



Did clickbait crack the code on virality?

Prithwiraj Mukherjee¹ · Souvik Dutta² · Arnaud De Bruyn³

Received: 18 August 2020 / Accepted: 23 November 2021 / Published online: 29 January 2022
© Academy of Marketing Science 2021

Abstract

Although clickbait is a ubiquitous tactic in digital media, we challenge the popular belief that clickbait systematically leads to enhanced sharing of online content on social media. Using the Persuasion Knowledge Model, we predict that clickbait tactics may be perceived by some readers as a manipulative attempt, leading to source derogation where the publisher may be perceived as less competent and trustworthy. This, in turn, may reduce some readers' intention to share content. Using a controlled experiment, we confirm that high-emotional headlines are shared more and show evidence that clickbait often leads to inferences of manipulative intent and source derogation. We then use a well-known secondary data set containing 19,386 articles from 27 leading online publishers. We supplement it with Twitter share data, sentiment analysis, topic modeling, and additional control variables. We confirm that, on average, clickbait articles elicit far fewer shares than non-clickbait articles. Our results are stable, with large effect sizes even after controlling for endogenous selection.

Keywords Social media · Clickbait · Persuasion Knowledge Model · Source derogation · Sharing · Topic modeling · Sentiment analysis · Propensity score matching

Introduction

With the proliferation of online content, the competition for readers' attention is fierce (Teixeira, 2014). In this context, online publishers often use a tactic called clickbait to induce readers to click on their content. Coined by a blogger (Geiger, 2006), this term refers to headlines that *bait* the user to *click* on their web links because of the way they are phrased.

Today, the term is common parlance amongst social media users and almost synonymous with online virality. For example, Bazaco et al. (2019) term clickbait as “a strategy of viral

journalism.” BuzzFeed, a major player in digital content, often associated with clickbait, is also regularly credited for having cracked the formula for shareable content and viral news (Madhavan, 2017; Rowan, 2014; Webics, 2014). Similarly, CoSchedule, a popular headline optimizing service used widely by bloggers, online journalists, and digital marketers, also asserts that clickbait is shared more. All of this suggests the common assumption that clickbait leads to more online sharing.

Virality is often a goal of online marketers. It is thus intuitive that clickbait specialists, such as BuzzFeed, are cited as examples to emulate and copy. For instance, the IsItWP Headline Analyzer suggests including uncommon, emotional, or power words to “make your headline irresistibly clickable” and “drive traffic and shares.”

Yet, a leaked internal BuzzFeed document indicates that, between 2011 and 2014, the company spent about three-quarters of its editorial budget on buying traffic from Facebook (Trotter, 2015). BuzzFeed's “viral” content may not be as viral as many believe and might instead be an artifact of paid traffic. The assumption that clickbait is associated with sharing, a relationship that BuzzFeed is often cited as exemplifying,¹ appears questionable.

In this article, and contrary to common belief in the industry, not only do we argue that clickbait tactics do not

Jenny van Doorn served as Area Editor for this article.

✉ Prithwiraj Mukherjee
pmukherjee@iimb.ac.in

Souvik Dutta
souvik@iiitd.ac.in

Arnaud De Bruyn
debruyn@essec.edu

¹ Indian Institute of Management Bangalore, Bannerghatta Road, Bangalore 560076, India

² Indraprastha Institute of Information Technology Delhi, Okhla Industrial Estate Phase-III, New Delhi 110020, India

³ ESSEC Business School, Avenue Bernard Hirsch, 95000 Cergy, France

¹ See for example Tandoc Jr (2018, p. 202) or Stringer (2020).

contribute to more shares and word of mouth, we demonstrate that they might often impede them. This claim has profound consequences for online marketers, copy editors, and web marketing services who may have been wrongfully convinced that clickbait tactics have cracked the code on virality.

We organize this article as follows. In “Theoretical framework,” we lay the theoretical foundations of the paper and show that past research, indeed, predicts that clickbait will generate more shares. We rely on the sharing literature and its antecedents, as well as on the concept of curiosity gap, and demonstrate that there is a plausible and theoretically justified behavioral pathway to predict that clickbait headlines will be widely shared indeed. However, we nuance these predictions and envisage the idea that some consumers may be increasingly aware of these manipulative tactics: they form theories and beliefs about the intents of the publishers. Following the Persuasion Knowledge Model (Friestad & Wright, 1994) and other related research (Abelson & Miller, 1967; Zuwerink Jacks & Cameron, 2003), we predict that this belief about the publishers may backfire and generate mistrust and defiance, which may, in turn, impede virality.

In “Study 1: Controlled experiment,” while we confirm the role of emotions and curiosity in online sharing, we also validate our theory that some consumers are aware of the manipulative tactics of the publishers. We show that this belief creates a negative halo effect around the publisher and generates a perception of untrustworthiness, incompetence, and lack of sincerity, which translates into a lower likelihood of sharing.

In “Study 2: Field study,” armed with a better understanding of the underlying mechanisms at play, we investigate clickbait vs. non-clickbait shares using 19,386 articles from the Webis Clickbait Challenge (Potthast et al. 2018a, b)—a publicly available data set from a machine learning contest inviting novel approaches to detect clickbait. These articles are randomly sampled from the Twitter feeds of 27 prominent English-language publishers and manually rated for their degree of “clickbaitiness.” We augment this data set by extracting the number of times each article was shared on Twitter and the number of likes each garnered. To account for the possibility that clickbait may be an endogenous treatment prevalent for certain types of content, we employ propensity score matching. In addition, we control for content characteristics through sentiment analysis and topic modeling. We show that clickbait is both liked less and shared less than non-clickbait. These results hold even after we introduce numerous controls such as the number of Twitter followers of each publisher (Yoganarasimhan, 2012), sentiment scores, readability scores (Berger, 2011; Berger & Milkman, 2012; Berger & Schwartz, 2011), political readership, and publisher dummies.

Our work questions the widespread belief that clickbait has “cracked the code” on virality, both theoretically and empirically. Given the ubiquity of clickbait tactics and the importance of digital content sharing (Zubcsek & Sarvary, 2011), this research has important managerial implications. It opens interesting avenues for future research, which we discuss in the conclusions.

Theoretical framework

Definition of clickbait

The phenomenon of clickbait is as prevalent as it is loosely defined. Merriam-Webster officially added the word in 2015, defining it as: “Something (such as a headline) designed to make readers want to click a hyperlink, especially when the link leads to content of dubious value or interest.” This definition is not without merit but remains problematic. All editors, journalists, and bloggers write headlines hoping that readers would click on them, which makes the Merriam-Webster definition too vague to be useful for our purpose. Potthast et al. (2018b) propose no less than four different definitions of clickbait, whereas it focuses on (a) the result of a headline-optimization process, (b) the intention of the publisher, (c) the effect obtained on clickthrough, or (d) the perception of “bait” from its readers. Since “proving malicious intent of individual journalists or publishers as a whole [...] is virtually impossible” (Potthast et al., 2018b, p. 1501), they settle on the last definition and define clickbait as “teaser messages perceived by (some) readers as bait to click a link.” However, they remain vague as to which specific editorial tactic gives rise to that perception of “bait” among readers. Sanders (2017) clarifies that ambiguity and specifies that clickbait is a “style of headline designed to entice consumption by strategically withholding information” (see also Munger et al., 2020, p. 49). We embrace that distinction and define clickbait as: “A headline that strategically withholds information to entice the reader to click on a link.”

Since the amount of information withheld in a headline—and the inferred strategic intent of doing so—are both continuous constructs, the definition of what constitutes clickbait is itself a matter of degree. We reflect this continuity in our theoretical developments and our manipulation of clickbait in our controlled experiment. However, for reasons we will discuss later in greater detail, clickbait is often dichotomized in practice (e.g., in detection algorithms).

McNeal (2015) reminds us that “storytelling by its very nature encourages sensationalism and attention-grabbing tactics.” What differentiates clickbait from other editorial tactics is the strategic withholding of information to create an artificial “curiosity gap” (Loewenstein, 1994). A clickbait

headline highlights a critical piece of missing information that the readers could unveil only by clicking on the link and reading more. For instance, the headline “CEO Samantha Jones unexpectedly fired from XYZ Corp” is good journalism: it is informative yet attention-grabbing. The same headline completed with “You won’t believe how she learned the news” is a pure clickbait tactic. The publisher could have specified “by email” but instead decided to withhold that information to entice the audience to click. Readers are influenced into believing that the way she was fired is sensational and noteworthy and that learning more about it is well worth a click.

The marketing literature provides plausible pathways to explain why clickbait headlines might be shared to a greater extent than non-clickbait headlines. We present these arguments next and then conclude with a dissenting view about clickbait’s virality.

The case for emotions

Marketing researchers have extensively investigated online sharing and its antecedents. Berger (2011) and Berger and Milkman (2012) find that physiological arousal can induce content to be shared; high arousal content causing awe and anger tends to be shared more than content with low arousal like sadness. Akpınar and Berger (2017) study online ads to find that content with emotional appeal is shared more than content with informative appeal. Schulze et al. (2014) establish that viral campaign ads for hedonic but not utilitarian goods are successful in being rebroadcast on social media. Araujo et al. (2015) demonstrate that informational cues like product details and brand URL information drive retweets of commercial content on Twitter; they also find that emotional cues reinforce informational cues to drive sharing. Tellis et al. (2019) present several factors driving rebroadcasting of videos on social media. They find that positive emotions enhance sharing while prominent brand placement works adversely. They also find that emotional ads are shared more on general social media platforms like Facebook, Google +, and Twitter, while informational ads are shared more on professional platforms like LinkedIn. Zhang et al. (2017) establish that message-user fit heavily influences sharing. Jalali and Papatla (2019) find that tweets that start with more topic-related words get retweeted more. Berger (2014) provides a comprehensive review of sharing motivations, noting that motivations could range from (a) impression management where individuals convey positive impressions of themselves, (b) emotion regulation where individuals manage their emotions, (c) information acquisition where individuals seek inputs from others, (d) social bonding where individuals seek to connect with others, and (e) persuading others. Based on extent research, we expect to replicate the finding that:

H1: Headlines with a high-emotional appeal are more likely to be shared than headlines with a low-emotional appeal.

Content information value is a strong determinant of sharing behaviors (Araujo et al., 2015). Clickbait headlines, characterized by their purposeful withholding of information (Munger et al., 2020; Sanders, 2017), cannot rely on the information value they provide; the more information the headline contains, the smaller the curiosity gap, hence the less teasing the article. If information value must be kept at a minimum, clickbait headlines must, therefore, maximize their emotional value instead, another significant antecedent of sharing behaviors (Akpınar & Berger, 2017; Tellis et al., 2019). Clickbait is built upon the idea of heightened emotions and curiosity. For instance, the non-clickbait headline “Donald Trump’s son-in-law Jared Kushner appointed as senior White House adviser” is designed to be neutral and informative. Inversely, the clickbait headline “You won’t believe whom Donald Trump appointed as senior White House adviser” has been carefully designed to elicit emotions such as anger, curiosity, or even disgust. Some professional copy editors even propose online tools to maximize the number of emotional words in a headline (e.g., The Emotional Marketing Value Headline Analyzer by Advanced Marketing Institute). To compensate for their by-design low information value, we expect that:

H2: Clickbait headlines are more likely to be crafted as high-emotional appeal than low-emotional appeal.

The case for the curiosity gap

Clickbait is often portrayed as a bait and switch tactic designed to deceive rather than inform. Consequently, even news outlets strongly associated with clickbait have tried to distance themselves from the practice. Ben Smith, the former editor of BuzzFeed News (the news arm of BuzzFeed) says (Smith, 2014), “[...] many people in the media industry confuses what we do with true clickbait. We have admittedly (and at times deliberately) not done a great job of explaining why we have always avoided clickbait at BuzzFeed. In fact—and here is a trade secret I’d decided a few years ago we’d be better off not revealing—clickbait stopped working around 2009.”

Despite these protestations, the general sentiment regarding clickbait among journalists and in the general population is negative. “Put simply, [clickbait] is a headline which tempts the reader to click on the link to the story. But the name is used pejoratively to describe headlines which are sensationalized, turn out to be adverts or are simply misleading” (Frampton, 2015).

Still, some marketing tactics are annoying but effective. Despite the bad press, there is a case to be made for the

purposeful withholding of information that defines clickbait headlines, which might explain its widespread success and omnipresence online. Clickbait is designed to induce curiosity in the reader, increasing the likelihood that they will click. Digital marketers refer to clickbait in the context of a “curiosity gap” publishing model based on the information gap theory (Golman & Loewenstein, 2018; Loewenstein, 1994). This theory posits curiosity as an important driver of human behavior. Such a curiosity gap can be induced by omitting important information about a news piece or entertainment feature in its headline.

Note that while—by their very definition—clickbait headlines “*strategically withhold information*” (emphasis on strategically), not all headlines that omit relevant information qualify as clickbait. For instance, in our research (cf. Study 1), the headline “After Going Vegan For 10 Weeks, Olivia Petter Reports Health Benefits” rated low on clickbait but high on omitted information. The concepts of *clickbait* and *omitted information* are closely related but remain conceptually distinct. While omitting important information is not the appanage of clickbait headlines, we still expect that:

H3: Clickbait headlines are more likely to omit important information than non-clickbait headlines.

Headlines that omit essential information create novelty, suspense, excitement, and intrigue, which has been shown to be a strong determinant of sharing behaviors. T. Teixeira (2012) finds that content featuring an “emotional roller coaster” is more likely to capture attention and be shared, “much the way a movie generates suspense by alternating tension and relief” (p.26), such as when a curiosity gap is artificially created and then relieved. Epistemic curiosity, also dubbed as the “drive to know” (Berlyne, 1954, p. 187), is a strong human trait that motivates individuals to eliminate information gaps (Litman, 2008). By sharing a clickbait headline, the source provides to its audience both the excitement of the unknown and the solution to relieve that tension. Therefore, we hypothesize that:

H4: Headlines that omit important information are more likely to be shared.

The case for resistance to persuasion

In a world flooded by information, companies have long understood that customers’ attention is the new currency (Davenport & Beck, 2001). But customers are increasingly aware of it as well. Most consumers now understand that companies go to great lengths to capture their attention, track their behaviors, capture their data, and profile their habits and preferences. It is now common knowledge that, on social networks, “we’re not the customers. We are the

product” (Rushkoff, 2011). Academics have made the headlines—and spurred controversy—by demonstrating the ability of social networks to manipulate the emotions of unaware (and unwilling) individuals at a large scale (Kramer et al., 2014).

Clickbait headlines try to manipulate online readers by creating an artificially exacerbated curiosity gap, sometimes to the point of ridicule. An infamous headline from the San Francisco Globe read: “When You Find Out What These Kids Are Jumping Into, Your Jaw Will Drop!”, with the subtitle “This is unbelievable! I have NEVER seen anything like THIS in my entire life! Wow.” The big revelation was that the kids in the picture were jumping into a swimming pool. In such a context, it is likely that some online visitors are now aware of the constant and repeated persuasion attempts targeting them. Inferences of manipulative intent (IMI) are reflexive processes by which consumers think that a market agent “is attempting to persuade [them] by incongruent, unfair, or manipulative means” (Campbell, 1995, p. 228). IMI are subjective: readers might infer manipulative intent when there is none and infer none when there clearly is. Consequently, IMI need to be measured at the reader’s level. We hypothesize that:

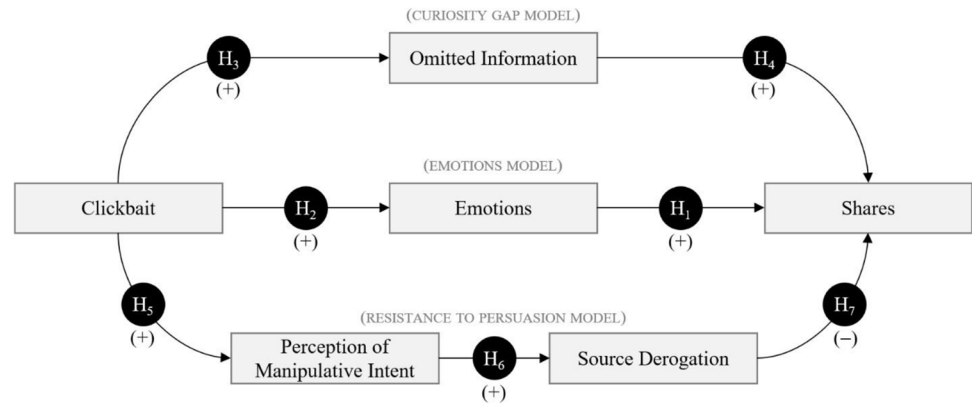
H5: When exposed to a clickbait headline, readers may infer the publisher’s manipulative intent.

IMI have been shown to influence customers’ attitudes and, ultimately, behaviors across a wide variety of contexts. For instance, the influence of IMI has been studied in servicescape and co-creation (Lunardo et al., 2016), customer service (Warren et al., 2020), comparative advertising (Kalro et al., 2017), advertising disclosure (V. L. Thomas et al., 2013), and even retail store atmospherics (Lunardo & Mbenque, 2013).

The Persuasion Knowledge Model (Friestad & Wright, 1994) predicts that readers will develop and use persuasion knowledge to cope with persuasion attempts, and in doing so, will refine their attitude towards the publishers themselves. Although readers who face and resist persuasion attempts may follow different strategies, such as avoidance, processing, or empowerment, a particular strategy is likely to be triggered when facing concerns of deception: *contesting strategies* (Fransen et al., 2015). Facing a deception attempt, a frequently used resistance strategy is to question the trustworthiness of the message (P. Wright, 1975; Zuwerink Jacks & Cameron, 2003) or the source, also known as *source derogation* (Abelson & Miller, 1967; Zuwerink Jacks & Cameron, 2003).

Source derogation has been studied extensively in advertising appeals or political marketing tactics (e.g., Belch, 1981; Kamins & Assael, 1987; Meirick, 2002). In this research, source derogation occurs instead as a psychological

Fig. 1 (Study 1) Summary of hypotheses



reactance to high-pressure tactics (Brehm, 1966). In essence, it is a defense mechanism invoked when facing a persuasion attempt aimed at limiting one’s freedom of choice. It requires minimum cognitive effort and pushes the target of the manipulation attempt to question the expertise or trustworthiness of the source (P. Wright, 1975; P. L. Wright, 1973). Consequently, we conjecture that:

H6: After identifying a manipulation intent, readers may resist a persuasion attempt by engaging in a source derogation strategy.

Zuwerink Jacks and Cameron (2003) note that “source derogation involves insulting the source, dismissing his or her expertise or trustworthiness, or otherwise rejecting his or her validity.” Extant research reports that source credibility and trustworthiness are significant in word-of-mouth perceived value and influence (e.g., Bansal & Voyer, 2000; Gilly et al., 1998; Tkaczyk & others, 2016). Source credibility in word of mouth largely depends on the characteristics of the source, such as its professionalism (Wangenheim & Bayón, 2007), and influences the perceived informational value of the message being shared (Liang & Yang, 2015; Martin & Lueg, 2013). Therefore, we expect that online content published by derogated sources will be shared to a lesser extent:

H7: Headlines from derogated sources are less likely to be shared than headlines from non-derogated sources.

Summary

Past research has extensively studied the pathway *content* → *emotion* → *more shares* (we refer to it as the “emotions model”). Since clickbait headlines are more likely to elicit emotional appeal than informational appeal, emotions are likely to play a role in clickbait’s sharing, indeed.

More centrally, clickbait headlines rely heavily on crafting an artificial curiosity gap that makes the headlines

“irresistibly clickable” and creates a sense of excitement and suspense, such that *clickbait* → *omitted information* → *more shares* (the “curiosity gap” model).

However, we hypothesize that, especially with today’s well-informed consumers, clickbait headlines might also trigger readers’ psychological reactance, such that *clickbait* → *perception of manipulative intent* → *source derogation* → *fewer shares* becomes plausible as well (the “resistance to persuasion” model). We summarize our hypotheses in Figure 1.

In the next section, we report a controlled experiment where we tested and compared the underlying mechanisms at play. Whether clickbait headlines are shared to a lesser extent in today’s online environment, beyond the confines of a controlled experiment, is an empirical question that we later explore with a field study in Study 2.

Study 1: Controlled experiment

Study design and measurements

Starting from nine actual headlines identified in the press and spanning various topics, we created four variants for each headline (for a total of 36), from “least clickbait” to “most clickbait.” We list the 36 headlines in Table 1.

We described what constitutes a clickbait headline to three independent judges. We then asked them to grade each of the 36 headline variants on a scale from 1 (“Not clickbait at all”) to 5 (“Extremely clickbait”). The inter-rater agreement was high, with weighted Cohen-Kappa’s correlation coefficients of 0.76, 0.88, and 0.83, respectively. The manipulations were successful, with an average clickbait rating of 1.11 (N = 27, standard deviation of 0.32) for the “least clickbait” manipulation, 1.52 (0.75) for the second, 2.89 (1.55) for the third, and 4.56 (0.75) for the last category (“most clickbait”). All categories are statistically different from one another at $p < 0.01$.

Table 1 (Study 1) List of 9 topics manipulated into four categories, from “least clickbait” to “most clickbait.”

Least clickbait	→		Most clickbait
Wikibuy Compares Prices From 235 Ecommerce Websites, Helped Customers Save More Than \$70 Million Last Year		Wikibuy Helps You Compare Hundreds Of Ecommerce Websites To Obtain Best Prices, Helped Customers Save Millions Last Year	Before You Renew Amazon Prime, You Need To Compare Prices With Wikibuy
Women Who Drink 2.5 Or More Glasses Of Milk A Day Have A Higher Risk Of Fractures Than Women Who Drink Less Than One Glass A Day		For Grown-Ups, Consuming Too Much Milk And Dairy Can Actually Be Harmful	Before You Drink Your Milk, You Need To Read This Article. It's Not As Good For Your Health As You Think
Health Experts Explain How Ebola Could Mutate And Start Being Transmitted By Cough And Sneeze		Health Experts: Ebola Could Mutate And Be Transmitted In The Air	A New Strain Of Ebola In The Air? A Nightmare That Could Happen
After Going Vegan For 10 Weeks, Olivia Peter Reports Better Digestion, More Energy, Clearer Skin		After Going Vegan For 10 Weeks, Olivia Peter Reports Health Benefits	I Went Vegan For 10 Weeks. And You Won't Believe What Happened To My Body
President Trump Pushed Out John R. Bolton, His Third National Security Adviser, Amid Fundamental Disputes Over North Korea		President Trump Pushed Out National Security Adviser Amid Fundamental Disputes Over International Politics	You Are Fired! This Is Whom President Trump Has Sacked Yesterday
A Maryland Baby With Heart Defect Has Survived Since March Despite Being Taken Off Life Support		A Maryland Baby Survives Unexpectedly After Being Taken Off Life Support	A Maryland Baby Taken Off Life Support Was Not Expected To Survive. You Won't Believe What Happened To Her
Canadian Community Is Mobilizing To Offer Hockey Equipment To Yaman, An 8-Year Old Syrian Refugee		Canadian Community Is Mobilizing To Offer Sports Equipment To Young Syrian Refugee	Canadians Welcome Syrian Refugee Family. What Happened Next Will Melt Your Heart

Consistent with the notion of curiosity gap, we asked the same judges to rate whether the headlines were omitting important information on a scale from 1 (“None at all”) to 5 (“A great deal”), regardless of the inferred intent of the publisher. The inter-rater agreement was high as well, with weighted Cohen-Kappa’s correlation coefficients of 0.86, 0.72, and 0.81.

For the main study, we recruited 150 respondents from Prolific. Respondents were 62% females between the age of 18 and 63 (average 33.2, s.d. 10.8). We limited our sample selection to native speakers in English-speaking countries (U.K., U.S., Ireland). The population’s sample was diverse, with 60 full-time employees, 32 part-time employees, 22 unemployed, 14 homemakers or retirees, and 22 “others” (including students). For each respondent, we displayed a unique headline picked randomly from the pool of 36 headlines described above.

For the emotions model, we replicated the Berger and Milkman (2012) dimensions. We asked each respondent to rate the headline positivity, emotionality, awe, anger, and sadness. We also included relevant control variables such as practical utility, interest, and surprise. It should be noted that Berger & Milkman initially relied on independent judges to rate the 6,956 articles under their consideration. However, the same political headline might enrage one reader but delight another. Likewise, the perceived interest of a news article is likely subjective. Since emotions and interests are individual-specific, we found it more appropriate to ask readers to rate headlines themselves rather than rely on external judges. This also decreased the likelihood of spurious correlations between clickbait and the other measures.

We asked respondents to rate whether they perceived a manipulative intent using the five-item version of the Inferences of Manipulative Intent (IMI) scale developed by Campbell (1995). Cronbach’s alpha is high at $\alpha = 0.89$.

For source derogation, we relied on the six-item scale developed by Reser (1972) to measure the influence of manipulation intent on the perception of the source. The scale includes likable, pleasant, sincere, trustworthy, competent, and well-informed. They are averaged to measure the overall perception of the “manipulator” (Reser, 1972, p. 38). If respondents infer manipulative intents, we expect the source to suffer from a negative gestalt and accordingly be evaluated badly across various dimensions (McCornack & Ortiz, 2021). This negative “halo effect” is indeed confirmed by a high Cronbach’s alpha of 0.94.

Finally, we asked respondents to indicate how likely they would share this article on their social media feed on a scale from 1 (“Extremely unlikely”) to 7 (“Extremely likely”). We report all the scales used, along with their descriptive statistics, in Online Appendix A.

Table 2 (Study 1) Factor loadings of a principal component analysis after varimax rotation. Factor loadings above 0.4 are in bold; those below 0.1 are not reported. The constructs used by Berger and Milkman (2012) are summarized along three orthogonal dimensions: positive emotion, negative emotion, and utility

	<i>PC1</i>	<i>PC2</i>	<i>PC3</i>
	Positive emotion	Negative emotion	Utility
Positivity	0.436	-0.376	
Emotionality	0.565	0.235	-0.114
Awe	0.547		0.112
Anger		0.639	
Sadness	0.117	0.622	
Utility	-0.119		0.789
Interest	0.123		0.588
Surprise	0.380		
Variance explained	30.8%	25.4%	11.5%
Eigenvalues	2.467	2.030	0.922

The emotions model

Initially, our attempt to replicate the findings from Berger and Milkman (2012) was not met with great success, and only one parameter barely achieved statistical significance. Some factors might be blamed for this failed replication. The authors’ dependent variable is binary and focuses on outliers (i.e., article belongs to the “most shared articles” list); in contrast, ours is continuous and covers the whole spectrum of sharing intentions (on a 1–7 scale). Also, the sample size is markedly different ($N = 150$ vs. $N = 6,956$), and the authors’ data were entirely keyed in by professional judges. More importantly, it appears that our data suffer from multicollinearity. Unsurprisingly, factors such as sadness and anger ($r = 0.650$), awe and emotionality ($r = 0.489$), or anger and positivity ($r = -0.427$) are highly correlated. In our survey, the median inter-item correlation is 0.266; it is 0.065 in Berger & Milkman.

To circumvent this issue, we run a principal component analysis (PCA) and find that the data can be summarized along three dimensions² that we label “positive emotion,” “negative emotion,” and “utility.” We report the factor loadings (after varimax rotation) in Table 2.

After reducing the data’s dimensionality through principal component analysis, a much clearer picture emerges (see Table 3 “Model 1”): as reported in Berger &

² A scree plot analysis points to a clear 3-factor solution. An analysis of the eigenvalues leads to a less clear-cut diagnosis, with the third factor having an eigenvalue of 0.9216, below the traditional cutoff value of 1. However, the third dimension captures the control variables, and we retain it in the analyses for that reason. Replicating the analyses with 2, 4, and 5 factors do not lead to substantially different results.

Table 3 (Study 1) Parameter estimates for the full model and various nested specifications

	<i>MODEL 1</i>		<i>MODEL 2</i>		<i>MODEL 3</i>		<i>MODEL 4</i>	
	Emotion model only (linear)		Emotion model only (quadratic)		Clickbait model only		Full specification	
Sharing:								
Intercept	2.107	****	1.475	****	4.205	****	2.953	****
Positive emotion	0.342	****	0.322	***			0.217	**
(Positive emotion) ²			0.125	**			0.103	*
Negative emotion	0.012		0.016				0.096	
(Negative emotion) ²			0.079				0.473	
Utility	0.352	***	0.386	****			0.202	
(Utility) ²			0.162	**			0.146	**
Omitted information					0.172	**	0.129	
Source derogation					-0.868	****	-0.575	****
<i>R</i> ² (adjusted)	0.185		0.234		0.224		0.289	
<i>RMSE</i>	1.513		1.467		1.476		1.413	
Omitted information:								
Intercept					0.447	****	0.447	****
Clickbait					0.891	****	0.891	****
<i>R</i> ² (adjusted)					0.931		0.931	
<i>RMSE</i>					0.366		0.366	
Manipulative intent:								
Intercept					2.364	****	2.364	****
Clickbait					0.186	****	0.186	****
<i>R</i> ² (adjusted)					0.070		0.070	
<i>RMSE</i>					0.978		0.978	
Source derogation:								
Intercept					0.883	****	0.883	****
Manipulative intent					0.729	****	0.729	****
<i>R</i> ² (adjusted)					0.608		0.608	
<i>RMSE</i>					0.593		0.593	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Milkman, positive emotion is a strong predictor of shares, even after controlling for interest and utility.

Beyond the valence of the emotions (i.e., positive vs. negative), Berger & Milkman also report that articles that elicit high-arousal emotions are more likely to be shared than those eliciting low-arousal emotions. To go beyond emotion valence and capture emotion strength, we introduce quadratic effects in the model (see Table 3, “Model 2”). We find that the relationship between positive emotion and shares is not linear. Extreme-emotion headlines are disproportionately likely to be shared. The nonlinear influence of negative emotions is also better captured, although it still fails to achieve significance. Introducing quadratic effects improve model fit markedly, even after taking into account the increased number of parameters (adjusted R^2 increases from 0.185 to 0.234, in line with Berger & Milkman’s reported R^2 of 0.28—their full model includes many more control variables than ours).

If we assume that the quadratic terms capture high-arousal emotions, we can confidently report that we replicate Berger & Milkman’s results: “positive content is more viral than negative content [...]. Content that evokes high-arousal [...] emotions is more viral. [...] These results hold even when the authors control for how surprising, interesting, or practically useful content is (all of which are positively linked to virality)” (Berger & Milkman, 2012, p. 192). In the context of this research, we find strong empirical evidence to support H1; headlines with a high emotional appeal are more likely to be shared than headlines with a low-emotional appeal.

For H2, however, supporting evidence is weak at best. Clickbait headlines are not strongly related to increased emotional appeals in the headlines (details not reported in the interest of space). The only (barely) significant parameter is between clickbait and negative emotion ($\beta = 0.142$, $p = 0.065$). The quadratic term is not significant, and

clickbait headlines are not perceived higher in terms of positive emotional appeal or utility. Therefore, we reject H2. We discuss this finding in greater detail at the end of this section.

The curiosity gap and the resistance to persuasion models

When it comes to predicting shares, especially when it comes to clickbait headlines, the role of emotions might only be a part of the story. We calibrate a model that focuses on the curiosity gap and resistance to persuasion models (Table 3, “Model 3”).

For the curiosity gap model, clickbait is indeed a strong predictor of the amount of omitted information in a headline. The parameter is positive and strongly significant ($\beta=0.891$, $p<0.001$), hence supporting H3. Consistent with the curiosity gap model developed by (Loewenstein, 1994), the amount of omitted information is, in turn, a predictor of how likely a headline will be shared ($\beta=0.172$, $p=0.053$). However, the relationship appears weaker than expected (see full specification model hereafter). H4 is only partially supported.

Clickbait is strongly associated with a perception of manipulation intent ($\beta=0.186$, $p<0.001$), hence providing empirical evidence for H5. Perception of manipulative intent, in turn, causes source derogation ($\beta=0.729$, $p<0.001$), confirming H6. A full mediation analysis using the PROCESS macro (*processR* by Moon, 2021) confirms that the effect is fully mediated (indirect effect: $p<0.001$; direct effect: $p=0.562$).

Source derogation of the publisher causes a significant drop in the likelihood of sharing its article ($\beta=-0.868$, $p<0.001$). The influence of inference of manipulative intent on shares is fully mediated by source derogation (indirect effect estimated using 200 bootstrap draws: $p=0.001$; direct effect: $p=0.724$). We find strong evidence in support of H7 as well.

The model specification that focuses exclusively on the curiosity gap and resistance to persuasion pathways has a similar explanatory power to Berger & Milkman’s original emotions model ($R^2=0.224$ vs. $R^2=0.234$).

Full model

The preceding specifications can be seen as nested versions of a full model that incorporates all possible paths to sharing: emotions, curiosity gap, resistance to persuasion, and additional control variables. The results of this fully-specified model are reported in “Model 4” in Table 3. The results are summarized in Fig. 2.

A few key results are worth noting. First, after controlling for positive and negative emotions and utility, the amount of omitted information—which was only marginally

contributing to shares in the nested curiosity gap model—becomes insignificant ($\beta=0.129$, $p<0.137$). Second, the emotions model does not predict that clickbait headlines will be largely shared because they do not seem to either trigger high positive emotions or convey high utility value (the most important determinants of sharing in the emotions model). Third, even after controlling for emotions and utility, the resistance to persuasion model is strongly confirmed. Actually, source derogation is the strongest predictor of sharing behaviors.

While it is partly conjecture, and we do not have longitudinal data to back up such a claim, it seems that clickbait tactics have been used—and sometimes overused—by online publishers to a point where online readers’ became partially immune to their lure. While clickbait headlines purposely withhold information to create a curiosity gap, online customers have long learned that there is not much value behind the artificially-crafted mysteries. Likewise, the hyped headlines, exciting promises, and abuse of exclamation points do not seem to trigger significant emotional responses anymore. At the other end of the spectrum, many online readers are now attuned to the manipulative tactics employed and derogate publishers who use them, hence impeding shares.

Whether these findings hold at a large scale, beyond the confine of a controlled experiment, is an empirical question that we examine in Study 2.

Study 2: Field study

Introduction

In our second study, we validate with real-life data whether, as predicted, clickbait headlines are shared to a lesser extent than non-clickbait headlines. This question raises a significant challenge in terms of disentangling the effect of clickbait framing on the one hand and the nature of the articles that are typically framed as clickbait on the other hand. For instance, if clickbait tactics are more widely used for celebrity-related news than for politics-related news, and everything else being equal, celebrity headlines are more likely to be shared, then there will be a confounding effect between the clickbait treatment and the nature of the articles this treatment is more likely to be applied to.

We address this looming endogeneity challenge as follows. First, we include a wide array of control variables in the model, including topic analysis, emotion classification, headline characteristics, publisher characteristics (e.g., number of followers, political stance, audience reach, topic specialization), and audience characteristics. For completeness, we also test a model with as many dummy variables as publishers. We detail the data and control variables below.

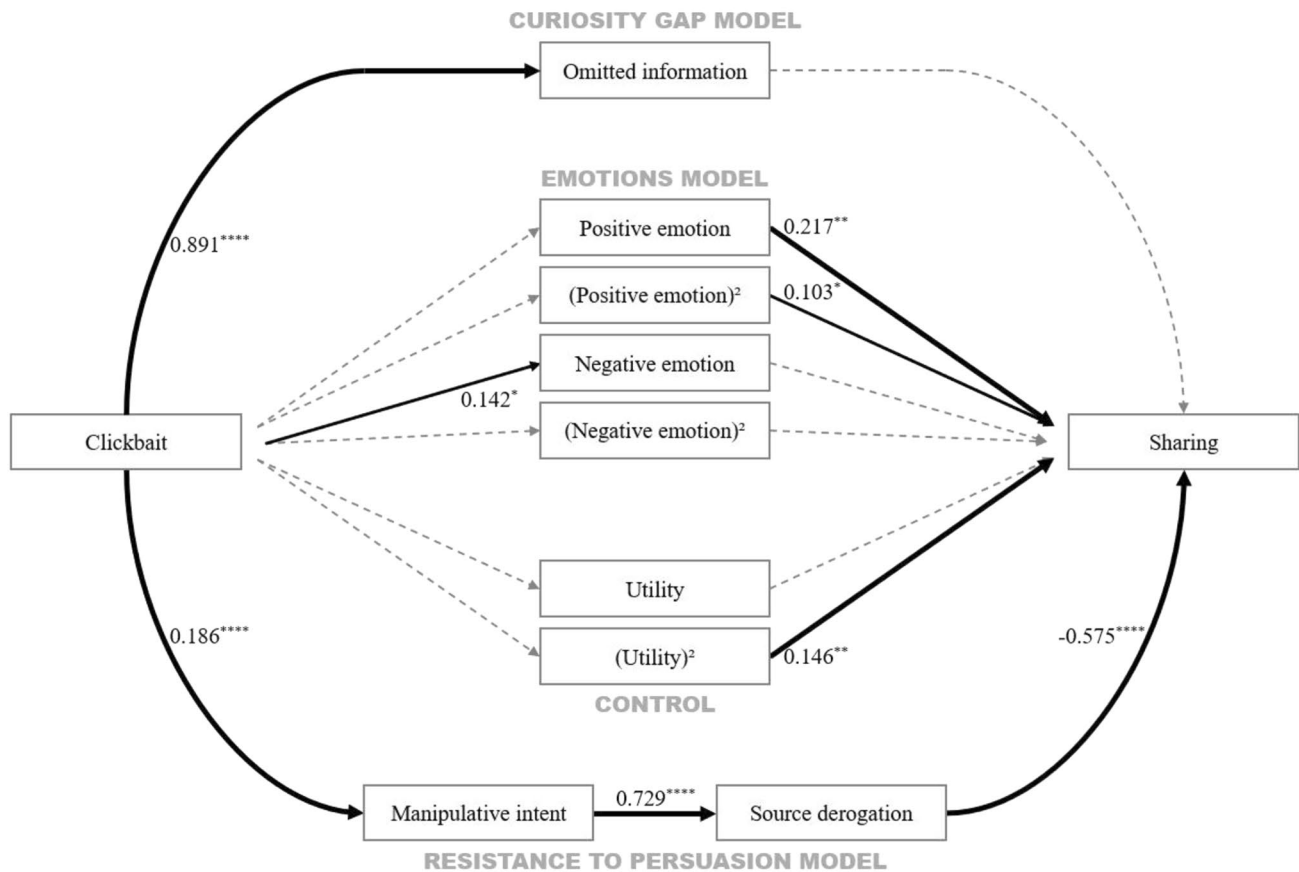


Fig. 2 (Study 1) Key results for the fully-specified model, accounting for all paths to sharing behaviors: emotions, curiosity gap, resistance to persuasion, and control variables

Second, our modeling approach consists of creating two distinct groups of articles: the control and the treatment group, using propensity score matching (PSM). The control group's headlines are non-clickbait, whereas the treatment group's headlines are clickbait, but are otherwise comparable along all other dimensions. Therefore, the difference in sharing between the two groups can be attributed to clickbait alone. We detail the methodology hereafter.

Data

Webis Clickbait Challenge data set

The Webis Clickbait Challenge 2017 was a machine learning contest instituted to develop novel methods to detect clickbait from online content (Potthast et al., 2018a, b), and is today the best-known training set for machine learning researchers interested in developing clickbait detection algorithms. Potthast et al. (2018b) divided their data set into two parts: a publicly available training set and a privately hosted validation set to validate contestants' responses. In our work, we use only the publicly available training set

for our analysis. To the best of our knowledge, this is the first application of this data set in mainstream marketing research. Potthast et al. (2018b) present their data collection protocol, which we summarize here.

Potthast et al. (2018b) selected the 27 most shared English language media handles on Twitter. For each of these publishers, they collected every news article published on Twitter for four months, between December 1, 2016, and April 30, 2017. From this massive corpus of over half a million tweets, they randomly sampled 38,517 tweets, keeping the number of tweets per publisher similar. They published 19,518 of those tweets, which we use in our analysis. These tweets are all of the format “headline + URL.” The date, time of posting, and full text of each article are available in the Webis Clickbait Challenge 2017 data set. The data set does not directly provide the publisher's identity. However, we could infer the names of the publishers from their URL (e.g., The New York Times uses *nyt.ms*) or by pinging each shortened URL.

Potthast et al. (2018b) presented each article in the clickbait corpus to five human judges on Amazon's Mechanical Turk. Each judge was presented with articles' headlines and

Table 4 (Study 2) Percentage of clickbait by publisher. ABC News (Australia) is used as the base publisher in regressions henceforth

Publisher name	# of Obs	% Clickbait	Publisher name (cont'd)	# of Obs	% Clickbait
Breitbart	724	83.98	The Washington Post	766	20.50
BuzzFeed	734	62.53	The New York Times	775	20.12
Forbes	737	38.53	The Telegraph	738	16.67
Indiatimes	745	37.72	Complex	762	16.53
Mashable	698	36.82	CNN	717	15.06
Yahoo	720	36.81	The Guardian	788	14.59
The Independent	720	33.89	The Wall Street Journal	734	12.81
Business Insider	686	31.49	CBS News	722	12.19
BBC	739	24.36	ABC News (Australia)	744	11.83
ESPN	360	23.06	NBC News	790	11.52
Daily Mail	732	22.40	Billboard	762	10.63
Huffington Post	776	21.39	Bleacher Report	548	8.21
Bloomberg	739	20.84	Fox News	757	6.07
			ABC News (US)	673	5.05
Total	19,386	24.31			

Table 5 (Study 2) Examples of clickbait and non-clickbait headlines from our data

Headline	Publisher	Clickbait?
A Superior Chicken Soup	The New York Times	Yes
Leah Remini's Reddit AMA reveals juicy secrets of Scientology	Mashable	Yes
Can Work Life Balance Be a Reality? This Company Makes it Possible	NBC News	Yes
Visit Myanmar's Capital Now! There's Still a Lot Not to See	The Wall Street Journal	Yes
Panama Papers: Europol links 3,500 names to suspected criminals	The Guardian	No
100 Women 2016: On the frontline with the women policing the peace in Afghanistan	BBC	No
Older Viewers and Conservatives Are Watching Less NFL, Survey Finds	The Wall Street Journal	No
5 Dead, 7 Injured After Tornadoes in Alabama and Tennessee	ABC News (U.S.)	No

their URL, which they could click on to view the main content. Respondents rated each article on the degree of “clickbaitiness” on a 4-point Likert scale (0 = “Not clickbaiting”; 0.33 = “Slightly clickbaiting”; 0.66 = “Considerably clickbaiting”; 1 = “Heavily clickbaiting”). The data set provides five individual clickbait ratings for each headline.

Potthast et al. (2018a) define the binary variable *clickbait* as 1 if the *mode* of the five responses is greater than 0.5 and as 0 otherwise. Although, as we discussed earlier, clickbait is best viewed as a continuous construct, we retain this dichotomous classification for several reasons:

1. This dichotomous conceptualization is widely used in practice, both in academic research (e.g., Grigorev, 2017; Kumar et al., 2018; Papadopoulou et al., 2017; Potthast et al., 2018a; Thomas, 2017; Wiegmann et al., 2018) and in the industry (for clickbait detection). Therefore, it ensures comparability between our analyses and those reported in the literature.
2. From a methodological perspective, standard propensity score matching requires a dichotomous classification of

the dependent variable. A continuous scale would make controlling for endogeneity—a pressing concern in this setting—much more arduous.

3. The results of various tests (cf. Online Appendix D) confirm that results are robust to a mean-split operationalization of clickbait as a dichotomous variable as well.

Based on their mode, 4,713 of the 19,386 articles³ are classified as clickbait, illustrating the prevalence of clickbait tactics in the publishing industry. We report statistics of clickbait by publisher in Table 4, and selected examples of headlines in Table 5.

Twitter data

The dependent variable of our model is whether clickbait articles are less likely to be shared than non-clickbait

³ The final sample size is 19,386 due to some deleted tweets, some non-English articles inadvertently included in the original data source and other parsing errors.

articles. We took each URL in the corpus, and using a Python script, inferred how many times it had been liked and shared on Twitter (including versions using URL shorteners). This was done in May 2019, two years after the last article was posted. For each of these tweets, we extracted the total number of shares and likes using a browser automation script. We chose Twitter to augment our data set with share and like counts because: (a) the original Webis Clickbait Challenge 2017 corpus itself was sourced from Twitter, (b) a much larger fraction of Twitter profiles are public as compared to Facebook, (c) unlike Facebook whose algorithm explicitly suppresses the propagation of clickbait (Babu et al., 2017),⁴ Twitter has no such policy, and (d) to the best of our knowledge, none of the posts were tagged as “promoted tweet” by Twitter.

Topic modeling

Since clickbait tactics may be more prevalent in some news categories (e.g., fashion, entertainment) than others (e.g., politics), it is important to control for the headline topics. We performed topic modeling on the entire article corpus via Latent Dirichlet Allocation (Blei et al., 2002, 2003; Tirunilalai & Tellis, 2014) using R's *tm* package (Hornik & Grün, 2011). To determine the optimal number of clusters, we used the procedure of Zhang et al. (2017), which involves performing Latent Dirichlet Allocation on the corpus from 2 to N clusters, performing a probit regression on the selection equation (see Online Appendix B), and determining where the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) start to taper off after a steep decline. In our data, we determined this optimal number of clusters to be 12.

After determining the optimal number of clusters, we examined the top 20 words of each cluster and manually inferred each one's theme from its keywords. For example, cluster 1 contains top word stems like *show*, *year*, *star*, *film*, and *music*, and we manually inferred its topic to be ENTERTAINMENT. In contrast, cluster 2 contains top word stems like *polic*, *offic*, *year*, *told*, and *report*—words associated with CRIME reporting. Interestingly, we found several clusters related to politics. On closer inspection, we found that they were easily identifiable as international politics, U.S. politics, and geopolitics. As the number of clusters increases, it is natural to find multiple clusters with very similar themes. The cluster names we assigned to each topic are commonly used keywords in digital media. As an

additional check, we manually verified that clusters with the appropriate labels came from corresponding sources. For example, it is natural for ESPN and Bleacher Report to carry articles primarily affiliated to the SPORTS cluster, while general publishers like The New York Times and BBC carry articles across a larger spectrum of topics. Online Appendix B presents each cluster's top 20 keywords along with its label. We recorded the probabilities of an article belonging to each cluster as covariates in our analysis.

Sentiment analysis

As we have shown in Study 1, emotions play an important part in content sharing, and they need to be included as control variables in the model. We performed sentiment analysis and emotion detection on the entire article corpus, separately on each article's headline and its main body text. We implemented this using R's *syuzhet* package (Jockers, 2017), which uses the widely used NRC lexicon (Mohammad & Turney, 2010) to infer a positive valence score, a negative valence score, and scores on eight discrete emotions (anger, anticipation, disgust, fear, joy, sadness, surprise, and trust) on each article's headline and main body.

Other control variables

Some articles might be easier to read than others, and this characteristic may influence how likely they might be shared. We computed the Flesch Readability Score for each article headline and body using R's *quanteda* package (Benoit et al., 2018) and included them as control variables. The Flesch readability score is a well-known measure to assess a text body's readability and complexity and is widely used in academic research (Flesch, 1948; McLaughlin, 1969).

We also tagged articles published during the weekend (dummy variable *weekend*), as the day of the week a headline is published might affect its likelihood of being shared.

Finally, the Webis Clickbait Challenge data set covers a wide array of publishers, from fashion and sports magazines to news outlets and media companies. These publishers vary in terms of topics covered, audience, and reach. We list below the various control variables that we include at the publisher level:

1. Number of Twitter followers (as headlines from publishers with more followers are, everything else being equal, more likely to be shared).
2. We inferred the political stance of the publisher using the Media Bias Fact Check portal. Two publishers, Billboard and Complex, were not listed. We present later analyses both with and without this variable. The variable (*left*) is dichotomous.

⁴ This update was released in May 2017, just after the time period of the Webis Clickbait Challenge 2017 articles. They did not have such an algorithm in 2014 or earlier, which is the time period when Trotter (2015) accessed BuzzFeed's internal memo.

3. In terms of audience provenance, we used the Similar-Web portal to infer the percentage of traffic coming from direct, referral, search, and social media.
4. The topic coverage of the publisher (*general*) is equal to 0 for narrowly-defined, domain-specific publications (e.g., Billboard) and to 1 for general-purpose, wide-ranging topic publishers (e.g., The Guardian).
5. The binary variable *U.S.* indicates whether the media house is headquartered in the U.S. (e.g., NBC News) or not (e.g., BBC).
6. Some publishers are Web-only (e.g., BuzzFeed). In contrast, others have some offline counterparts (e.g., The Washington Post), and this characteristic might impact the average profile and typical (sharing) behaviors of their target audience. We include the control variable *webonly* to control for that.
7. We also define the binary variable *paid* if the publisher has any pay-walled content (e.g., The New York Times) and 0 if all content is free (e.g. Huffington Post).

Despite our best efforts to include as many relevant control variables as possible, some key publishers' characteristics may not be properly captured, such as reputation, general trustworthiness, audience composition, or other factors potentially influencing sharing behaviors. To avoid any confounding effects, we calibrate different model versions where we include as many dummy variables as there are publishers. As we report later, the results are markedly similar.

The names and descriptions of each variable in the data, along with their respective means and standard deviations, are available in Online Appendix B.

Model-free evidence before propensity score matching

In line with our main hypothesis, it appears that clickbait headlines are liked and shared to a lesser extent than non-clickbait headlines. The average number of likes and shares for non-clickbait headlines are 329.22 and 155.65, respectively. The same figures for clickbait headlines shrink to 194.27 (-41%) and 80.61 (-48%). Clickbait articles appear to be both liked and shared less than non-clickbait articles. However, these results do not correct for endogeneity and should not be interpreted at face value.

Methodology

We use propensity score matching (PSM), a well-established method to reduce the bias caused by confounding variables, namely, the effect of clickbait on shares on the one hand and the higher propensity of some articles to be framed as clickbait on the other hand. We generate a sample of non-clickbait (control) articles that closely match

Table 6 (Study 2) Probit model results. Dependent variable: clickbait. Publisher dummies are included in regression but excluded here for conciseness. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

VARIABLES	Parameter	St. Dev
headline:anger	-0.0347	0.0285
headline:anticipation	0.00487	0.022
headline:disgust	0.0238	0.0342
headline:fear	-0.0275	0.0249
headline:joy	-0.0174	0.0307
headline:sadness	-0.0305	0.028
headline:surprise	-0.0463*	0.0246
headline:trust	0.0176	0.02
headline:negative	-0.0265	0.0218
headline:positive	0.00275	0.0191
headline:wordcount	-0.0273***	0.00355
headline:readability	0.00678***	0.000474
topic:crime	-0.214	0.171
topic:sports	-0.106	0.175
topic:health	0.970***	0.147
topic:socialmedia	1.325***	0.184
topic:lifestyle	2.136***	0.164
topic:people	2.201***	0.177
topic:intlpolitics	-0.519**	0.186
topic:travel	0.285*	0.159
topic:USpolitics	-0.597***	0.154
topic:economy	0.354***	0.143
topic:geopolitics	-0.442**	0.174
weekend	0.0637***	0.024
Constant	-1.779***	0.126
Observations	19,386	
AIC	17,242.18	
BIC	17,643.67	

clickbait (treatment) articles based on observable characteristics in our data (Dehejia & Wahba, 2002; Hofmann et al., 2017; Kumar et al., 2016; Rishika et al., 2013; e.g., Rosenbaum & Rubin, 1983). We relegate the methodology details to Online Appendix B and provide the salient points here.

We conduct the propensity score analysis using a probit model. The control variables in this selection equation, along with the estimated coefficients of the probit model, are listed in Table 6. For instance, headlines covering politics (U.S., international, or geopolitics) are less likely to be framed as clickbait. Conversely, clickbait headlines are more likely to be published during the weekend, covering topics such as social media, people, or health, and with shorter word counts and higher readability scores.

We match the articles in the treatment (*clickbait* = 1) and control (*clickbait* = 0) groups using these calculated propensity scores by the nearest neighbor (with one neighbor)

matching with replacement. As a result, the original 4,712 clickbait articles (one clickbait article is dropped due to lack of common support) are matched to an equivalent number of non-clickbait articles. In the nearest neighbor matching with replacement, an article in the control group can be used more than once as a match, which increases the average quality of matching and decreases bias (Caliendo & Kopeinig, 2008). The matched sample consists of 7,576 unique articles: 4,712 clickbait articles, and 2,864 non-clickbait articles (some matched more than once), for a balanced sample (50% control, 50% treatment) of 9,424 articles.

Next, we check if the underlying assumptions of the PSM methodology hold. We need to ensure that there is substantial overlap in the characteristics of clickbait and non-clickbait articles so that the common support condition holds (Rosenbaum & Rubin, 1983). As per Lechner (2002)'s recommendation, we plot distributions of propensity scores before matching and after matching as box plots and histograms in Online Appendix B, offering some visual confirmation of the common support condition. We also use the minima-maxima approach (Caliendo & Kopeinig, 2008), where we delete all observations whose propensity score is smaller than the minimum and larger than the maximum in the opposite group. The results do not change substantially.

The quality of matching can be assessed by comparing the means of the conditioning variables for clickbait and non-clickbait articles before and after matching. Online Appendix B provides details of covariate balance between the clickbait and non-clickbait groups post-matching. We also report another indicator, “standardized bias” (*S.B.*), to assess whether the difference in means is large. The *S.B.* approach does not have a clear indication for the success of the matching procedure, even though in most empirical studies, an *S.B.* below 3% or 5% after matching is considered sufficient (Caliendo & Kopeinig, 2008). As evident from Table 9, except for only one variable (*topic: crime*), the standardized bias after matching is below 5%. The median (mean) standardized bias for all covariates is 12.5 (14.2) before matching and 1.3 (1.6) after matching.

The *pseudo-R*² before and after matching is an alternative measure of matching effectiveness. Matching reduces *pseudo-R*² from 0.203 to 0.005. The hypothesis of the joint insignificance of all the regressors cannot be rejected after matching (*p*-value = 0.132) (Sianesi, 2004). Thus, matching does a good job in making the treatment and control groups comparable.

Finally, we conduct an analysis to check for any “hidden” bias. We aim at analyzing how strongly an unmeasured variable may influence the decision to frame an article as clickbait, thereby potentially undermining the results of the matching analysis (DiPrete & Gangl, 2004; Imbens, 2003). We check for this “hidden” bias by deleting some

variables used for estimating the propensity scores, followed by the matching process, and measuring how the results are affected. The results do not change substantially. Our results do not change substantially either, after using an augmented set of variables. We report the details in Online Appendix C.

All the aforementioned tests and validation procedures lead to the same conclusion: all the PSM assumptions are fulfilled. We can assume with reasonable certainty that the PSM approach we employed satisfactorily reduces the possible bias caused by confounding variables.

Model-free evidence after propensity score matching

Before PSM, the clickbait headlines were shared significantly less than non-clickbait headlines (80.61 for treatment vs. 155.65 for control). After PSM and controlling for selection effects, the difference shrinks but remains strongly significant; a clickbait headline elicits on average 48.58 fewer shares than a non-clickbait headline (*p* < 0.01). It also commands 119.17 fewer likes than non-clickbait (*p* < 0.01).

Regression analyses

Even after PSM, clickbait articles appear to be shared less than non-clickbait articles. However, this result does not control for publishers' Twitter followers, which may play a role in sharing outcomes (e.g., Yoganarasimhan, 2012). Certain publisher and audience characteristics might also affect sharing, above and beyond the inferred characteristics of the article being shared. To control for the possible impact of such confounding factors, we run five ordinary least squares regressions on the matched sample:

- Models 1 and 4 do not control for articles' characteristics, such as sentiment, topic, word count, published during the weekend, and readability score. Models 2, 3, and 5 do.
- Models 1 to 3 include numerous controls for the publishers' idiosyncratic characteristics: Model 1 includes the number of Twitter followers; Model 2 adds additional controls such as generalist, Web-only publisher, U.S. origin, traffic source, or presence of a paid wall (see the “Data” section for details); Model 3 incorporates political stance. However, this latter indicator is missing for two publishers (Complex and Billboard). Model 3 is, therefore, calibrated on a smaller sample of articles.
- To account for the possibility that the full list of controls about the publishers and their audiences may still ignore unidentified confounding effects, Models 4 and 5 replace all the publisher-related control variables with as many

Table 7 (Study 2) Model versions. The variable left is only available for some publishers, hence Model 3 is applied to a subsample. As soon as individual publisher dummies are included, all other publisher-specific indicators become redundant and need to be removed (Models 4 and 5)

VARIABLES	Model 1	Model 2	Model 3	Model 4	Model 5
Article characteristics:					
– Headline: sentiment analysis (10)		✓	✓		✓
– Headline: word count, readability (2)		✓	✓		✓
– Article: sentiment analysis (10)		✓	✓		✓
– Article: word count, readability (2)		✓	✓		✓
– Topic modeling (11)		✓	✓		✓
– Weekend (1)		✓	✓		✓
Publisher characteristics:					
– Twitter followers (1)	✓	✓	✓		
– U.S., paid, general... (8)		✓	✓		
– Left (1)			✓		
– Publisher dummies (26)				✓	✓

dummy indicators as there are publishers (minus one for identification purposes).

Table 7 summarizes the predictors of each model, and Tables 8 and 9 present the results for sharing and like, respectively.

All models confirm that, even after controlling for headline or article corpus sentiment, Twitter followers, topic categories, and publishers' characteristics, the negative impact of clickbait on likes and shares is both large and statistically significant. Clickbait impedes sharing both directly (fewer shares) and indirectly (fewer likes, leading online platform algorithms to feature these articles less aggressively in readers' feed), offering strong externally valid evidence to our assertion that clickbait is shared on average less than non-clickbait on social media.

Discussion

Summary findings

The key findings of this research can be summarized as follows. First, the use of clickbait is often seen by the reader as a publisher's manipulative tactic. Second, this perception may lead to the reader resisting the said manipulation by engaging in a source derogation strategy, which in turn may reduce the publisher's perceived competence and trustworthiness. Third, readers appear less likely to share links from such a source-derogated publisher. All the above are confirmed in our controlled experiment. Finally, we show with actual secondary data that clickbait is indeed, *on average*, shared much less than non-clickbait on social media, with large effect sizes, even after controlling for endogeneity and other covariates.

Managerial implications

Our results are relevant to online marketing in general and online media and news portals in particular, especially as clickbait is an ever-present phenomenon today. While many believe that clickbait has “cracked the code” on shareable content, our results show this claim may not be warranted. Clickbait may be counterproductive, especially if a publisher relies on it to increase its reach via sharing. Thus, clickbait as a headline framing treatment represents a tradeoff. While it may enhance direct readership (several headline-optimizing services confirm that clickbait headlines are clicked more), it also impedes organic reach via likes and sharing. The usage of clickbait tactics, therefore, necessitates higher expenditures on social media platforms to propagate the same content. A profit-maximizing online publisher needs to take this into account, especially as word-of-mouth cascades are an important driver of reach (e.g., Zubcsek & Sarvary, 2011).

We acknowledge that some organizations, notably BuzzFeed—which is famous for its analytics and A/B tests on content (Wang, 2017)—may be fully aware that clickbait impedes sharing. However, they may persist with clickbait for two possible reasons: (a) the optimal resource allocation for click-based revenue dictates the use of clickbait along with large expenses on sponsored content propagation, and (b) clickbait acts as a teaser to get a reader on to their portal, after which they browse more articles there. Our data do not allow for such structural analysis.

Avenues for further research

Clickbait is ubiquitous in digital media today and has many facets of interest to marketing and adjacent disciplines. While our work questions the widespread notion of clickbait as viral content, its omnipresence in the more

Table 8 (Study 2) Ordinary Least Squares Regression. Dependent variable: shares and primary independent variable: clickbait. Publisher dummies are included in regressions in models 4 and 5, but excluded here for conciseness. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

VARIABLES	Model 1	Model 2	Model 3	Model 4	Model 5
clickbait	-42.78**	-48.99***	-51.92***	-50.26***	-48.51***
headline:anger		-0.0553	-0.0161		-2.494
headline:anticipation		-14.47	-14.20		-17.12
headline:disgust		-11.34	-11.22		-14.63
headline:fear		12.50	11.92		15.68
headline:joy		4.873	4.717		4.876
headline:sadness		5.418	2.996		3.737
headline:surprise		16.73	17.84		16.73
headline:trust		5.492	5.479		6.435
headline:negative		3.388	4.430		2.687
headline:positive		-5.891	-5.009		-3.569
headline:word count		2.097	1.983		0.658
headline:readability		0.309	0.262		0.469
article:anger		-0.391	-0.602		-0.563
article:anticipation		0.818	0.802		0.857
article:disgust		-1.422	-1.581		-1.58
article:fear		1.704*	2.078**		1.852**
article:joy		0.0358	0.172		0.396
article:sadness		-1.225	-1.328		-1.38
article:surprise		0.461	0.641		0.99
article:trust		0.0836	0.0824		0.0661
article:negative		0.101	0.115		0.365
article:positive		-0.119	-0.0600		-0.301
article:word count		-0.0183	-0.0263		-0.0277
article:readability		-0.544	-0.619		-1.196
topic:crime		12.00	88.75		-31.1
topic:sports		63.30	181.0*		-259.9***
topic:health		-85.27	-12.80		-86.66
topic:socialmedia		-66.26	2.339		-44.49
topic:lifestyle		-100.7	-22.16		-107.3
topic:people		-52.81	29.66		-78.25
topic:intlpolitics		-36.50	30.78		-45.12
topic:travel		-104.4	-11.69		-82.37
weekend		15.39**	14.74**		19.29***
Twitter followers	4.396***	4.827***	5.544***		
U.S		28.74	29.42		
paid		7.171	12.08		
general		7.409	20.13		
web only		27.30	48.21		
percent direct		5.599	8.167		
percent referral		18.36	19.92*		
percent search		4.694	6.988		
percent social		6.712	10.24		
left			-35.19		
Constant	84.03***	-514.3	-836.5	130.0***	212.2***
Publisher dummies	NO	NO	NO	YES	YES
Observations	7,576	7,576	7,188	7,576	7,576
R ²	0.034	0.079	0.087	0.126	0.143

Table 9 (Study 2) Ordinary least squares regression. Dependent variable: likes and primary independent variable: clickbait. Publisher dummies are included in regressions in models 4 and 5, but excluded here for conciseness. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

VARIABLES	Model 1	Model 2	Model 3	Model 4	Model 5
clickbait	-89.61**	-115.0**	-124.1**	-112.1**	-113.0***
headline:anger		17.81	17.31		8.671
headline:anticipation		-39.57	-41.34		-49.3
headline:disgust		-6.270	-7.418		-17.68
headline:fear		-7.104	-9.046		4.992
headline:joy		23.06	22.50		25.56
headline:sadness		-5.053	-9.052		-7.725
headline:surprise		35.76	37.11		38.26
headline:trust		5.294	7.073		9.562
headline:negative		4.530	7.995		4.146
headline:positive		-2.079	0.160		2.291
headline:word count		3.479	2.785		0.223
headline:readability		1.092	1.064		1.551
article:anger		-1.846	-2.559		-2.119
article:anticipation		2.200*	2.361*		2.133
article:disgust		-1.419	-1.778		-1.678
article:fear		2.751	3.553		2.902
article:joy		1.358	1.734		2.305**
article:sadness		-0.813	-0.979		-1.228
article:surprise		0.118	0.669		1.654
article:trust		-0.464	-0.487		-0.487
article:negative		-0.969	-0.780		-0.0906
article:positive		-0.655	-0.770		-1.097
article:word count		-0.0359	-0.0488		-0.058
article:readability		-1.621	-1.833		-3.861
topic:crime		-252.6	-87.59		-440.3
topic:sports		407.1	715.5**		-633.4**
topic:health		-387.1	-234.0		-418.3*
topic:socialmedia		-305.2	-184.6		-186.9
topic:lifestyle		-157.4	3.304		-254.4
topic:people		4.020	168.4		-124.8
topic:intlpolitics		-256.0	-120.4		-306.6**
topic:travel		-387.3	-190.9		-368.1
topic:USpolitics		-226.4	-83.25		-180
topic:economy		-403.6*	-185.8		-347.2*
topic:geopolitics		-217.4	-50.36		-301
weekend		29.54	27.51		41.60*
Twitter followers	9.489***	11.71***	13.75***		
US		111.7	110.6		
paid		-37.37	-38.34		
general		44.69	65.67		
webonly		96.13	150.9		
percentdirect		6.101	15.27		
percentreferral		36.47	42.41		
percentsearch		3.142	11.66		
percentsocial		14.19	27.33		
left			-106.6		
Constant	199.0**	-467.3	-1,502	302.1***	687.5***
Publisher dummies	NO	NO	NO	YES	YES
Observations	7,576	7,576	7,188	7,576	7,576
R ²	0.018	0.062	0.066	0.122	0.133

technologically sophisticated media houses could surely be explained with more granular data (clickstream, advertising, revenues, etc.) that allows for structural modeling of the actual editorial decisions going beyond clickbait. With more granular data, we envision the potential of a study analogous to Van den Bulte and Lilien's (2001) landmark study of medical referrals to establish the effects of a firm's marketing efforts (in this case, paid online traffic) over word of mouth. Additionally, data from large-scale field experiments such as Matias and Munger (2019) can be exploited to decompose the role of direct clicks as well as shares on the actual reach of online clickbait and non-clickbait content.

One particular aspect of clickbait tactics, unexplored in this research, is the actual *consumption* of the articles under consideration. In Study 1, while source derogation had a strong impact on *sharing* an article, it had little to no impact on *clicking* on it ($\beta = -0.163$, $p = 0.428$), and neither did omitted information ($\beta = -0.056$, $p = 0.601$). Still, positive emotion ($\beta = 0.350$, $p = 0.006$) and perceived utility ($\beta = 0.511$, $p < 0.001$) remained strong determinants of article consumption (full results not reported here in the interest of space, but available from the authors). It should be noted that the sharing of an article on social media is not necessarily conditional on its prior consumption: readers can—and regularly do—share articles they have not read or even clicked on, and social media user interfaces have been designed to facilitate such behaviors (e.g., prominent Facebook's share and Twitter's retweet buttons). In Study 1, several respondents indicated they were likely to share an article they were otherwise unlikely to click on. Clicking on clickbait headlines or consuming cheesy articles on the Internet has far fewer social implications than sharing them. Therefore, sharing antecedents might be markedly different from consumption antecedents. Unfortunately, the available data for Study 2 does not allow us to address the second aspect of the research problem. The distinction between sharing and consumption—and the potential disconnect between the two—may warrant further research.

Along with content publishers, brands too use clickbait in their online ads. This opens up a wide array of questions of interest to behavioral researchers on the efficacy of clickbait advertising on several managerially important outcomes like attitude to the brand, brand recall, awareness, and more. If clickbait triggers source derogation, as Study 1 suggests, could this “negative gestalt” effect spill over the brand behind the article and impact its reputation?

While we find the existence of clickbait in multiple topical domains, it is often associated with fake news (e.g., Munger, 2020) and right-wing politics (e.g., Luca et al., 2021; Munger et al., 2020). Indeed, events like Brexit, the 2020 U.S. Presidential elections, and the COVID-19 pandemic have brought both these phenomena into global prominence. Our data show that clickbait is not exclusive to these

domains. While it is true that the extremely right-wing Breitbart is a regular user of clickbait, the highly conservative Fox News employs little to no clickbait at all. Meanwhile, BuzzFeed, known for its liberal stance, is a major user of it. Highly reputable media houses with a liberal slant, like The New York Times, BBC, and The Washington Post all use more clickbait than Fox News, as Table 4 indicates. We posit the question “what are the organizational antecedents of clickbait?” as a promising avenue of investigation.

Our results are also encouraging for policy-makers. They highlight that citizens could be (partly) shielded from the harmful consequences of manipulative tactics as long as they can properly identify publishers' manipulative intent. While stopping the spread of fake news and extremist propaganda—which erode the foundations of our democracies—might be an elusive goal in today's online environment, educating citizens about online manipulative tactics might prove a nobler and far more effective public policy.

The domain of clickbait is of interest to researchers at the intersection of marketing and information systems. Especially of interest are the reactions of publishers to changes in social media platforms' policies to propagate clickbait—of note is Facebook's momentous decision to reduce the visibility of clickbait—leading to significant drops in Upworthy's online traffic (Sanders, 2017). Ongoing research by Sen and Yildirim (2015) uncovers a “clicks bias” in a major Indian online publisher, where a subject's initial traffic affects its future coverage. It is of interest to fully understand the economics of clickbait in the context of such publisher-platform dynamics.

In our controlled experiment, the publisher's name is unknown, and the reader has no prior exposition to the source. Hence, in that peculiar experimental context, source derogation is solely influenced by *one* headline. Yet, in reality, publishers may *routinely* use clickbait tactics. It would be interesting to see the long-term effects of said tactics on source derogation. Does it worsen over time, or do readers get accustomed to it, expect little of clickbait-heavy publishers, and hence do not feel manipulated as much?⁵

Finally, while this research focuses on the average effects of clickbait on sharing, it might be interesting to explore how different population substrates react to such tactics. For instance, sophisticated readers might be more likely to discern the manipulative intent of the publisher, while

⁵ In our field study, clickbait has a negative impact even after introducing dummy variables to control for each publisher's idiosyncratic characteristic. The headlines' clickbaitness still significantly influences shares. Since some sources publish more clickbait headlines than others, and since our model controls for publishers, one could argue that the effects we find likely underestimate clickbait's true effect on shares. Consequently, the results we report should be seen as conservative estimates.

other readers might remain oblivious to them. This seems to suggest that clickbait tactics might be more effective with specific customer profiles. Younger generations are also of particular interest. On the one hand, they are more likely to suffer from the fear-of-missing-out (FOMO, Metz, 2019), hence being particularly sensitive to artificially-created curiosity gaps. On the other hand, as the consumption of clickbait is highly age and ideology-dependent (Luca et al., 2021), we believe that a deeper dive into the socio-demographics of “consumption” of clickbait might be an interesting avenue for future research.

Conclusions

Clickbait is synonymous with “viral” content. However, we demonstrate this claim is not as grounded as previously believed. A controlled experiment indicates that clickbait usage may cause the publisher to be derogated in the eyes of the reader, leading to lowered intention to share. We back this up with a rigorous field study, demonstrating that clickbait articles are indeed shared much less than non-clickbait articles on social media.

To the best of our knowledge, this is the first systematic study of clickbait going beyond machine learning detection algorithms. Knowing that, among the 19,386 randomly selected articles from 27 leading online publishers, a quarter of them have been classified as clickbait (indicating the prevalence of the phenomenon in online marketing and publishing), this gap in the marketing literature is surprising. Our study adds to the digital and interactive marketing literature, especially in the domain of unintended consequences of interactivity (Deighton & Kornfeld, 2009). It also fits well in research agendas laid out in word-of-mouth (Berger, 2014), digital marketing (Kannan & Li, 2017), and social media in marketing (Appel et al., 2020) research.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11747-021-00830-x>.

Acknowledgements We are thankful to Albert Bemmaor, Joel Bothello, Pradeep Chintagunta, Manish Gangwar, Reetika Gupta, Pranav Jindal, Sreelata Jonnalagedda, Gilles Laurent, Vivek Kaushal, Girish Mallapragada, Dalhia Mani, Ayse Onculer, Sonja Prokopec, Varun Ramachandra, Vithala Rao, Madhu Viswanathan, Sudhir Voleti, Sai Yayavaram, and participants of the Chicago Booth-India Quantitative Marketing Conference 2019, the ISB Conference on the Digital Economy 2019, the ISMS Marketing Science Conference 2020 and the Interactive Marketing Research Conference 2020 for their helpful comments and suggestions. We thank Kiran Jonnalagedda for helping us implement the scripts to capture data from Twitter.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

- Abelson, R. P., & Miller, J. C. (1967). Negative persuasion via personal insult. *Journal of Experimental Social Psychology*, 3(4), 321–333.
- Akpinar, E., & Berger, J. (2017). Valuable virality. *Journal of Marketing Research*, 54(2), 318–330.
- Appel, G., Grewal, L., Hadi, R., & Stephen, A. T. (2020). The future of social media in marketing. *Journal of the Academy of Marketing Science*, 48(1), 79–95.
- Araujo, T., Neijens, P., & Vliegenthart, R. (2015). What motivates consumers to re-tweet brand content?: The impact of information, emotion, and traceability on pass-along behavior. *Journal of Advertising Research*, 55(3), 284–295.
- Babu, A., Liu, A., & Zhang, J. (2017). New updates to reduce clickbait headlines. *Facebook Newsroom*.
- Bansal, H. S., & Voyer, P. A. (2000). Word-of-mouth processes within a services purchase decision context. *Journal of Service Research*, 3(2), 166–177.
- Bazaco, Á., Redondo, M., & Sánchez-García, P. (2019). Clickbait as a strategy of viral journalism: Conceptualisation and methods. *Revista Latina De Comunicación Social*, 74, 94.
- Belch, G. E. (1981). An examination of comparative and non-comparative television commercials: The effects of claim variation and repetition on cognitive response and message acceptance. *Journal of Marketing Research*, 18(3), 333–349.
- Benoit, K., Watanabe, K., Wang, H., Nulty, P., Obeng, A., Müller, S., & Matsuo, A. (2018). quanteda: An R package for the quantitative analysis of textual data. *Journal of Open Source Software*, 3(30), 774.
- Berger, J. (2011). Arousal increases social transmission of information. *Psychological Science*, 22(7), 891–893.
- Berger, J. (2014). Word of mouth and interpersonal communication: A review and directions for future research. *Journal of Consumer Psychology*, 24(4), 586–607.
- Berger, J., & Milkman, K. L. (2012). What makes online content viral? *Journal of Marketing Research*, 49(2), 192–205.
- Berger, J., & Schwartz, E. M. (2011). What drives immediate and ongoing word of mouth? *Journal of Marketing Research*, 48(5), 869–880.
- Berlyne, D. E. (1954). A theory of human curiosity. *British Journal of Psychology. General Section*, 45(3), 180–191.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan), 993–1022.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2002). Latent dirichlet allocation. *Advances in Neural Information Processing Systems*, 601–608.
- Brehm, J. W. (1966). *A theory of psychological reactance*.
- Caliendo, M., & Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1), 31–72.
- Campbell, M. C. (1995). When attention-getting advertising tactics elicit consumer inferences of manipulative intent: The importance of balancing benefits and investments. *Journal of Consumer Psychology*, 4(3), 225–254.
- Davenport, T. H., & Beck, J. C. (2001). The attention economy. *Ubiquity*, 2001(May), 1-es.
- Dehejia, R. H., & Wahba, S. (2002). Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and Statistics*, 84(1), 151–161.
- Deighton, J., & Kornfeld, L. (2009). Interactivity’s unanticipated consequences for marketers and marketing. *Journal of Interactive Marketing*, 23(1), 4–10.
- DiPrete, T. A., & Gangl, M. (2004). Assessing bias in the estimation of causal effects: Rosenbaum bounds on matching estimators

- and instrumental variables estimation with imperfect instruments. *Sociological Methodology*, 34(1), 271–310.
- Flesch, R. (1948). A new readability yardstick. *Journal of Applied Psychology*, 32(3), 221.
- Frampton, B. (2015). Clickbait: The changing face of online journalism. *BBC*.
- Fransen, M. L., Smit, E. G., & Verlegh, P. W. (2015). Strategies and motives for resistance to persuasion: An integrative framework. *Frontiers in Psychology*, 6, 1201.
- Friestad, M., & Wright, P. (1994). The persuasion knowledge model: How people cope with persuasion attempts. *Journal of Consumer Research*, 21(1), 1–31.
- Geiger, J. (2006). Definition of clickbait. *Jay Geiger's Blog*. <http://www.jaygeiger.com/index.php/2006/12/01/definition-of-clickbait/>. Accessed 20 Mar 2021
- Gilly, M. C., Graham, J. L., Wolfinger, M. F., & Yale, L. J. (1998). A dyadic study of interpersonal information search. *Journal of the Academy of Marketing Science*, 26(2), 83–100.
- Golman, R., & Loewenstein, G. (2018). Information gaps: A theory of preferences regarding the presence and absence of information. *Decision*, 5(3), 143.
- Grigorev, A. (2017). Identifying clickbait posts on social media with an ensemble of linear models. *ArXiv Preprint*.
- Hofmann, J., Clement, M., Völkner, F., & Hennig-Thurau, T. (2017). Empirical generalizations on the impact of stars on the economic success of movies. *International Journal of Research in Marketing*, 34(2), 442–461.
- Hornik, K., & Grün, B. (2011). topicmodels: An R package for fitting topic models. *Journal of Statistical Software*, 40(13), 1–30.
- Imbens, G. W. (2003). Sensitivity to exogeneity assumptions in program evaluation. *American Economic Review*, 93(2), 126–132.
- Jalali, N. Y., & Papatla, P. (2019). Composing tweets to increase retweets. *International Journal of Research in Marketing*, 36(4), 647–668.
- Jockers, M. (2017). Package syuzhet. <https://Cran.r-Project.Org/Web/Packages/Syuzhet>. Accessed 20 Mar 2021
- Kalro, A. D., Sivakumaran, B., & Marathe, R. R. (2017). The ad format-strategy effect on comparative advertising effectiveness. *European Journal of Marketing*, 51(1), 99–122.
- Kamins, M. A., & Assael, H. (1987). Two-sided versus one-sided appeals: A cognitive perspective on argumentation, source derogation, and the effect of disconfirming trial on belief change. *Journal of Marketing Research*, 24(1), 29–39.
- Kannan, P. K., & Li, H. “Alice.” (2017). Digital marketing: A framework, review and research agenda. *International Journal of Research in Marketing*, 34(1), 22–45.
- Kramer, A. D., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24), 8788–8790.
- Kumar, A., Bezawada, R., Rishika, R., Janakiraman, R., & Kannan, P. (2016). From social to sale: The effects of firm-generated content in social media on customer behavior. *Journal of Marketing*, 80(1), 7–25.
- Kumar, V., Khattar, D., Gairola, S., Kumar Lal, Y., & Varma, V. (2018). Identifying clickbait: A multi-strategy approach using neural networks. *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 1225–1228.
- Lechner, M. (2002). Some practical issues in the evaluation of heterogeneous labour market programmes by matching methods. *Journal of the Royal Statistical Society: Series A (statistics in Society)*, 165(1), 59–82.
- Liang, J., & Yang, M. (2015). On spreading and controlling of online rumors in we-media era. *Asian Culture and History*, 7(2), 42.
- Litman, J. A. (2008). Interest and deprivation factors of epistemic curiosity. *Personality and Individual Differences*, 44(7), 1585–1595.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1), 75.
- Luca, M., Munger, K., Nagler, J., & Tucker, J. A. (2021). You Won't Believe Our Results! But They Might: Heterogeneity in Beliefs About the Accuracy of Online Media. *Journal of Experimental Political Science*, 1–11.
- Lunardo, R., & Mbengue, A. (2013). When atmospherics lead to inferences of manipulative intent: Its effects on trust and attitude. *Journal of Business Research*, 66(7), 823–830.
- Lunardo, R., Roux, D., & Chaney, D. (2016). The evoking power of servicescapes: Consumers' inferences of manipulative intent following service environment-driven evocations. *Journal of Business Research*, 69(12), 6097–6105.
- Madhavan, A. (2017). I reverse-engineered BuzzFeed's most viral posts and the truth is shocking! *Hacker Noon*. <https://www.webics.com.au/blog/content-marketing/upworthy-buzzfeed-viral-marketing/>. Accessed 20 Mar 2021
- Martin, W. C., & Lueg, J. E. (2013). Modeling word-of-mouth usage. *Journal of Business Research*, 66(7), 801–808.
- Matias, J. N., & Munger, K. (2019). *The Upworthy Research Archive: A Time Series of 32,488 Experiments in US Advocacy*.
- McCornack, S., & Ortiz, J. (2021). *Choices & Connections: An Introduction to Communication*. Bedford/St. Martin's.
- McLaughlin, G. H. (1969). SMOG grading—a new readability formula. *Journal of Reading*, 12(8), 639–646.
- McNeal, M. (2015). One writer explored the marketing science behind clickbait. You'll never believe what she found out. *Marketing Insights*, 27(4), 24–31.
- Meirick, P. (2002). Cognitive responses to negative and comparative political advertising. *Journal of Advertising*, 31(1), 49–62.
- Metz, J. (2019). FOMO and Regret for Non-Doings. *Social Theory and Practice*, 451–470.
- Mohammad, S. M., & Turney, P. D. (2010). Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, 26–34.
- Moon, K.-W. (2021). *ProcessR: Implementation of the “PROCESS” Macro*. <https://CRAN.R-project.org/package=processR>. Accessed 20 Mar 2021
- Munger, K. (2020). All the news that's fit to click: The economics of clickbait media. *Political Communication*, 37(3), 376–397.
- Munger, K., Luca, M., Nagler, J., & Tucker, J. (2020). The (null) effects of clickbait headlines on polarization, trust, and learning. *Public Opinion Quarterly*, 84(1), 49–73.
- Papadopoulou, O., Zampoglou, M., Papadopoulou, S., & Kompatsiaris, I. (2017). A two-level classification approach for detecting clickbait posts using text-based features. *ArXiv Preprint*.
- Potthast, M., Gollub, T., Hagen, M., & Stein, B. (2018). The clickbait challenge 2017: Towards a regression model for clickbait strength. *ArXiv Preprint*.
- Potthast, M., Gollub, T., Komlossy, K., Schuster, S., Wiegmann, M., Fernandez, E. P. G., Hagen, M., & Stein, B. (2018). Crowdsourcing a large corpus of clickbait on twitter. *Proceedings of the 27th International Conference on Computational Linguistics*, 1498–1507.
- Reser, J. P. (1972). *Perception and Awareness of Manipulative Intent*.
- Rishika, R., Kumar, A., Janakiraman, R., & Bezawada, R. (2013). The effect of customers' social media participation on customer visit frequency and profitability: An empirical investigation. *Information Systems Research*, 24(1), 108–127.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55.

- Rowan, D. (2014). How BuzzFeed mastered social sharing to become a media giant for a new era. *Wired*. <https://www.wired.co.uk/article/buzzfeed>. Accessed 20 Mar 2021
- Rushkoff, D. (2011). Does Facebook really care about you? *CNN*. <https://edition.cnn.com/2011/09/22/opinion/rushkoff-facebook-changes/>. Accessed 20 Mar 2021
- Sanders, S. (2017). Upworthy Was One Of The Hottest Sites Ever. You Won't Believe What Happened Next. *NPR*.
- Schulze, C., Schöler, L., & Skiera, B. (2014). Not all fun and games: Viral marketing for utilitarian products. *Journal of Marketing*, 78(1), 1–19.
- Sen, A., & Yildirim, P. (2015). Clicks bias in editorial decisions: How does popularity shape online news coverage? *Available at SSRN 2619440*.
- Sianesi, B. (2004). An evaluation of the Swedish system of active labor market programs in the 1990s. *Review of Economics and Statistics*, 86(1), 133–155.
- Smith, B. (2014). Why BuzzFeed doesn't do clickbait. *BuzzFeed*. https://www.buzzfeed.com/bensmith/why-buzzfeed-doesnt-do-clickbait?utm_term=.lv1o9x7Ge#.byNoYW3DN. Accessed 20 Mar 2021
- Stringer, P. (2020). Viral media: Audience engagement and editorial autonomy at buzzfeed and vice. *Westminster Papers in Communication and Culture*, 15(1).
- Tandoc, E. C., Jr. (2018). Five ways BuzzFeed is preserving (or transforming) the journalistic field. *Journalism*, 19(2), 200–216.
- Teixeira, T. (2012). The new science of viral ads. *Harvard Business Review*, 90(3), 25–27.
- Teixeira, T. (2014). The rising cost of consumer attention: Why you should care, and what you can do about it. *HBS Working Paper*.
- Tellis, G. J., MacInnis, D. J., Tirunillai, S., & Zhang, Y. (2019). What Drives Virality (Sharing) of Online Digital Content? The Critical Role of Information, Emotion, and Brand Prominence. *Journal of Marketing*, 1–20.
- Thomas, P. (2017). Clickbait identification using neural networks. *ArXiv Preprint*.
- Thomas, V. L., Fowler, K., & Grimm, P. (2013). Conceptualization and exploration of attitude toward advertising disclosures and its impact on perceptions of manipulative intent. *Journal of Consumer Affairs*, 47(3), 564–587.
- Tirunillai, S., & Tellis, G. J. (2014). Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent dirichlet allocation. *Journal of Marketing Research*, 51(4), 463–479.
- Tkaczyk, J., et al. (2016). The Importance of Similarity and Expertise of the Information Source in the Word-Of-Mouth Communication Process. *International Conference on Marketing and Business Development Journal*, 2(1), 61–71.
- Trotter, J. (2015). Internal Documents Show BuzzFeed's Skyrocketing Investment in Editorial. *Gawker*. <http://tktk.gawker.com/internal-documents-show-buzzfeed-s-skyrocketing-investm-1709816353>. Accessed 20 Mar 2021
- Van den Bulte, C., & Lilien, G. L. (2001). Medical innovation revisited: Social contagion versus marketing effort. *American Journal of Sociology*, 106(5), 1409–1435.
- Wang, S. (2017). Adaptation, A/B testing and analytics: How BuzzFeed optimizes the news for its audience. *International Journalists' Network*. <https://ijnet.org/en/story/adaptation-ab-testing-and-analytics-how-buzzfeed-optimizes-news-its-audience>. Accessed 20 Mar 2021
- Wangenheim, F. V., & Bayón, T. (2007). The chain from customer satisfaction via word-of-mouth referrals to new customer acquisition. *Journal of the Academy of Marketing Science*, 35(2), 233–249.
- Warren, N., Hanson, S., & Yuan, H. (2020). Feeling Manipulated: How Tip Request Sequence Impacts Customers and Service Providers? *Journal of Service Research*, 24(1), 66–83.
- Webics. (2014). How Upworthy and BuzzFeed are Masters of Viral Marketing. *Webics*. <https://www.webics.com.au/blog/content-marketing/upworthy-buzzfeed-viral-marketing/>. Accessed 20 Mar 2021
- Wiegmann, M., Völske, M., Stein, B., Hagen, M., & Potthast, M. (2018). Heuristic Feature Selection for Clickbait Detection. *ArXiv Preprint*.
- Wright, P. (1975). Factors affecting cognitive resistance to advertising. *Journal of Consumer Research*, 2(1), 1–9.
- Wright, P. L. (1973). The cognitive processes mediating acceptance of advertising. *Journal of Marketing Research*, 10(1), 53–62.
- Yoganarasimhan, H. (2012). Impact of social network structure on content propagation: A study using YouTube data. *Quantitative Marketing and Economics*, 10(1), 111–150.
- Zhang, Y., Moe, W. W., & Schweidel, D. A. (2017). Modeling the role of message content and influencers in social media rebroadcasting. *International Journal of Research in Marketing*, 34(1), 100–119.
- Zubcsek, P. P., & Sarvary, M. (2011). Advertising to a social network. *Quantitative Marketing and Economics*, 9(1), 71–107.
- Zuwerink Jacks, J., & Cameron, K. A. (2003). Strategies for resisting persuasion. *Basic and Applied Social Psychology*, 25(2), 145–161.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.