



All the News That's Fit to Click: The Economics of Clickbait Media

Kevin Munger

To cite this article: Kevin Munger (2020) All the News That's Fit to Click: The Economics of Clickbait Media, Political Communication, 37:3, 376-397, DOI: [10.1080/10584609.2019.1687626](https://doi.org/10.1080/10584609.2019.1687626)

To link to this article: <https://doi.org/10.1080/10584609.2019.1687626>



Published online: 03 Dec 2019.



Submit your article to this journal [↗](#)



Article views: 920



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)



All the News That's Fit to Click: The Economics of Clickbait Media

KEVIN MUNGER

The news media industry has changed as the internet and social media have matured and become integral to modern life. I describe these changes through a theoretical analysis of the economic structure of the industry and explore the implications for scholars of online media and politics. The crux of my argument is that social media simultaneously serves as a distribution platform and reputation builder, as social recommendations take the place of expensive investments in high-quality journalism. This development rendered crucial portions of previous models of the market for news inaccurate due to the declining importance of firm reputation. This mechanism interacts with the massive heterogeneity in digital literacy and growing animosity toward the news media among conservatives to create “credibility cascades,” which I argue are a necessary condition for Fake News to flourish.

Keywords online media, media economics, Clickbait Media

Introduction

Trust in the news media has been declining steadily ever since the 1970s (Ladd, 2011), especially among conservatives. This latter trend has driven the development of a parallel conservative news media defined by cults of personality and outrage (Berry & Sobieraj, 2013).

The discussion of “Fake News” as a significant factor in the 2016 election is a striking manifestation of the lack of confidence in the media. Although current estimates place the number of people who recalled seeing the average “Fake News” story at 1.2% (Allcott & Gentzkow, 2016), that the concept resonated so well across both sides of the political spectrum indicates that people fear that the media can no longer perform its role as “gatekeeper,” deciding which stories to disseminate to the public. Estimates of exposure to “Fake News” are considerably higher than of recall, and hugely moderated by age: Guess, Nyhan, and Reifler (2017) find that “only 7% of people age 59 or younger consumed one or more pro-Trump Fake News articles compared with 19% among those age 60 or older.”

The existence of “Fake News” is directly implied by the technological, regulatory and cultural structure of the online media industry; indeed, “Fake News” is less of an innovation than the culmination of existing trends in online media. I explore how the structure of the online media industry has evolved over the lifespan of the internet: entry

Kevin Munger is an Assistant Professor of Political Science and Social Data Analytics at Penn State University. He studies the internet and politics.

Address correspondence to Kevin Munger, Pond Lab, Penn State University, State College, PA 16801, USA. E-mail: kevinmunger@gmail.com

into the market is cheaper; the costs of the inputs of online news production (information, skilled labor, distribution) are lower; consumers of political news are increasingly partisan and have widely varying levels of technical sophistication. These factors combine to produce what I call “credibility cascades”, a novel mechanism for attracting an audience in the contemporary online media environment: stories acquire credibility as they are shared along social networks, becoming more desirable at the same time as they increase their potential audience.

Credibility cascades enable news content from firms that have not developed a reputation for expertise or trustworthiness to spread on social media. Firm reputation is still a strong signal of information quality, but the social information encoded in metrics such as the number of “likes” or “shares” on a post acts as an alternative source of credibility.

Underlying this macro-level phenomenon is the essential fact that each news story stands alone, competing for attention with all of the other news stories on the internet. In contrast to subscription-based newspapers (or magazines, or even websites), or to temporally linear formats like radio or television, media outlets creating content to be shared on Facebook are unable to bundle their news stories. The resulting pressure to attract attention affects both the format and subject matter of stories. During the heyday of nonpartisan media in the 20th century, journalists explicitly embraced their roles as curators and verifiers of information, roles which have today been subsumed by the structural logic of credibility cascades.

I call this media environment “Clickbait Media,” although in its current iteration it is primarily Facebook’s “Clickbait Media.” The importance of Facebook in particular to the contemporary media environment cannot be overstated; specific content and design-related decisions the company undertakes have the capacity to annihilate media companies around the globe. In October 2017, for example, Facebook rolled out an experimental redesign to its central Newsfeed feature that separated posts from individual users and posts by “pages” (like those run by media companies) into separate feeds. Local media companies in the targeted countries—including Slovakia and Cambodia—reported drops of up to 75% in their audience size, making their business models untenable (Cellan-Jones, 2017). The modern media environment has been subject to Facebook’s evolutionary pressures for years now, as the case of Upworthy discussed below illustrates.

The history of the media industry in the US suggests that its current state is not an aberration; the cultural and technological conditions of the postwar era actually make *that* period the anomaly. I describe the economic incentive structures that characterize Clickbait Media, as well as the specific cultural and technological conditions that make its current iteration distinct from earlier iterations.

The economic theory of “contestable markets” is particularly helpful here (Baumol, Panzar, & Willig, 1983). Under certain economic conditions (that I argue obtain in the contemporary market for online news), firms can rapidly enter and exit a marketplace, competing away any profit from existing firms. In equilibrium, the threat of this entry forces incumbent firms to adjust their business models before the widespread competition arrives. This can help explain the puzzle of why the largest newspaper companies have persisted in the current context, and why no new online news site has matched their position: the best model for media entrepreneurs is to start cheap online websites that can be profitable for a short time, then exit the marketplace when costs get too high.

In this paper, I attempt to combine institutional theory from economics that explains the structure of the media industry with behavioral theories from sociology,

communication and political science that model individual media preferences and effects. The former tends to take consumer preferences as given (or at least possible to formally model); the latter tends to take media supply as given (or at least attributable to elite preferences).¹ The decreased costs of news production described in sections 4 and 6 below indicate that such models may soon become necessary. I argue that the synthesis I perform here in defining “Clickbait Media” is necessary to understand and predict the dynamics of online news media, zero-credibility “Fake News” and otherwise.

Note that this article presents no novel empirical evidence. My task is instead to tie together disparate disciplines in order to open up new research questions for scholars of political communication to empirically test.

A Brief History of Media in the United States

Clickbait Media represents more of a return to business-as-usual than an unprecedented development. The baseline for media normalcy in our cultural consciousness is deceptive. The mid-1960s were the golden age of the institutional news media: a combination of the high barriers to entry but extensive reach of broadcast television, strict regulations about content, and low partisan polarization allowed the news media to become fully institutionalized and commit to norms of professionalism and objectivity (Ladd, 2011). The rise of outrage media and the internet mirror the rise of cheap newsprint and yellow journalism a century ago.²

Prior to Hearst and Pulitzer, newspapers were primarily financed directly by political parties (or even the US government). Objectivity was not generally seen as essential or even necessarily desirable. The proliferation of high-volume printing presses *and the decline in the cost of paper* in the late 19th century changed the business model of publishers in large cities, as subscription models waned in favor of on-the-spot sales. As a result, headlines became designed to grab as much attention as possible. There were even some papers that featured quizzes on the front page that purported to tell about the reader’s personality and predict their future. There truly is nothing new under the sun—not even BuzzFeed.

This structural shift made newspapers more entertainment-focused, and part of that entertainment meant taking political positions. The big change, though, was *away* from being officially associated with major parties. So the papers took up populist issues that appealed to their readers’ sensibilities, sometimes at the expense of sound reasoning or sober policy advice—most famously, the incitement to war against the Spanish in 1898.

This media-economics state was one in a which a technology (newsprint) developed to the point where the marginal cost of production was driven to zero, and in which opting into reading any particular article (as hawked by a newsboy, at least in cities) could be done at a whim. I argue that this represents the mature stage of a given medium, in which competition and technological innovation reduce costs and give consumers maximal choice.

Radio and television followed the same pattern, from a near-monopoly due to high startup costs to the eventual democratization of production. These technologies were different from the Yellow Journalism equilibrium of the turn of the century, though, because of their fundamentally linear nature: unlike newspapers, consumption could not be begun or ended on demand, limiting the actual choice available to the consumer. Interestingly, we may be seeing the final maturation of these forms only because of the internet, which allows radio and television to be consumed on demand. As Mutz (2015)

points out, the internet as a distribution platform reduces the mechanical distinctions between these mediums.

The period when broadcast television was the primary outlet for news represents *an anomaly* in the economics of media; the proliferation of cable news channels, talk radio stations, blogs and social media seem to be the norm given mature technologies and a moderate regulatory burden.

However, the current cultural and technological environment produces a novel media industry-political news dynamic: contemporary Clickbait Media. The term “clickbait” is clearly native to the computer; Merriam-Webster, which added the word to its dictionary in 2015, defines it as “something (such as a headline) designed to make readers want to click on a hyperlink especially when the link leads to content of dubious value or interest,” and claims that it was first used in 2010 (Editors, 2015).

Clickbait is the essential manifestation of Clickbait Media, but the specific form it takes has varied over time, even in the brief period that the term has existed. The example of Upworthy is illustrative. The “fastest-growing media site of all time,” Upworthy implemented a new style of clickbait headline designed to entice consumption by strategically withholding information (Karpf, 2016). Less than two years after its founding in March 2012 by Eli Pariser and Peter Koechley, Upworthy had over 80 million unique visitors each month—more than either the New York Times or Washington Post. In November 2013, however, Facebook announced that it would penalize deceptive headlines in their ranking algorithm, and within a year, Upworthy’s business collapsed. In November 2014, the site had only 20 million unique visitors (Karpf, 2016). These changes had massive implications for how people consumed political news in the early 2010s, as other sites quickly caught on to Upworthy’s success.

The flood of news sites inspired by Upworthy was enabled and entailed by the novel economic incentive structure of contemporary Clickbait Media. The rapid rise and fall of Upworthy points to the necessity of understanding that incentive structure to appreciate which elements of Clickbait Media are central and which are ephemeral. The following section aims to adapt the approach in Hamilton (2004) to account for the way that new media technologies have altered the market for news.

The motivating puzzle is this: the (remarkably prescient) Chapter 7 in Hamilton (2004) holds that “News on the Web is still an experience good; to know the good is to consume the good. This means brand names and reputations will be important as signals about the potential content of sites.” But the 2016 US Presidential Election served as an existence proof that this is no longer true: some portion of the audience was willing to consume news from firms with no brand or reputation.

In 2004, the social web basically did not exist, but in 2016 it was dominant. As Settle (2018) demonstrates, a news story shared on social media always contains both source information and social information about who has already read the story and what they think about it. These bundles are inseparable, and this fundamentally changes the market for news. In this context, the “information cascades” that are central to Hamilton (2004)’s theory are in fact even more powerful, as quantified social information supplants source reputation in vouching for the quality of an information good. I call these “credibility cascades.”

To model the market enabled by “credibility cascades” and a host of other important technological, cultural and regulatory changes, I leverage an old and well developed

theoretical framework from economics that applies remarkably well to contemporary Clickbait Media: contestable markets (Baumol et al., 1983).

The Economics of Contemporary Clickbait Media: Contestable Markets

Motivated by the example of the newly-deregulated airline industry, William Baumol and his coauthors described several economic conditions that would need to obtain to create a peculiar industry structure.

The primary feature of contestable markets relative to the standard model of firm competition is that the latter assumes that competition has to come from *actually existing* firms. In a contestable market, the threat of *potential entry* into the market is sufficient to produce market discipline in the form of low prices and/or quality goods.

In the standard model, an entrepreneur might notice that the airline industry was charging too much for flights and set up a competing airline. This process takes a long time, and might not return a profit for years. The costs associated with developing a brand with a good reputation and buying physical infrastructure would have to be spent up front, in the hopes of one day being recouped. And if incumbent firms have access to proprietary technology, the costs of developing an alternative version also have to be paid over years. These barriers make it so that the only competitive pressure to keep prices down comes from other firms already in the market.

In a contestable market, on the other hand, the entrepreneur must be able to rent a few airplanes and set up a small business very quickly, and with no up front capital costs. Products have to be undifferentiated, so consumers are indifferent to the brand of a given flight, preferring simply the cheapest option. And both incumbent firms and challengers have access to the same technology. In this case, the entrepreneur can quickly win market share by charging a lower price. If the incumbent firm responds by lowering their price, the entrepreneur can cheaply and easily exit the market, taking their temporary profits with them.

This is what Baumol described as “hit-and-run” competition: new firms entering the industry and charging sub-market rates in order to attract business from established firms, then exiting the industry once prices drop. In the long run, then, there will be very few firms in the industry, each of them earning close to zero profit. The threat of entry prevents any firm from charging a price higher than their marginal cost in order to earn a profit.

For normal goods, contestable markets have the desirable property of ensuring that consumers get low prices without requiring many firms to operate in the same industry. In fact, even a monopolist in a contestable market is forced to keep their prices as low as possible to stave off contestation from new entrants.

News, however, is an information good, so the normative implications are less clear. Some of the desirable properties of independent investigative journalism are in an economic sense inefficient—firms that invest in this expensive practice are unable to capture all of the benefits, and other firms can “free ride” off of the stories they produce. Marquee investigative journalism does not drive profits in the short run, but it does increase the reputation of a media firm. This “reputation” is anti-competitive, making it more difficult for new firms to enter the market.

As further illustration, consider the broadcast era of journalism discussed above. Technological and regulatory conditions essentially enforced an oligopoly run by 3 firms. Yet this time period was the high water mark for American trust in the media; the

“supranormal” (to use the term from economics) profits enabled by restricted competition were invested into the professionalism of the institution.

Baumol identified several aspects necessary and jointly sufficient to characterize a contestable market. In the real world in which entropy exists, none of these idealized conditions can actually occur, so it makes more sense to refer to the degrees of contestability in a market rather than contestability as a binary condition.

- **Entry/Exit:** There cannot be formal barriers to entry or exit from the market. Firms need to be able to rapidly set up shop without any kind of explicit licensing requirements, and they need to be able to easily exit the market—importantly, they cannot have long-term commitments like employee pensions.
- **Symmetric Information/Technology:** There cannot be any specialized technology or knowledge available to the incumbent firms but not the new entrants.
- **Sunk Costs:** There cannot be any capital investments (in either physical or intellectual capital) that cannot be recouped.

Baumol may have been wrong about the airline industry being a contestable market. Instead of deregulation leading to the rapid entry of many low-cost firms, the largest airlines have consolidated, and the industry has tended toward oligopoly (Martin, 2000). However, the three above conditions *do* describe contemporary *Clickbait Media*. We appear to be in the short term “hit-and-run” competition described by Baumol: the barriers to entry for new firms have been decreasing, the sunk costs associated with a media firm have been disappearing, and the primary technology of news creation—access to information—has been radically democratized. The current environment sees more (national) media firms directly competing with each other than ever before in history.

The crucial question is what will happen in the long run. The current situation does not appear sustainable, and indeed contestable markets theory predicts that it is not. The new entrants are playing a different game than legacy media outlets, producing online content at a fraction of the price. It may be the case that the steady state of this market is also what contestable markets theory predicts about the long-run equilibrium: very few firms making very low profits, constrained by the threat of other costless entrants. However, real-world conditions necessarily deviate from idealized economic theory. My theory of contemporary *Clickbait Media* describes the current state of online news media to explain how it deviates from an ideally contestable market and offers several alternative possible equilibria for the industry.

Entry/Exit

Until recently, it was very costly to create a national news organization. The logistical challenges of distributing a physical newspaper meant that from WWII to today, only one national newspaper (the *USA Today*, in 1982) entered the marketplace (Hindman, 2008). The broadcast television market was limited to three firms by a combination of regulation and massive capital requirements. The cable television market was less regulated and somewhat cheaper to enter; several new cable television firms were founded, but costs and regulations limited this number as well (Prior, 2007).

The spread of the internet (and the necessary hardware and software) allowed anyone with some degree of technical know-how to set up their own blog and create news content on their own. These blogs allowed many people to share their views outside the supervision of traditional media companies, but initial enthusiasm about the democratization of

information production was misplaced. Due to the link structure of the early web, attention was distributed according to what Hindman (2008) calls the “Googlearchy,” and all but the very most successful blogs received next to zero attention.

The online media market only became mature as internet use became common among a much broader swath of the population, a trend which co-developed with the ubiquity of social media use. In 2016, 68% of adults in the US used Facebook, and 74% of adult Facebook users use Facebook *every day* (Smith & Anderson, 2018), providing both an audience and a distribution platform for online content. Software for managing online content has become nearly free, so the cost of setting up a website with the potential to reach the millions of daily Facebook users has fallen to nearly zero.

Equally important to the cultural climate is the regulatory environment. During the postwar era, the power of the Federal Communications Commission has been continually eroded. The 1987 repeal of the Fairness Doctrine, the overturn of the obscenity provisions of the 1996 Telecommunications Act and the 2000 repeal of the “personal attack” and “political editorial” rules were all decisions that lifted restrictions on what the media could do (Berry & Sobieraj, 2013). The First Amendment provides robust protections to freedom of the press, and attempts to regulate internet content have been met with extremely negative reactions (Coleman, 2014). The generational makeup of Congress makes it unlikely that key members understand the technical challenges of regulating the internet.

Another form of regulation impacts firm entry: relative regulatory licensing restrictions. In the real world, there are a finite number of entrepreneurs who plan to start a company. One important consideration is the explicit regulatory hurdles in each industry; industries with fewer licensing requirements are, all else equal, more attractive.

In order to charge money for barbering services in California, you must pass a written and practical examination. Before you may attempt to do so, you must have logged 1,500 hours of (unpaid) training barbering; barbering without a license is subject to a \$1,000 fine (Department of Consumer Affairs and Cosmetology, 2016).

To set up a national media company based in California, you need to do exactly nothing; there are no licensing requirements and no potential fines. In 2013, Jestin Coler did just that: he established twenty or more media sites with next to zero editorial discretion, including the now-infamous Denver Guardian, which published stories with headlines like “FBI AGENT SUSPECTED IN HILLARY EMAIL LEAKS FOUND DEAD IN APPARENT MURDER-SUICIDE”. Coler—who considers himself an entrepreneur—reportedly makes hundreds of thousands of dollars a year. His employees are all freelancers who work completely anonymously; none of them face any regulatory burden whatsoever (Sydell, 2016).

Although this is an extreme case, online media companies are almost always structured so that they get the majority of their content from freelancers, or at least from nonunion workers.³ This institutional setup means that web native companies have minimal long-term commitments, giving them close to zero exit costs.

Technology/Skills

The necessary technology to create news content in contemporary Clickbait Media is unpatentable and (other than lengthy, in-person investigative reporting) within the reach of any incipient media company. The only inputs are internet access, computer hardware and digitally literate employees. The former two have decreased dramatically in price and require no special level of technical know-how to acquire and use. There is very little to

differentiate one news story from another when they are both shared on a Facebook Newsfeed; a link to the (fake) *Denver Guardian* contains a headline, photo and caption, as does a link to the *New York Times*.⁴

The biggest technological change is in the kind of training necessary to be a journalist. The ideal of the Journalism Schools founded by Joseph Pulitzer and William Randolph Hearst was to ensure that all journalists had rigorous training in journalistic ethics and practice—journalists were to be respected professionals, on a par with doctors and lawyers. There was never a bar exam or explicit certification process, but there was the expectation that everyone in a newsroom would—either through J-school or on-the-job mentorship—share a broad skill set and ethical standards.

The technological constraints of broadcast television and newsprint meant that the supply of J-school graduates and job openings for journalists were in a rough equilibrium. Employees with professional degrees are expensive, though, and web native companies have not found it necessary to require their entry-level news reporters to have such formal training. Indeed, each of these freelancers is competing directly with other freelancers, often working for multiple news outlets at the same time; they have little organizational loyalty, instead aiming to raise their own profile. This is precisely the mechanism for media bias proposed by Baron (2006):

[J]ournalists may bias their stories if their career prospects can be advanced by being published on the front page. News organizations can control bias by restricting the discretion allowed to journalists, but granting discretion and tolerating bias can increase profits if it allows journalists to be hired at a lower wage.

This leveling of journalistic skills is enabled by the medium of electronic distribution. Editors of print newspapers or cable news segments are space constrained: the inclusion of one story necessarily implies the exclusion of another. This is not the case on the web, which means that each story need not be held to the same standard of quality (in addition to the decreased importance of brands discussed above).⁵

The opportunity cost of each online story is zero, the marginal cost of distribution is zero, and the marginal cost of creation corresponds to the value of the journalists' time. Clickbait Media has taken advantage of the fact that there is a surplus of potential employees with the necessary skills: gathering information on the web, rapidly writing summaries of their findings, and possessing the social media/cultural savvy to promote their news content.⁶

These substantially decreased costs should cause new firms to enter the market for online news, and employment data from the US Bureau of Labor Statistics indicates that this is in fact the case. In contrast to constant or decreasing numbers of firms in newspaper, book, and magazine publishing and television and radio broadcasting, the number of internet firms has skyrocketed. In 2007, there were approximately 6,000 such firms; by 2018, that number had increased to 15,000.

Sunk Costs

Essential to operating a news media company is that consumers believe what you report. Developing a reputation for credible reporting is a pre-requisite for operating a serious national news outlet. This was not always the case; the establishment of professional journalistic practices was an explicit strategy of papers like the *New York Times* around

the turn of the century in order to differentiate themselves from less reputable competitors (Ladd, 2011). With expanding competition from cable news and talk radio, however, it became clear that these high standards were no longer essential.

There are still a large number of discerning news consumers who only consume content from reputable outlets; this product differentiation means that the industry is not a perfectly contestable market. But the august reputations of legacy news outlets are simply not important to many consumers; as I argue below, this reputation has actually become a liability for many (especially conservative) consumers who distrust established media outlets.

The theory motivating high-quality journalism was that the expensive and time-consuming process of developing a reputation was a sunk cost that traditional outlets hoped would scare off new entrants—when a newspaper goes out of business, there is no way to recoup these reputation-building costs—but the modern news industry does not require these sunk costs. Gentzkow and Shapiro (2008) argue that a diverse media environment leads to greater investment in high-quality investigative reporting because of market discipline, but they also claim that “this mechanism will only operate if firms value a reputation for reporting the truth.”

The mature online news industry represents the culmination of this de-valuing of reputation. Even with cable news or radio, reputation matters because consumers have to decide to change the channel to a specific media outlet. The same process occurred on the early web, where people had to decide which websites to visit. On social media, each individual piece of content competes with every other piece of content as individual users decide which pieces of content to share.

Reputation still could matter, though, to convince a potential reader that the reporting in a story is credible. The fact of social recommendation provides an alternative source of credibility. Scholars of political communication frequently cite experimental evidence which suggests that consumers prefer to share content with more anonymous recommendations on Facebook (Messing & Westwood, 2012). Social recommendations from friends or family members provide even more legitimacy. Furthermore, people are less likely to fact-check information they encounter on social media (Jun, Meng, & Johar, 2017).

The literature on alternative sources of credibility is more central to the field of communication. “Bandwagon effects” are a well-established phenomenon in which consumers use the decisions made by previous consumers as a heuristic, choosing more popular options (Schmitt-Beck, 2015). This has important implications for website design; online stores that show the popularity of each of their items make consumers more likely to select the most popular items (Sundar, Oeldorf-Hirsch, & Xu, 2008). These cues are not deterministic, but in some cases (like zero-credibility websites), they are all the consumer has to work with. In the context of selecting a movie to watch, for example, bandwagon effects have been found to be largest when the quality of the movie is more uncertain (Xu & Fu, 2014). This phenomenon has been shown to be particularly effective in the choice of YouTube videos, an information good just like news stories from unknown websites (Fu & Sim, 2011).

Communication theory tends to explain the bandwagon effect through the lens of psychological processing: a surfeit of choices drives consumers to use cognitive shortcuts. But this behavior also concords with theories more amenable to political scientists. Consider the rational case for information acquisition. The likelihood that following a news story will cause an individual to change their views sufficiently to change their vote choice and that this will in turn change the outcome of an election is minimal; if this is the only benefit to information acquisition, ignorance is rational (Downs, 1957).

A complementary strand of scholarship, most famously advanced by Campbell, Converse, Miller, and Donald (1960) and seeing renewed interest with Green, Palmquist, and Schickler (2002), conceptualizes partisanship as a social identity (Iyengar, Sood, & Lelkes, 2012; Mason, 2016).⁷ Instead of being rational truth-seekers aiming to make the best possible vote choice, partisans are team players who aim to follow social cues about the correct views and arguments for members of their social group.

This story comports with the explanations given in qualitative interviews reported in Berry and Sobieraj (2013). People who consume what Berry and Sobieraj call “outrage media” do so because they want to feel like a member of a morally righteous social circle, and to feel educated: they want to have talking points ready to go for the next time they have a political discussion. Because these people are likely to avoid political discussions with members of the opposing party, these talking points serve primarily to establish their legitimacy with co-partisans.

As a result, the current technological/cultural environment is ripe for market segregation, enabling some media companies to attract niche audiences while sinking minimal costs into investing in credibility (Stroud, 2011).⁸ The media to which people are exposed on social media is precisely the media that people in their social networks think is important, and because social networks are homophilious, this means that content should be spread among individuals who share a social identity.

This is the explicit strategy of web native media outlets like Gawker, Vice, and especially BuzzFeed. BuzzFeed makes content that people want to share. As a trivial example of how this works, consider the series of quizzes that purport to explain where the reader is from (a form of social identity):

- “Your Cheesecake Factory Order Will Tell Us Which State You’re From”
- “We Can Guess Where You’re From Based On Your Bagel Choices”
- “Can We Guess The State You Live In Based On Your Restaurant Choices?”
- “We Know Where In America You Actually Live”
- “Can We Guess Where In The USA You Actually Live?”

Even these outlets, though, decided to develop a brand identity of some kind, in order to get the initial consumers necessary to begin spreading their content. They also focused on the traditionally coveted (and digitally active and proficient) demographic of young, educated consumers. The 2016 US Presidential election demonstrated that even this minimal level of reputation-building sunk cost may no longer be necessary.

News reports in the wake of the election focused on the problem of “Fake News”—false or wildly misleading online content peddled by unknown media outlets and spread via social media. The most notorious example is the Denver Guardian, discussed above. In addition to employing anonymous freelancers, its business model was explicitly to publish outrageous stories that would appeal to partisan biases (Sydell, 2016).

The crucial innovation of these sites is to spread news that is verifiably false. This strategy cannot work if source reputation matters at all. Additionally, consumers must be unwilling or unable to consult other media sources and learn that they have been deceived.

In controlled settings, studies have shown various degrees of success in correcting misinformation, with partisans being less likely to accept that an ideologically congruent belief is false (Bode & Vraga, 2015; Garrett, Nisbet, & Lynch, 2013; Nyhan & Reifler, 2010; Nyhan, Reifler, Richey, & Freed, 2014). These studies do an excellent job of delineating the necessary and sufficient conditions for the correction of misinformation, but they cannot speak to how often those conditions obtain in different media consumption contexts.

An exception is Jun et al. (2017), who find that people are less like to fact-check in social situations, like “platforms that are inherently social (e.g., Facebook) or ... features of online environments such as “likes” or “shares.” “People tend to “let their guard down” when consuming news obtained in social settings, such as social media.

There is a large literature on the “digital divide” between people with internet access (or social media accounts) and those without (Chadwick, 2006; Mossberger, Tolbert, & McNeal, 2007). This *interest* in the technological aspect of internet use has sometimes been too focused on the digital access binary, rather than the fact that skills are unevenly distributed among internet users (DiMaggio & Hargittai, 2001).

Hargittai (2001) calls this the “second-level digital divide”: the wide disparity in the accuracy and speed at which internet users can perform even a standard task like information retrieval. The OECD performs the most comprehensive research on adults’ skills and finds evidence of a massive disparity in the skill sets of the digital elite (a category to which essentially all producers of online media belong) and the majority of internet users (Kankaraš, Montt, Paccagnella, Quintini, & Thorn, 2016). The OECD’s survey instrument is designed to measure skills related to “problem solving in technology-rich environments” (PSTRE), which they define as “Ability to use digital technology, communication tools and networks to acquire and evaluate information, communicate with others, and perform practical tasks.” This measure is not explicitly designed to capture adults’ capacity to determine the authenticity of a piece of news content on Facebook, but the skills it measures are closely related.

The extent of this “second-level digital divide” bears emphasizing. According to results published in 2016, US adults fall into one of five skill levels, each defined by the most sophisticated computer task people at that level can complete:

- Can’t use computers (26%)
- Can delete an e-mail (14%)
- Can use “reply-all” to send an e-mail to three people (29%)
- Can “find a sustainability-related document that was sent to you by John Smith in October last year” (26%)
- Can calculate “what percentage of the emails sent by John Smith last month were about sustainability” (5%) (Nielsen, 2016)

These examples are all related to e-mail and not directly relevant to evaluating a news story on Facebook; they are presented merely to demonstrate the level of difficulty of the tasks the OECD uses to calculate its PSTRE measure. Most of the tasks have to do with extracting information from specifically curated (and simple) web pages.

These numbers from the OECD are aggregates of all “adults”: people aged between 16 and 65. However, there is a massive heterogeneity across age cohorts: in the United States, 39% of people aged 25–34 scored in the top two categories, but only 20% of people aged 55–65 did. It is overwhelmingly likely that the level of computer skills among those over 65 is even lower than for those in this age range. Combining these data with those concerning social media use leads to a startling conclusion:

68% of adults (this figure includes people over 65) use Facebook, but only 60% of adults (ages 16 to 65) are able to reply-all to an e-mail.

The combination of these low-skill social media users and the power of social recommendation are necessary conditions for the zero-credibility firms that operate in the Clickbait Media context. By buying Facebook ads promoting their articles or pages, they are able to get the initial exposure they need. This process can happen without the majority of news

consumers (the “mainstream audience”) ever being aware because of the capacity for these social media sites to sell unprecedentedly well-targeted ads. Facebook touts this ability: in marketing material for advertisers, it claims that “[w]ith our powerful audience selection tools, you can target the people who are right for your business” (Facebook, 2016).

Zero-credibility firms target precisely these social media users who are least willing or able to intuit or ascertain the accuracy of their content. These are the people who have access to the internet and who use social media and yet have low levels of digital literacy. A strong proxy for digital literacy, at least in the contemporary United States, is age (Hargittai, Piper, & Morris, 2018).

The other decision facing zero-credibly news firms is what kind of content to produce. It makes little sense to cover mainstream news topics. Instead, the ideal strategy is to create content of interest to a hyper-partisan, low digital literacy audience, preferably one with less trust in the reputations of mainstream news outlets. Evidence from both the personal accounts of fake news “entrepreneurs” and academic research indicates that the ideal audience during the 2016 US Presidential Election was older conservatives on Facebook.

Using web tracking data during the 2016 election, Guess et al. (2017) find that “only 7% of people age 59 or younger consumed one or more pro-Trump Fake News articles compared with 19% among those age 60 or older,” and that “58.9% of all visits to fake news websites came from the decile of news consumers with the most conservative information diets.” For the purposes of my theory, sharing articles is more important than viewing them, and here the results are even more stark (for age, and a bit less stark for ideology): “Users over 65 shared nearly 7 times as many articles from fake news domains as the youngest age group” (Guess, Nagler, & Tucker, 2019).

These results are the product of a dynamic market in which the supply of zero-credibility conservative news adjusted to meet the demand. A memorable example of this flexible supply was reported by Silverman and Alexander (2016): teenagers in Veles, Macedonia who switched from producing zero-credibility websites about health-related topics to American politics. With very little information about the context of the US election, they found the audience for zero-credibility conservative news the most receptive: “some in Veles experimented with left-leaning or pro-Bernie Sanders content, but nothing performed as well on Facebook as Trump content.”

If media sites are able to convince these less digitally literate and highly partisan people to consume and then share this inflammatory but fictitious content, their audience expands *and* the fact of the social recommendation lends legitimacy to their content.

The precise relationship between the digital literacy of the consumer, the ideological affinity between consumer and news item, and the consumer’s lack of trust in the reputations of mainstream media firms is unclear, but the three are likely to be mutually reinforcing. Crucially, in the 2016 US election, these three co-occurred in the biggest audience for zero-credibility news: older Trump supporters.

The mechanism connecting ideological consonance and digital literacy can usefully be thought of as *cost*. People who are less pleased with the ideological content of a story may be more likely to expend the effort to verify its veracity, and their level of digital literacy moderates the difficulty of this task. Firm reputation serves as a shortcut—people who sufficiently value the reputation of established media firms will never share zero-credibility news; bandwagon effects are dramatically weaker in the presence of strong prior preferences or beliefs.

For people who do *not* value media firm reputation, including a growing number of conservatives, the shortcut comes more in the form of peer recommendation. News stories

shared on social media are bundles of source cues and social information, including the total number of people who liked or shared the story and who among them are your friends. Both indicate that other people, faced with the same consumption decision you currently face (but with fewer social recommendations), made the decision to consume and share the story. Because social networks tend to be homophilous, it is likely that these readers have at least somewhat similar tastes. And if the recommendation comes from a close friend, it should be even more informative because you know more about their ideological preferences and how they relate to yours.

A hyperpartisan news story from a zero-credibility media outlet can initially only appeal to people who do not value firm reputation and who have decided that the cost of verifying a story is too great (from a combination of high ideological congruence or low digital literacy). Once they have decided to share the story, however, other people facing that same consumption decision have access to the bundled social recommendation information, giving them a shortcut as to the stories' quality.

In this way, people with lower ideological congruence with a story or with greater digital literacy are enticed to consume and possibly share content that they might otherwise not. They become aware that other people with whom they share their partisan (social) identity are consuming this information, and they thus have reason to read it themselves, even though the media organization that produced it has sunk absolutely nothing investing in the credibility of their brand.

As the story continues to be read and shared, the audience grows exponentially but also changes qualitatively as an incrementally more ideologically diverse and more digitally literate news consumers find the story credible. Updating the concept of the "information cascade" at the center of Hamilton (2004)'s theory, I call this process a *credibility cascade*.

Credibility cascade: Social recommendation provides credibility to news stories as they spread along online networks, cascading through layers of increasingly sophisticated digital news consumers.

Once a story accumulates enough social recommendations, it acquires sufficient legitimacy that more traditional news outlets cover it. This process represents a parallel credibility cascade, as increasingly high reputation media outlets cover the story, based on the "reporting" done by less established outlets. These higher-level credibility cascades might helpfully be considered an extension of the "media storms" described by Boydston, Hardy, and Walgrave (2014).

"Media storms" are a specialized media time cycle in which an external event temporarily lowers newsworthiness thresholds and media firms imitate each other's choice of news. Higher-level credibility cascades take these behaviors as intertwined; the newsworthiness is defined by imitation. Even in the case of completely inane "news," firms are rewarded for creating their own story if the "news" has demonstrated the ability to flow through online social networks.

Consider the nonpolitical case of "Laurel or Yanny." A Twitter user tweeted an audio clip that different listeners heard differently, a point of minor interest. The initial tweet spread quickly, indicating consumer interest in the story. Then, hundreds of online media firms created articles discussing the clip. The *New York Times* alone published four separate articles on the phenomenon over the course of three days in May 2018: an initial story, an interview with experts, a report on President Trump's contribution to the debate, and an innovative web tool that allowed listeners to hear either pronunciation (Salam & Victor, 2018).

The half-life of this "news" story, however, was very short. The stickiness described in Boydston (2013) no longer obtains; no news organization is sending reporters to cover

this nonevent, so there are no institutional forces prolonging media attention to the topic. The desire to avoid being the first news organization to switch off of a popular topic described by Boydston et al. (2014) is muted as new topics arise so frequently in a world where newsworthiness is determined by imitating popular topics already in circulation and the topics are not necessarily of substantive import. Clickbait Media is comprised of a never-ending sequence of media storms where the top outlets are simply the best at imitating other firms in a way that encourages credibility cascades.

Clickbait Media only describes a portion of the online news ecosystem; the more the conditions for a contestable market obtain, the better Clickbait Media characterizes a given portion of the media/audience sphere. Legacy media firms that have sunk massive costs to develop their reputation through high-quality journalistic practice are able to offer a *differentiated* product, but only to the audience that appreciates that differentiation. These commitments to journalistic practice are sticky, and inhere in the culture of a given institution. Usher (2014) documents the status games within the *New York Times*: veteran journalists competed to make the front page of the print edition even while the internet was remaking the news industry under their feet. The capacity to produce novel, well-reported news continues to set these institutions apart, and for them to set the agenda among a large portion of the audience. But even the Gray Lady had to report on Laurel or Yanny.

Gaming Information Verification Systems: Past, Present and Future

Creating zero-credibility news stories that can start credibility cascades is a successful strategy because it “games” a central system by which information on social networks is evaluated: social recommendation. An earlier form of information verification on the internet also suffered from a flood of hit-and-run entrants who took advantage of this informational market structure, and it is illustrative to recall how this cycle played out.

A decade ago, the best way to make money in the online information sphere was to appear near the top of Google’s search results and sell ads based on clicks. This ecosystem led to the proliferation of “content farms”: rapidly constructed websites that spammed newly popular search terms to ride the wave of search attention. The rise and fall of a company called Demand Media almost perfectly parallels the story of Upworthy described above.

Demand purchased over a million web domains and attracted traffic to, say, 3dblur-ayplayers.com, because of how Google’s search operated. They also gamed the incoming-links metric that is an important component of the essential PageRank algorithm that powers Google search. The centerpiece of the operation was “eHow.com,” which paid “cheap freelancers who don’t have the expertise” to produce inane articles about every conceivable activity (Wallenstein & Spangler, 2013).

But in 2011, Google changed their search ranking algorithm to penalize “junk” websites (Miller, 2011). Within two weeks, Demand had lost almost half of its valuation, and continued to bleed traffic.⁹

The internet is a marketplace for information, and has consistently been plagued with profit-seeking fraudsters trying to game the systems by which information is deemed to be valuable. As social recommendation has become increasingly central, the technique of creating zero-credibility, hyper-partisan news becomes more powerful. Of course, this is only useful insofar as it can be monetized.

For both the content farms of the early 2010s and the zero-credibility websites of the latter half of the decade, the primary source of revenue has been displaying ads. Immediately after the 2016 election, Google AdSense—currently the dominant ad

service—announced a change to their policy to ban fake news websites from using the service to sell ads (Nicas, 2016). While this surely decreased the returns to the business model, it has proven difficult to enforce comprehensively, and there are other ad services these websites can use.

The makeup of the social web is constantly shifting, as different groups of people adjust their usage patterns between platforms. Each of these platforms, in turn, is constantly tinkering with and sometimes overhauling its affordances. The specific techniques used by zero-credibility firms to create credibility cascades may be different in the future. But as long as the fundamental components of Clickbait Media described in the previous three sections continue to obtain, these firms will always be waiting to enter and contest the market for online news.

With that caveat in place, it is still illustrative to describe how this played out in the most robustly studied context of zero-credibility news to date: the 2016 US Presidential Election.

A major limitation of Facebook from a researcher's perspective is the extent to which most of the data on the platform are difficult to collect at scale. Worse, some crucial news dissemination takes place in "closed groups" that can only be accessed by members (Silverman, Lytvynenko, & Vo, 2018). One technique popular during the 2016 campaign was for owners of zero-credibility firms to infiltrate nonpolitical groups and then spam them with their political content. Jestin Coler, the fake news entrepreneur, reports that this is a common tactic:

Joining Facebook groups ... then spamming links throughout those groups using various aliases, is quite effective ... Members of the group then essentially become "bots" and often share content to their network (outside of the group) creating a more organic-looking campaign.

Shares of the news story within the group are not recorded outside of the group, and so do not contribute to the story's credibility, but this tactic can produce many distinct seeds of credibility cascades, any of which could take off.

Twitter, although far less popular overall, is central to the online news ecosystem, and makes all of their data public. Bovet and Makse (2019) use this to their advantage to track the networks through which mainstream and zero-credibility news spread during the election. Their findings concord with the established theory of agenda setting among the mainstream news, but an inverted, bottom-up network for spreading fake news:

top news spreaders of center and left leaning news outlets are the ones driving Twitter activity while top news spreaders of fake news are in fact following Twitter activity, particularly Trump supporters activity.

The "top spreaders" of fake news content represent the users who provide the biggest boost to the credibility of a given story. They find stories that are cascading up their networks and tweet about them, increasing their own fame in the process. But the divided audience described above limits the absolute reach of these stories. Many people will never share a news story that is not from a reputable media firm.

Some of these top spreaders have successfully parlayed this ability into sustainable fame and acquired valuable reputations of their own—but only among the audience who rejects the reputation of legacy media. The top two Twitter accounts identified by Bovet

and Makse (2019) belong to Paul Joseph Watson and Alex Jones, the Editor and Publisher, respectively, of the fake news website Infowars.

Infowars does not sell ads. Instead, their revenue comes from selling their own branded products. Particularly high margin products include male enhancement pills and dietary supplements like “DNA Force Plus.” While many major social media platforms have banned InfoWars and Alex Jones, and despite a lack of ads, this alternative business model—premised on a smaller but more devoted audience than ad-based media firms aim at—suggests the next evolution of online information fraud.¹⁰

Discussion

The synthetic nature of this article means that I have not demonstrated anything conclusively. Several points in my argument require theoretical leaps across disciplines and the application of empirical results to contexts somewhat different from where they came. I encourage future research on both of these dimensions.

In particular, the theory of credibility cascades in online news has not yet been proven. On the micro-level, I have argued for their existence by combining insights from the theories of digital literacy, bandwagon effects and homophily in online networks. On the macro-level, they help explain the economic model of online media firms that forgo previously essential investments by media firms in reputation and distribution. But both of these roles needs to be more explicitly tested.

On the other hand, the applicability of the theory depends on several key features of the contemporary market for online news; if conditions change, Clickbait Media might not persist. Even today, these features are not constant around the globe; I developed the theory in the US context.

My normative goal in describing Clickbait Media is to help interested parties in media, academia and government understand the conditions that got us here. I conclude this essay by proposing a number of ways to enable the marketplace for online news to transcend Clickbait.

Of course, many of the *fundamental* conditions described in this article are unlikely to change: the technological ability to rapidly set up an online media firm without a physical presence is here to stay, as is the desire for millions of Americans to consume online news. And the underlying cultural and political conditions that produced the affective polarization and decreased trust in media reputation that are necessary for Clickbait Media have deep roots that are causally prior to the contemporary online media environment.

Most of my suggestions, then, involve imposing some kind of friction on the production or sharing of online news, either at the level of the individual piece of content or the level of the online news industry. Economists are generally loath to impose frictions, correctly identifying them as a deadweight loss on economic activity. The logic behind my suggestions comes from the fact that news is not a normal good; it is an *information good*, requiring nonstandard economic models.

It follows that one area in which we should decrease frictions involves industries creating normal goods. The marketplace for online information is just one of many marketplaces, but it has attracted so many “entrepreneurs” because it is completely unregulated while, say, California barbers are legally required to complete 1,500 hours of unpaid training. The Federal government needs to rebalance the regulatory burden across industries.

Regulating online media directly poses a host of issues related to the freedom of the press. Furthermore, the rapidly changing nature of information verification system online means that it will be challenging to keep regulation up to date. *Something* should be done, but in the meantime, states can change the decision calculus of an aspiring entrepreneur by lowering the barriers to entry into non-Clickbait Media options.

Non-governmental actors also have the power to change the conditions that produce Clickbait Media. Facebook (and Google, and a few other big platforms) could embrace a more editorial role and give greater weight in their algorithmic systems to content from “reputable” outlets. Currently, they prefer to optimize for “engagement,” for obvious economic reasons. Downweighting niche/novel sources of online news would stem the proliferation of firms, but might well cost them the membership of the users who would prefer to consume that kind of news. However, Facebook has memorably not shied away from user complaints when overhauling their interface, correctly assuming that users will adapt. From a long-term perspective, the company might also come to realize that excluding some people is good for the overall health of the platform. The company has already taken steps in this direction.

This approach is not trivial. A version of this system already exists in the form of “verification” via a blue check mark, used first by Twitter and now by other platforms to indicate that an account has been deemed important enough to be individually verified as legitimate. Initially designed to prevent impersonation, it is now a binary status symbol. Twitter has essentially ceased to verify any new users (and YouTube recently un-verified thousands). The same anti-media sentiment that enables zero-credibility websites has led some users to employ “blue check” as a slur, particularly in the plural (“the blue checks are rallying around the media establishment”).

If these platforms implement a more assertive system for verifying information, then, they face two risks. Users dissatisfied with a heavy hand might decamp for freer (or at least other-handed) platforms, replicating Clickbait Media. To date, despite consistent predictions to the contrary, the forces of centralization have dominated the forces of fragmentation online (Hindman, 2018).

Another risk comes from a governmental response. These platforms have thus far insisted that they not be treated as media companies despite the growing implausibility of this claim; if they officially certify news content from certain outlets, their position becomes even weaker. And regardless of the legal reality, both ends of the political spectrum are increasingly disenchanted with the major online platforms, and may find it useful to argue that any of their editorial choices are biased—as they inevitably must be. Earlier eras of political media were much more heavily regulated, and again, some form of regulation will prove necessary, but the tech companies who have the capacity to make immediate changes to the online media environment will not make these changes if it increases their exposure to regulation that harms their business.

Another possibility is internet regulation at a much more fundamental level. Near-instantaneous one-to-many communication has brought many benefits, but it has been temporally bundled with the spread of smartphones and other digital technology, so the true value of that speed is difficult to gauge. While still allowing small-group videoconferencing, regulation could impose frictions in the form of either micro-payments or small delays on broadcast social media. This would hamstring the power of credibility cascades while leaving most of the rest of internet communication untouched.

The risk here is concomitantly more fundamental. A government with a precedent of regulating internet communication at this level would likely find other rationales for

regulations, regardless of their value to the public. The imposition of micro-payments would also destroy the democratic nature of the internet—but as Hindman (2008) argues, this “digital democracy” has always been a myth.

The demand side of the market for online news provides another avenue to transcend Clickbait Media, albeit a slow one. Credibility cascades require the existence of a digital literacy gradient among internet users. The US is in a period of disequilibrium, with older users (who are lowest in digital literacy) getting online for the first time. These people will gain internet skills over time, and future generations will not have the issue of first encountering the internet at age 75. A higher average level—and lower variance—of digital literacy among internet users would reduce the power of credibility cascades and thus the economic viability of new online media firms.

The long term case is thus optimistic. The past two decades have been a maelstrom of change for information and communication technology, with the US as the epicenter of technological innovation but lacking the bureaucratic capacity for effective regulatory response. This will not continue; in fact, some observers argue that the rate of change for online communication has already slowed (Karpf, 2019). But this rapid rate of change has also made it difficult for media scholars to accumulate enough empirical knowledge to develop new theories explaining contemporary online media. My approach, synthesizing theories from a variety of disciplines to describe the supply and demand factors that comprise Clickbait Media, is second-best. But I believe that it provides a useful starting place for media scholars, technologists and regulators hoping to accelerate the move away from Clickbait Media.

Notes

1. This latter point relates to an important trend in contemporary media that the theory developed in this paper does *not* directly address: nonprofit maximizing firms. Recent evidence suggests that in both the case of national newspapers (Archer & Clinton, 2018) and local television stations (Martin & McCrain, 2018), being acquired by wealthy conservatives increases national politics coverage and the ideological slant of that coverage past the theoretical ideal point for profit maximization. This specific empirical strategy is less applicable to online news because of the ease of establishing new media firms, but a potentially analogous case is the decision by the family of conservative billionaire Robert Mercer to invest \$10 million in the far-right website Breitbart.com in 2011 (Mider, 2016). Because Breitbart is a privately held company, its finances are not publicly available, but given the Mercer’s political donations to other far-right causes, it is plausible that they would be willing to subsidize some losses to increase the supply of far-right online news. Many economic models of the market for media include firms who forgo some amount of profit for ideological reasons, but I am unaware of any in which there exist some firms who are willing to operate *at a loss* to promote an ideological agenda.
2. Much of the following section is developed in more detail in Ladd’s excellent *Why Americans Hate the Media, and Why it Matters* (2011).
3. Employees at Gawker media, one of the most successful of the first wave of web native media companies, voted to unionize early in 2016. Although this is not necessarily related, the company was mere months away from being sued into bankruptcy.
4. This was true at the time of the 2016 election. Since then, Facebook has tested several display and algorithmic changes to Newsfeed that were designed to ameliorate this problem.
5. In addition to the three major conditions outlined in above, Baumol specified that a contestable market needed to be one in which products were undifferentiated—the classic economics example points out that each bushel of wheat is indistinguishable from each other bushel. In one sense, the online media industry is maximally differentiated, as each piece of content is unique. In practice, though, each piece of content is just another combination of

“words on a screen.” If any piece of content is particularly popular, scores of imitators from other media companies will create versions of it that are only trivially different; often, these outlets will write a more inflammatory headline and introductory paragraph followed by a block quote from the original publication. These pieces of content, then, are close to undifferentiated.

6. Indeed, it as if the US higher education system were designed to produce a surplus of graduates with precisely these skills.
7. For an excellent intellectual history of the group theory of voter behavior, see Chapter 2 of Achen and Bartels (2016).
8. These effects are asymmetric. Republicans are more distrustful of the mainstream media and far more unified by symbolic ideology (Grossmann & Hopkins, 2016). In the course of the 2016 election, Allcott and Gentzkow (2016) estimate that the average American remembered .92 pro-Trump and .23 pro-Clinton stories from a zero-credibility outlet.
9. Focusing on the largest firms like Demand and Upworthy is helpful for exposition, but it also emphasizes that these markets are never *purely* contestable. These companies still exist, although both have been acquired by other firms, because any market in reality has *some* sunk costs and friction on entry/exit.
10. Another model that bears mention are streaming video creators on platforms like YouTube and Twitch. These “streamers” aim to create a devoted “community” of viewers who are willing to watch hours of their videos every day. Successful “streamers” accept donations or “tips” in exchange for reading a viewer’s message on their video.

Acknowledgements

I thank two anonymous reviewers, the editors, Jay Hamilton, Amber Boydston, Drew Millard, Scott de Marchi, Jonathan Nagler, Josh Tucker, Ben Goldhaber, Andy Guess, Naomi Zucker, Jim Bisbee, Gabor Simonovits, and Michael Munger. I also thank the working journalists and editors who agreed to be interviewed on background, providing much of the intuition that underlies this piece. All remaining errors are mine.

Disclosure statement

No potential conflict of interest was reported by the author.

References

- Achen, C. H., & Bartels, L. M. (2016). *Democracy for realists: Why elections do not produce responsive government*. New York, NY: Princeton University Press.
- Allcott, H., & Gentzkow, M. (2016). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2), 211–236.
- Archer, A. M., & Clinton, J. (2018). Changing owners, changing content: Does who owns the news matter for the news? *Political Communication*, 35(3), 353–370. doi:10.1080/10584609.2017.1375581
- Baron, D. P. (2006). Persistent media bias. *Journal of Public Economics*, 90(1), 1–36. doi:10.1016/j.jpubeco.2004.10.006
- Baumol, W. J., Panzar, J. C., & Willig, R. D. (1983). Contestable markets: An uprising in the theory of industry structure: Reply. *The American Economic Review*, 73(3), 491–496.
- Berry, J. M., & Sobieraj, S. (2013). *The outrage industry: Political opinion media and the new incivility*. London, UK: Oxford University Press.
- Bode, L., & Vraga, E. K. (2015). In related news, that was wrong: The correction of misinformation through related stories functionality in social media. *Journal of Communication*, 65(4), 619–638. doi:10.1111/jcom.12166

- Bovet, A., & Makse, H. A. (2019). Influence of fake news in Twitter during the 2016 US presidential election. *Nature Communications*, 10(1), 7. doi:10.1038/s41467-018-07761-2
- Boydston, A. E. (2013). *Making the news: Politics, the media, and agenda setting*. New York, NY: University of Chicago Press.
- Boydston, A. E., Hardy, A., & Walgrave, S. (2014). Two faces of media attention: Media storm versus non-storm coverage. *Political Communication*, 31(4), 509–531. doi:10.1080/10584609.2013.875967
- Campbell, A., Converse, P. E., Miller, W. E., & Donald, E. (1960). *The American Voter*. New York, NY: University of Chicago Press.
- Cellan-Jones, R. (2017). *Facebook's news feed experiment panics publishers*. BBC.com.
- Chadwick, A. (2006). *Internet politics: States, citizens, and new communication technologies*. London, UK: Oxford University Press.
- Coleman, G. (2014). *Hacker, hoaxer, whistleblower, spy: The many faces of anonymous*. New York, NY: Verso Books.
- Department of Consumer Affairs and Cosmetology. (2016). Frequently asked questions. *CA Website*.
- DiMaggio, P., & Hargittai, E. (2001). From the 'digital divide' to 'digital inequality': Studying internet use as penetration increases. *Princeton: Center for Arts and Cultural Policy Studies, Woodrow Wilson School, Princeton University*, 4(1), 4–2.
- Downs, A. (1957). An economic theory of political action in a democracy. *The Journal of Political Economy*, 135–150. doi:10.1086/257897
- Facebook. (2016). Choose your audience. *Facebook Advertising Website*.
- Fu, W. W., & Sim, C. C. (2011). Aggregate bandwagon effect on online videos' viewership: Value uncertainty, popularity cues, and heuristics. *Journal of the American Society for Information Science and Technology*, 62(12), 2382–2395. doi:10.1002/asi.21641
- Garrett, R. K., Nisbet, E. C., & Lynch, E. K. (2013). Undermining the corrective effects of media-based political fact checking? The role of contextual cues and nave theory. *Journal of Communication*, 63(4), 617–637. doi:10.1111/jcom.12038
- Gentzkow, M., & Shapiro, J. M. (2008). Competition and truth in the market for news. *The Journal of Economic Perspectives*, 22(2), 133–154. doi:10.1257/jep.22.2.133
- Green, D., Palmquist, B., & Schickler, E. (2002). *Partisan hearts and minds: Political parties and the social identities of voters*. New York, NY: Yale University Press.
- Grossmann, M., & Hopkins, D. A. (2016). *Asymmetric politics: Ideological republicans and group interest democrats*. London, UK: Oxford University Press.
- Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, 5(1), eaau4586. doi:10.1126/sciadv.aau4586
- Guess, A., Nyhan, B., & Reifler, J. (2017). Inside the fake news bubble? Consumption of online fake news in the 2016 U.S. election. *Working Paper*.
- Hamilton, J. (2004). *All the news that's fit to sell: How the market transforms information into news*. London, UK: Princeton University Press.
- Hargittai, E. (2001). Second-level digital divide: Mapping differences in people's online skills. *arXiv preprint cs/0109068*.
- Hargittai, E., Piper, A. M., & Morris, M. R. (2018). From internet access to internet skills: Digital inequality among older adults. *Universal Access in the Information Society*, 1–10.
- Hindman, M. (2008). *The myth of digital democracy*. New York, NY: Princeton University Press.
- Hindman, M. (2018). *The internet trap: How the digital economy builds monopolies and undermines democracy*. New York, NY: Princeton University Press.
- Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology a social identity perspective on polarization. *Public Opinion Quarterly*, 76(3), 405–431. doi:10.1093/poq/nfs038
- Jun, Y., Meng, R., & Johar, G. V. (2017). Perceived social presence reduces fact-checking. *Proceedings of the National Academy of Sciences* 114(23), 5976–5981.
- Kankaraš, M., Montt, G., Paccagnella, M., Quintini, G., & Thorn, W. (2016). *Skills matter: Further results from the survey of adult skills. OECD skills studies*. London, UK: OECD Publishing.

- Karpf, D. (2016). *Analytic activism: Digital listening and the new political strategy*. New York, NY: Oxford University Press.
- Karpf, D. (2019). Something I no longer believe: Is internet time slowing down? *Social Media+ Society*, 5(3), 2056305119849492.
- Ladd, J. M. (2011). *Why Americans hate the media and how it matters*. New York, NY: Princeton University Press.
- Martin, G. J., & McCrain, J. (2018). Local news and national politics. *American Political Science Review*, 113(2), 372–384.
- Martin, S. (2000). The theory of contestable markets. *Bulletin of Economic Research*, 37(1), 65–68.
- Mason, L. (2016). A cross-cutting calm: How social sorting drives affective polarization. *Public Opinion Quarterly*, 80(S1), 351–377.
- Messing, S., & Westwood, S. J. (2012). Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online. *Communication Research*, 41(8), 1042–1063.
- Mider, Z. (2016). What kind of man spends millions to elect Ted Cruz? *Bloomberg Politics*.
- Miller, C. C. (2011). Seeking to weed out drivel, Google adjusts search engine. *New York Times*.
- Mossberger, K., Tolbert, C. J., & McNeal, R. S. (2007). *Digital citizenship: The internet, society, and participation*. Cambridge, MA: MIT Press.
- Mutz, D. C. (2015). *In-your-face politics: The consequences of uncivil media*. New York, NY: Princeton University Press.
- Nicas, J. (2016, November 16). Google to bar fake-news websites from using its ad-selling software. *Wall Street Journal*.
- Nielsen, J. (2016). The distribution of users' computer skills: Worse than you think. *Nielsen Norman Group*. Retrieved from <https://www.nngroup.com/articles/computer-skill-levels/>
- Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, 32(2), 303–330. doi:10.1007/s11109-010-9112-2
- Nyhan, B., Reifler, J., Richey, S., & Freed, G. L. (2014). Effective messages in vaccine promotion: A randomized trial. *Pediatrics*, 133(4), e835–e842. doi:10.1542/peds.2013-2365
- Prior, M. (2007). *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. London, UK: Cambridge University Press.
- Salam, M., & Victor, D. (2018). Yanny or Laurel? How a sound clip divided America. *New York Times*.
- Schmitt-Beck, R. (2015). Bandwagon effect. *The International Encyclopedia of Political Communication*, 1–5.
- Settle, J. (2018). *Frenemies: How social media polarizes America*. London, UK: Cambridge University Press.
- Silverman, C., & Alexander, L. (2016). How teens in the Balkans are duping Trump supporters with fake news. *Buzzfeed News*, p. 3.
- Silverman, C., Lytvynenko, J., & Vo, L. T. (2018). *How Facebook groups are being exploited to spread misinformation, plan harassment, and radicalize people* (Technical report BuzzFeed). Retrieved from <https://www.buzzfeednews.com/article/craigsilverman/how-facebook-groups-are-being-exploited-to-spread>
- Smith, A., & Anderson, M. (2018). Social media use in 2018. *Pew Research Center*. Retrieved from http://assets.pewresearch.org/wp-content/uploads/sites/14/2018/03/01105133/PI_2018.03.01_Social-Media_FINAL.pdf
- Stroud, N. J. (2011). *Niche news: The politics of news choice*. London, UK: Oxford University Press on Demand.
- Sundar, S. S., Oeldorf-Hirsch, A., & Xu, Q. (2008). The bandwagon effect of collaborative filtering technology. In *CHI'08 extended abstracts on human factors in computing systems* (pp. 3453–3458). ACM.
- Sydell, L. (2016). We tracked down a fake-news creator in the suburbs. Here's what we learned. *NPR.com*.
- The Editors. (2015). *A growth spurt*. Retrieved from <http://unabridged.merriam-webster.com/blog/2015/05/a-growth-spurt/>

- Usher, N. (2014). *Making news at the New York Times*. Ann Arbor, MI: University of Michigan Press.
- Wallenstein, A., & Spangler, T. (2013). Epic fail: The rise and fall of demand media. *Variety*. Retrieved from <https://variety.com/2013/biz/news/epic-fail-the-rise-and-fall-of-demand-media-1200914646/>
- Xu, X. W., & Fu, W. (2014). Aggregate bandwagon effects of popularity information on audiences' movie selections. *Journal of Media Economics*, 27(4), 215–233. doi:10.1080/08997764.2014.963229