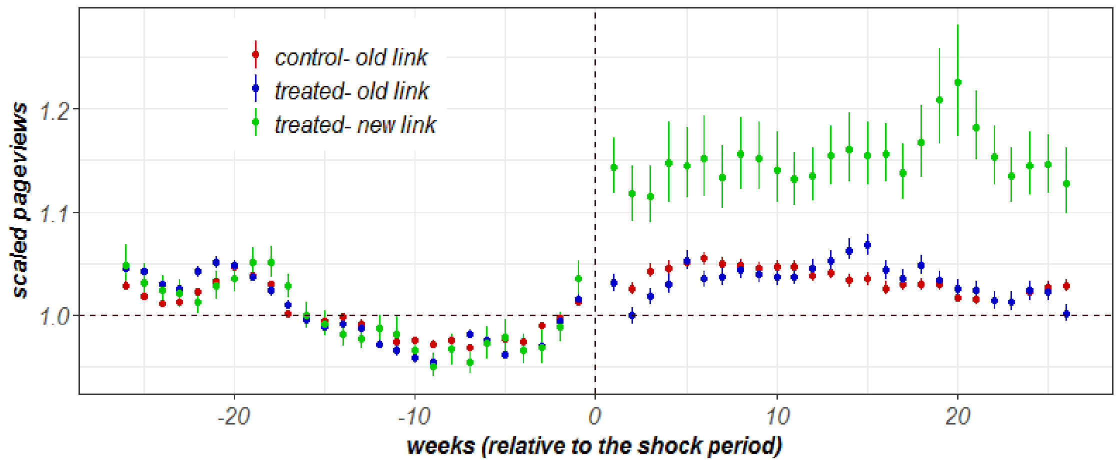

Figure 4 displays the pageviews dynamics for hyperlink articles based on whether the downstream article is connected through a new link or an old link. The time period is from 26 weeks prior to the contribution shock to 26 weeks after. Dots represent mean value of scaled pageviews in each bin and whiskers represent the corresponding standard deviation.

The model-free plot of the spillover effect for new links confirms our reasoning. Spillover of attention across newly created links is clearly significant and the temporal pattern of spillover closely follows the pattern of the post-shock pageviews of directly treated articles. Compared to an old link, a new link can convey an additional 15% pageviews to target articles on average.

Editing behaviors

Prior research has suggested that content contributions are self-promoting – that, in addition to boosting future attention (consumption), they also drive future contributions. We examine model-free evidence to determine whether the exogenous content contribution to articles leads to future contributions to those articles. We retrieved the full revision history of all articles in our sample and constructed two measures of editing behavior, the number of total edits and the number of unique editors in the six months prior to and after the shock period for each article. Because contribution behavior is relatively rare, we collapse the time series into a “pre” and “post” period. For each article, we look at the editing behavior before and after the content shock and their difference across treatment and control groups.

Table 2: Editing Behavior Before and After the Shock Period

	Total edits			Unique editors		
	Before	After	Δ	Before	After	Δ
Control	11.2	11.3	0.1	6.2	6.5	0.2
Treatment	11.7	15.4	3.7	6.7	9	2.2
t-test (p-value)	0.45	-	<1e-9	0.36	-	<1e-16

Note: The values display under the columns “Before” and “After” are counts of total edits and unique editors in the 6 months before and after the shock period. Δ = After - Before. The values in the row “t-test” are p-values from a two-sided t-test of the null hypothesis that control and treatment group have the same mean.

Editing behavior is similar across treatment and control groups during the pre-shock period, as expected: t-tests fail to reject the null hypothesis that the treatment and control group have the same mean number of total edits ($p = 0.45$) and number of unique editors ($p = 0.36$) prior to the shock. For treated articles, in the 6-month period after the contribution shock, the number of total edits increased by 3.7 (p<1e-9) and the number of unique editors increased by 2.2 persons (p<1e-16). In contrast, control group articles did not experience any significant increase in number of total edits or number of unique editors. These results confirm that exogenous content shocks significantly drive future editing behavior.

Overall, the model-free evidence confirms that exogenous content contributions drive future attention and editing behavior and that spillover of attention occurs significantly for newly added links. To capture the impact of varying intensity of treatment and heterogenous treatment impact, we turn to econometric modeling.

Empirical Methods

Direct Impact of Contribution Shock

In this section, we use econometric models to infer how differing intensities of content shocks affected treated articles contingent on article characteristics, in terms of future content consumption and future editing behavior. We further investigate the source of attention increases to treated articles by analyzing the internal and external inbound traffic to treated pages.

Content Consumption

We estimate the average treatment effect on the treated (ATT) for content consumption using the following simple specification as the baseline model:

$$Pageviews_{it} = \alpha PostShock_{it} + \gamma_i + \delta_t + e_{it} \quad (eq\ 3)$$

where i is a Wikipedia article and t indexes the week. The dependent variable $Pageviews_{it}$ is the scaled pageviews for article i at week t as defined in eq 1. For brevity, we have defined $PostShock_{it} = PostShockPeriod_t * Treatment_i$, a dummy variable equal to 1 if the period t is after shock and the article i is a treated article, and 0 otherwise. We include article and week fixed effects (γ_i and δ_t) to account for article level heterogeneity and common pageviews trends over time on the platform. Equation (3) estimates a simple Difference-in-Difference model of the impact of exogenous content contribution.

However, content contribution may have different impacts on articles with different characteristics. For example, less popular articles (with less average attention prior to the shock) may have been more or less affected. Article characteristics include article length, tenure and popularity (defined as average pageviews over the 6 months period before the shock). Moreover, not all treated pages received equal contributions during the shock period. Actual contributions varied significantly across treated articles, ranging from hundreds to tens of thousands of characters added through the course of student edits. To account for varying treatment intensity and to allow for heterogeneous treatment effects, we estimate the following model:

$$Pageviews_{it} = \beta_1 PostShock_{it} * \log(charCount_i) + \beta_2 PostShock_{it} * X_i$$
$$+\gamma_i + \delta_t + e_{it} \quad (eq\ 4)$$

where log $(charCount_i)$ is the logarithm of number of characters added to article i by a student during the shock period[13]. It represents the variation of treatment intensity. X_i is a vector of article characteristics measured before the content shock, including article tenure, size, and popularity. To provide better interpretability of model estimates and to avoid the assumption of linearity, we bin these three continuous variables to low and high levels by their median value and include dummy variables that are equal to 1 when the value is high and 0 otherwise (i.e., older article, longer article, and more popular article) in the vector X_i. Diagnostic tests show that two bins for our continuous variable is a reasonable choice (see Appendix for more detail). The interaction term of $PostShock_{it}$ and X_i allows us to investigate heterogeneous treatment effects. We retain article fixed effects and week fixed effects. The parameters of interest are β_1 and β_2.

[13] For articles in the control group, the value of log $(charCount_i)$ is set to zero.

We use linear regression to estimate the above models and results are reported in Table 3. Because we scale the pageviews of each article with respect to its average pageviews over the six months prior to the shock, all estimates can be conveniently interpreted as the percent changes of pageviews relative to their pre-shock average. Following the suggestion of Bertrand et al.(2004), all reported standard errors allow for arbitrary serial correlation across time and heteroscedasticity across articles to properly gauge the uncertainty around the estimates for serially correlated outcomes in panel data.

Overall, we find post-shock pageviews for treated article increased by 12% on average. The magnitude of the treatment effect is positively correlated with treatment intensity and the impact is stronger for articles that are younger and less popular. The effect is both economically and statistically significant. Based on the model estimates in (3), a relatively young and less popular article with 6000 characters added (the average number of characters added for treated articles in our sample) during the shock period experienced a 25% boost in post-shock pageviews. The impact is even larger for similar articles that received a more intense treatment.

Table 3: The Impact of Content Contribution on Consumption

	Scaled pageviews		
	(1)	(2)	(3)
PostShock	0.119***		
	(0.017)		
PostShock*log(char count)[14]		0.035***	0.065***
		(0.005)	(0.008)
PostShock*old article			-0.041*
			(0.024)
PostShock*popular article			-0.142***
			(0.025)
PostShock*long article			-0.015
			(0.025)
Article fixed effect	Yes	Yes	Yes
Time fixed effect	Yes	Yes	Yes
Observations	287,664	287,664	287,664
Adjusted R^2	0.122	0.122	0.124

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

We perform diagnostics to assess our modeling assumptions in terms of linear interaction effects and common support. Results show that both assumptions are satisfied. For robustness, we also estimated alternative specifications. Using linear regression, we drop article fixed effects γ_i and retain only a simple treatment indicator, and all estimates are similar (see the Appendix for more details).

[14] Note that in models 2 and 3, we include PostShock*log(char count) and exclude a bare PostShock term because log(char count) captures the intensity of a treatment (and every article that received a contribution as a consequence of treatment had some number of characters added).

Editing Behavior

Beyond the impact on attention, we are also interested in whether exogenous content contributions spur future editing behavior. Because editing behavior is typically sparse for a Wikipedia article, for modeling purposes, we collapse the time series into just “pre” and “post” periods for the 6 months prior to and after the contribution shock. For each article, this yields two 6-month time periods during which we count the number of total edits and number of unique editors and these comprise the dependent variables. Compared to alternative approaches (such as multistage, zero-inflated models), this transformation permits a simpler linear model which retains interpretability and avoids more restrictive modeling assumptions (such as distributional assumptions on the error term that are required by Poisson or Negative Binomial regression). In addition, as suggested by (Bertrand et al. 2004), the “pre” and “post” time series collapse allows us to obtain a consistent estimator for the standard errors of the treatment effect in the Difference-in-Difference model. The models estimated here are similar to models in equation (3) and (4) for content consumption, apart from the time period collapse and the exchange of the dependent variable for editing behavior. For the sake of interpretability, we report the results from a linear regression, but results from Poisson regression and Negative Binomial regression are qualitatively similar (see Appendix for details).

Table 4: The Impact of Contribution Shock on Future Editing behavior

	Number of total edits			Number of unique editors		
	(1)	(2)	(3)	(4)	(5)	(6)
PostShock	3.596***			1.996***		
	(0.855)			(0.243)		
PostShock *log(char count)		1.173***	1.186***		0.640***	0.606***
		(0.229)	(0.234)		(0.068)	(0.065)
PostShock *old article			1.446			0.691**
			(0.957)			(0.339)
PostShock *long article			-1.840**			-0.829***
			(0.926)			(0.305)
PostShock *popular article			0.241			0.333
			(0.856)			(0.326)
Article Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,964	10,964	10,964	10,964	10,964	10,964
Adjusted R^2	0.63	0.63	0.63	0.82	0.82	0.82

Notes:
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

As we can see from Table 4, the contribution shock has a significant impact on future editing behavior in terms of both number of total edits and number of unique editors. Based on model estimates from column (1) and (4) in Table 4, an article that received content contribution in the shock period had approximately 3.6 more edits and 2 more unique editors in the 6 months after the shock period, compared to articles that did not receive exogenous content contribution. Similar to our findings for content consumption, the magnitude of the treatment effect increases with treatment intensity. Based on the estimates from column (2) and (4), an article with 6000 characters added during the shock period attracts 4.5 more edits and 2.5 editors in the 6 months post-shock period. As for heterogeneous treatment effects, the most significant factor we weaker impact for articles that already have a substantial amount of content.

Sources of Increased Attention

Both model-free results and estimates from statistical models confirm that exogenous contributions to articles drive future attention. But from where does this increased attention originate? In general, articles can receive attention directly from external sources (e.g., traffic arriving to an article from outside of the information network, such as through search engine discovery or links from external websites) and internal sources (e.g. traffic flowing to an article from another upstream article). This distinction is interesting and meaningful from a policy perspective as some articles may act to pull attention into the information network from external sources, thereby increasing the overall attention to the platform. Articles also play a role in the redistribution of attention throughout the platform, which is relevant from the standpoint of information equity. An article's role in the flow of attention on the information network is illustrated in Figure 5.

Figure 5 Attention Flow on Wikipedia Network

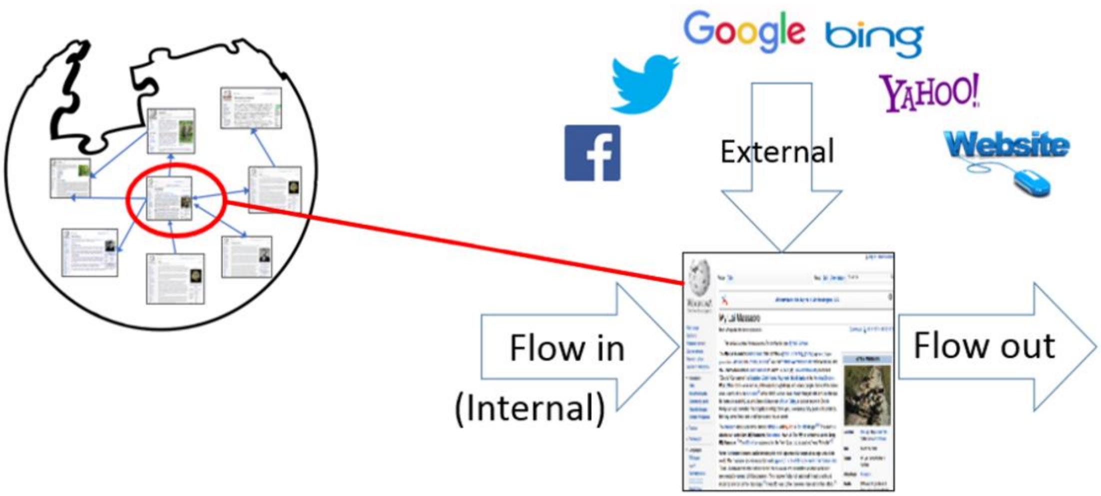

Figure 5 illustrates the flow of attention on information networks with respect to a particular article in terms of flow in (internal and external) and flow out.

For many large-scale real-world information systems, we cannot directly observe the detailed flow of attention (traffic). However, recently released data of monthly Wikipedia clickstream[15] snapshots provide exactly this level of detail for all Wikipedia articles. The clickstream data show how users arrive at an article and what links they click on within the article over the course of a given month, aggregated at the article level. They contain counts of (referrer, resource) pairs extracted from the Wikipedia HTTP request logs, where a referrer is an HTTP header field that identifies the address of the webpage that linked to the resource being requested. In other words, the clickstream data gives a weighted network of articles and external sites, where the weight of each edge corresponds to the traffic flow along that edge. These counts are aggregated at the monthly level and any (referrer, resource) pair with greater than 10 observations in a month are included in the dataset. To give a sense of the scale of the data, the August

---

[15] https://meta.wikimedia.org/wiki/Research:Wikipedia_clickstream

2016 release contains 25.8 million (referrer, resource) pairs from a total of 7.5 billion requests for about 4.4 million English Wikipedia articles. Figure 6 displays an example from the Wikimedia website, which illustrates incoming and outgoing traffic to the page "London" on English Wikipedia.

Figure 6 Incoming and Outgoing Traffic for an Example Page: London

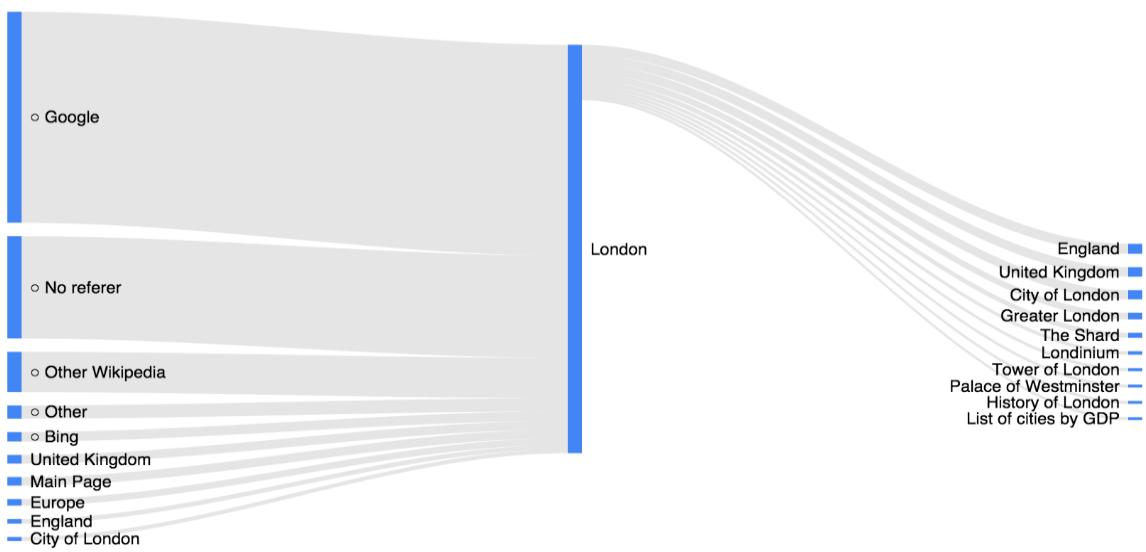

Figure 6 displays the sources of incoming and outgoing traffic for the "London" Wikipedia article, as determined from the clickstream monthly data snapshots provided by the Wikimedia foundation.

We leverage this data to shed light on the sources from which increased attention originate. The clickstream data snapshots are only available for a limited number of months during the period of our natural experiment. To look at the change of traffic flow, we need to compare snapshots before and after the shock period. Fortunately, the Wikimedia Foundation released clickstream snapshots for both August 2016 and January 2017, which are just before and after articles were treated in the fall semester of 2016.

For each article, we calculate its total inbound traffic (combined internal and external traffic arriving at the article), total outbound traffic (traffic leaving the article), internal inbound traffic[16] (traffic flow to the article from other articles in the network) and external inbound traffic (traffic flow to the article from a search engine or other external website). We use CEM to ensure that articles in the treatment group and control group are comparable across all traffic measures prior to the start of the natural experiment (i.e. in the August 2016 snapshot). The k-to-k CEM procedure leaves us with 1,017 articles in both the treatment and control group (see Appendix for distribution and balance checks for clickstream data).

First, we look at changes in network structure in terms of newly created incoming links. During the shock period, it is likely that links to articles in either the treatment or control group were created, either by student editors or as part of the natural evolution of the information network. Matching the 2,024 treatment and control articles in our sample with the clickstream data snapshots (for August 2016 and January 2017), we find that the number of active incoming links[17] for treated articles grew significantly faster as compared to control group articles. As we see in Table 5, articles in the treatment group received on average 0.9 more active links during the shock period (compared to 0.4 for articles in control group). New incoming links make an article more discoverable by creating new channels to capture attention flow within the network. These increased channels may explain how contributions ultimately drive attention.

---

[16] The link traffic only includes links from other Wikipedia articles. The link traffic from other website outside of the ecosystem of Wikipedia were classified under the “external traffic” category.
[17] We define an active incoming link as one that conveys at least 10 pageviews in a month. The monthly clickstream data snapshots filter out any (referrer, resource) pairs that do not meet this criterion.

Table 5: Number of Incoming Links

	Number of incoming links per articles		
	Before	After	Δ
Control	6.6	7.0	0.4
Treatment	6.6	7.5	0.9
t-test (p-value)	0.96	-	< 1e-15

Notes: The values display under the columns "Before" and "After" are the average number of incoming links per articles in the 6 months before and after the shock period. Δ = After - Before. The values in the row "t-test" are p-values from a two-sided t-test of the null hypothesis that control and treatment groups have the same mean.

Attention from external sources can also explain the attention increases we observed. To determine the extent to which observed attention increases derive from internal or external sources, we compare pre/post shock changes in internal, external, and total incoming traffic across treatment and control articles in Table 6. The control group serves as a counterfactual to account for natural fluctuations arising from seasonal or other pageview trends, leading to a simple DID style estimator:

Table 6: Incoming traffic breakdown

	Total incoming traffic			internal traffic ($T^{internal}$)			external traffic ($T^{external}$)		
	Before	After	Δ	Before	After	Δ	Before	After	Δ
Control	45.4	53.6	8.2	10.2	12.2	2.0	35.2	41.4	6.0
Treated	44.7	59.3	14.6	10.2	14.0	3.8	34.4	45.2	10.8
t-test (p-value)	0.85	-	0.01	0.97	-	0.05	0.80	-	0.03

Notes: The values display under the columns "Before" and "After" are the average traffic per article per day in the 6 months before and after the shock period. Δ = After - Before. The values in the row "t-test" are p-values from a two-sided t-test of the null hypothesis that control and treatment groups have the same mean.

From Table 6, we see that the total incoming traffic increased by 14.6 pageviews per article per day for the treatment group relative to 8.2 for the control group. The extra 6.4 pageviews can be interpreted as the Average Treatment Effect on the Treated (ATT), which is about a 14% increase relative to the pre-shock average. This result is consistent with our prior estimates, which were based on article-level pageviews data. Hence, we demonstrate the impact of content shock using two different data sources (clickstream data and pageviews data) and find similar effect sizes. We can also see that both internal and external sources conveyed increased attention, indicating that content contributions yield attention gains from within the information network and from without. We suggest that attention gains from external sources are likely the result of increased visibility of the articles in search engine results[18]. Modern search engine algorithms are clearly sensitive to the recency of content changes. Though we do not know the actual details of search engine ranking algorithms (proprietary information), more incoming hyperlinks to a page convey a higher ranking in ordinary PageRank. We define the ratio of internal to external traffic as $R(T) = T^{internal}/T^{external}$. New traffic has a higher ratio, ($R(\Delta T) = 0.4$) relative to the pre-shock ratio ($R\left(T_{Before}\right) = 0.3$), indicating that new traffic originates slightly more from internal sources.

Attention Spillover

The impact of content shocks is not limited to directly treated articles. Attention resulting from the shock can also spillover onto other downstream articles through the

[18] Search engines traffic dominates other external sources such as external websites in external traffic.

hyperlink network. Conceptually, we can think of the spillover as a dyadic relationship between each source (directly treated or control) and target article. As our consideration of model-free evidence showed, new links, which build bridges between source and target articles, seem to play a critical role in facilitating spillover. It also seems plausible that the popularity of source and target articles may moderate the extent of the spillover. We test these hypotheses with the following model:

$$Pageviews_{it} = \beta_0 PostShock_{it} + \beta_1 PostShock_{it} * stPopularity_i + \beta_2 PostShock_{it} * newLink_i + \beta_3 PostShock_{it} * stPopularity_i * newLink_i + \gamma_i + \delta_t + e_{it} \quad (eq\ 5)$$

Where i is a target article and t is the week. $stPopularity_i$ is a 2-dimension vector $(sourcePopularity_i, targetPopularity_i)$, representing the popularity of the source article (i.e., the treated article that received an exogenous content contribution) and the target article (that was linked to from the treated article), respectively. The indicator $newLink_i$ is equal to 1 if the link between source article and target article was added during the treatment period, 0 otherwise. The parameters of interest are $\beta_1, \beta_2, \beta_3$. We include each term in successive models gradually to investigate how they parcel out the overall spillover effect. The results are displayed in Table 7.

Table 7: The Attention Spillover of Contribution Shock

	Scaled pageviews			
	(1)	(2)	(3)	(4)
PostShock	0.008***	0.027***	-0.006	-0.005
	(0.003)	(0.006)	(0.004)	(0.007)
PostShock*popularTargetArticle		-0.013**		-0.004
		(0.005)		(0.005)
PostShock*popularSourceArticle		-0.016**		0.000
		(0.007)		(0.007)
PostShock*newLink			0.129***	0.148***
			(0.012)	(0.018)
PostShock*popularTargetArticle*newLink				-0.138***
				(0.023)
PostShock*popularSourceArticle*newLink				0.073***
				(0.023)
Article fixed effect	Yes	Yes	Yes	Yes
Time fixed effect	Yes	Yes	Yes	Yes
Observations	6,862,648	6,862,648	6,862,648	6,862,648
Adjusted R^2	0.104	0.104	0.104	0.104

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

We can see from column (1) of Table 7 that the overall effect (i.e., when averaged over all articles) is small but significant. This result is consistent with the model-free evidence and our intuition given the large heterogeneity across articles. Column (2) of Table 7 shows how the treatment effect varies with the popularity of source and target articles. Evidently, spillover from low popularity source articles to low popularity target articles yielded a 2.7% increase in pageviews ($p<0.01$). While this effect size may initially seem small, it is measured with respect to a single outgoing link from the treated article to one target article. In general, treated articles link to multiple downstream target articles, suggesting that the overall collective effect of spillover can be quite substantial.

Interestingly, spillover is enhanced when both source and target articles are less popular, which is a typical scenario for underdeveloped pages, particularly in informationally impoverished regions in the Wikipedia network.

A more interesting insight emerges when we consider whether the link between source and target articles was new. Surprisingly, for new links, the impact of the spillover can be as large as around 13%, which is close in magnitude to the average direct effect. As illustrated in our discussion of model-free evidence, the rationale is that a new link can "open the valve" between source and target article and convey both the preexisting and increased attention from the source to the target. We note that old links clearly convey attention (as the clickstream data illustrate). However, they convey only increased attention from the source to the target and we lack the statistical power to see it directly in this model. Finally, the attention spillover is even larger (14.8%) for new links between less popular source and target articles. As underdeveloped regions of information networks likely satisfy all these criteria (i.e. low popularity of articles and lack of preexisting link structures between articles), policies that focus on promoting such regions can benefit from strategies that harness spillover.

Policy Simulation of Attention Contagion

Our spillover results indicate that attention shocks in Wikipedia have a local network effect. Articles in the system benefit when upstream articles receive attention. Some spillovers direct attention to downstream articles that already receive significant exposure. On the other hand, some of this attention may increase exposure to underdeveloped articles. This begs the question: By focusing attention on connected sets of underdeveloped articles, can we optimally harness spillovers in order to redirect attention to articles that would benefit the most from increased exposure?

To better understand this question, we conduct policy simulations in which we integrate our findings from the econometric estimates into an empirically-calibrated attention diffusion model and to guide policy decisions through the analysis of counterfactuals. We propose a policy in which editors are encouraged to focus their editorial efforts on a set of targeted underdeveloped articles that are intimately related to one another, in order to harness attention contagion and maximize joint exposure. Targeted sets of related articles will be well-connected either at the outset (i.e., a set of stub articles that are already well-connected but remain underdeveloped) or will become well-connected as a consequence of directed editorial efforts. That is, the links between sets of related articles need not exist prior to being edited but can arise as a consequence. The rationale is that attention spillovers to underdeveloped articles are more valuable to the platform (in terms of the information equity that they convey) than spillovers to articles that are already well-developed.

Intuition – a Mean-field Estimation

We begin by providing an intuition for how network structure can impact attention spillover using a mean-field estimation. To represent a set of related and highly connected articles in a manner that is simple, we consider network cliques, defined as a set of n completely connected nodes in a network. To demonstrate our intuition, we analytically calculate the spillover in cliques of size n using mean-field assumptions.

For an n-clique, assuming each node receives direct traffic T and where spillover over a single step is given by $T_{spillover} = fT$, the total spillover exposure gain is given by: $\sum_{\boldsymbol{k=2}}^{\boldsymbol{n}} \frac{\boldsymbol{n}!}{(\boldsymbol{n-k})!} f^{\boldsymbol{k-1}}$. The summand represents all partial permutations of a set of at k nodes, describing the paths of length $k-1$ that successive spillovers take (each contributing a multiplicative factor of f) from each starting node to each other ending node. Figure 7 displays the total spillover gain for all articles in the clique (i.e., the total additional exposure gained from spillover from each article in the clique onto all other articles).

Figure 7 Mean Field Estimate of Total Spillover to a Clique

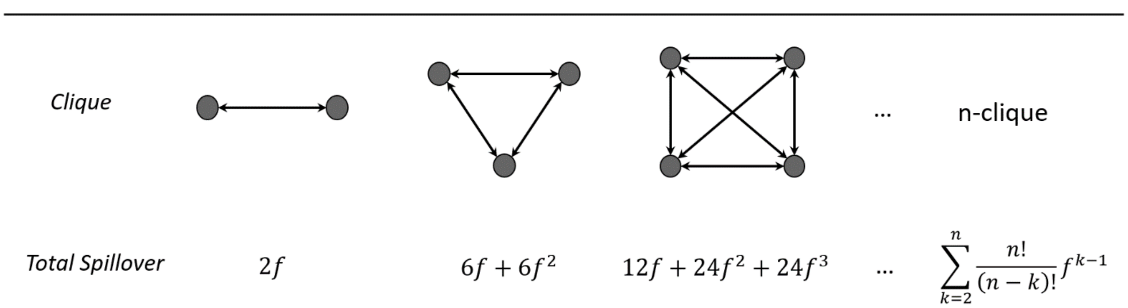

Note. For each clique shown, we calculate the mean field estimate of the total spillover to all nodes in the clique under the dynamic process described in the text.

For example, for a mean spillover of $f = 0.10$ and for cliques of sizes n=3, 4, 5, the total spillover exposure gain is 0.66, 1.46, and 2.73, respectively, as measured in units of proportion of incident direct traffic. This estimate assumes constant spillover (f), and equal traffic from any node in the clique to any other, which is unlikely to hold in the real world. Fortunately, we can relax these assumptions by using exact and fine-grained data on traffic flowing on all links in Wikipedia and traffic to all pages from external sources (e.g., traffic from search engines that arrive at Wikipedia pages) from the monthly Clickstream snapshots[19]. We leverage this data to estimate spillover and assess policies designed to capture spillover through empirically-calibrated simulations.

Diffusion Simulation

Our mean-field estimation is useful to obtain stylized estimates of policies that focus attention on clusters of well-connected articles and to develop an intuition about why this might work, but it does not account for real-world heterogeneity in actual traffic flow on the links between articles. To address this, we test policies more realistically and comprehensively through simulations of traffic flow that arise from attention perturbations. We define perturbations as increases in incident traffic from external sources. These policy simulations make use of highly detailed clickstream data for calibration, to ensure that traffic flow changes follow pathways in proportion to real-world patterns on Wikipedia. To accomplish this, we use a generalization of the

[19] Ellery Wulczyn, Dario Taraborelli (2015). Wikipedia Clickstream. https://meta.wikimedia.org/wiki/Research:Wikipedia_clickstream

personalized PageRank algorithm[20]. PageRank is widely recognized as one of the most important algorithms used for network-based information retrieval. It represents traffic flow as a random walk process on the information network, and is given in the iterative form by:

$$\vec{r}_{t+1} = (1-\alpha)\vec{r}_0 + \alpha G \cdot \vec{r}_t \qquad (eq\ 6)$$

Where $\vec{r}_t$ is a vector of the traffic (attention) landing on article i for the t-th iteration of the diffusion process; $\vec{r}_0$ is a vector of the initial distribution of traffic or whenever the process involves "hopping" rather than following a hyperlink from an article to a downstream article. The "hopping" occurs with probability $(1-\alpha)$ – the so-called damping factor. G is a matrix of normalized out-flow of traffic from any article i that hyperlinks to an article j. Convergence of the iterative form of PageRank is achieved for some $\vec{r} \equiv \vec{r}_{t+1}$ when $|\vec{r}_{t+1} - \vec{r}_t| < \epsilon$, for a small choice of ϵ. The converged vector $\vec{r}$ represents the normalized accumulated traffic to each article i that results from the simulated random walk process. We represent this simulation process functionally as: $\vec{r} = PR(\vec{r}_0, G, \alpha, \epsilon)$.

Ordinary PageRank assumes an equal initial distribution of traffic, $\vec{r}_0 = 1/N$, and equal probability of out-flow along all links, $G_{ij} = A_{ij}/k_j$ where A_{ij} is the adjacency matrix and k_j is the degree of article j. The damping factor is conventionally chosen as $(1-\alpha) = 0.15$. Personalized PageRank relaxes the assumption of equal initial distribution of traffic for an arbitrary normalized $\vec{r}_0$. To guarantee realism, we relax these

[20] Personalized PageRank has recently been *formally* related to the task of community detection in networks (Kloumann et al. 2016)

assumptions even further and leverage the clickstream data (see section 3, *Sources of Increased Attention* for a description) to empirically calibrate internal and external traffic flows in the simulation[21]. In personalized PageRank, we set the vector $\vec{r}_0$ to the normalized empirical distribution of external incident traffic on each article i, and the matrix G to the normalized empirical distribution of out-flow traffic from article i to article j. Having defined the simulation process, we are now in a position to assess how perturbations to attention (i.e., increases in incident traffic from external sources—for example, arising from content contribution shocks) drive accumulated attention to all articles in the network. We represent a general perturbation to some set of articles S as $\vec{r}_{0p}^{\,S} = \vec{r}_0 + \overrightarrow{\delta r}_{0p}^{\,S}$ and set the perturbation according to:

$$\left(\delta r_{0p}\right)_i = (r_0)_i \begin{cases} p, & \text{for } i \in \mathrm{S} \\ 0, & \text{otherwise} \end{cases} \qquad (eq\ 7)$$

where $p > 0$ represents a constant percentage increase of attention shock to affected articles (those in the chosen perturbed set S). In other words, we create relative perturbations of attention that are correlated across a set S of chosen articles. For each perturbation, we calculate the resultant PageRank vector $\vec{r}_p^{\,S} = PR(\vec{r}_{0p}^{\,S}, G, \alpha, \epsilon)$ and compare it to the unperturbed PageRank vector $\vec{r} = PR(\vec{r}_0, G, \alpha, \epsilon)$. Specifically, we are interested in the resultant *excess attention* (EA) received by underdeveloped articles which comprise the articles in the perturbed set:

$$EA(S, p) = \sum_{i \in S} \frac{r_{p,i}^S - r_i}{r_i} \qquad (eq\ 8)$$

[21] In prior research, others have calibrated PageRank with internal traffic from Wikipedia clickstream data (Dimitrov et al. 2017), but have not accounted for variation in external traffic.

Because any perturbation of a set of articles will result in those articles receiving excess attention, we compare excess attention across two different policies: i. an *Attention Contagion Policy* (ACP) where editorial efforts are focused on clusters of well-connected, underdeveloped articles; ii. an *Undirected Attention Policy* (UAP) where editorial efforts are focused on randomly chosen underdeveloped articles that are not necessarily (but may incidentally be) connected to one another. The random selection of underdeveloped articles under this latter UAP policy will lead to contributions to articles that are more spread out across the information network as compared to the ACP policy.[22] The two policies are illustrated in Figure 8. The UAP policy represents a simple and useful baseline for comparison. It may be that without guidance editors already cluster their editorial focus to some extent. However, we do not parametrize clustering under UAP to avoid introducing unnecessary assumptions and additional complexity.

[22] In fact, because UAP spreads out editorial focus through the network, it conveys excess attention to more unique articles. But, under ACP more articles receive a larger share of excess attention. For more details see Fig A11 and the related discussion in the Appendix.

Figure 8 Concentration of Attention Across Network Communities or Cliques for the Two Policies

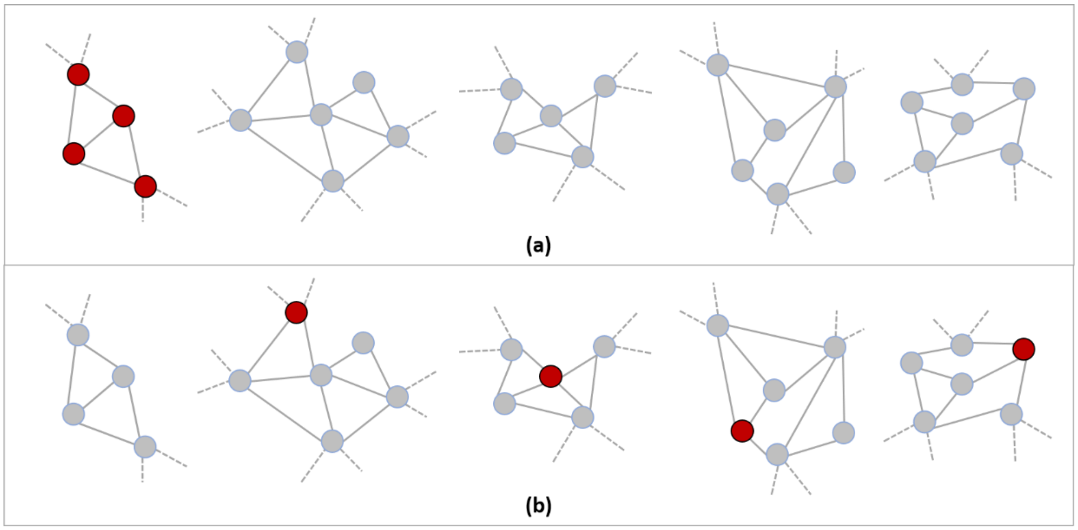

Figure 8 illustrates concentration of attention across network communities or cliques for the two policies. Red nodes receive increased attention (perturbed). Panel (a) illustrates the Attention Contagion Policy (ACP), where attention to red nodes (which constitute the perturbed set S_c^{ACP} for a given clique or community, c) is clustered within a community or clique. Panel (b) illustrates the Undirected Attention Policy (UAP), where attention is spread out randomly across communities or cliques in the network. To compare policies fairly, red nodes in panel (a) are matched one-to-one to red nodes in panel (b), (comprising the set $S_{m_c}^{UAP}$, as described in the text).

To compare these two policies, we first need to identify sets of well-connected articles in Wikipedia that appear in clickstream data and are good empirical proxies for underdeveloped articles. Importantly, many actual sets of related, underdeveloped articles will likely lack the linking structure that would naturally arise from directed editorial focus. That is to say, while these underdeveloped pages are related to one another, they do not *yet* possess the linking structure to connect them. To avoid making unnecessary and potentially ill-informed assumptions about unobserved network structure and its relationship to content, we instead focus only on actual links that appear in the clickstream data and that experienced actual traffic flow. To accomplish this, we use the weighted directed graph of traffic flow between articles and seek tightly connected sets of nodes in the form of both cliques and communities. To find cliques, we computed a large sample of maximal cliques via depth-first-search with Bron-Kerbosh style pruning (Tomita et al. 2006). To find communities, we modify the well-known label propagation algorithm (LBA) (Raghavan et al. 2007): to address the instability of the original LBA, we perform the algorithm 200 times and assign articles to the same community if and only if they were assigned to the same community in at least 95% of the runs. This approach produces stable, tightly connected communities with minimal noise. It is also efficient, fast and able to cope with networks of millions of nodes. We filter maximal cliques and communities and retain only those of small to moderate size ($2 \leq n \leq 6$). For each such clique or community, we match each article to another article in a different clique or community with the closest external incident traffic. This yielded a set of well-connected articles to perturb according to the *Attention Contagion Policy*, S_c^{ACP}, and a corresponding matched set of articles to be used in the *Undirected Attention Policy*, $S_{m_c}^{UAP}$, where c labels the clique or community and m_c labels the matched set. Note that the articles in S_c^{ACP} belong to the same clique or community (c), whereas articles in $S_{m_c}^{UAP}$ can belong to many different cliques or communities. Because testing large numbers of perturbations is computationally intense, we select a random subset of 600 cliques and communities and, for each clique or community, we simulate the perturbations for both policies and compare the distribution of excess attention $EA(S_c^{ACP}, p)$ to $EA(S_{m_c}^{UAP}, p)$. The results are displayed in figure 9 for cliques (panel a) and communities (panel b) for simulation with $p = 0.1$.

Figure 9 Distribution and Kernel Density Estimates of Excess Attention for Perturbative Simulations (p = 0.1) of the ACP and UAP for 600 Cliques and Communities

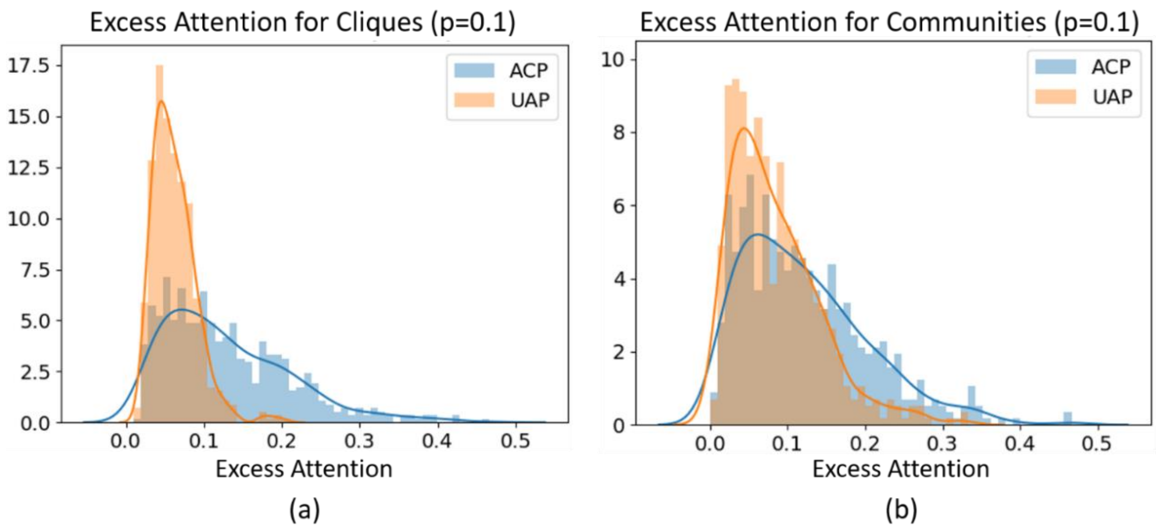

Figure 9 displays the distribution and kernel density estimates of Excess Attention for perturbative simulations (p=0.1) of the *Attention Contagion Policy* (ACP) and *Undirected Attention Policy* (UAP) for 600 cliques (a) and communities (b). The ACP policy leads to significantly more excess attention.

The *Attention Contagion Policy* clearly leads to significant excess attention directed towards underdeveloped pages as compared to the *Undirected Attention Policy*, yielding a relative increase of mean excess attention (ACP over UAP) of 106% for cliques and 44.2% for communities (p<1e-71 from two-sided t-test)[23]. Because editors may already cluster their editorial attention to some extent even without a guidance policy, our results should be interpreted as an upper bound to the value conveyed by the Attention Contagion Policy. Excess attention scales linearly with the size of the perturbation, which follows from the definition of excess attention and the expansion of the iterative perturbed PageRank equation. The shape of the distributions of excess attention for either policy is determined entirely from the network structure around the perturbation set,

[23] Alternatively, two sample KS-tests reject the null hypothesis that the distributions are equal with p<1e-63

implying that the results are identical up to a scale factor (p) for different choices of perturbation size. Results are also robust to different random samples of cliques or communities (see Appendix for details).

Conclusion

Open collaborative platforms have fundamentally changed the way that knowledge is produced, disseminated and consumed in the digital era. This study directly contributes to our understanding of the interaction between production and consumption of information and the phenomenon of attention contagion on Wikipedia, arguably the largest and most successful example of such platforms. To conduct valid causal inference so that we can inform policy with high confidence, we employed a battery of methods including natural experiment, matching, econometric modeling, and empirically-informed simulation. We found that real-world exogenous contributions increase future attention by 12% on average with stronger impact for more significant contributions. They also increase future contribution by 3.6 more edits and 2 more unique editors to affected articles over a 6-month period. This impact is both economically significant and persists for a long time. In addition, we obtained causal estimates of the extent of spillover impact and identified characteristics of articles and links between them that receive the most benefit from spillovers. Specifically, we find that spillover is greatest across new links that point to less popular target articles, yielding an impact as high as 22% for new links from popular source articles to unpopular target articles and 15% for new links from less popular source articles to less popular target articles.

Overall, our results confirm the existence of positive feedback loops of production and consumption of information on Wikipedia. This, unfortunately, also implies that underdeveloped articles experience a poor-get-poorer phenomenon and are therefore naturally disadvantaged in the cumulative development process. This observation is deeply troubling because it suggests that impoverished regions in collaborative information systems will remain impoverished in the absence of policies that are specifically designed to address this problem. More importantly, because information poverty is often correlated with economic poverty (Forman et al. 2012, Graham et al. 2014, Norris 2001, Yu 2006), this phenomenon can act to exacerbate economic, social, political, and cultural inequalities. Fortunately, our findings suggest that less developed regions of information networks can benefit substantially from spillovers. We carry this insight further and propose and compare policies that drive editorial attention using diffusion simulations that are based on real-world traffic flows on Wikipedia. We evaluate the Attention Contagion Policy that leverages spillovers to stimulate development of impoverished regions. We find that this policy can yield up to a twofold increase in excess attention relative to the baseline Undirected Attention Policy. These results are directly relevant to concerns of information equity and have managerial implication for collaborative information platforms. Although we focus on Wikipedia, our findings are relevant to the uneven coverage problem that exists in many platforms that facilitate collaborative content production in domains such as open-source software creation (e.g., GitHub), knowledge markets (e.g., Stack Overflow or Quora), and product reviews (e.g., Amazon or Steam).

Our results suggest that two policies can be effective for encouraging the development of underdeveloped articles or impoverished regions in the information network. First, editors may be encouraged to identify popular articles that should naturally (semantically) link to a focal underdeveloped article. Our results show that creating such a link can harness the largest attention spillover (as much as 22%). Although care should be taken to ensure that added links are semantically meaningful. Second, and perhaps more importantly, Wikipedia should consider encouraging coherent development of impoverished regions. Our results show that underdeveloped regions, which typically lack both attention and the linking structure to connect related articles, are precisely positioned to benefit from attention contagion policies. Currently, the quality and importance of Wikipedia articles is assessed through a tagging system implemented on talk pages. Tools exist that use these metrics to allow editors to search for specific articles that are both important and in need of attention. Additional features could be added to these tools to encourage a coherent focus for individual editors or even groups of editors.

This work is not without limitations. This work tackles causality by leveraging a natural experiment, matching, econometric techniques and empirically-informed simulation. However, cleaner causal inference could be achieved in future work through controlled randomized experiments. As we examine attention spillover due to a second order shock to attention (that itself is driven by a contribution shock), we may miss subtle heterogeneous spillover effects. Future work could consider perturbations to link structure and real-world experimental tests of attention contagion policies. Furthermore, Wikipedia is subject to other natural experiments that may be discoverable. In particular, examination of clickstream data may permit the discovery of natural experiments that can help us better understand attention flow in information networks.

CHAPTER TWO

How Media Ownership Impacts Information Skews: A Study of Televised News Using Massive-Scale Text Transcripts

Introduction

Broadcast TV in the United States is an information system comprised of hundreds of local television stations that both produce their own information and syndicate information from other sources (e.g., major networks). In this system, a few media conglomerates produce and disseminate a vast of this information. As media consolidation has increased in recent years, a very real concern has emerged: that conglomerate owners have both the means and motivation to skew information.

The reality of this threat was exemplified in 2018 when the Sinclair Broadcast Group forced dozens of local news anchors to recite the same script verbatim. When conglomerate owners act coherently, they can skew information to emphasize certain views and perspectives, as well as frame and cover information in ways that align with these views and perspectives. Unfortunately, we know little about how ownership affects information skew, and we lack a systematic, empirical evaluating of content and ownership at a sufficient scale and level of detail. In this study, I quantify and investigate the consequences of information system ownership, specifically in terms of diversity of information and political polarization, in one of our most important mass information systems—broadcast televised news.

Broadcast media has a dramatic impact on political and social outcomes and undeniably shapes the national dialogue surrounding important issues. Extensive research from communication and political science has shown that television significantly impacts a wide range of real-world outcomes, including voting behavior (DellaVigna and Kaplan, 2007; Enikolopov et al., 2011; Campante and Hojman, 2013; DellaVigna et al., 2014; Schroeder and Stone, 2015; Philippe and Ouss, 2018), disaster relief efforts (Eisensee and Stromberg, 2007), politicians' resignations (Garz and Sorensen , 2017), terrorism (Durante and Zhuravskaya, 2018), fertility decisions (Kearney and Levine, 2015), divorce rates (Chong and Ferrara, 2009), and economic and social development (La Ferrara, 2016; also see Kleemans and Vettehen, 2009, for an overview of the causes and effects of television sensationalism).

A controversial event in 2018 underscores why media ownership deserves critical attention. That year, dozens of news anchors on many local TV networks recited the same script word-for-word in their respective broadcast news segments. The events captured in the broadcast news video clips were striking, as local viewers expect to receive information primarily about local affairs, rather than discover that their local, respective anchors were all repeating the same news stories verbatim in different media markets. It turns out this event was not a coincidence. All the stations in the video clips were owned by a single media company, the Sinclair Broadcast Group (Martin and McCrain, 2019). This discovery led to a huge backlash on social media, as well as from a range of professionals, news producers and journalists. Many quit or announced their resignation at the end of their respective contracts. As former KHGI news producer Justin Simmons said, "It wasn't long before we started receiving the segments that our parent company said we had to air during our local news broadcasts . . . the 'must-runs' were something I was willing to get fired over." This controversial event reflects recent policy changes with respect to the media industry. In 2017, the Trump administration relaxed the FCC's restrictions in terms of TV ownership, which allowed even more consolidation. In turn, the extent to which a single company may own even more TV stations remains a question, if not a threat, fraught with concern.

Figure 10

While the Sinclair example underlined this threat in a very public way, we know little about how ownership affects information skew. What is needed is a systematic empirical evaluation of broadcast news media content and impact of ownership at a sufficient scale and level of detail. How can we evaluate the health and quality of such large-scale information conveyed through mass media? What does it mean for only a few corporate entities to control broadcast news coverage—both local journalism and reporting—throughout a single country? Those questions are particularly challenging to address. Local TV news is not as transparent or easily evaluated as national news (i.e., local TV news purports to be "local"). Moreover, quality of content has many important aspects, some of which are difficult to measure at scale and in an automated fashion. For instance, both "accuracy" and "importance" are key aspects when it comes to evaluate the quality of news. However, it is difficult to quantify "importance" in an objective way as it is inherently a subject concept and "accuracy" speaks to factual characteristic of news and can only be evaluated manually on a story-by-story basis. In this study, we operationalize quality of news through diversity of information and political polarization of the news content. Those are also important aspect of news content and are related to "accuracy" and "importance": Slanted news is unlikely to be accurate on its factual basis and it is desirable to have a diverse and balance news diet.

The goal of this study is to quantify and understand the consequences of information system ownership, specifically in terms of diversity of information and political polarization, in one of our most important mass information systems—broadcast televised news. My research questions are: (1) Can we reliably measure the diversity of information and political polarization in news content (and how)? (2) How and to what extent is information diversity and political polarization of local TV news correlated with a station's network affiliation and ownership? (3) How does media ownership impact information diversity and political polarization in TV news? To examine the impact of ownership, I focus on the natural experiments of station acquisitions by new owners.

Using data of hundreds of acquisitions over a six-year period, I examine changes of ownership, which allow me to disentangle the impact of a TV station's owner from other characteristics that may also be correlated with the information produced by a station.

Background

I first provide some background about television systems in the United States. The U.S. is divided into 210 local media markets, which are called Designated Market Areas (DMAs). People who live in the same DMA will be exposed to the same set of local channels. Historically, these channels corresponded to areas of broadcast reception; however, the shapes of DMAs today make less sense in terms of cable transmission possibilities.

Figure 11 Designated Market Area Map of the United States

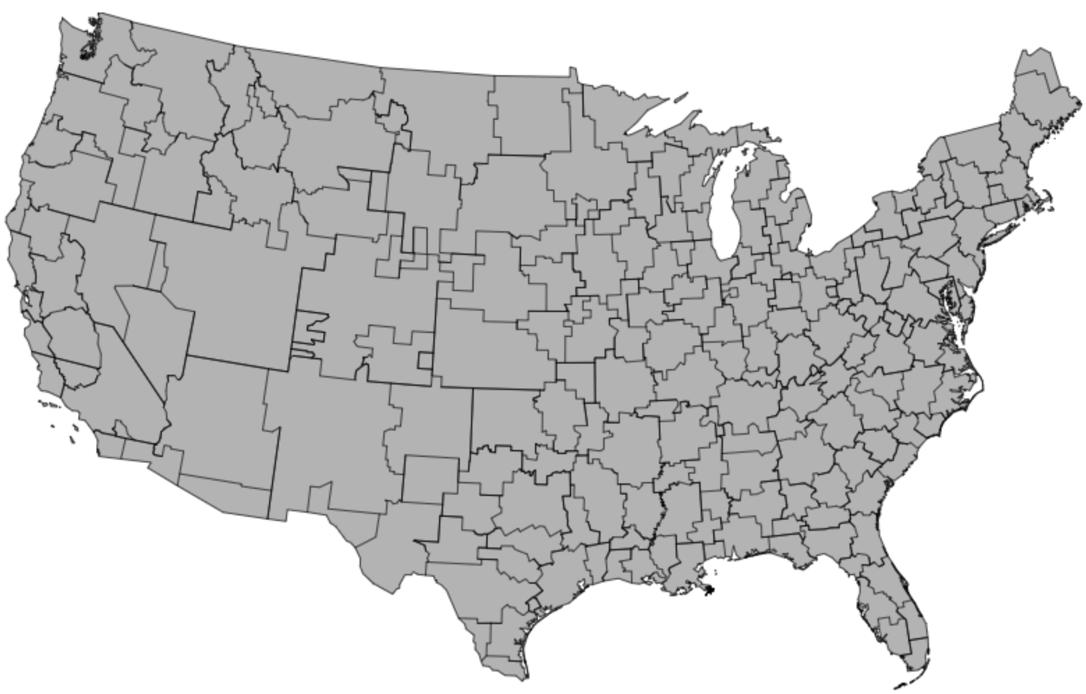

For example, Boston's DMA is displayed in Figure 12. The shape of this particular DMA does not align with state or county boundaries; rather, the DMA covers part of Massachusetts as well as some regions of Rhode Island, Vermont, and New Hampshire. The Boston DMA has its own set of local channels, each of which is affiliated with a network (Figure 11). TV networks occasionally choose to own some TV stations, but the vast majority of TV stations in the U.S. are not owned by networks. Instead, a local TV station is usually owned by a media company. Some major media companies in the industry include the Nexstar Media Group, the Sinclair Broadcast Group, Gray Television, Raycom Media, and Tribune, among others. Local channels receive some content from the networks with which they are affiliated, and this information is called syndicated network content. Syndicated content is the same across DMAs: the content typically includes network news, TV series, and talk shows. Local TV stations also produce their own content, which is known as locally produced content.

A local TV station has control over both its locally produced content and its network syndicated content. It is fairly obvious how a local TV station influence its local content, that which it produces for its own purposes. For network syndicated content, however, even though programs are produced at a centralized, network level and are broadcast across designated market areas, a local TV station usually negotiates with its affiliated network to determine what programs local TV stations will carry and how much network programming they would like to broadcast. Because of this arrangement, local TV stations have some control over the precise syndicated content they will broadcast, even though they are not directly involved in producing that content.

Figure 12 Illustration of Boston DMA

Data and Measurement

To answer my research questions, I leverage computational techniques in natural language processing and machine learning to analyze structured and unstructured data from multiple sources. First, to obtain the news content itself, I use a massive scale text corpus of television program transcripts. The corpus is comprised of complete transcripts of roughly 800 television channels across all 210 DMAs in the U.S. from January 2013 – March 2018. This time period spans one presidential election and two mid-term elections. The complete set of transcripts has a rich variety of content, including news, commentary, sports, entertainment, and advertising (commercial and political). In total, the transcripts contain over 10 million hours of broadcasting and over 60 billion words, which is about 27 times that of the entire Wikipedia in English. In this study, I focus on televised news programs.

The size of this text corpus is several terabyte data in total (200GB per month). The television news transcripts were divided into "pages" or short chunks of text. Each page is about 616 words or 4.2 minutes of broadcasting, on average. To better analyze the content, I combine "pages" into 15-minute time blocks. This is consistent with how the television industry analyzes content. For instance, Nielsen TV ratings are also measured in 15-minute increments. I then processed this content extensively by using machine learning and natural language process, including program categorization, phrase detection, word and document embeddings, and topic classifications, among others. In this study, I focus on news content only. I will describe in detail how I classified news and how I classified syndicated news and locally produced news.

In addition to the massive-scale text transcripts of news broadcasting, I also collected data from multiple other sources, including: (1) network affiliations, ownership status, and shared service agreements of all local TV stations in the data set from official Federal Communication Commission filings; (2) a variety of demographic data from the U.S. Census, aggregated at the DMA level; (3) measures of political leanings of congressional voting districts (e.g., the Cook Partisan voting index); (4) Congressional records that contain all the congressional speeches in the U.S. Senate and House from 2008 to 2017; and (5) party affiliation and DW-NOMINATE ideological scores for each congressional representative in the Congressional Records from 2008 to 2017.

The ownership and acquisition information is of special importance. To obtain that information, I use ownership data from the FCC. Unfortunately, these FCC records of local station ownership are limited, as the FCC does not maintain a consistent record of ownership change. Thus, I examine all electronic paper trails to capture local station acquisition. These data include FCC Forms 302, 314, 315, 316, and 345 from the FCC Media Bureau CDBS Public Access and Biennial Ownership Declarations from the FCC OPIF (Public Inspection Files) API. These electronic paper trails are not sanitized. In fact, because the records are spotty, there could exist multiple media holding companies, personal and family trusts. Because of these issues, determining who owned what proved challenging. Much manual effort was required to clean records and investigate the relationship between different entities. Finally, I produced a verified set of records of local station ownership and acquisition. This is an important step for my further analysis. Moreover, beyond this research, public accountability and policy transparency cannot be achieved with bad data. This is a problem to be fixed on its own for better public governance.

Classification of Syndicated News Content

I first needed to extract relevant news content from the raw corpus, as I was interested in measuring the informational content from the news. The raw data provide a continuous text stream with information regarding broadcast time. To classify news, I pair the text corpus with the EPG TV guide data, which marks television programs and associated meta information with time periods. As a result, I can ascribe any text in the corpus to a particular TV program.

Next, I faced the challenging task of defining what TV programs are classified as news. Syndicated news and locally produced news have different production processes and, thus, different implications; as a result, I study them separately. In this section, I describe how I identified syndicated news (or network news). To do that, I took as a starting point the set of programs that were labeled by the EPG TV guide as NEWS. I note that this programming title from the EPG TV guide is complex and often has many program titles listed for the same underlying program (e.g., “CBS Evening News” and “CBS Evening News with Jeff Glor” both appear in the data and are the same show). As a result of this complexity, there exists thousands of different program titles, even though the real number of programs is much lower. Moreover, there are many one-off news specials (e.g., NBC News Special Report: Inauguration Day). Thus, a much smaller list of program titles accounts for the vast majority of viewing time. I develop a procedure in which I count the number of stations a given title appears on. As I am interested in syndication news in this section, syndicated content—by its very definition—should appear on multiple stations.

There are 232 program titles that have been broadcast on more than 5 TV stations. I examined this list and combined titles that indicated the same program, but that had minor differences in their names. Hence, I obtain a relatively clean set of program titles of syndicated programs that are labeled by the EPG TV guide as News. Unfortunately, its classification of news is too broad and includes many programs that traditionally would not be deemed as news (e.g., crime investigation program, news magazines). Two experienced coders with good domain knowledge about TV news independently reviewed the list and classified the programs into five categories: Hard News, Journalism (not formal news), Newsmagazine, Election-related, and other content. In my analysis, I only include the text transcript from programs in the Hard News category, which includes programs such as CBS This Morning, Good Morning America, CBS Evening News, Meet the Press, and ABC World News. An analysis of soft news like that in the Journalism and Newsmagazine categories could potentially prove interesting, as such programs are also part of citizens' broader news consumption. That said, I restrict my analysis to only programs in the Hard News category, so I may interpret my findings more clearly. In particular, this strategy allows me to draw conclusions about the impact on informational content derived from traditional news reporting (e.g., morning news and evening news).

The complete list of syndicated news programs is listed below: Today, CBS This Morning, Good Morning America, Good Morning America: Weekend Edition, CBS

Overnight News, Meet the Press, Sunday Today With Willie Geist, Fox News Sunday With Chris Wallace, Nightline, ABC World News Tonight With David Muir, CBS Evening News, America This Morning, CBS Weekend News, ABC World News, NBC Nightly News With Lester Holt, ABC World News With David Muir, CBS Morning News, Early Today, Up to the Minute, NBC Nightly News, World News Now, CBS Evening News With Scott Pelley, NBC Nightly News With Brian Williams, ABC World News With Diane Sawyer, Nightline Prime, NBC News Special, CBS Evening News with Jeff Glor, This Week With George Stephanopoulos, Today With Kathie Lee and Hoda, Face the Nation, Today – Kathie Lee & Hoda, FOX News Sunday, This Week, ABC News, and CBS News.

Classification of Locally Produced News Content

In addition to syndicated news content, locally produced news is of particular interest in this study. The owners of local TV stations have direct control of the management and editorial policies that impact local TV stations' news production. Because owners impact news content so directly, the impact will shape locally produced content.

However, identifying locally produced news content from the massive text corpus is not as straightforward as it seems. Neither the transcript corpus nor the EPG data have labels to designate what content was produced locally. Each channel may air locally produced content for different schedules. More broadly, these factors make it challenging to even define "local production." To address this challenge, I devise a procedure to identify locally produced content based on broadcast patterns. Specifically, I pool program titles across all 800 stations and count the number of stations for which each title has appeared. By definition, local content should not appear on multiple DMAs; they are locally produced, have a local focus, and thus will only appeal to local audiences. To complete this process, I choose a threshold of five: any program that appears on more than 5 stations will no longer be considered as local content. In setting this threshold of 5, I assumed that some news programs may be shared between stations in the same DMA. The most common case is that of a DMA having four local TV stations. By using this strategy, I circumvent the need to identify detailed production information about a program while, at the same time, can proceed to select programs that intuitively are likely to be produced locally.

With this procedure, I compile a list of 9,717 news programs that are locally produced. For brevity's sake, I do not provide the complete list. Sample program titles include WZZM 13 News @ 6pm, WYFF News at 6am, WTOC News at 5P, Good Morning Wyoming, and Missouri Viewpoints.

Measure of Partisan Slant

In this study, I measure the ideological slant of news by adapting and building upon the seminal work by Gentzkow and Shapiro (Gentzkow and Shapiro 2010), which investigates political slant by examining the characteristics of language in a text. They identify a list of partisan phrases by comparing the different language use patterns of Democrat and Republican members of Congress in their respective Congressional speeches. In using this approach, I assume that elected officials' political stance can be derived solely by observing their word choices and language patterns. This approach leverages the fact that legislators across the political spectrum talk about different policies and issues or tend to use different phrases to speak to their constituents. For example, conservatives tend to use the phrases “illegal immigrants” or “death tax” frequently, while liberals tend to use the phrases “undocumented immigrants” or “estate tax” frequently.

To implement this strategy, I follow two steps. First, I use a customized chi-squared statistic to select phrases that have distinguishing power between the speech of two groups (e.g. Republicans and Democrats). Next, I then use counts of those phrases to score the political slant of text. I can use more sophisticated methods, such as lasso or other machine learning methods, by treating this as a prediction problem. That is, I wish to predict the political orientation of text content using a list of pre-identified phrases as features. The essence of this approach is that the counts of those polarized phrases are the distinguishing features that predict political orientation.

This approach has been applied to a wide variety of large-scale news corpuses, including newspapers (Gentzkow and Shapiro 2010), online news (Gentzkow and Shapiro, 2011), cable news (Martin and Yurukoglu, 2017), and Wikipedia (Greenstein and Zhu, 2012, 2016, 2017). In the literature, when researchers use this approach, they identify polarized phrases from Congressional records or simply use the original list of phrases identified by Gentzkow and Shapiro (2010). Congressional records are the records of Congressional speeches delivered by U.S. Senate and House representatives. One potential issue that arises from using this strategy is that news anchors and TV personnel do not use language in the ways that members of Congress do. Previous research has shown that cross-domain ideological classification can yield poor results (Yan et al., 2017), which indicate that polarized phrases identified in Congressional speeches may not be the most informative feature of television news in terms of predicting political orientation. Thus, using this approach to analyze TV news corpuses may not yield accurate classifications.

Therefore, to obtain more accurate results, I identify polarized phrases by using "labeled" TV data and then use those phrases to supplement the phrases identified from Congressional records, using the same procedure. Arguably, FOX is the most conservative cable news channel, and MSNBC is the most liberal. Instead of comparing the different language patterns of Democrats and Republicans in their respective Congressional speeches, I compare language patterns in news reporting and commentary content from FOX and MSNBC.

Republican -> FOX Democrat -> MSNBC

I used this approach for both Democrats' and Republicans' Congressional speeches and content from MSNBC/FOX cable news broadcasts to identify two sets of polarized phrases that appear in Congressional Records and cable news channel broadcasts, respectively. To accomplish this, I define customized χ^2 statistics to select the phrases that have high distinguishing power to contrast ideological viewpoints. There is about a 20% overlap of the Top 1000 identified polarized phrases between the two corpuses. There are indeed numerous phrases that I identified on TV broadcasts, but do not emerge from this procedure that applied to Congressional records. The phrases I identified on TV broadcasts tend to be colloquial and topic specific.

Below are examples of the polarized phrases that have high χ^2 value:

Table 8: Examples of Polarized Phrases

TV (2012–2018)		Congressional Record (2008–2016)	
MSNBC	FOX	Democrats	Republicans
Right wing	Terrorist attack	Voting Rights Act	Healthcare law
Marriage equality	Obamacare	Immigration reform	Strong work ethic
Voting Rights Act	Mainstream media	Gun violence	National media
Background checks	IRS scandal	Tax cuts	Balanced budget amendment
Minimum wage	Al Qaeda	Civil rights	May God bless
Immigration reform	Middle East	Paycheck Fairness Act	Behind closed doors
Same-sex marriage	Healthcare law	Working families	Job creators
Gun violence	Merry Christmas	Social safety net	Illegal immigrants
Undocumented immigrants	Illegal immigrants	Unemployment benefits	Tax increase
Medicaid expansion	First amendment	People with disabilities	Small business
Middle class	Fair and balance	Equal Pay Act	Boy Scouts of America
George Washington bridge	Benghazi attack	Human rights	US Air Force
Gay rights	North Korea	Undocumented immigrants	Obama administration

To verify these results, I first apply this method to three cable channels: CNN, MSNBC, and FOX. CNN falls in between MSNBC and FOX (Figure 13). This is consistent with our prior expectation. In addition, I calculate the ideological slant for 10 popular news programs from the three cable channels. The results see Figure 14. “The

Rachel Maddow Show" is the left-most slanted program among the ten, based on the measure I use, while "Fox and Friends," "Hannity," and "America's News HQ" all fall to the far-right side of the spectrum. These results are consistent with subjective judgement from domain experts and confirm the method I use.

Figure 13 Political Slant for Cable News Channel between 2012-11 and 2014-08

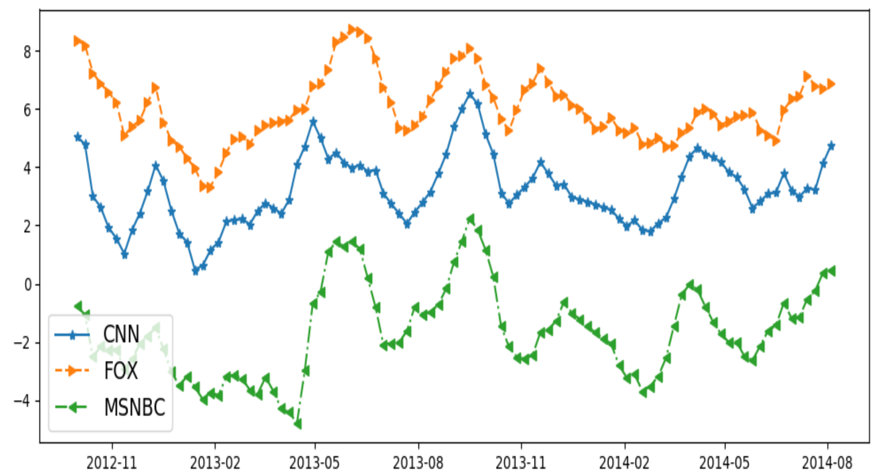

Figure 14 Political Slant at News Program Level

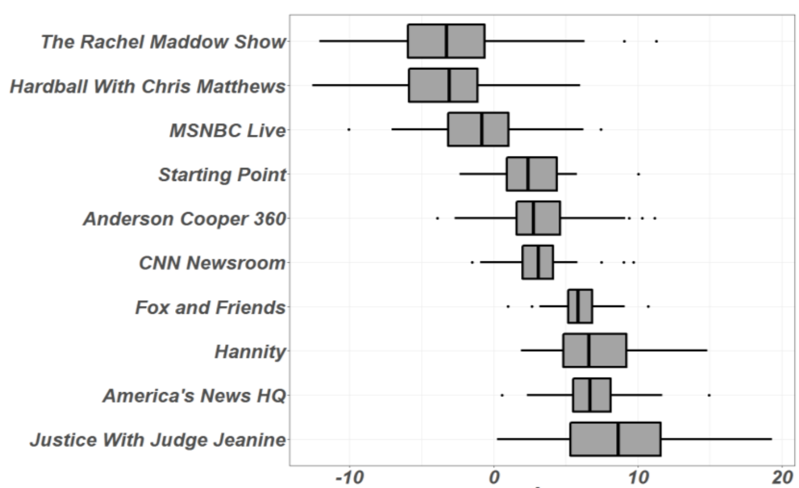

It is worth noting some downsides of using this Gentzkow and Shapiro (2010) approach. This methodology is based on a bag-of-phrases strategy and can only classify content using pre-identified features. The identified polarized phrases are only around 1,000 in the original list in Gentzkow and Shapiro (2010). With this strategy, a good deal of content does not include any of these phrases and, therefore, remains unclassified. A question then emerges: is such content neutral? Or we did not have the right features to identify the political slant in those content? The problem with the approach in Gentzkow and Shapiro (2010) is that phrases have no relationship or association with other phrases, even if their semantic meanings are similar.

To address this limitation, I use word embeddings, a technique that maps words and phrases to a vector space by encoding the semantic meaning and association between the words and phrases. Word embedding and, more broadly, representation learning is very powerful and have been proven to capture semantic meaning, support additivity, and provide meaningful projections, all of which help relate words, phrases, and even documents to one another. Thus, I combine polarized phrases with word embeddings to determine an axis of ideology that can be used to score the slant of text.

To demonstrate this idea, I use Figure 15 to illustrate a stylized example. The classic example of word embedding capturing semantic relationship is depicted in the direction mapped from king to queen and is parallel to the direction mapped from man woman in the embedding space. Similarly, with respect to ideology, an ideological axis in the embedding space might exist that maps the direction from liberal to conservative in ways that is parallel with the direction mapped from undocumented immigrants to illegal immigrants. Thus, the vector representation of any set of concepts (e.g., words, phrases, or documents) can be projected onto this ideological axis to measure political slant.

Figure 15 Gender and Ideology Axis in Embedding Space

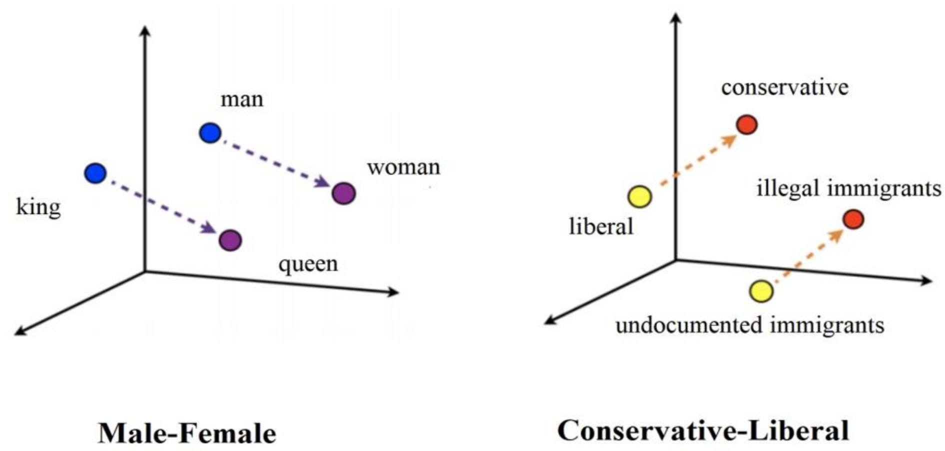

Moreover, these ideological axes might be topic specific. As a result, there are likely many instead of just one. Authors of content from either side of the political aisle frequently talk about particular topics or issues (i.e., coverage bias), but words or phrases belonging to these topics are not necessarily slanted. Intuitively, there are different types of phrases: absolutely slanted, slanted conditional on topic, and neutral topic phrases. For instance, a Congressional representative from Iowa may discuss agriculture in his or her Congressional speech while also holding a set of relatively conservative views with respect to economic and political issues. However, one should not conclude that anyone who talks about agriculture is necessarily conservative.

For this reason, I can identify ideological axes that are independent of topics by using deep learning and neural networks. This intuitive approach builds a neural network that splits the embedding of a document into two parts: topic and slant. The model punishes topic for predicting slant labels and then learn the most slanted phrases for each topic using some labeled data (e.g., Congressional records). Following this approach, I create axes from slanted topic phrases that can be applied to any pre-trained embedding model (e.g., GloVe, word2vec)

Measure of Information Diversity

In addition to political slant, I am also interested in the impact of ownership change on the diversity of news content. I also seek to advance the method used to measure Information Diversity from text data. In the previous work, the best idea so far is to use a conventional measure of diversity (https://en.wikipedia.org/wiki/Diversity_index) and extend this measure to text data. Among several possible diversity indices, Shannon entropy has been the most popular diversity index in the ecological literature, where it is also known as the Shannon diversity index. It is a non-parametric measure to quantify the uncertainty of a set of states and is calculated as follows:

$$D = -\sum_{i}^{R} p_i ln p_i$$

in which p_i is the proportional abundance of the ith type and R is the total number of types.

When the index is applied to text data, we can adopt a bag-of-words approach and define this diversity for words themselves, in which case p_i is just the fractional occurrence of a word i, and it is divided by some baseline (e.g., inversed document frequency). The normalizing factor is necessary as some words are more common than others in particular languages. we can also extend this bag-of-words approach by representing each document with a set of topics. In such cases, p_i would be the fractional distribution of topic i. We then divide by some baseline-like average proportion of a given topic for the collection of documents.

In addition to the Shannon diversity index, other diversity measures can also be adapted to text data. For example, Aral and Van Alstyne measure “novel information” in employees’ email communication in a workplace by using an average distance from the mean (i.e., center-of-mass) $D = 1/N \sum_i^N (1 - cos(d_i, \overline{d}))^2$ (Aral and Van Alstyne, AJS 2011). In this measure, d_i is a vector representation of a document (either bag-of-words or distribution over topics), and $\overline{d}$ is the mean of all N vector d_i (center of mass).

However, bag-of-words representations neither properly capture variations in word choice for the same semantics, nor account for correlation between topics. Thus, I use word embedding to capture the semantic meaning of words and measure the spread of document vectors in the embedding space. Briefly, I represent content in the following way:

$$\overrightarrow{d_i} \sim \text{embedding vec of a 15-min document}$$
$$C = \left(\overrightarrow{d_1} \;\; \overrightarrow{d_2} \;\; \ldots \;\; \overrightarrow{d_N}\right) \sim \text{stack them for each day}$$

Given the neural embedding representation of text content, I measure information diversity using the trace of the covariance matrix of C or the sum of eigenvalues of the singular value decomposition of C. Both measures are tightly correlated. Below, I show the information diversity measure of CNN as an example.

Figure 16 Information Diversity for CNN Over Time

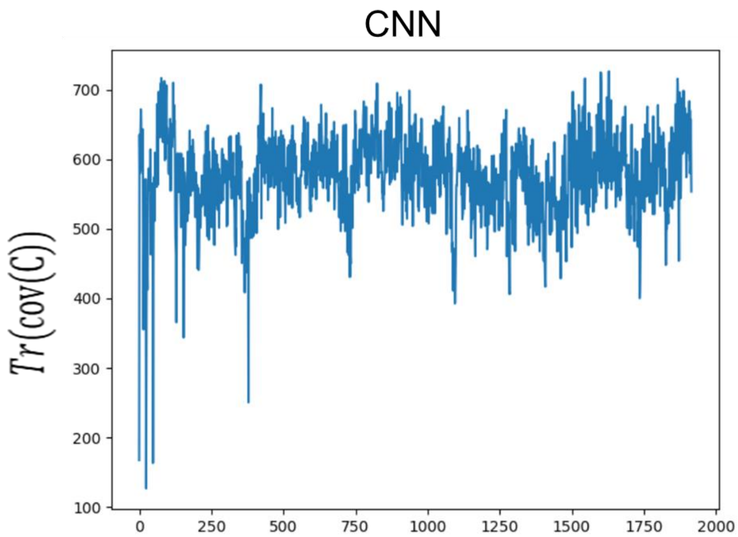

Analysis of Local Channels across the United States

To begin my analysis, I first apply the measure of political slant I developed to both network news and locally produced news from around 800 local TV channels across all 210 DMAs in the U.S. This is the first-time political slant has been measured for all local channels on a national scale using massive-scale text data. I aggregate the results by their network affiliations (i.e., FOX, NBC, ABC, CBS), which allows us to identify interesting patterns related to networks, as well as identify temporal patterns.

Figure 17 Political Slant of Local Channels Aggregated by Network Affiliation

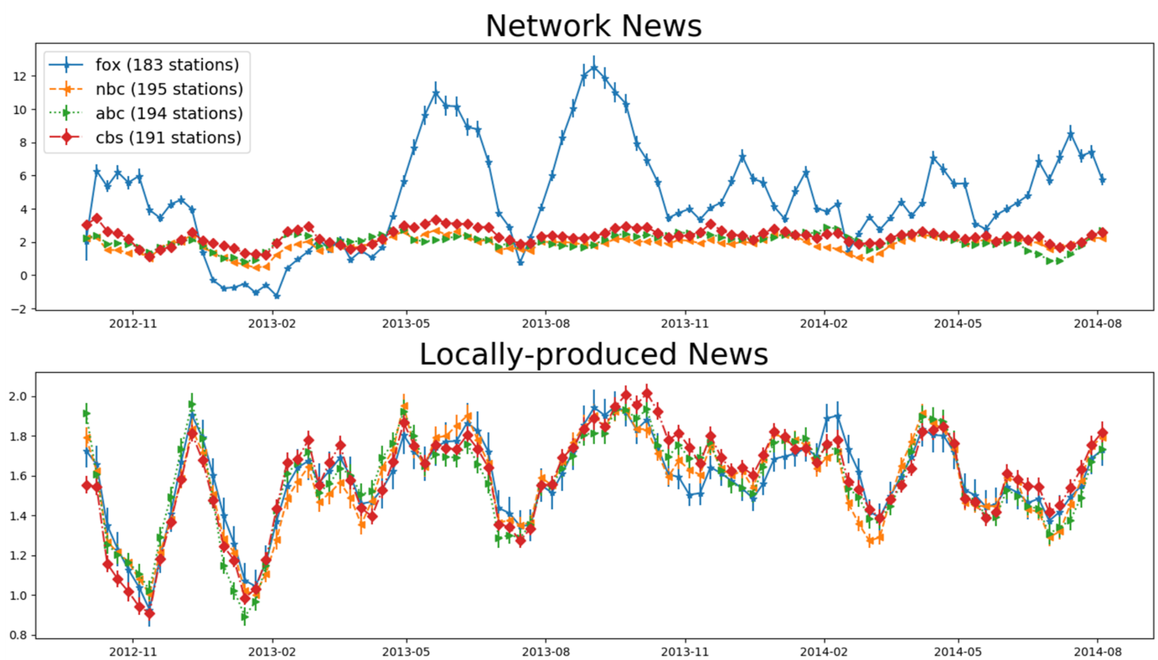

Several key findings emerge from this analysis. First, broadcast news has its cycle, so does their political slant. This is true, regardless of network affiliation. In Figure 15, the index of political slant trends up and down over time for all networks, although to different degrees. Second, a striking pattern exists in the top panel: with respect to network news, FOX-affiliated local channels experience a much larger fluctuation than all other local channels, partially because FOX offers less network news (e.g., Fox News Sunday with Chris Wallace) to local channels; as a result, it is easier for those channels to experience changes in their average slants over time. Nonetheless, the difference is statistically significant and stark in magnitude. It implies that there is something particular about news production operations for FOX: the network seems to have a way to spin things to the right when there are potential underlying events happening. Lastly, I do not observe a significant difference in local news in terms of political slant across different networks, which suggests that network-affiliation may not have an impact on the content of locally produced news.

Next, before examining the impact of ownership change, I examine the correlation between station ownership and informational content in terms of news coverage, so I may determine the extent to which stations owned by different media companies have different political slant. To confirm, I find significant variations in terms of political slant across media groups. This along speak to why we should care about the ownership of the news outlets where we get our information every day. In the next section, I will dive into identify whether the driving force behind this association is ownership when we examine the impact of ownership change.

I use a simple additive model, in which the slant of station is explained by its owner, its network affiliation, a time fixed effect, and an idiosyncratic term. All factors that are not captured by owner, network affiliation, and the date will be left in the idiosyncratic error term. With this modeling strategy, we are interested in capturing the extent to which political slant in news is associated with station ownership in some ways.

$$\text{Slant}_{it} = Owner_i + Network_i + Date_t + \epsilon_{it}$$

I find that significant variations exist in terms of political slant of the news across different media companies. Specifically, media ownership is highly correlated with political slant. Also, stations owned by different media companies have different political slant with respect to both network news and locally produced news.

In terms of network news, although local stations are not involved in the production process of syndicated network news, it is nevertheless possible for an owner to have control over the network news. Local channels and their respective network affiliates must negotiate over how much and what network news to syndicate on local channels. The results of these negotiations point to a more subtle way for influencing information broadcast to audiences, especially in ways that previously did not attract much attention. This analysis indeed indicates that these negotiated decisions significantly affect the political slant of network news broadcast on local stations. On the right-most part of Figure XX, we observe that stations owned and operated by FOX (FOX O&O) have the highest political slant toward conservative perspectives. Presumably, this slant is due to those stations having the largest amount of network news syndicated from FOX.

Next, I analyze the level of political slant in local news across different owners of local stations. In local news, there is an even bigger variation in political slant across different media companies. Companies with the highest political slant have an index almost twice that of companies with the lowest political slant. In between those two extremes, there exists a full spectrum of different political slant. With respect to locally produced news, local stations have direct control over the news production process. The significant variation in political slant across different media companies echoes a primary concern: the extent to which large owners skew the information viewers receive, so owners may emphasize certain views, perspectives, and coverage that they deem more coherent with their respective worldviews.

Figure 18 Political Slant of Local Channels by Media Company

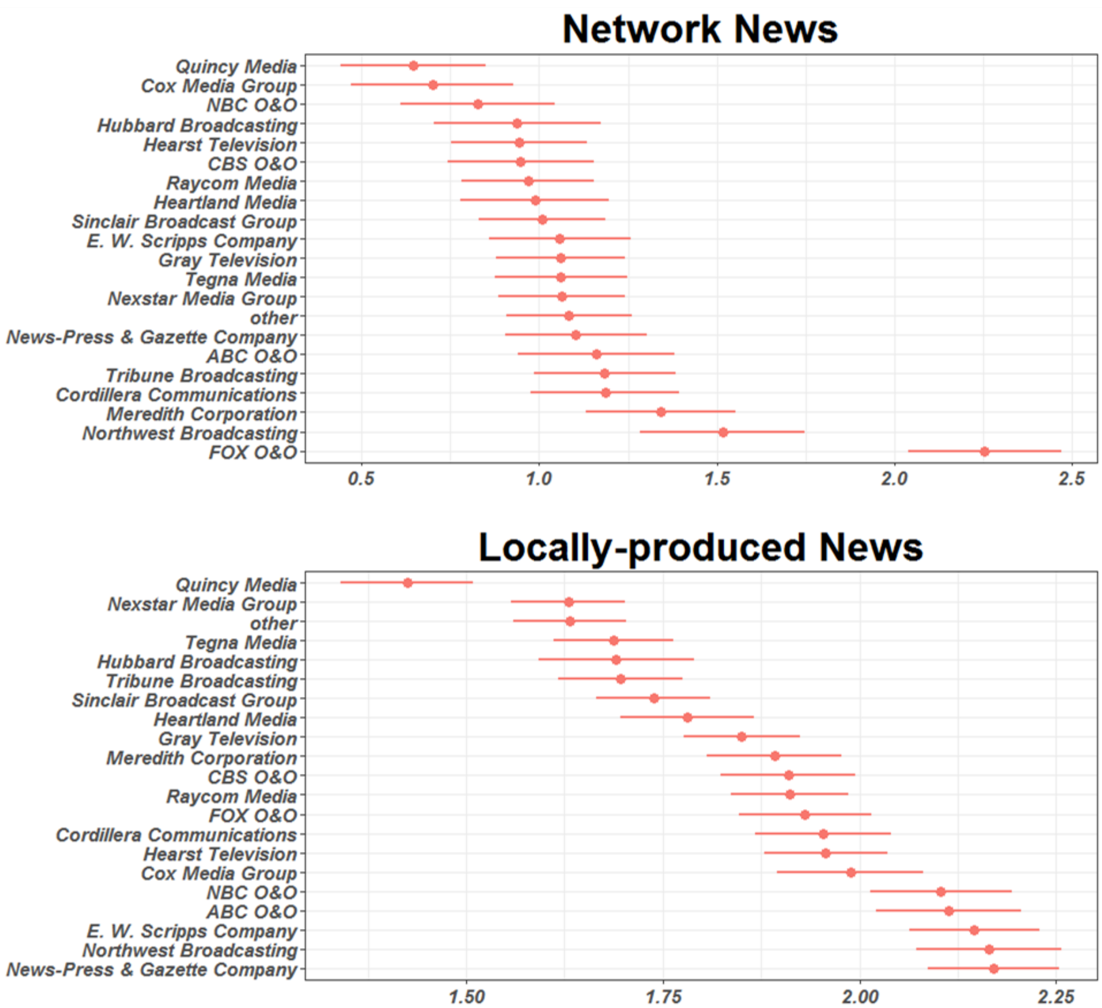

However, the model above only provides a correlational interpretation of the relationship between ownership and political slant; as a result, it is reasonable to suspect that some confounding factors exist that drive this relationship, but that are not directly attributable to ownership per se. Moreover, this correlation could result from a selection process: a media company could choose to only acquire stations that are relatively more conservative, which would result in this particular media company owning a conservative portfolio of TV stations. However, in this case, we would not expect the owners to change the behavior of news reporting at local news organizations, which is the effect this study seeks to identify. Therefore, to identify a causal impact of media ownership, I next focus on the impact of station acquisition in the next section.

Causal Inference using Synthetic Control Methods

To identify the causal impact of ownership, I now focus on the many natural experiments of station acquisitions by a new owner. To do so, I examine whether a station acquired by a new owner experiences political slant before and after the acquisition and benchmark that difference to a baseline. This strategy allows us to construct reasonable counterfactuals (i.e., what would happen if the station acquisition had not happened) and identify a causal impact by comparing the quantity of interest from the real world and its counterfactuals.

Despite the intuitive nature of this approach, it actually presents a challenging modeling problem. To obtain a concrete, complete understanding of the impact of ownership, it is necessary to conduct a systematic investigation into many station acquisitions by different owners at different time periods between 2013 and 2017. In this setting, each station is unique in terms of the content it produced, and each acquisition could have a heterogeneous impact, based on the media company involved and the media market in which it is located. Traditional approaches for causal inference with panel data, such as two-way fixed effects (FE) and difference-in-differences (DiD), are limited and have several drawbacks in this setting: 1) the treatment effect heterogeneity leads to bias in two-way FEs and DiD approaches; 2) it is difficult to evaluate the strong assumption of strict exogeneity; and 3) it requires strong assumptions about functional forms, such as the additivity of fixed effects and linearity in covariates.

Recent advances in causal inference for panel data can help overcome these limitations. In particular, I use a Bayesian structural time-series modeling approach

(Brodersen et al., 2015) to model the impact of ownership change. This approach compares observation under the post-treatment (after acquisition) to imputed counterfactuals using information about outcomes and covariates from the past and other units (during the pre-treatment period). Intuitively, it uses data from the pre-treatment period to construct a synthetic control for each treated unit. In this way, this approach allows for more flexibility in modeling the counterfactual of each station and provides a greater ability to select reference groups (i.e., comparable stations based on domain knowledge). Using different strategies, my approach is to select TV stations in the same or adjacent markets. In addition to allowing heterogeneity in terms of treatment effect, this approach further eliminates unobserved confounding factors by accounting for local trends, seasonality, and other dynamic patterns that the model is able to learn from the reference groups. Below is a summary of the model:

$$y_t = Z_t^T \alpha_t + \varepsilon_t \qquad \text{(observation eq)}$$

$$\alpha_{t+1} = T_t \alpha_t + R_t \eta_t \qquad \text{(latent state eq)}$$

state vector *Transition matrix* *matrix of controls (e.g., seasonality)*

State vector can account for:

- local trends
- seasonality
- contemporaneous static / dynamic cov's learned from the reference group

The model predicts outcomes for treated units based on fitted control and provides a confidence measure of the prediction. By comparing predicted values with actual values of the treated unit in the post-treatment period, I draw inferences about whether the treatment has a significant impact on the treated unit. Below are two examples of acquisitions where this method is applied: the Sinclair purchase of KPTM and the Nexstar purchase of KWKT. The left column reveals a significant shift in political slant for the Sinclair purchase of KPTM, and in the right column, no such shift is present (null effect) for the Nexstar purchase of KWKT.

Figure 19 Example analysis of KPTM and KWKT

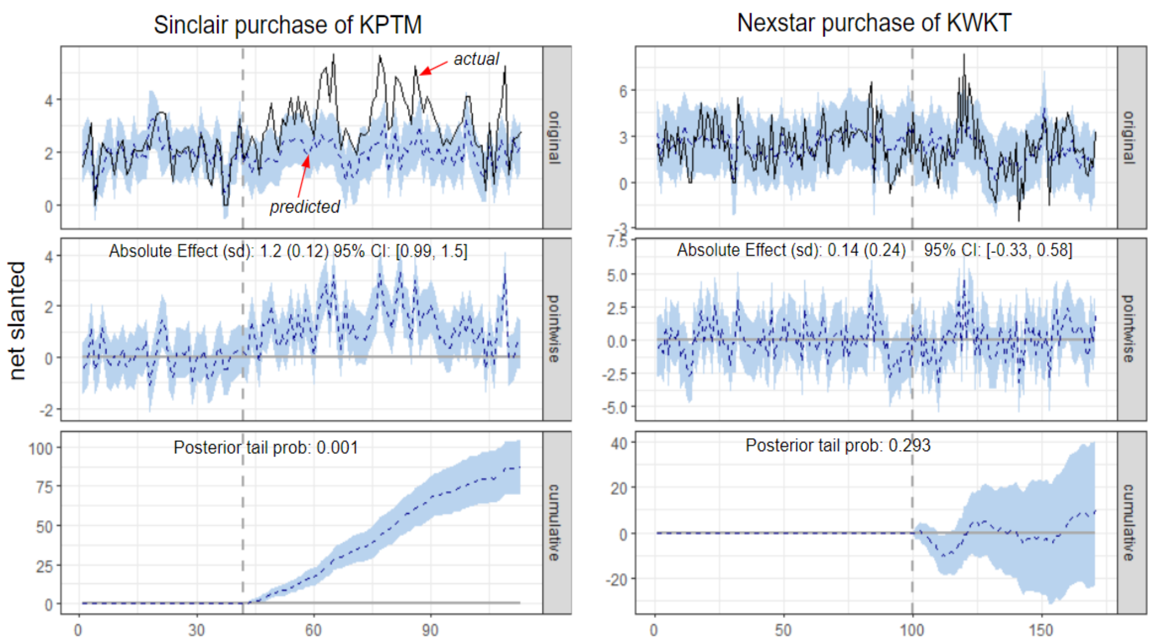

This method is applied to all ownership acquisitions between 2013 and 2017. Below are the results for the impact of ownership change on political slant. When this methodology is applied to many station acquisitions in Figure 18, each point and confidence interval comes from a natural experiment involving a station acquired by a new owner, similar to the analysis captured in the acquisitions of KPTM and KWKT in Figure 19. The overall effect can be summarized with one point and confidence interval. I then group and color new owners and omit the stations that did not experience a significant change in slant after acquisition.

As shown in Figure 20, a significant pattern of change in political slant occurs after acquisitions. Many stations experienced significant change in political slant in both directions after being acquired by a new owner relative to their own stance before the acquisition, even after the model accounted for an expected baseline change. This analysis presents a more complex picture of ownership's impact on political slant in local news. While some stations significantly slant to the right, there are also stations that slant more left. This finding suggests that the impact of acquisition may not be a simple mechanism, and that different media companies may have different political agendas. When acquisitions from the same media company are combined and the average impact across multiple acquisitions is calculated, evidence suggests that different owners impact political slant in the news differently. In particular, Media General (-0.67) and Hearst (-0.65) acquisitions are associated with the most left slant, and Sinclair (0.57) and Gannett (0.56) acquisitions are associated with the most right slant.

Figure 20 Impact of Ownership Change on Political Slant

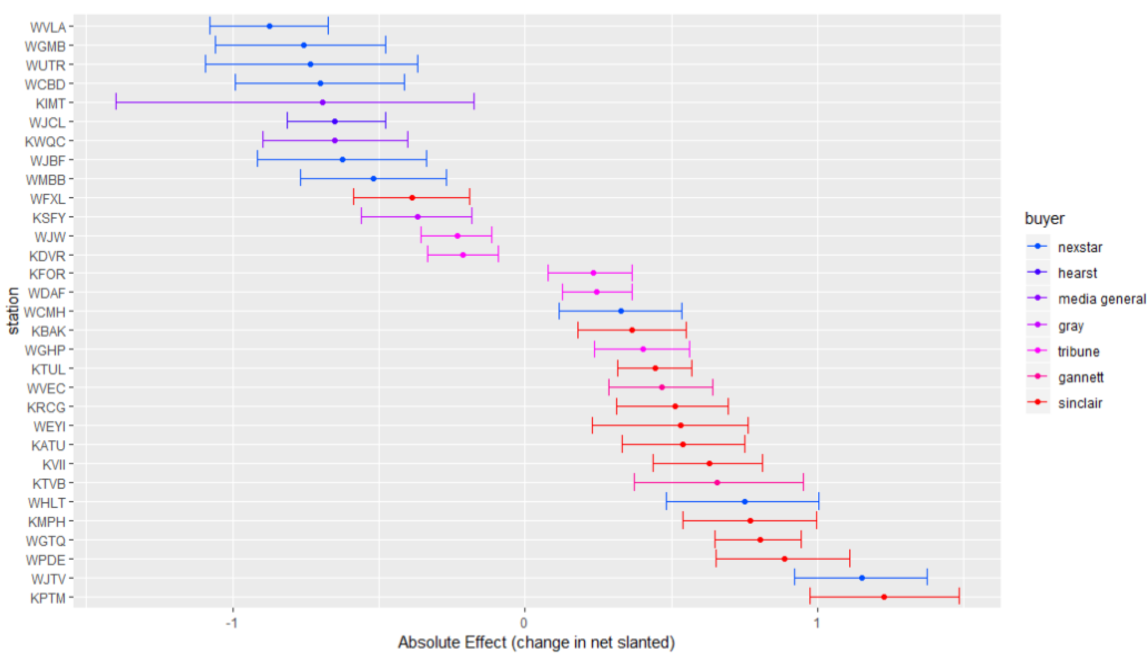

I conduct a similar analysis of impact of stations acquisition in terms of information diversity. The results are shown in Figure 21. Once more, an impact in both directions is observed, and a striking pattern is detected: some stations experienced a sharp drop in their information diversity after been acquired by a new owner. There is clearly asymmetry between stations that had an increase in information diversity and stations that had a decrease in information diversity. The increases are usually very mild, while the decreases could be dramatic. Aggregated to the level of media ownership, Media General, Sinclair, and Tribune acquisitions are associated with the largest decreases in information diversity, while Gannett, Nexstar, and Gray acquisitions are associated with the largest increases in information diversity.

One hypothesis is that a new owner may want to leverage economies of scale and cut local production budget and local newsroom staff. Hence, local news may lose its local focus and become more and more nationalized at the same time, as information diversity drops. This could be a dangerous trend, as local news is critical to local governance and democracy at the community level. Future studies of the mechanism that explains changes in information diversity in local news are needed. Moreover, a large-scale study of how the content and focus of local news evolves over time across different geographic areas at the national level is especially needed.

Figure 21 Impact of Ownership Change on Information Diversity

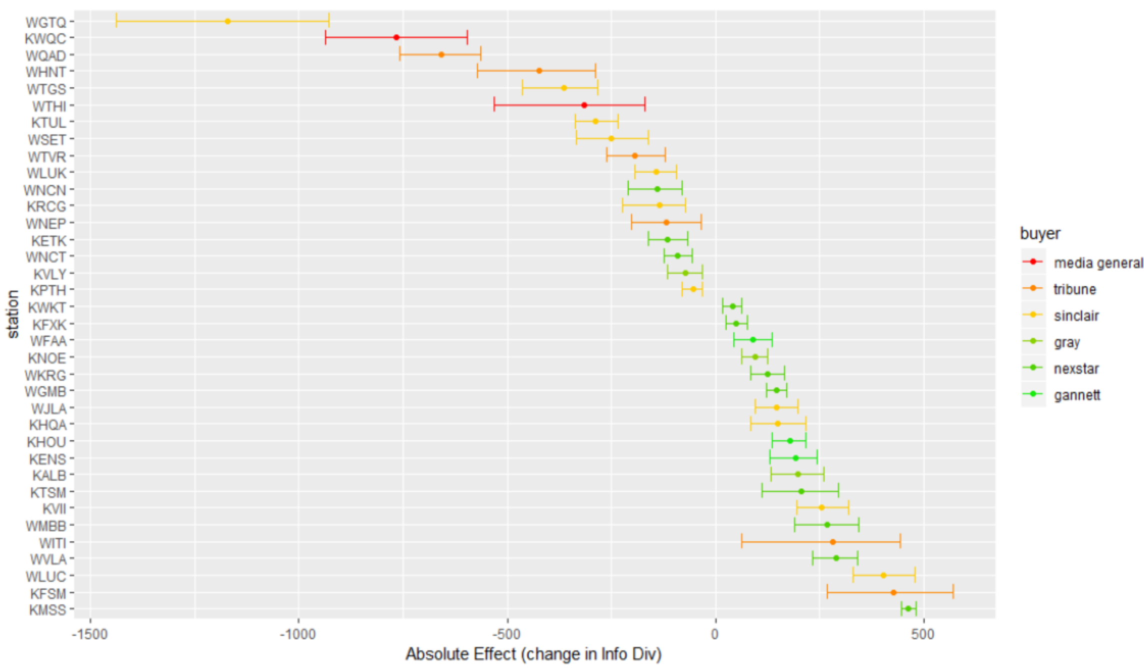

Conclusion

Broadcast TV in the U.S. is an information system comprised of hundreds of local television stations that both produce their own information and syndicate information from other sources (e.g., major networks). In this system, a few media conglomerates are behind the vast amount of information produced and disseminated. There has been an increasing trend of media consolidation in recent years, which raises a very real concern that conglomerate owners have both the means and motivation to skew information. In this study, I explore how media ownership impacts information diversity and the political polarization in TV news. To examine the impact of ownership, I focus on the natural experiments of hundreds of station acquisitions by new owners over a six-year time span. The changes of ownership allow us to disentangle the impact of owner of a TV station from other characteristics that may also be correlated with the information produced by a station. I find that ownership has a significant and causal impact on both the political slant in news and its information diversity. The impact is more complex than previously expected. Specifically, I observe that changes in political slant could be in either direction for a station after an acquisition. When aggregated at the owner level, different owners seem to have different political agendas. In terms of information diversity, the most worrisome finding is that some stations witness a dramatic drop in information diversity of its news content after acquisition. The implication of this drop in information diversity is profound. Overall, I find evidence that suggests large owners of local news stations can skew information to emphasize particular views, perspectives, and coverage. This is particularly concerning, as media consolidation has increased in recent years. More attention is needed from both researchers and citizens in terms of how this trend impacts the quality and comprehensiveness of local journalism and local news content.

CHAPTER THREE

A Warped Mirror? Skewed Coverage of Gun Violence in the Mass Media

Introduction

Gun violence in the United States has reached epidemic proportions, as guns have been identified as the primary cause for more than 30,000 fatalities and 80,000 injuries annually. Dozens of high-profile gun violence events (e.g., mass shootings, school shootings) over the past few years have provoked a national conversation and a heated debate about gun violence, its regulation, and related public policy. However, the vast majority of gun violence incidents that happen daily—including suicides, accidental shootings, and domestic violence shootings, among others—rarely attract the national media's attention, despite the fact that these incidents account for a much larger proportion of gun deaths and gun injuries.

The public's perception and understanding of the gun violence epidemic in the U.S. is driven by their exposure to media coverage of gun violence incidents, predominantly through televised news. Mass media, like TV news, has the power to set national agendas (Behr & Iyengar, 1985; McCombs & Shaw, 1972), activate public discussions (King, Schneer, & White, 2017), influence citizens' views and perceptions of national affairs (Gilliam & Iyengar, 2000; Iyengar, 1987), and has a real-world consequence with respect to a wide range of social outcomes (DellaVigna & Kaplan, 2007; Eisensee & Strömberg, 2007; Kearney & Levine, 2015). In an increasingly polarized world in which citizens are more likely than ever to view media that only reinforces their particular worldview. Shared narratives and mutual agreements about the basic framework of a society are necessary for a society to function, if not flourish. Therefore, broadcast media plays a particularly important role in helping citizens who share different perspectives nevertheless appreciate that a democracy allows space for these different perspectives to coexist. To this end, the public's understanding of gun violence is important, as public opinion will impact the pressure that citizens exert on policy makers who influence gun policy (See Figure 22).

Despite the attention that gun violence receives in the media, an empirical understanding of the nature of this coverage and its impact is still remarkably limited. Most importantly, does the media offer citizens an accurate depiction of the gun violence problem? Or do citizens instead "look into a warped mirror" when they view media coverage on gun violence? To answer these questions, I conducted a nationally representative survey and found that the public is generally ill-informed regarding the prevalence of different types of gun violence in the U.S. Notably, when asked about the relative prevalence of different types of gun violence, survey respondents' responses were worse than random. In addition, I find that differences in levels of informedness with respect to gun violence prevalence reflects respondents' gender, race, and education level, both in terms of their general awareness of the epidemic itself and their accurate sense of of the specific types of gun violence that contribute to the epidemic.

Figure 22: Social process of gun violence incidents to gun policy

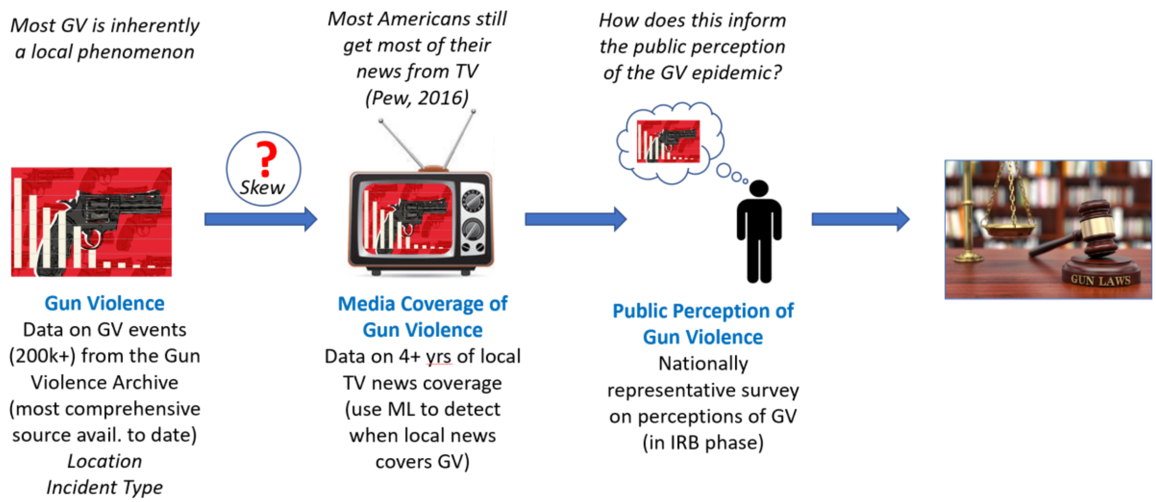

Media coverage of gun violence incidents plays an important role in informing the general public about the daily events happening in their local areas. A 2016 Pew survey revealed that the majority of Americans obtain their news from televised cable, network, and local nightly coverage (Pew Research Center, 2016), which suggests that an understanding of gun violence issues stems from this coverage. That said, the majority of gun violence incidents in the U.S. are inherently local events that do not receive as much attention as mass shootings or school shootings, for example. In short, Americans largely learn about gun-related crime through local media coverage. How has gun violence been covered and reflected in local TV news broadcasts? This question is challenging to answer, as a systematic evaluation of local information on a national scale proves equally challenging. However, such a systematic evaluation is much needed for two reasons: 1) without an empirically-informed approach to this question, individual beliefs about gun violence are subject to potential cognitive bias, due to views developed early in life and to the use of availability heuristics; and 2) a consensus for a policy solution to address gun violence issues cannot be reached without a shared understanding of a set of facts, around which conversation about gun issues can unfold.

A systematic investigation of the media's coverage of gun violence is needed to provide a holistic view of this epidemic. Such a study would require a comprehensive record of gun violence incidents that occur throughout the country as well as thorough digital records of media reporting at the local level on incidents of a national scale. Both sources of data have not been available until very recently. Thus, this data now allows for an accurate study of this issue, thanks to the digitization of both gun violence incident records and televised news, as well as to recent advancements in Natural Language Processing. Using the most comprehensive records to date of individual gun violence incidents across the U.S. over a multi-year period (2013-2018), I examine the media's coverage of gun violence by analyzing a large-scale TV transcript data of around 800 local TV stations in all U.S. media markets over the same time period with machine learning and natural language processing (NLP) techniques.

I found systematic skews in media coverage based on incident types. Specifically, workplace shootings, school shootings, mass shootings, and the shooting of police officers were systematically covered more frequently; in contrast, accidental shootings, suicide, and domestic violence events were systematically covered less frequently. In addition, ***Officer Victim*** is covered twice as much as ***Officer Shooter***. Assault weapon incidents gets much more coverage than more prevalent incidents that are committed with handguns. I also find that U.S. regions with high vs. low gun ownership rates receive different types of media coverage, which has the potential to divide an already divided public. Specifically, concealed carry victims are positively associated with gun ownership rates while concealed carry perpetrators are negatively associated with gun ownership rates. School and workplace shootings are both negatively associated with gun ownership rates.

The findings of this study cannot be explained by a simple, single factor. After all, A complex function of numerous factors impact journalist’s decision and hence media coverage, particularly when it comes to violence and crime. Some studies suggest that "statistical unusualness" is newsworthy, while others do not. Also, other factors may include the impact of gender, incidents that occur multiple times, crimes committed by strangers, the notion of an "ideal victim," or a victim portrayed as "more innocent" than other victims (e.g., drug users). A deeper investigation of the mechanisms behind skewed coverage is beyond the scope of this study. However, regardless of the mechanisms, this skewed coverage phenomenon is important, as it has consequences and implications for citizens. The human judgement of prevalence is affected by heuristic cues and, hence, could be biased by an availability of information (Lichtenstein, Slovic, Fischhoff, Layman, & Combs, 1978; Tversky & Kahneman, 1973, 1974). In addition, numerous studies over the past few decades have demonstrated that skewed coverage in the media impacts collective opinion and public behavior. (Andrews & Caren, 2010; Martin & Yurukoglu, 2017). This is particularly salient in criminal justice and violence incidents, shaping public beliefs and fears that would lead to downstream policy pressure beyond its impact, especially in terms of how individuals gauge the prevalence of incidents (Beale, 2006; Chyi & McCombs, 2004; Haider-Markel & Joslyn, 2001).

Lastly, using a nationally representative survey on public perceptions of gun violence, I investigate the impact of media coverage of gun issues. I find that several factors influence how people allocate resources when trying to solve gun violence problems; of these factors, the amount of news coverage has a very significant impact. By linking these factors to budget allocation, I reveal the different ways that respondents assign importance to different types of gun violence.

Public's Misguided Perceptions of Gun Violence Prevalence

I conduct a nationally representative survey to determine the public's attitude toward different types of gun violence. Using an online national representative panel with quotes for state, race, age, gender, income, and education, I surveyed 1,018 individuals across all 50 states and the District of Columbia from January 10–20, 2020.

To examine the accuracy of respondents' perceptions with respect to the relative prevalence of different types of gun violence, I asked respondents to rate the frequency of each type of gun violence in their local area on a scale from 1 to 10. I then computed errors of their ratings by comparing their ratings to the real frequencies of different types of gun violence in their respective states (after the transformation in which the most frequent type is normalized to 10). To gauge how accurate their ratings of gun violence prevalence were, I calculated absolute errors with which I compared the true errors of their ratings to random errors, for which I generated random numbers between 1 and 10 and compared them to real frequencies as simulated errors. When I did so, I found that the true error was generally greater than the random error. As shown in Figure 23, the distribution of the true error (with a median value of 3.8) spread in higher values in support when compared to that of the random error (with a median value of 3).

Figure 23: True error is greater than random error

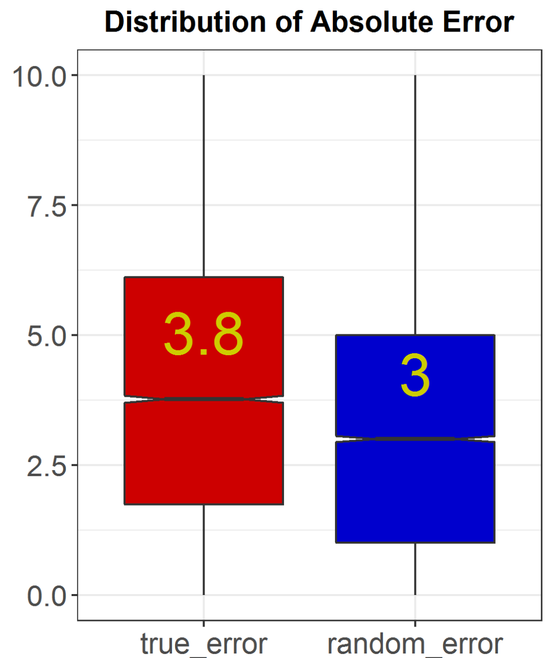

In addition, I explored this finding further by conducting a t-test of the means of the two distributions. This test result confirmed that the mean absolute error of the survey answer (mean=4.1) was worse than a random guess (mean=3.3) and was highly significant (p=1e-15). In sum, respondents judged the relative frequency of different types of gun violence poorly, which reflects how ill-informed respondents proved to be in terms of this topic.

Furthermore, I explored what characteristics are associated with errors of rating. When I examined the relationship between the demographics of respondents and their error of rating at the individual level, I found that female and racial minorities tend to have higher errors (the base group in the model for gender is **Female** and for race is **White**). Higher income and education are generally associated with lower errors (the base group for income is **Less than $50,000** and for education is **Less than High**

School).

Table 9: Characteristics associated with more errors in rating when determining which type of gun violence is more common

	Dependent variable:
	common_true_error
Male	-0.275***
	(0.056)
age_between35And54	0.008
	(0.064)
age_aboveAnd55	-0.065
	(0.069)
income_between50kAnd100k	-0.200***
	(0.062)
income_between100kAnd200k	-0.117
	(0.082)
income_above200k	-0.316**
	(0.124)
education_someCollege	-0.276***
	(0.075)
education_assocaiteOrBachelorDegree	-0.263***
	(0.069)
education_graduateDegree	-0.108
	(0.098)
race_african_american	0.519***
	(0.086)
race_asian	-0.055
	(0.120)
race_hispanic	0.273***
	(0.078)
race_other	0.372***
	(0.129)
Incident Type Fixed Effect	Yes
Observations	10,764
R^2	0.175
Note:	*p<0.1; **p<0.05; ***p<0.01

Skewed Coverage of Gun Violence Incidents in the Mass Media

Our perception of gun violence comes from the lens of the media, particularly the media's coverage of gun violence incidents. As most citizens do not have direct and personal exposure to gun violence, poor perceptions of gun violence's relative prevalence begs an important question: are these perceptions skewed because citizens relate to this topic through "a warped mirror of reality" that reflects skewed coverage of different types of gun violence in the media? I explore this question further by examining the local news coverage that reflects a comprehensive record of gun violence in the U.S. over a four-year period.

For each shooting incident, I match the incident to local news coverage based on the location of the incident and the time it occurred. To determine the extent of coverage of gun violence incidents in local televised news transcripts, I adopt a robust procedure that leverages natural language processing and hundreds of thousands of reports of gun violence incidents to develop a lexicon of gun violence features. Further details are provided in the SI Appendix. Importantly, while this method can detect coverage of gun violence in absolute terms, I adopt a modeling strategy that further eliminates potential noise in this signal by considering only the change in detected coverage for stations in a market where an incident has occurred in a short time period following an incident, relative to coverage in a short time period prior to the incident's occurrence. I first identify local TV channels in the Designated Market Area where the incident occurred and look at a 2-day window before and after the incident date. In so doing, I attribute the difference in the amount of gun violence coverage to a specific incident. The strategy is illustrated in Figure 24.

Figure 24: Attribute coverage to a gun violence incident

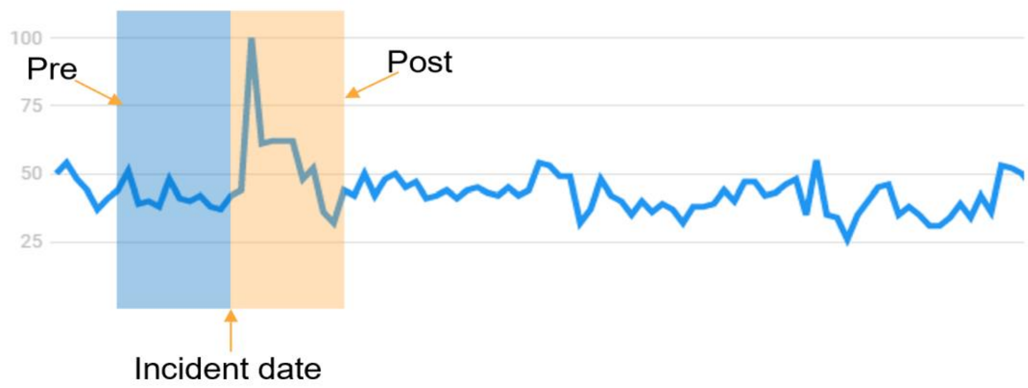

I model the lift in news coverage with a linear fixed effect model in which the dependent variable is the difference in news coverage of gun violence in the pre- and post- window. The coefficients of interest are vector β, which represent how each incident characteristic is associated with changes in the amount of news coverage. The model also includes $channel$ and $year \times month$ fixed effects to account for market heterogeneity as well as seasonality and common time trends.

$$\Delta Y_{ij} = \beta X_i + \gamma_{channel} + \delta_{year\times month} + \epsilon_{ij}$$

where ΔY_{ij} is a lift in gun violence coverage for incident i on channel j, and X_i is a vector of dummy variables that represents incident characteristics. Figure 22 visualizes the estimated coefficients of β in this model. Each coefficient represents the amount of change in news coverage of gun violence in the pre- and post-window that is associated with the corresponding incident characteristics. I rank incident characteristics by the size of their point estimates in a descending fashion.

Despite their infrequency, **Workplace Shootings** and **School Shootings** receive the most coverage. In contrast, suicide by gun, which is the most common cause of gun death, systematically receives the least coverage. **Domestic Violence** and **Accidental Shootings** also receive less coverage, even though they are two relatively common types of gun violence. In my data set, I identify 10,762 domestic violence events involving guns and 8,158 accidental shootings, which are 62 times and 47 times more common than workplace shooting (173 incidents), respectively (see the SI Appendix for a full description of the frequency of incident types in our data). **Officer-involved Shootings** (officer victim, bystander shot or killed, officer shooter) receive greater coverage in local news. Interestingly, shooting incidents in which an officer was the victim receive significantly more coverage than incidents in which an officer was the shooter. Incidents involving children (**Child Shooter** or **Child Victim**) are also associated with greater news coverage. Finally, **Assault Weapon** incidents receive more coverage than more prevalent incidents that are committed with handguns.

Figure 25: How each incident's characteristics are associated with changes in the amount of news coverage

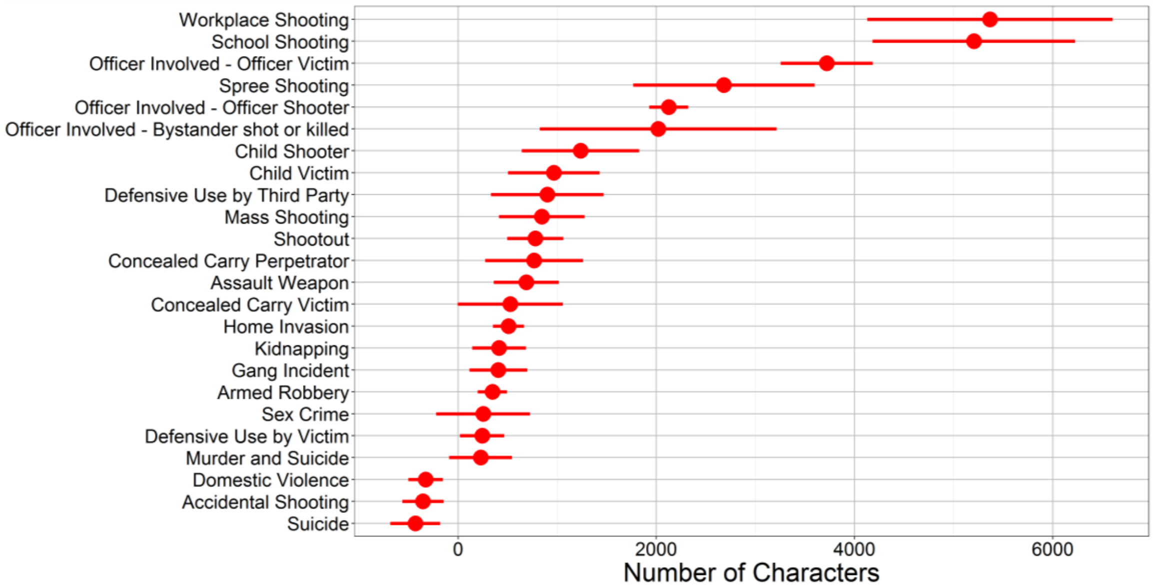

Different populations are exposed to different media coverage

The market areas reached by televised local news in the U.S. differ in their population makeup. As a result, local news coverage of gun violence may differ across these areas, providing a different impression of gun violence for residents of some areas as compared to others. One interesting area characteristic is that of a population's gun ownership rate. Residents of areas with high- and low gun ownership likely share different beliefs about gun violence and take different stances in debates on gun control policy. As such, a set of important questions emerge: does media coverage provide the same portrayal of gun violence to residents across all market areas? Or does the media instead exacerbate disagreement between these populations through information availability bias, which creates very different impressions with respect to the nature of gun violence? To answer these questions, I analyze how gun violence coverage varies with respect to gun ownership rates.

Accurate statistics on state-level gun ownership across the U.S. are not available. Therefore, I use an up-to-date proxy for state-level gun ownership that is determined from both the ratio of firearm suicides to all suicides and the per capita number of hunting licenses. This measure has been validated and shown to correlate strongly with survey-measured gun ownership in years when such surveys were available (Siegel, Ross, & King, 2014). To understand whether gun ownership is related to local news coverage of gun violence incidents, I incorporate the proxy measure for gun ownership in the model and interact it with the vector of incident characteristics to investigate the heterogeneity of this relationship across different types of gun violence. I then classify a region as either a high or low gun ownership region if it is either one standard deviation above or below the mean, respectively.

$$\Delta Y_{ij} = \beta_1 X_i + \beta_2 X_i \times high\_gun\_ownership\_region$$
$$+ \beta_3 X_i \times low\_gun\_ownership\_region$$
$$+\gamma_{channel} + \delta_{year \times month} + \epsilon_{ij}$$

Here, $high\_gun\_ownership\_region$ and $low\_gun\_ownership\_region$ are two dummy variables that represent whether an incident occurs in a region that has a gun ownership rate that is greater or smaller than one standard deviation above or below the average gun ownership rate. Figure 23 visualizes the differential coverage for high and low gun ownership regions for each incident type. For brevity's sake, I omit characteristics for which estimates do not vary significantly from zero (see the SI Appendix for complete estimates). Interestingly, I find that regions with low gun ownership rates have a higher level of media coverage than regions with high gun ownership rates for many gun incident types. As exceptions, **Gang Incidents** and **Defensive Use by Third Party** reverse this relationship: they tend to be covered more frequently in regions with high gun ownership rates. In addition, I observe the asymmetric relationship between **Concealed Carry Perpetrator** and **Concealed Carry Victim**, for which a **Concealed Carry Perpetrator** has higher levels of media coverage in regions with low gun ownership rates, whereas the opposite is true for **Concealed Carry Victim** (although not significant). In summary, high and low gun ownership regions receive different levels of media coverage of gun violence, which has a tendency to further divide an already divided public.

Figure 26: Media Coverage of gun violence in high and low gun ownership regions.

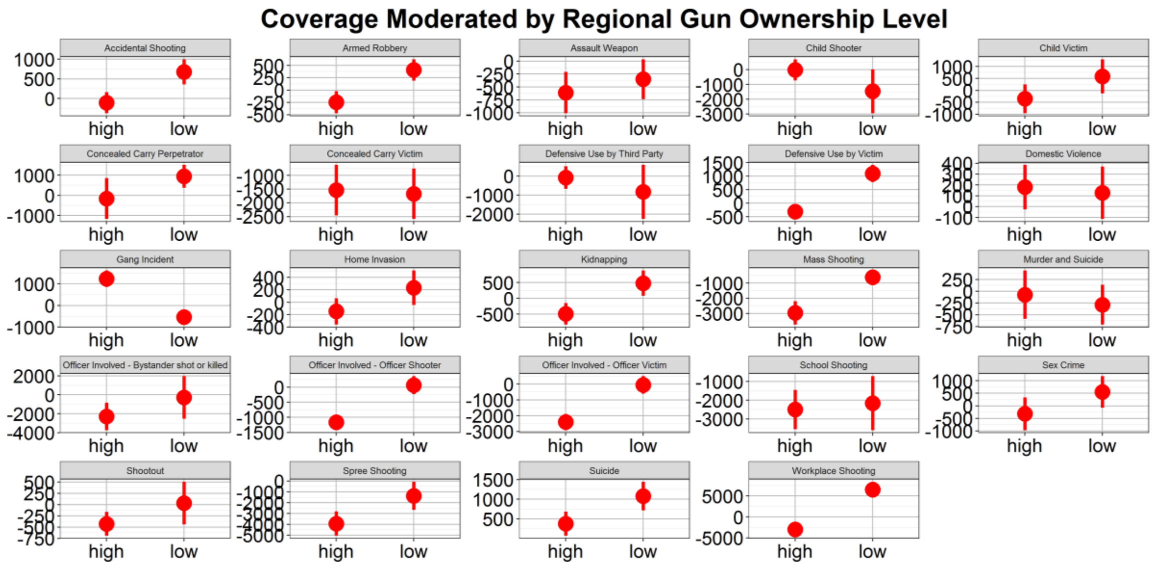

Note: Error bars indicate a 95% confidence interval.

Does skewed media coverage of gun violence incidents matter? To answer this question, I designed and conducted a nationally representative survey to analyze perceptions of gun violence. In the survey, I asked respondents to allocate a limited budget (i.e., 100 points in total) to help mitigate different types of gun violence as well as to share their opinions on those type of gun violence. In Table 9, I show that various factors lead to their allocation. This includes how informed they believe they need to be about such incidents, how important it is that the media report such incidents, and how common and how severe they rate the incidents themselves. Most importantly, I then incorporate my model estimate of skewed coverage in equation (1) in the regression and find that it is significant and has the highest magnitude among the factors associated with a respondent's budget allocation decision. This finding indicates that media coverage of gun violence incidents plays an important role in how people perceive the relative importance of different incident types.

Table 10: Various factors are associated with how people allocate resources to judge different types of gun violence incidents

	Linking Factors to Budge Allocation		
	Budget Allocation		
	(1)	(2)	(3)
Report rating	1.025***	0.914***	0.901***
	(0.071)	(0.072)	(0.071)
Informed rating	0.581***	0.476***	0.465***
	(0.072)	(0.072)	(0.072)
Commonality rating	0.660***	0.717***	0.725***
	(0.055)	(0.055)	(0.054)
Severity rating	0.749***	0.675***	0.632***
	(0.066)	(0.066)	(0.065)
Coverage estimate		0.922***	
		(0.084)	
Coverage average			0.957***
			(0.066)
Respondent Fixed Effect	Yes	Yes	Yes
Observations	10,764	10,764	10,764
R^2	0.192	0.202	0.209

Note: *p<0.1; **p<0.05; ***p<0.01

Discussion

Public perception of gun violence is important, as public opinion will impact the pressure citizens exert on policy makers (e.g., the elected officials), which ultimately impacts gun policy. This perception stems from our exposure of media's coverage of gun violence incidents. To agree upon on what solution ought to be, we need to first reach a census of what the problem is. By analyzing patterns of how the media covers these incidents, I highlight potential blind spots in journalistic practices and media operations and hence help us become more aware of our biases when thinking about policy solutions for gun violence. A well-recognized shared reality is essential to the democracy process. Exposing the bias and skewness in media coverage of gun violence to internal stakeholders ((journalists, news producers) and external stakeholders (general public, policy makes) via rigorous empirical evidence will force them to confront the world as it is rather than how it is being projected from a warped mirror. Providing this national-scale systematic empirical evidence, I hope to facilitate the forming of a common ground upon which we can begin a meaningful and constructive conversation around gun violence in American society and its media coverage.

APPENDIX

Data and Descriptive Statistics

Table A1: Distribution of Article Characteristics

Statistic	N	Min	Pctl(25)	Median	Pctl(75)	Max
Popularity	5,532	2	96	392	1,617	48,850
Size	5,532	19	3,678.8	8,788.5	22,122	147,469
Tenure	5,532	18	422.8	538	627	808
CharCount	2,766	501	1,331	3,180	7,292	159,912

Note: 1. We display the min, max and each 25 percentile values for popularity, size, and tenure which are all pre-shock article characteristics (in both treatment and matched control groups). Binary variables used in main analysis are binned by corresponding median values; 2. CharCount for control group is defined as zero and the distribution displayed above is only for articles in the treatment group.

Table A2: Distribution of Traffic Flow of Matched Articles in Clickstream Data

Statistic	N	Min	Pctl(25)	Median	Pctl(75)	Max
Inbound	2,034	14	251.2	869	2,329.8	53,088
Outbound	2,034	3	3	40	219.8	10,778
Link Count	2,034	0	1	3	7	149
Link Share	2,034	0.001	0.054	0.114	0.214	0.647

Note: We display the min, max and each 25 percentile values for Inbound traffic, Outbound traffic, Link Count (number of incoming links), and Link Share (the proportion of link traffic in the inbound traffic).

Table A3: Balance Checks for Clickstream Matching

	Inbound	Outbound	Link Count	Link Share
Control	1916	192	5.5	0.15
Treatment	1915	192	5.6	0.15
t-test (p-value)	0.99	0.99	0.78	0.99

Note: The table illustrates the quality of our matching procedure for clickstream data. “Inflow” and “Outflow” are traffic per month per article. “Link Count” is number of incoming links per article. “Link Share” is the proportion of link traffic in the total inflow traffic for each article. T-tests indicated that we cannot reject the null hypothesis that articles in treatment and control group have the same mean across all four characteristics.

Matching on Article Topics

The model-free plot in Figure 2 indicates that the control articles matched on pre-shock article characteristics closely mirror the treated articles in seasonality and natural time trends prior to the shock. This constitutes strong evidence that the controls serve as good counterfactuals for treated articles and capture what would have happened had they not received exogenous content contributions. Despite this, one may still have the concern that the topic distribution of treated and control groups may not be exactly the same. In this section, we reproduce our analysis with an alternative control sample that matches with treated articles on topics in addition to the other pre-shock article characteristics of popularity, size, and tenure. The results are qualitatively similar with only very small differences in magnitudes compared with the results presented in the main analysis of the paper, giving us strong confidence that our results are insensitive to matching procedure choices. The major challenge of topic matching is that each Wikipedia article is associated with multiple topics or categories and collectively the topic distribution of all treated articles resides in a high dimensional space. No traditional matching method is designed to deal with this problem. We adopt a novel two-step procedure to tackle this unconventional matching problem and ensure that we can match reasonably well on topics.

In the first step, we construct a pool of candidate control articles to use for matching through a random sampling procedure that leverages the "Category" information associated with each article. Each Wikipedia article has a set of "Category" labels added by its editors. Because category labels are user-defined, they are prone to errors and not subject to sanitization, e.g. very few articles have exactly the same set of labels and very few labels appear multiple times in a randomly selected set of articles. We adopt a strategy to leverage category information in our sampling procedure that avoids potential pitfalls. To do this, we randomly draw articles only from the categories to which our treated articles belong, where the number of draws from each category is proportional to the category frequency in the treated sample. This sampling procedure can achieve sample-level matching on categories but does not guarantee a direct correspondence between each individual control and treated article. For simplicity, we refer to this category-matched sample as "the control sample" in the remainder of this discussion.

In the second step, to better account for direct topic matching, we turn to the popular text-mining technique of Latent Dirichlet Allocation (LDA) topic modeling. Topic modeling is a frequently used machine learning tool for discovering hidden semantic structures in a corpus of text. We use LDA topic modeling to discover the latent topics from the text of each article in an unsupervised fashion and then match each treated article with a control article in the latent topic space. We train our topic model with the complete text of English Wikipedia (about 5.3 million articles and 15 GB) extracted from the October 2018 Wikipedia data dump. The number of topics is set to 100, though our method is robust to different choices. Manual inspection of word distribution of each topic indicates that our model captures coherent latent topics from the texts. Some example topics from our topic model are displayed in Figure A1. Next, we apply the topic model to treated and control articles to obtain their topic distribution in the latent topic space. Using this, we generate a tailored pool of candidate control articles C_i for each treated article T_i by searching for all articles in the control sample that are sufficiently similar on topic, according to the cosine similarity $\cos\left(\overrightarrow{Topic}[C_i], \overrightarrow{Topic}[T_i]\right) > 0.6$. We experimented with different cosine similarity thresholds and the results are robust to the choice of threshold; naturally, the size of the matched sample monotonically decreases with a stricter similarity requirement. Finally, we use Coarsen Exact Matching on the treated and topically similar control samples to further match on the other pre-shock article characteristics of popularity, size, and tenure. We opt for a k-to-k matching solution by choosing the closest matched control article in terms of Euclidean distance. The above procedure yields 2,747 pairs of matched treated and control articles. For each article, we construct a panel of weekly pageviews from 26 weeks before the shock to 26 weeks after (excluding the shock period itself). Our final sample consists of a balanced panel of 52 periods for 5,494 articles, or 285,688 observations at the article-week level. Finally, we redo our analysis of direct effect on this new sample -- the results are displayed in the table A4.

Figure A1: Some Example Topics

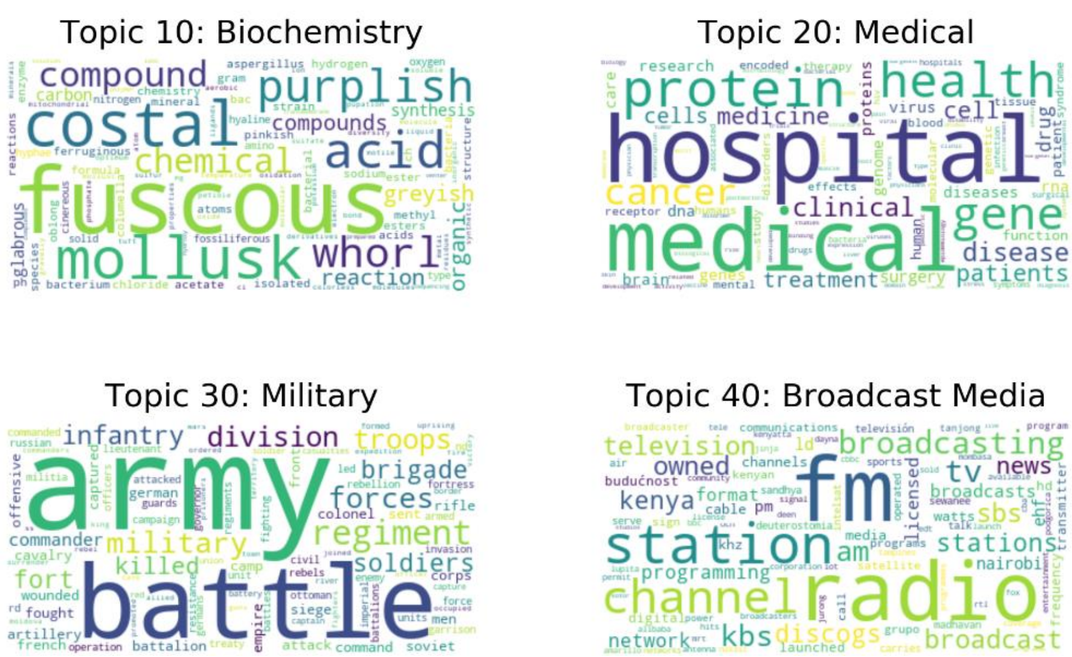

Table A4: Direct Effect with Matching on Topic

	Scaled pageviews		
	(1)	(2)	(3)
PostShock	0.106***		
	(0.016)		
PostShock*log(char count)		0.032***	0.060***
		(0.004)	(0.006)
PostShock*old article			-0.060**
			(0.024)
PostShock*popular article			-0.116***
			(0.024)
PostShock*long article			-0.030
			(0.025)
Article fixed effect	Yes	Yes	Yes
Time fixed effect	Yes	Yes	Yes
Observations	285,688	285,688	285,688
Adjusted R^2	0.172	0.172	0.174

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Matching on Network Characteristics

In this section, we replicated our analysis for direct effect of content contribution based on a sample that is matched on network characteristics. We note that the so-called curse of dimensionality affects every matching method - as the number of covariates over which we match grows, the chance of finding matches with similar values of all covariates rapidly goes to zero (King and Nielsen, 2019). Hence, we were very careful about selecting matching variables in the main analysis because adding variables to the matching procedure comes at a cost of lowering the chance of finding good matches and reducing the size of matched sample. We think the most important variable to match on is pre-popularity of an article as it conveys the information about how much attention an article receives prior to the treatment period and we want to compare the impact for treated and control articles that receive the similar amount baseline attention.

Some of the network characteristics, e.g., in-degree, incoming internal traffic, or incoming external traffic, carry information about how the attention arrives at an article, not the amount of attention arriving, which is already accounted for by pre-popularity. We therefore regarded these to be less relevant to the matching procedure and analysis. Still, matching on in-degree and in-traffic might be appropriate, as it allows us to compare treated articles to control articles that receive attention at the same proportion through channels both internal and external to Wikipedia. We would like to demonstrate that our results are also robust to matching based on network characteristics.

We conducted additional analysis and show that the results are very similar when we match on some of the network characteristics of the articles. We conducted Coarsened Exact Matching (CEM) on three network characteristics of an article, i.e., in-degree, average incoming external traffic, average incoming internal traffic. The matched sample consists of 2,058 treated articles and control article, respectively. We did a balance check and it shows that we cannot reject the null hypothesis that articles in treatment group and control group have the same mean values across number of incoming links, internal traffic from other pages in Wikipedia and external traffic (See Table A5). We replicated our analysis of direct effect with this new matched sample and the results are very similar as in the original (See Table A6). The original model, however, is preferable as: 1) we are matching on characteristics on which we evaluate heterogeneous treatment effects; and 2) pre-popularity is already a very good control for incoming traffic and in-degree.

Table A5: Balanced Check for Matched Sample of Network Characteristics

	Number of Incoming links	Internal traffic/month	External traffic/month
Control	7.20	447	2014
Treatment	7.20	446	2015
t-test (p-value)	0.96	0.90	0.96

Table A5 illustrates the quality of our matching procedure. It compares pre-shock network characteristics of articles in the matched groups. T-tests indicate that we cannot reject the null hypothesis that articles in treatment and control group have the same mean across all three network characteristics.

Table A6: Direct Effect with Matching on Network Characteristics

	Normalized pageviews		
	(1)	(2)	(3)
PostShock	0.118***		
	(0.015)		
PostShock*log(char count)		0.037***	0.068***
		(0.004)	(0.006)
PostShock*old article			-0.074***
			(0.026)
PostShock*popular article			-0.140***
			(0.027)
PostShock*long article			-0.023
			(0.024)
Article fixed effect	Yes	Yes	Yes
Week fixed effect	Yes	Yes	Yes
Observations	214,032	214,032	214,032
Adjusted R^2	0.138	0.139	0.141

Notes:
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Checks of Modeling Assumptions for Multiplicative Interactions

Binning Estimates

The plots below serve as a diagnostic tool for two main modeling assumptions: common support and linear interaction effect. The distribution of the covariate presented at the bottom of each plot demonstrates that the assumption of common support, which is needed for a multiplicative interaction model, is satisfied. Moreover, the number of bins in the plot is two and equal-sized bins are created based on the distribution of each covariate. The plots confirm that using two bins to represent low/high values for the covariates is a reasonable choice. We provide a set of diagnostic statistics to further justify that choice in Table A7.

Figure A2

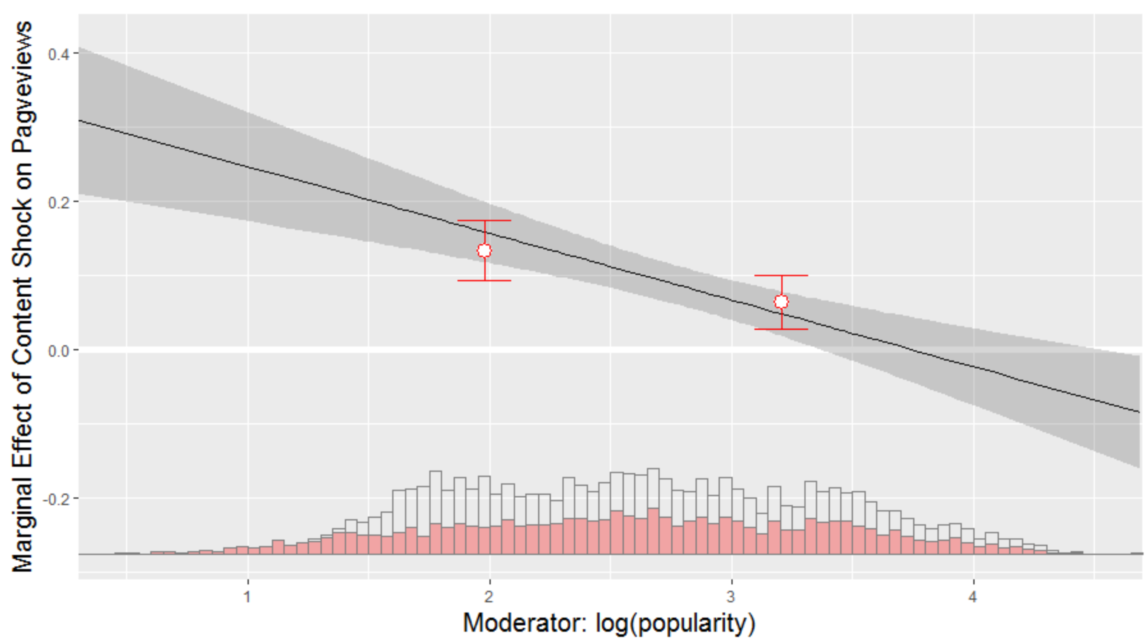

Figure A3

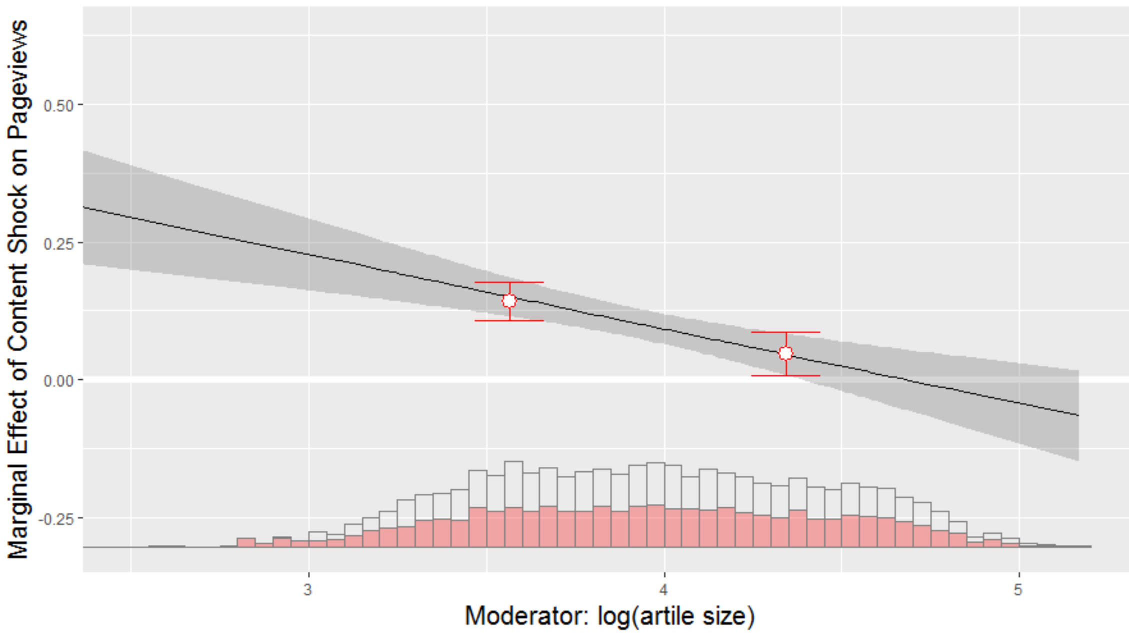

Figure A4

Table A7: Model Estimates and Test Statistics of Binning Estimators

	range	median	coef	se	CI_lower	CI_upper	t-test
log(popularity):high	[0.23, 2.59]	1.980	0.133	0.021	0.093	0.173	0.012
log(popularity):low	(2.59, 4.69]	3.209	0.064	0.019	0.027	0.099	
log(article size):high	[1.28, 3.94]	3.566	0.143	0.018	0.108	0.177	3e-04
log(article size):low	(3.94, 5.17]	4.345	0.048	0.020	0.009	0.087	
articlue tenure:high	[18,538]	424	0.109	0.021	0.068	0.149	0.09
articlue tenure:low	(538,808]	627	0.063	0.017	0.029	0.097	

Note: 1. The binning estimates for the three article characteristics correspond to the above three plots; 2. The column "t-test" displays the p-value of t-test for the two binning estimates.

Kernel Estimates

The kernel method produces non-linear marginal effects that are much more flexible and closer to the effects implied by the true data-generating process. The standard errors are produced by a non-parametric bootstrap. The below kernel plots show that covariates exhibit linear behavior over most of their range provide further evidence that our linear interaction model with two bins well approximates the more flexible models while also maintain good interpretability.

Figure A5

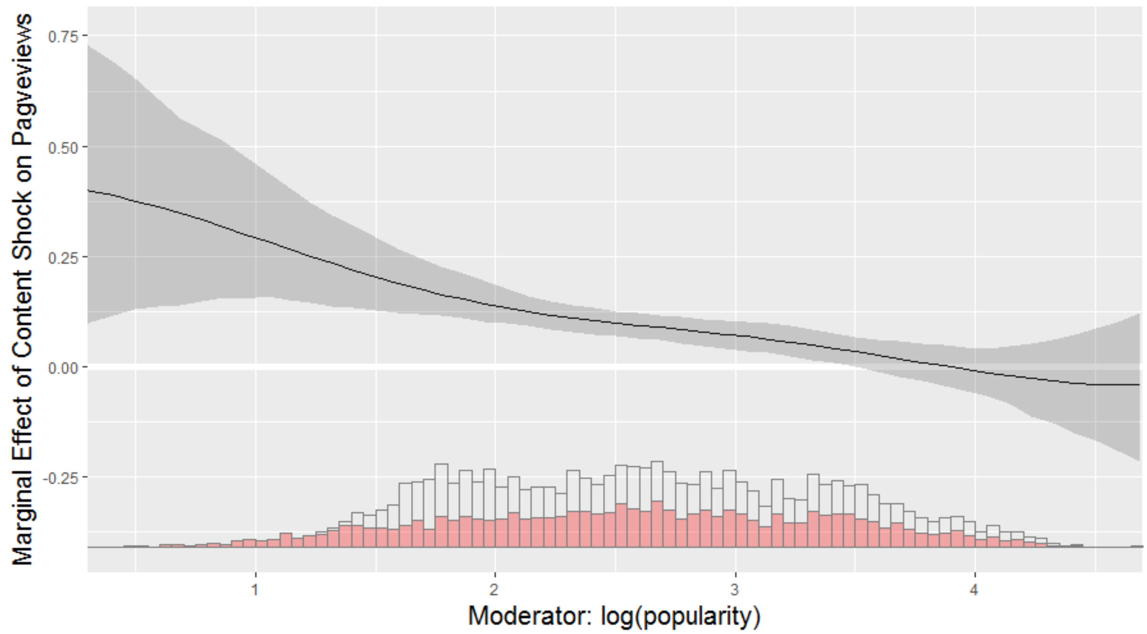

Figure A6

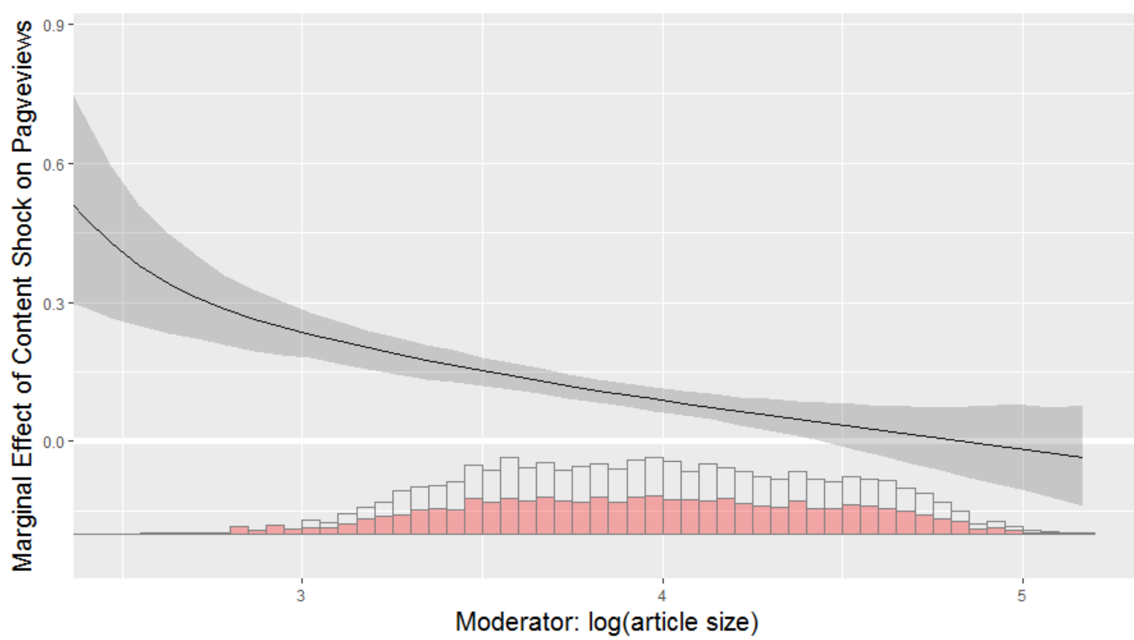

Figure A7

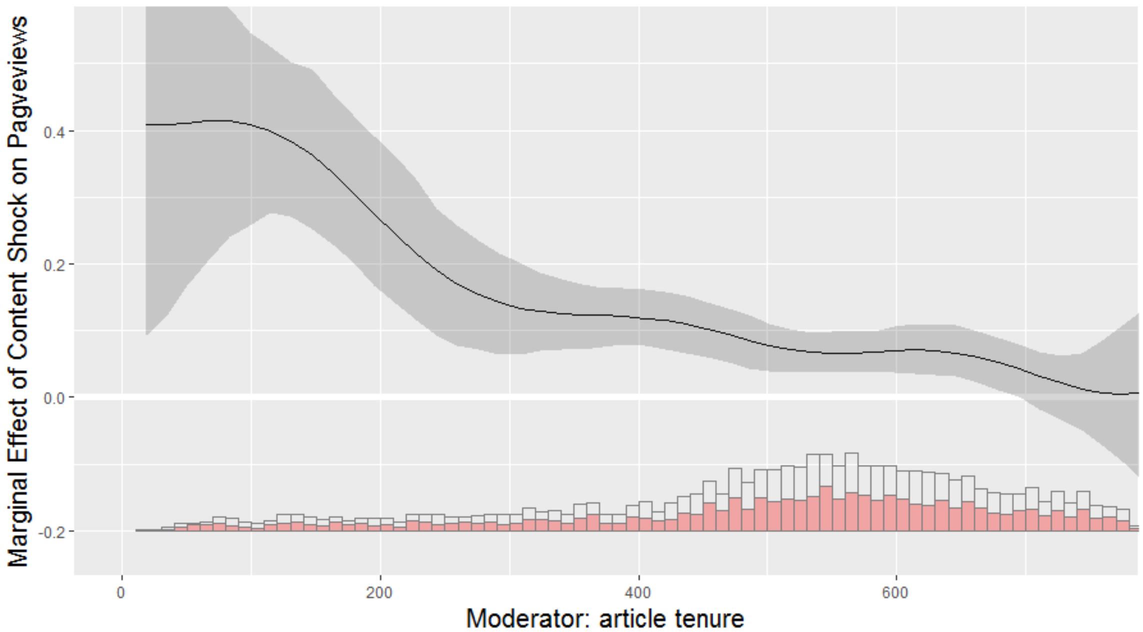
Marginal Effect of Content Shock on Pagveviews
0.4
0.2
0.0
-0.2
0
200
400
600
Moderator: article tenure

Checks of Model Specification

No Fixed Effect

Because they do not account for heterogeneity across articles, models without fixed effects tend to overestimate effect sizes. However, we find that they lead to qualitatively similar results. We provide the model estimates as robustness checks in Tables A8-A9.

Table A8: Direct impact of content shock -- No Fixed Effect

	Scaled pageviews		
	(1)	(2)	(3)
Post	0.062***	0.056***	0.062***
	(0.004)	(0.004)	(0.004)
Treated	0.000	-0.006	0.000
	(0.004)	(0.004)	(0.004)
PostShock	0.111***		
	(0.005)		
PostShock*log(charCount)		0.035***	0.062***
		(0.002)	(0.002)
PostShock*old article			-0.030***
			(0.006)
PostShock*popular article			-0.150***
			(0.007)
PostShock*long article			-0.010
			(0.006)
Article fixed effect	No	No	No
Week fixed effect	No	No	No
Observations	287,664	287,664	287,664
Adjusted R^2	0.009	0.010	0.013

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table A9: Spillover Effect – No Fixed Effects

	Scaled pageviews			
	(1)	(2)	(3)	(4)
Post	0.039***	0.039***	0.039***	0.039***
	(0.001)	(0.001)	(0.001)	(0.001)
Treated	0.000	-0.000	-0.000	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)
PostShock	0.008***	0.014***	-0.002	-0.002
	(0.001)	(0.002)	(0.002)	(0.003)
PostShock*popularTargetArticle		-0.015***	-0.022***	-0.005***
		(0.002)	(0.002)	(0.002)
PostShock*popularSourceArticle		-0.002	0.029***	0.014***
		(0.002)	(0.002)	(0.003)
PostShock*newLink			0.128***	0.143***
			(0.002)	(0.005)
PostShock*popularTargetArticle*newLink				-0.137***
				(0.005)
PostShock*popularSourceArticle*newLink				0.069***
				(0.005)
Article fixed effect	No	No	No	No
Week fixed effect	No	No	No	No
Observations	6,862,648	6,862,648	6,862,648	6,862,648
Adjusted R^2	0.001	0.001	0.001	0.001

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Spillover models with treatment intensity

Multiway interaction models require more restrictive modeling assumptions and are not easily interpretable. We did not incorporate the treatment intensity in our main spillover models for this reason. As a robustness check, we present the result for the spillover model that accounts for both treatment intensity and new link indicators. We did not estimate a model that simultaneously incorporates treatment intensity, new link indicators, and target and source popularity, as these would involve complex four-way interactions that are difficult to interpret. Estimates for this model are displayed in Table A10.

Table A10: Spillover Model with Treatment intensity

	Scaled pageviews		
	(1)	(2)	(3)
PostShock*log(charCount)	0.002^{***}	0.008^{***}	-0.002
	(0.001)	(0.002)	(0.002)
PostShock*log(charCount)*popularTargetArticle		-0.003^{**}	
		(0.002)	
PostShock*log(charCount)*popularSourceArticle		-0.005^{***}	
		(0.002)	
PostShock*log(charCount)*newLink			0.032^{***}
			(0.004)
Article fixed effect	Yes	Yes	Yes
Week fixed effect	Yes	Yes	Yes
Observations	6,862,648	6,862,648	6,862,648
Adjusted R^2	0.104	0.104	0.104
Notes:	***Significant at the 1 percent level.		
	**Significant at the 5 percent level.		
	*Significant at the 10 percent level.		

Poisson and Negative Binomial Regression for Editing Behavior

Editing behaviors (i.e.. the number of total edits and number of unique editors) in a certain period (6 months) are counting processes. We show below that Poisson regression and Negative Binomial regression produce qualitatively similar results as the linear regression that we use in the main analysis.

Table A11: Number of Total Edits

	Number of Total Edits					
	Poisson Regression			Negative Binomial Regression		
	(1)	(2)	(3)	(4)	(5)	(6)
Postshock	0.278***			0.526***		
	(0.007)			(0.020)		
Postshock*log(charCount)		0.093***	0.240***		0.157***	0.191***
		(0.003)	(0.006)		(0.006)	(0.010)
Postshock*old article			0.105***			0.093**
			(0.020)			(0.046)
Postshock*long article			-0.454***			-0.290***
			(0.022)			(0.046)
Postshock*popular article			-0.292***			0.002
			(0.025)			(0.051)
Article Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,964	10,964	10,964	10,964	10,964	10,964
Log Likelihood	-43,202.5	-43,060.7	-42,497.8	-35,368.3	-35,326.3	-35,303.4

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table A12: Number of Unique Editors

	Number of Unique Editors					
	Poisson Regression			Negative Binomial Regression		
	(1)	(2)	(3)	(4)	(5)	(6)
Postshock	0.290***			0.411***		
	(0.016)			(0.015)		
Postshock*log(charCount)		0.090***	0.178***		0.123***	0.158***
		(0.003)	(0.007)		(0.004)	(0.008)
Postshock*old article			0.059**			0.087**
			(0.026)			(0.034)
Postshock*long article			-0.301***			-0.234***
			(0.027)			(0.035)
Postshock*popular article			-0.142***			-0.040
			(0.031)			(0.039)
Article Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,964	10,964	10,964	10,964	10,964	10,964
Log Likelihood	-30,997.0	-30,928.4	-30,789.9	-30,230.1	-30,184.3	-30,154.1

Notes:
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Robustness of Simulation

The results on distribution of excess attention for the ACP and UAP policies are similar for different choices of cliques or communities. While we perturbed all 600 cliques that met our size criteria, there are significantly more communities that do so. We repeated the analysis for an alternate set of communities. Results are displayed in Figure A7 and are qualitatively similar to the main results. Differences in excess attention arise from differences in network structure, though ACP consistently captures more attention than UAP on average.

Figure A8

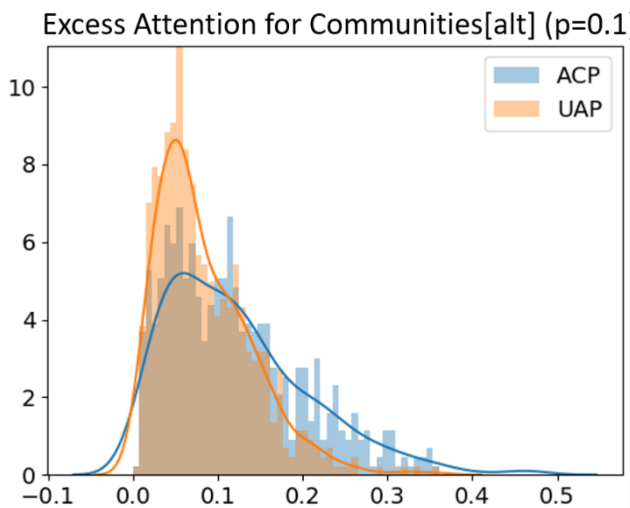

As described in the text, the shape of the distribution of excess attention is entirely a consequence of the network structure around the perturbation set, where the size of perturbation p acts as a simple scaling factor. This can be seen by iteratively expanding the PageRank equation and examining only the elements of the PageRank vector that correspond to the nodes of the perturbed set. For this set of nodes, p is a common factor which can be factored out. We verify that our distributions are consistent with this reasoning by examining two other choices perturbation size $p = 0.25, 0.5$ for the same set of chosen cliques or communities, as displayed in Figures A8, A9.

Figure A9

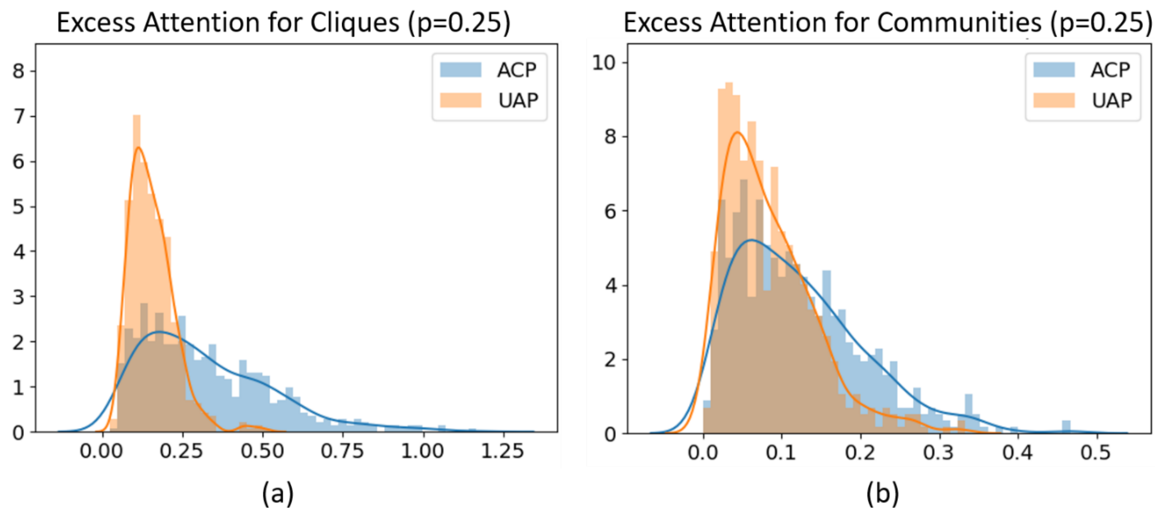

Figure A10

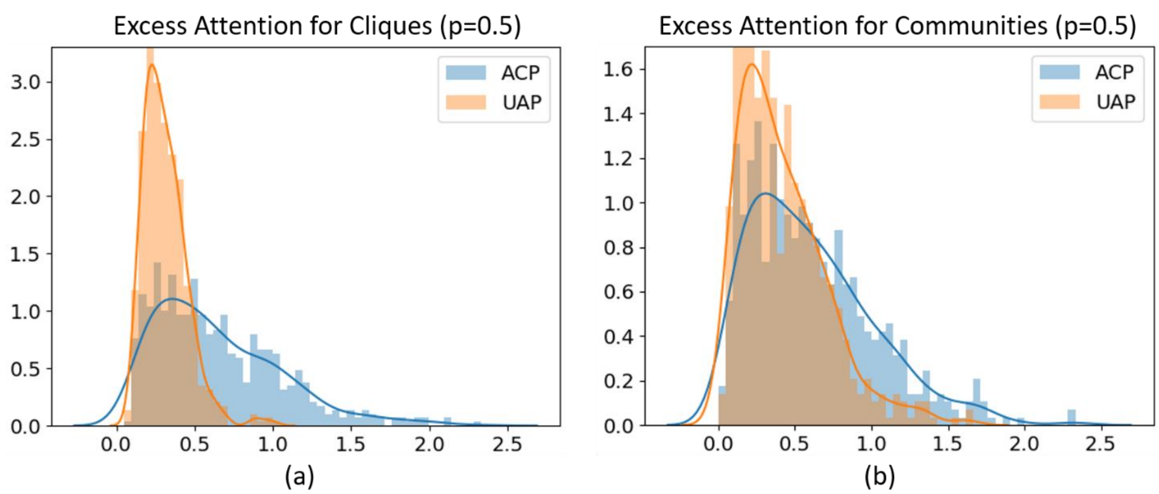

Excess Attention Conveyed by ACP vs. UAP

The aim of the Attention Contagion Policy (ACP) is to bring attention to specific (underdeveloped) regions of the information network that could benefit from it the most. When editors cluster the focus of their editorial attention under ACP, spillovers compound, conveying excess attention locally. In contrast, the undirected - or essentially random – editorial focus of the Undirected Attention Policy (UAP) will convey excess attention more widely across the information network. In other words, UAP will convey significant excess attention to more unique articles overall. But, under ACP, more articles receive a larger share of excess attention. This is illustrated in Figure A11, which shows the distribution of percentage increase in attention across *all* articles under both policies. It is clear that ACP shifts the weight of the distribution to the right relative to UAP.

Figure A11

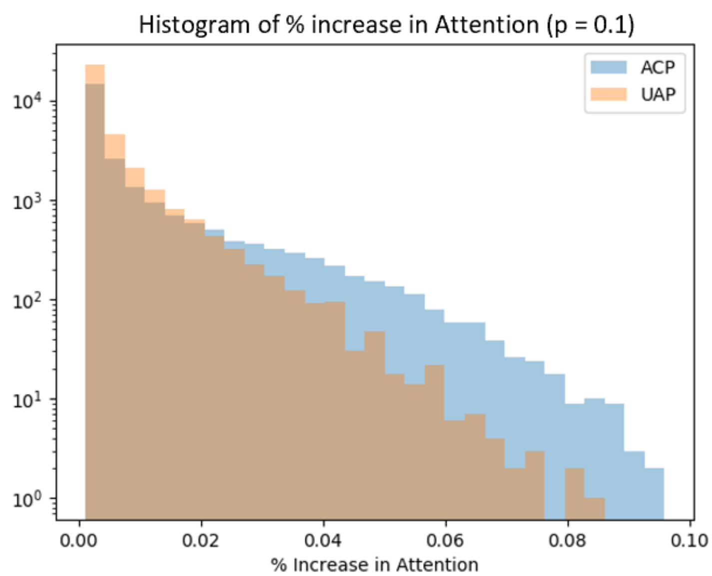

Proxy measure for state-level gun ownership

Quote from abstract from Siegel et al. (2014): “ Survey measures of gun ownership in all 50 states, however, are only available for 3 years, and no state-level data have been collected since 2004. Consequently, the National Research Council has declared the development of a valid proxy that can be constructed from accessible, annual, state-level data to be a priority. While such a proxy does exist (the proportion of suicides in a state committed using a gun (FS/S), its correlation with state estimates of gun ownership in recent years is only 0.80. Using state-level data for the years 2001, 2002 and 2004, we developed an improved proxy for state-level gun ownership that uses FS/S (firearm suicides divided by all suicides) and also the per capita number of hunting licenses. We validated this measure using data from surveys of gun ownership conducted in 48 states during 1996 and 1999, and in 21 states during 1995– 1998. Adding per capita hunting licenses to the proxy increased its correlation with survey-measured gun ownership from 0.80 to 0.95. The correlations of the new proxy with gun ownership in the two validation studies were 0.95 and 0.97. We conclude that the combination of FS/S and per capita hunting licenses improves substantially upon FS/S alone. This new proxy is easily computed from data that are available annually by state and may be useful for investigating the effect of gun prevalence on firearm-related morbidity and mortality.”

	state	year	firearm_suicide_ratio	gun_ownership_proxy
1	Montana	2017	68.489	62.634
2	Wyoming	2017	63.057	58.814
3	North Dakota	2017	60.390	53.305
4	Idaho	2017	62.245	51.975
5	South Dakota	2017	43.455	50.944
6	West Virginia	2017	67.939	50.500
7	Alabama	2017	68.541	50.176
8	Alaska	2017	59.500	48.769
9	Mississippi	2017	66.292	47.605
10	Arkansas	2017	62.124	46.420
11	Tennessee	2017	60.978	44.874
12	Oklahoma	2017	58.069	43.647
13	Louisiana	2017	61.667	43.007
14	Maine	2017	55.109	42.641
15	Missouri	2017	60.904	42.138
16	Kentucky	2017	61.039	41.968
17	Vermont	2017	55.357	41.840
18	Georgia	2017	62.578	40.862
19	Kansas	2017	57.143	40.279
20	South Carolina	2017	63.365	39.361
21	Wisconsin	2017	46.976	37.660
22	Arizona	2017	59.683	37.569
23	Utah	2017	51.885	36.243
24	North Carolina	2017	55.227	35.789
25	Texas	2017	57.120	35.497
26	Minnesota	2017	46.616	35.483
27	Oregon	2017	53.212	35.282
28	Virginia	2017	56.997	34.387
29	Indiana	2017	55.311	34.258
30	Nebraska	2017	45.818	34.080
31	Pennsylvania	2017	48.916	34.016
32	New Mexico	2017	53.157	33.906
33	Michigan	2017	46.809	32.287
34	Colorado	2017	50.127	32.200
35	Iowa	2017	46.347	31.928
36	Ohio	2017	52.759	31.881
37	New Hampshire	2017	49.811	31.009
38	Nevada	2017	52.472	30.514
39	Florida	2017	53.393	29.583
40	Washington	2017	49.268	28.727
41	Delaware	2017	44.643	25.235
42	Maryland	2017	43.016	24.399
43	Illinois	2017	39.145	22.447
44	California	2017	37.338	19.452
45	New York	2017	25.531	14.424
46	Connecticut	2017	26.914	13.372
47	New Jersey	2017	25.660	12.329
48	Rhode Island	2017	25.581	12.252
49	Massachusetts	2017	22.287	10.229
50	Hawaii	2017	13.216	4.532

Model Results

	tag	estimate	std.error	statistic	p.value	conf.low	conf.high
1	Workplace shooting by disgruntled employee	3,732.989	479.669	7.782	0	2,792.855	4,673.123
2	School Shooting	3,566.686	567.857	6.281	0	2,453.707	4,679.665
3	Officer Involved - Officer Victim	2,936.538	191.927	15.300	0	2,560.368	3,312.707
4	Officer Involved - Bystander shot or killed	2,413.333	503.187	4.796	0	1,427.105	3,399.560
5	Spree Shooting	2,080.113	404.608	5.141	0	1,287.096	2,873.130
6	Officer Involved - Officer Shooter	1,411.683	95.193	14.830	0	1,225.108	1,598.259
7	Mass Shooting	1,023.105	201.340	5.081	0	628.486	1,417.725
8	Child Shooter	774.365	245.566	3.153	0.002	293.064	1,255.665
9	Shootout	630.675	111.223	5.670	0	412.682	848.668
10	Defensive Use by Third Party	595.809	229.209	2.599	0.009	146.567	1,045.050
11	Child Victim	595.655	247.651	2.405	0.016	110.268	1,081.041
12	Murder/Suicide	534.108	177.117	3.016	0.003	186.964	881.252
13	Assault weapon	526.417	125.562	4.192	0	280.319	772.514
14	Concealed Carry Victim	471.237	215.167	2.190	0.029	49.519	892.956
15	House party	458.520	166.711	2.750	0.006	131.773	785.266
16	Attempted Murder/Suicide	341.682	283.578	1.205	0.228	-214.121	897.485
17	Car Jacking	339.323	86.934	3.903	0	168.935	509.710
18	Bar Incident	338.805	79.846	4.243	0	182.308	495.301
19	Sex Crime	297.482	180.446	1.649	0.099	-56.186	651.149
20	Concealed Carry Perpetrator	268.565	220.224	1.220	0.223	-163.065	700.196
21	Gang Incident	260.148	87.291	2.980	0.003	89.060	431.236
22	Road rage	223.711	122.156	1.831	0.067	-15.709	463.132
23	Home Invasion	200.328	46.683	4.291	0	108.832	291.825
24	Illegal Posession	169.054	41.271	4.096	0	88.165	249.943
25	Kidnapping/abductions/hostage	136.705	90.377	1.513	0.130	-40.431	313.842
26	Armed robbery	127.283	39.885	3.191	0.001	49.110	205.457
27	Drive-by	94.012	46.709	2.013	0.044	2.464	185.561
28	Toy Gun	32.709	129.033	0.253	0.800	-220.192	285.610
29	Drug involvement	-28.853	43.622	-0.661	0.508	-114.351	56.644
30	Defensive Use by Victim	-74.854	122.225	-0.612	0.540	-314.411	164.704
31	Stolen Gun	-77.225	59.754	-1.292	0.196	-194.341	39.892
32	Under the influence of alcohol or drugs	-136.453	104.368	-1.307	0.191	-341.011	68.106
33	Domestic Violence	-211.237	54.452	-3.879	0	-317.961	-104.513
34	Gun range/gun shop/gun show shooting	-249.654	225.806	-1.106	0.269	-692.226	192.918
35	Accidental Shooting	-311.905	63.564	-4.907	0	-436.488	-187.322
36	Suicide	-1,129.704	99.589	-11.344	0	-1,324.893	-934.514

Frequency by Incident type

	tag	frequency	percentage
1	Illegal Posession	44,289	0.186
2	Armed robbery	19,674	0.083
3	Drug involvement	16,987	0.071
4	Drive-by	13,541	0.057
5	Stolen Gun	12,930	0.054
6	Domestic Violence	10,762	0.045
7	Home Invasion	10,590	0.044
8	Officer Involved - Officer Shooter	8,815	0.037
9	Accidental Shooting	8,158	0.034
10	Suicide	5,879	0.025
11	Gang Incident	5,608	0.024
12	Car Jacking	4,928	0.021
13	Bar Incident	4,530	0.019
14	Defensive Use by Victim	4,378	0.018
15	Under the influence of alcohol or drugs	3,463	0.015
16	Kidnapping/abductions/hostage	2,663	0.011
17	Murder/Suicide	2,451	0.010
18	Shootout	2,308	0.010
19	Road rage	2,127	0.009
20	Assault weapon	2,043	0.009
21	Toy Gun	1,709	0.007
22	Child Victim	1,535	0.006
23	Mass Shooting	1,383	0.006
24	Officer Involved - Officer Victim	1,278	0.005
25	Child Shooter	786	0.003
26	House party	761	0.003
27	Sex Crime	653	0.003
28	Concealed Carry Victim	621	0.003
29	Concealed Carry Perpetrator	556	0.002
30	Defensive Use by Third Party	537	0.002
31	Attempted Murder/Suicide	463	0.002
32	Gun range/gun shop/gun show shooting	434	0.002
33	Spree Shooting	394	0.002
34	School Shooting	391	0.002
35	Workplace shooting by disgruntled employee	173	0.001
36	Officer Involved - Bystander shot or killed	169	0.001

Tag-tag co-occurrence heat map

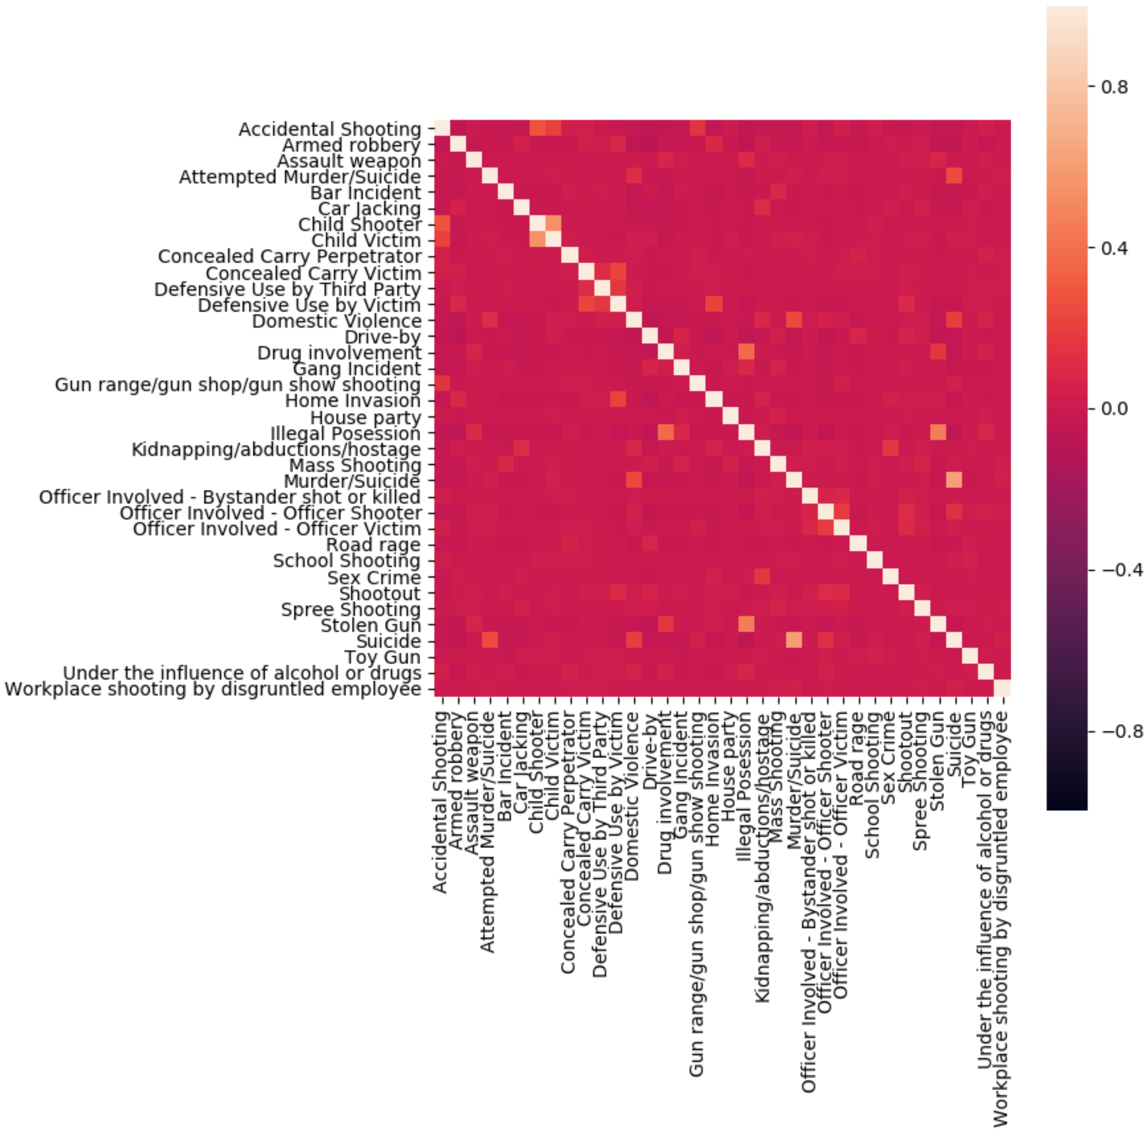

BIBLIOGRAPHY

Aaltonen A, Seiler S (2016) Cumulative Growth in User-Generated Content Production: Evidence from Wikipedia. Management Science 62(7), 2054–2069.

Alesina, A., Brioschi, B. and Ferrara, E.L., 2016. Violence against women: A cross-cultural analysis for Africa (No. w21901). National Bureau of Economic Research. Retrieve from: https://www.nber.org/papers/w21901

Andrews, K. T., & Caren, N. (2010). Making the News: Movement Organizations, Media Attention, and the Public Agenda. American Sociological Review, 75(6), 841–866.

Aral, S., Van Alstyne, M. (2011). The diversity-bandwidth trade-off. American Journal of Sociology, 117(1), 90–171.

Aral S, Muchnik L, Sundararajan A (2013) Engineering social contagions: Optimal network seeding in the presence of homophily. Network Science 1(2):125–153.

Barabási AL, Albert R (1999) Emergence of scaling in random networks. Science 286(5439), 509–512.

Beale, S. S. (2006). The news media's influence on criminal justice policy: How market-driven news promotes punitiveness. William & Mary Law Review, 48, 397.

Behr, R. L., & Iyengar, S. (1985). Television News, Real-World Cues, and Changes in the Public Agenda. Public Opinion Quarterly, 49(1), 38–57.

Bertrand M, Duflo E, Mullainathan S (2004) How Much Should We Trust Differences-In-Differences Estimates? The Quarterly Journal of Economics 119(1), 249–275.

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.

Carmi E, Oestreicher-Singer G, Stettner U, Sundararajan A (2017) Is Oprah Contagious? the Depth of Diffusion of Demand Shocks in a Product Network. MIS Quarterly 41(1), 207–221.

Chong, A. and Ferrara, E.L., 2009. Television and divorce: Evidence from Brazilian novelas. Journal of the European Economic Association, 7(2–3), 458–468.

Chyi, H. I., & McCombs, M. (2004). Media Salience and the Process of Framing: Coverage of the Columbine School Shootings. Journalism & Mass Communication Quarterly, 81(1), 22–35.

DellaVigna, S. and Kaplan, E., 2007. The Fox News effect: Media bias and voting. The Quarterly Journal of Economics, 122(3), 1187-1234. Retrieved from https://academic.oup.com/qje/article-abstract/122/3/1187/1879517

DellaVigna, S., Enikolopov, R., Mironova, V., Petrova, M. and Zhuravskaya, E., 2014. Cross-border media and nationalism: Evidence from Serbian radio in Croatia. American Economic Journal: Applied Economics, 6(3), 103–132.

Dimitrov D, Singer P, Lemmerich F, Strohmaier M (2017) What Makes a Link Successful on Wikipedia? Proceedings of the 26th International Conference on World Wide Web – WWW ’17. (ACM Press, New York, NY, USA), 917–926.

Durante, R. and Zhuravskaya, E., 2018. Attack when the world is not watching? US news and the Israeli-Palestinian conflict. Journal of Political Economy, 126(3), 1085–1133.

Eisensee, T. and Strömberg, D., 2007. News droughts, news floods, and US disaster relief. The Quarterly Journal of Economics, 122(2), 693–728.

Enikolopov, R., Petrova, M. and Zhuravskaya, E., 2011. Media and political persuasion: Evidence from Russia. American Economic Review, 101(7), 3253–3285.

Forman C, Goldfarb A, Greenstein S (2012) The Internet and local wages: A puzzle. American Economic Review 102(1), 556–575.

Gallus J (2016) Fostering Public Good Contributions with Symbolic Awards: A Large-Scale Natural Field Experiment at Wikipedia. Management Science (October):mnsc.2016.2540.

Garz, M. and Sörensen, J., 2017. Politicians under investigation: The news Media's effect on the likelihood of resignation. Journal of Public Economics, 153, 82–91.

Gentzkow M, Shapiro JM (2010) What Drives Media Slant? Evidence From U.S. Daily Newspapers. Econometrica 78(1), 35–71.

Gilliam, F. D., & Iyengar, S. (2000). Prime Suspects: The Influence of Local Television News on the Viewing Public. American Journal of Political Science, 44(3), 560–573.

Graham M, Hogan B, Straumann RK, Medhat A (2014) Uneven Geographies of User-Generated Information: Patterns of Increasing Informational Poverty. Annals of the Association of American Geographers 104(4), 746–764.

Haider-Markel, D. P., & Joslyn, M. R. (2001). Gun Policy, Opinion, Tragedy, and Blame Attribution: The Conditional Influence of Issue Frames. The Journal of Politics, 63(2), 520–543.

Halavais A, Lackaff D (2008) An analysis of topical coverage of Wikipedia. Journal of Computer-Mediated Communication 13(2), 429–440.

Harhoff D, Henkel J, Von Hippel E (2003) Profiting from voluntary information spillovers: How users benefit by freely revealing their innovations. Research Policy 32(10), 1753–1769.

Hinnosaar, M, Hinnosaar, T, Kummer, M, and Slivko, O. Wikipedia Matters (July 14, 2019). Journal of Economics & Management Strategy, 2021, 1–13. https://doi.org/10.1111/jems.12421, Available at SSRN: https://ssrn.com/abstract=3046400 or http://dx.doi.org/10.2139/ssrn.3046400

Ho DE, Imai K, King G, Stuart EA (2007) Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. Political Analysis 15(3), 199–236.

Iacus SM, King G, Porro G (2012) Causal inference without balance checking: Coarsened exact matching. Political Analysis 20(1), 1–24.

Iyengar, S. (1987). Television News and Citizens' Explanations of National Affairs. The American Political Science Review, 81(3), 815–831.

Kämpf M, Tessenow E, Kenett DY, Kantelhardt JW (2015) The Detection of Emerging Trends Using Wikipedia Traffic Data and Context Networks Rocchini D, ed. PLoS One 10(12):e0141892.

Kämpf M, Tismer S, Kantelhardt JW, Muchnik L (2012) Fluctuations in Wikipedia access-rate and edit-event data. Physica A: Statistical Mechanics and its Applications 391(23), 6101–6111.

Kane GC (2009) It's a Network, Not an Encyclopedia: A Social Network Perspective on Wikipedia Collaboration. Academy of Management Proceedings 2009(1), 1–6. https://doi.org/10.5465/ambpp.2009.44243222

Kane GC, Ransbotham S (2016) Research Note—Content and Collaboration: An Affiliation Network Approach to Information Quality in Online Peer Production Communities. Information Systems Research (August):isre.2016.0622.

Kearney, M.S. and Levine, P.B., 2015. Media influences on social outcomes: The impact of MTV's 16 and pregnant on teen childbearing. American Economic Review, 105(12), 3597-3632.

Keegan B, Gergle D, Contractor N (2013) Hot Off the Wiki: Structures and Dynamics of Wikipedia's Coverage of Breaking News Events. American Behavioral Scientist 57(5), 595–622.

King, G., Schneer, B., & White, A. (2017). How the news media activate public expression and influence national agendas. Science, 358(6364), 776–780.

Kittur A, Chi EH, Suh B (2009) What's in Wikipedia? Proceedings of the 27th international conference on Human factors in computing systems – CHI 09. (ACM Press, New York, New York, USA), 1509.

Kleemans, M. and Hendriks Vettehen, P.G.J., 2009. Sensationalism in television news: A review.

Kloumann I, Ugander J, Kleinberg J (2016) Block Models and Personalized PageRank. Proceedings of the National Academy of Sciences of the United States of America 114(1), 33–38.

Kummer, M, Spillovers in Networks of User Generated Content: Pseudo-Experimental Evidence on Wikipedia (December 29, 2014). ZEW - Centre for European Economic Research Discussion Paper No. 14-132, Available at SSRN: https://ssrn.com/abstract=2567179 or http://dx.doi.org/10.2139/ssrn.2567179

Kummer ME (2019) Attention in the Peer Production of User Generated Content - Evidence from 93 Pseudo-Experiments on Wikipedia (July 11, 2018). Available at SSRN: https://ssrn.com/abstract=3431249 or http://dx.doi.org/10.2139/ssrn.3431249

Kummer ME, Saam M, Halatchliyski I, Giorgidze G (2016) Centrality and content creation in networks - The case of economic topics on German wikipedia. Information Economics and Policy 36, 36–52.

Kuznetsov S (2006) Motivations of Contributors to Wikipedia. SIGCAS Computers and Society. 36(2).

Lampe C, Obar J, Ozkaya E, Zube P, Velasquez A (2012) Classroom Wikipedia participation effects on future intentions to contribute. Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work – CSCW '12:403.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). Judged frequency of lethal events. Journal of Experimental Psychology. Human Learning and Memory, 4(6), 551–578.

Lin Z, Goh KY, Heng CS (2017) The Demand Effects of Product Recommendation Networks: An Empirical Analysis of Network Diversity and Stability. MIS Quarterly 41(2), 397–426.

Martin, G. J., & Yurukoglu, A. (2017). Bias in Cable News: Persuasion and Polarization. The American Economic Review, 107(9), 2565–2599.

McCombs, M. E., & Shaw, D. L. (1972). The Agenda-Setting Function of Mass Media. Public Opinion Quarterly, 36(2), 176–187.

Muchnik, L., Aral S (2013) Social influence bias: a Randomizaed Experiment. Science 341(6146), 647–651.

Muchnik L, Pei S, Parra LC, Reis SDS, Andrade Jr JS, Havlin S, Makse HA (2013) Origins of power-law degree distribution in the heterogeneity of human activity in social networks. Scientific Reports 3(1), 1783.

Norris Pippa (2001) Digital divide : civic engagement, information poverty, and the Internet worldwide (Cambridge University Press).

Nov O (2007) What motivates Wikipedians? Communications of the ACM 50(11), 60–64.

Pew Research Center. (2016, July 7). How Americans get their news. Retrieved October 27, 2020, from Pew Research Center website: https://www.journalism.org/2016/07/07/pathways-to-news/

Philippe, A. and Ouss, A., 2018. “No hatred or malice, fear or affection”: Media and sentencing. Journal of Political Economy, 126(5), 2134-2178

Raghavan UN, Albert R, Kumara S (2007) Near linear time algorithm to detect community structures in large-scale networks. Physical Review E 76(3).

Ransbotham S, Kane G, Lurie NH (2012) Network Characteristics and the Value of Collaborative User Generated Content. Marketing Science 31(3), 387–405.

Rosenbaum PR, Rubin DB (1983) The Central Role of the Propensity Score in Observational Studies for Causal Effects. Biometrika 70(1), 41–55.

Salganik MJ, Dodds PS, Watts DJ (2006) Experimental study of inequality and unpredictability in an artificial cultural market. Science 311(5762), 854–856.

Schroeder, E. and Stone, D.F., 2015. Fox news and political knowledge. Journal of Public Economics, 126, 52–63.

Siegel, M., Ross, C. S., & King, C., 3rd. (2014). A new proxy measure for state-level gun ownership in studies of firearm injury prevention. Injury Prevention, 20(3), 204–207.

Taylor SJ, Eckles D (2018) Randomized Experiments to Detect and Estimate Social Influence in Networks. In Lehmann S, Ahn YY, eds. Complex Spreading Phenomena in Social Systems. (Springer International Publishing, Cham), 289–322.

Thompson NC, Hanley D (2017) Science is Shaped by Wikipedia: Evidence from a Randomized Control Trial Evidence From a Randomized Control Trial. MIT Sloan Research Paper No. 5238-17, Available at SSRN: https://ssrn.com/abstract=3039505 or http://dx.doi.org/10.2139/ssrn.3039505

Tomita E, Tanaka A, Takahashi H (2006) The worst-case time complexity for generating all maximal cliques and computational experiments. Theoretical Computer Science 363(1), 28–42.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 5(2), 207–232.

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. Science, 185(4157), 1124–1131.

Walker D, Muchnik L (2014) Design of Randomized Experiments in Networks. Proceedings of the IEEE 102(12), 1940–1951.

West R, Leskovec J (2012) Human wayfinding in information networks. Proceedings of the 21st international conference on World Wide Web – WWW ’12. (ACM Press, New York, New York, USA), 619.

Wilkinson D, Huberman B (2007) Assessing the value of coooperation in wikipedia. Proceedings of the 2007 International Symposium on Wikis. ACM.

Wu F, Huberman BA (2007) Novelty and collective attention. Proceedings of the National Academy of Sciences of the United States of America 104(45), 17599–17601.

Xiaoquan Z, Lihong Z (2015) How Does the Internet Affect the Financial Market? an Equilibrium Model of Internet-Facilitated Feedback Trading. MIS Quarterly 39(1), 17-A10.

Xu SX, Zhang X (Michael) (2013) Impact of Wikipedia on Market Information Environment: Evidence on Marnagement Disclosure and Investor Reaction. MIS Quarterly 37(4), 1043–1068.

Yu L (2006) Understanding information inequality: Making sense of the literature of the information and digital divides. Journal of Librarianship and Information Science 38(4), 229–252.

Zhang X, Zhu F (2011) Group Size and Incentive to Contribute: A Natural Experiment at Chinese Wikipedia. The American Economic Review 101(June), 1–17.

Zhu H, Zhang A, He J, Kraut RE, Kittur A (2013) Effects of Peer Feedback on Contribution : A Field Experiment in Wikipedia. CHI '13 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: 2253–2262.

CURRICULUM VITAE

[redacted]

[redacted]

[redacted]