The Economics of Online News Media

July 11, 2017

Abstract

The news media industry has changed as the internet and social media have matured and become integral to modern life: there are orders of magnitude more stories with a national reach, distribution is free and instantaneous, and the context of media choice is more spontaneous and related more to the story than the media firm. These differences have important implications for the study of media choice and the ensuing effects of media consumption on information retention and partisan affect. I explain these aspects of online media through an analysis of the economic structure of the industry and explore the implications for scholars of online media and politics. My central claim 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 mechanism interacts with the massive (and unappreciated) heterogeneity in digital literacy to create "credibility cascades", which I argue are a necessary condition for Fake News to flourish.

1 Overview

The 2016 US Presidential election revealed that the mainstream news media are in a dire place. Trust in the institution has been declining steadily ever since the 1970s (Ladd, 2011), and this distrust among conservatives have driven the development of a parallel news media defined by cults of personality and outrage (Berry and 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 and Gentzkow, 2017), the fact 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.

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 existince trends in online media. In this article, I explore how the structure of the online media industry has evolved: entry 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", the mechanism which drives the online news industry: stories acquire credibility as they are shared along social networks, becoming more desirable at they same time as they increase thier potential audience.

This theory applies only to the mature online media environment; it makes no sense to analyze the impact of the "the internet" on news consumption because this relationship has changed rapidly over the past twenty years. The world in which a casual Facebook user reads the headline of a Breitbart article shared by a high school acquaintance on her smartphone is very different from the world in which high-information users read specific and esoteric political blogs on their expensive and difficult-to-use desktop computers. Similarly, we cannot understand the mature online media industry solely by analyzing the technological affordances of the internet; this technology has co-evolved with political culture and government regulation, and the interaction of these components can have unintuitive and nonlinear effects.

Political scientists know all of this to be true of television; although the fundamental capacity to transmit an image from one location to millions of other locations has existed for nearly a decade, the impact of television on news consumption has shifted with the cultural, regulatory and technological envionment. As Prior (2007) and other scholars have conclusively demonstrated, the proliferation of cable news had the opposite effect on media consumption than did broadcast television: incidental exposure to news content fell with the availability of entertainment-only television changes, leading to higher informational inequality and an smaller but increasingly polarized electorate.

The evolution of social media out of the early web appears to be having the converse effect. The early web was more similar to cable news, with specialized web sites that consumers had to explicitly seek out, producing little opportunity for incidental exposure. Social media–and especially the broadest and most captivating social media platform, Facebook–aims to be a one-stop shop for internet users, who can be exposed

to anything on any other part of the internet, including political news. It seems to be succeeding in that role, as more than half of US adults log on every day (Greenwood, Perrin, and Duggan, 2016).

On the other hand, if social media were directly analogous to broadcast television, it would be producing a high information, low polarization electorate. The general intuition is that the 2016 US Presidential Election was not characterized by a high information, low polarization electorate.

There are, of course, many specific ways in which the mature online media industry differs from the broadcast television industry. The type and number of news stories that are produced are radically different. These changes have heterogeneous effects on the media diets of different types of consumers, and on the way that they react to consuming that news. A Facebook user who prefers not to read news might still decide to click on a particularly sensational story shared by a close friend about the newest evil thing their political opponents are up to, and it might cause that user to feel even more alienated from politics.

On the other hand, it might not: there simply is not enough research on the political implications of media choice and media effects in the social media context. There are a dizzying array of physical (am I reading the story on a laptop or on a mobile device? but see Searles and Dunaway (2017); Searles et al. (2017)) and social (did my family member share the story, or an acquaintance?) differences in the way that people consume news on social media compared to television, radio or newspapers, none of which have been well studied by scholars of media effects. In a recent book, Mutz (2015) demonstrates the moderating impact of *camera distance* on the effects of cable news on partisan affect and information retention, suggesting that there might be large differences in the effect of reading eg the same New York Times op ed in the physical paper compared to reading it on an iPhone after your anarchist niece retweeted it with a caption deeming the article complicit in the murder of innocents.

Although this research is badly needed, it faces huge challenges: online media in the age of social media is not a stationary concept. The more any study of media effects attempts to capture the nuances of a specific kind of media exposure, the more quickly it will become obsolete. Consider the example of Upworthy. The "fastest-growing media site of all time," Upworthy implemented a new style of "clickbait" headline designed to entice consumption by strategically withholding information (Sanders, 2017). Less than two years after it was founded, 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 (**?**).

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. However, because of the timeline of academic publishing, the first peer-reviewed publication to contain the word "clickbait" was not published until 2015.[1]

Understanding the structural incentives underlying the online news industry may help media scholars focus their research agendas on the elements of online news most likely to remain important. To develop that understanding, I leverage an old and well developed theoretical framework from economics that I believe applies remarkably well to the online media industry: contestable markets (Baumol et al., 1982).

2 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. It was cheap to rent an airplane and operate a very small-scale business. This operation would easily be able to out-compete the existing firms, with their large organizational and legacy costs; airline flights are a relatively undifferentiated product, so price is the main way these firms compete. Baumol called this a *contestable market*, and made several predictions about how firms might behave in such an industry.

In the short run, new firms would engage in "hit-and-run" competition, entering the industry and charging sub-market rates in order to attract business from the established firms. Once the market price has adjusted to account for the increase in supply, these new firms will exit the industry. 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 any firm from charging a price higher than their marginal cost in order to earn a profit.

Baumol identified several characteristics necessary to characterize a contestable market. In the real world in which entropy exists, none of these idealized conditions can

[1] A Google Scholar search performed in June 2017 suggests that there have only been 8 peer-reviewed publications with at least 1 citation ever published; the majority of these are Computer Science articles in faster-moving Conference Proceedings.

actually occur, so it makes more sense to refer to degrees of contestability in a market rather than 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.

- Sunk Costs: There cannot be any capital investments (in either physical or intellectual capital) that cannot be recouped.

- Symmetric Information/Technology: There cannot be any specialized technology or knowledge available to the incumbent firms but not the new entrants.

Baumol may have been wrong about the airline industry being a contestable market; at least, it failed to be characterized by rapid entry and exit of extremely low-cost firms (Martin, 2000). However, the three above conditions do a disturbingly good job of describing the online media industry in 2016. We appear to be in the short run "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 sustainable. 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 the steady state of this market is also what constable 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 online news 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.

3 Entry/Exit

Until recently, it was very difficult 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 to technical know-how to set up their own blog and create news content on their own. Blogs are not media companies, however, and 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 use Facebook, and 51% of adults use Facebook *every day* (Greenwood, Perrin, and Duggan, 2016), 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 is the regulatory environment. During the postwar era, the power of the FCC 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 and 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 reactionsColeman (2014). The generational makeup of Congress makes it unlikely that they understand the technical challenges in regulating the internet.

Another form of regulation impacts firm entry: relative regulatory licensing restrictions. In the real world, there are finite number of entrepreneurs who plan to start a company. One important consideration are the explicit regulatory hurdles in each in-

dustry; 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 20 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 non-union workers.[2] In an interview with an editor at a prominent web native media organization, I learned that his role as "editor" was quite different from that at traditional media outlets.

Rather than a few editors and many reporters, the newsroom he worked in consisted almost entirely of "editors": people who were tasked with getting some number of stories (weighted by pageviews) published every day. They could write them themselves, but it was far more efficient to commission each story to be written by a freelancer. These freelancers rarely had any formal training, and rather than making any investments in developing their journalistic skills, the editor would make aggressive edits or merely refuse to publish a story that was too poorly written.

This institutional setup means that web native companies have minimal long-term commitments, giving them close to zero exit costs.

[2]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 Sunk Costs

Essential to operating a news media company is that consumers believe what you report. Developing a reputation for credible reporting was 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 the 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 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 potential reader that the reporting in a story is credible. The fact of social recommendation provides an alternative source of credibility; experimental evidence suggests that consumers prefer to share content with more anonymous recommendations on Facebook (Messing and Westwood, 2012). Social recommendations from friends or family members should provide even more legitimacy.

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).

An alternative strand of scholarship, beginning with Campbell et al. (1960) and seeing renewed interest with Green, Palmquist, and Schickler (2002), conceptualizes partisanship as a social identity (Iyengar, Sood, and Lelkes, 2012; Mason, 2016). 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 media environment is perfect for information segregation and media companies which need to sink zero costs into investing in credibility.[3] The media to which people are exposed on social media is explicitly the media that people in their social networks think is important, and because social networks are homophilious, this means that content should be spread on social media 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's strategy is to make 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?"

[3]These effects are likely to be asymmetric. Republicans are more distrustful of the mainstream media and far more unified by symbolic ideology (Grossmann and Hopkins, 2016). In the course of the 2016 election, Allcott and Gentzkow (2017) estimate that the average American remembered .92 pro-Trump and .23 pro-Clinton stories from a zero-credibility outlet.

- "We Know Where In America You Actually Live"
- "Can We Guess Where In The USA You Actually Live?"

This strategy applies to Buzzfeed's less asinine arm, Buzzfeed News. I interviewed an editor at Buzzfeed who explained that the company decided to begin covering poltical news when they realized that LGBT news was being under-provided. The topic fits in with the Buzzfeed model because of the built-in audience of people who share a social identity–the LGBT community.

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 is unnecessary.

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 biasesSydell (2016).

The crucial innovation of these sites is spread news that is verifiably false–for this to work, consumers must be unwilling or unable to consult other media sources and learn that they have been deceived.

This represents something of a blind spot for journalists and academics studying fact-checking. In controlled settings, studies have shown various degrees of success in correcting misinformation, with partisans being less likely to accept that a ideologically congruent belief is false (Bode and Vraga, 2015; Garrett, Nisbet, and Lynch, 2013; Nyhan and Reifler, 2010; Nyhan et al., 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 consuption contexts.

An exception is Jun, Meng, and Johar (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.

This study, like many studies involving online fact checking, recruited subjects via Amazon Mechanical Turk. The Mechanical Turk sample is far from nationally repre-

sentative in terms of age, race and gender (Huff and Tingley, 2015), but there is an even more pressing bias: *100% of MTurkers are digitally literate.*

On this dimension, this is precisely the wrong sample on which to test online fact-checking. Fact-checking is in fact beyond the technical capacity of a significant portion of social media users in the US.

I believe that there is an important sociological reason that this issue is not a major part of the conversation on online misinformation. Although the problem of internet use leading to balkanization/echo chambers/bubbles has been widely discussed in the partisan context, the class of people who use and study the internet has been structurally blinded to the existence of a bubble to which they necessarily belong. It is almost impossible to measure digital literacy with online samples, but there are millions of Americans who use social media yet lack the digital literacy required to perform the task of fact-checking.

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

Hargittai (2001) calls this the "second-level digital divide": there is a 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 nearly essentially all producers of online media belong) and the majority of internet users (Kankaraš et al., 2016). The OECD's survey instrument is desinged 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 email (14%)
- Can use "reply-all" to send an email 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 email and not directly relevant to the 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 delete an email.

The combination of these low-skill social media users and the power of social recommendation are necessary conditions for the Fake News business model. By buying Facebook or Twitter 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).

Fake News purveyors target precisely these social media users who are least 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. If they are able to convince these people to consume and then share their inflammatory but fictitious content, their audience expands *and* the fact of the social recommendations lends legitimacy to their content.

In this way, people with higher degrees of skepticism or with greater digital literacy are enticed to consume and share content that they might otherwise find questionable. 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 incrementally more and more sophisticated/digitally literate news consumers find the story credible. Borrowing from Kuran (1991), 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 tendencies in the traditional media have been well documented by other scholars of media and politics.

Boydstun (2013) describes a media industry characterized by path dependency and institutional conservationism If a particular story begins to attract media coverage, other media organizations are likely to cover it as well: it has demonstrated an ability to attract audience attention. Once a media company has devoted investigative resources to a story, there is tendency to follow up on related topics, meaning that media attention to a particular topic is likely to persist.[4]

A less benign comparison is to the symbiotic relationship between internet trolls and cable news described by Phillips (2015). Trolls' explicit aim was to upset as many people as possible; to that end, they would engage in acts of online cruelty (eg taunting the relatives of teen suicide victims on Facebook). Ostensibly out of moral concern, cable news channels (and especially Fox News) would devote hours of coverage to this phenomenon. This served to give trolls the attention they craved and to give the news

[4]This finding comports with a wealth of sociological evidence on the tendency for similar organizations to develop isomorphic institutional setups that decrease diversity of output (DiMaggio and Powell, 1983).

networks the viewers they wanted, and thus encouraged the trolls to continue.

This seems an apt description of the mainstream media in response to the emergence of the Fake News phenomenon in late 2016 (and, to a perhaps greater extent, the Republican Primary campaign of Donald Trump). Covering "Fake News" garnered a lot of attention for the mainstream news, but it also drove even more readers to the creators of Fake News, the culmination of the credibility cascade. In this way, these media organizations are able to attract widespread readership without needing to engage in the costly cost-sinking behavior of reputation building.

5 Technology/Skills

The necessary technology to create news content today is unpatentable and (other than lengthy, in-person investigative reporting) within the reach of any potential 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.

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.

In conversations with journalists and editors at Vice and Gawker, I found that the majority of people creating news content did not have J-school degrees, and indeed that such training was unnecessary and outdated. I spoke with an editor at Vice who

had achieved that position after 6 months as a reporter there, with no formal training. Each of these freelancers is competing directly with other freelancer, 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. Online media companies have 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.

Indeed, it as if the US higher education system were designed to produce a surplus of graduates with precisely these skills.

6 The Future

I have argued that online news media is currently a highly contestable market–new firms face minimal entry/exit costs, need not sink any costs into reputation building, and have access to the same technology as existing firms. The rapid entry (and sometimes exit)

[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.

of hundreds of new online media companies in the past decade is evidence of this fact. If this industry continues to conform to the theoretical expectations developed by Baumol et al. (1982), the long-term steady state may be one of far fewer firms and very low levels of profitability There is reason to believe that such an equilibrium–with little investigative reporting or choice among information sources–would have deleterious effects for democracy.

The theory offers clear implications for how to avoid this outcome: a market is more contestable the more these three conditions obtain.

The imposition of entry costs to the industry would straightforwardly discourage new entrants. If the US government imposed taxes or licensing restrictions on new media companies, there would be fewer new media companies.

This cure, however, could well be worse than the disease: an unregulated free press has long been seen as fundamental to the functioning of democracy. Furthermore, it would be very challenging to explicitly define a "media company" in an age of social media, where each individual can publish her own online content.

One important way in which the real world differs from ideal economic theory is the in the finite number of potential new entrants across all firms. The determining factor for each of these finite actors is not the entry cost to one industry in particular, but rather the relative entry costs among all potential industries. There is indeed a disparity in the level of government regulation between the media industry and most other industries, and it could be solved through deregulation.

In general, the past decade has seen an explosion in innovation in industries involving people sitting in front of a computer, but comparatively little elsewhere. To paraphrase Peter Thiel, there are many more barriers to entry in the flying car industry than in digital communication.[6]

Another possible avenue for making the online media market less contestable is to require media companies to sink costs into developing their reputations. It is tempting to think that a coalition of journalistic, academic, technological and political experts might be able to persuade consumers to only consume news from media companies with reputations for qualtiy; however, the rank failure of such a coalition to convince citizens to vaccinate their children, believe in anthropogenic climate change, or vote for Hillary Clinton suggests that this might not be a successful strategy. The late 2016 push for Facebook to label "Fake News" as such is a manifestation of this likely-to-fail reaction.

[6]Thiel does appear to be trying to impose barriers to entry to the digital media industry single-handedly: he funded the lawsuit that drove Gawker into bankruptcy.

An alternative approach would be to disrupt the credibility cascade that enables content from media companies without reputations to spread along social media. The challenge is to do so in a way that requires no editorial decision-making on the part of a social media platform or regulatory body.

One part of the process that might be disrupted is the speed at which these credibility cascades can occur. False information can be spread up the digital literacy chain faster than it can possibly be disproven–by the time it is read by people with the skills to fact check it, it may already have been read millions of times and mainstream news outlets might not able to resist covering it.

Social media outlets might give their users to *self-label* their posts as being legitimate news or not. If they decide not to, the *self-imposed* label is likely to carry far more weight than a "Fake News" label applied by distant platform administrators. If users do decide to label their post as news, there could be a 30-minute delay in posting it or sharing it once it has been posted.

This does not seem to be an unreasonable restriction of freedom of speech–rather, it seems like a return to the technological conditions of 1998. This proposal has clear parallels with proposals to regulate High Frequency Trading on the stock market: the speed of these phenomenon render them impossible for any individual human to understand or control, these institutions functioned tolerably well before the advent of near-infinite speed, and this speed has the clear and disastrous potential for adverse events.[7]

Even a simple reversion to the Facebook algorithm of June 29th, 2016, could have a significant impact. The company announced changes to their Newsfeed algorithm such that "Friends and family come first...To help make sure you dont miss the friends and family posts you are likely to care about, we put those posts toward the top of your News Feed...We dont favor specific kinds of sources or ideas. Our aim is to deliver the types of stories weve gotten feedback that an individual person most wants to see."(Mosseri, 2016). The update was specifically designed to prioritize posts from friends and family members at the expense of news outlets, and within news outlets, not to give any priority to established, reputable outlets. In analysis of Fake News trends, Silverman (2016) points out that "reactions" to the top Fake News stories (8.7 million)

[7]Twitter is distinct from Facebook in that has been used by dissidents and revolutions to spread information rapidly in the face of oppressive regimes, and to actually help organize real-world protests in real time. The problem is that Twitter is also the medium of choice for professional journalists. These are both valuable functions, but there is no reason that they both have to happen on the same platform.

eclipsed those of top mainstream news stories (7.3 million) in the months between that change and the election, but that the reverse was far and away the case before that change. Because these dynamics are so poorly understood, even these small algorithmic decisions can have a massive impact.

These suggestions are probably flawed and surely insufficient. The market for online media is in a precarious state in 2017, though, and it may continue to worsen unless one of the fundamental characteristics structuring it changes to make the market less contestable.

What is clear is that there are many questions that urgently need to be answered in the field of online media and politics. There are too many technological and social variables in flux for old models to fully capture the way that online media in the age of social media affect citizen beliefs and preferences. The specific styles, forms and technologies employed by online media companies will continue to change–potentially at an accelerating rate.[8] A reactionary, post hoc scholarship of these changes is unlikely to be sufficient, especially hindered by the rhythms of traditional academic publishing. For this reason, a structural understanding of the online media industry like the one I have attempted to provide in this paper is essential for scholars studying media and democracy in the 21st century.

[8]In particular, the price of video production technology and the bandwidth necessary to transmit it is rapidly decreasing, and many news outlets are eagerly creating video content designed for distribution on Facebook. At the same time, the technical capacity to create moderately convincing–convincing enough for consumers with low digital literacy–"fake" videos is rapidly expanding. Video is more emotionally effective and implies a greater degree of objectivity than text, so all of the above trends may be intensified as video content becomes the norm.

References

Allcott, Hunt, and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. Technical report National Bureau of Economic Research.

Baron, David P. 2006. "Persistent media bias." *Journal of Public Economics* 90 (1): 1–36.

Baumol, William J, John C Panzar, Robert D Willig, Elizabeth E Bailey, Dietrich Fischer, and Dietrich Fischer. 1982. "Contestable markets and the theory of industry structure.".

Berry, Jeffrey M, and Sarah Sobieraj. 2013. *The outrage industry: Political opinion media and the new incivility*. Oxford University Press.

Bode, Leticia, and Emily K. Vraga. 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.
URL: *http://dx.doi.org/10.1111/jcom.12166*

Boydstun, Amber E. 2013. *Making the news: Politics, the media, and agenda setting*. University of Chicago Press.

Campbell, Angus, Philip E Converse, Warren E Miller, and E Donald. 1960. "The American Voter.".

Chadwick, Andrew. 2006. *Internet politics: States, citizens, and new communication technologies*. Oxford University Press, USA.

Coleman, Gabriella. 2014. *Hacker, hoaxer, whistleblower, spy: The many faces of Anonymous*. Verso Books.

Department of Consumer Affairs, Board of Barbering, and Cosmetology. 2016. "Frequently Asked Questions." *CA Website* .

DiMaggio, Paul, Eszter Hargittai et al. 2001. "From the digital divideto 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.

DiMaggio, Paul, and Walter W Powell. 1983. "The iron cage revisited: Collective rationality and institutional isomorphism in organizational fields." *American Sociological Review* 48 (2): 147–160.

Downs, Anthony. 1957. "An economic theory of political action in a democracy." *The journal of political economy* pp. 135–150.

Facebook. 2016. "Choose your audience." *Facebook Advertising Website* .

Garrett, R Kelly, Erik C Nisbet, and Emily K Lynch. 2013. "Undermining the corrective effects of media-based political fact checking? The role of contextual cues and naïve theory." *Journal of Communication* 63 (4): 617–637.

Gentzkow, Matthew, and Jesse M Shapiro. 2008. "Competition and Truth in the Market for News." *The Journal of Economic Perspectives* 22 (2): 133–154.

Green, Donald, Bradley Palmquist, and Eric Schickler. 2002. "Partisan hearts and minds.".

Greenwood, S, A Perrin, and M Duggan. 2016. "Social Media Update 2016." *Washington, DC: Pew Internet & American Life Project. Retrieved November* 27: 2016.

Grossmann, Matt, and David A Hopkins. 2016. *Asymmetric Politics: Ideological Republicans and Group Interest Democrats*. Oxford University Press.

Hargittai, Eszter. 2001. "Second-level digital divide: mapping differences in people's online skills." *arXiv preprint cs/0109068* .

Hindman, Matthew. 2008. *The myth of digital democracy*. Princeton University Press.

Huff, Connor, and Dustin Tingley. 2015. ""Who Are These People?" Evaluating the Demographic Characteristics and Political Preferences of MTurk Survey Respondents." *Research and Politics* 2 (1): 1–12.

Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. "Affect, not ideology a social identity perspective on polarization." *Public opinion quarterly* 76 (3): 405–431.

Jun, Youjung, Rachel Meng, and Gita Venkataramani Johar. 2017. "Perceived social presence reduces fact-checking." *Proceedings of the National Academy of Sciences* p. 201700175.

Kankaraš, Miloš, Guillermo Montt, Marco Paccagnella, Glenda Quintini, and William Thorn. 2016. "Skills Matter: Further Results from the Survey of Adult Skills. OECD Skills Studies." *OECD Publishing* .

Kuran, Timur. 1991. "Now out of never: The element of surprise in the East European revolution of 1989." *World politics* 44 (01): 7–48.

Ladd, Jonathan M. 2011. *Why Americans hate the media and how it matters*. Princeton University Press.

Martin, Stephan. 2000. "The theory of contestable markets." *Bulletin of Economic Research* 37 (1): 65–68.

Mason, Lilliana. 2016. "A Cross-Cutting Calm How Social Sorting Drives Affective Polarization." *Public Opinion Quarterly* p. nfw001.

Messing, Solomon, and Sean J Westwood. 2012. "Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online." *Communication Research* p. 0093650212466406.

Mossberger, Karen, Caroline J Tolbert, and Ramona S McNeal. 2007. *Digital citizenship: The Internet, society, and participation*. MIt Press.

Mosseri, Adam. 2016. "Building a Better News Feed for You." *Facebookl* .

Mutz, Diana C. 2015. *In-your-face politics: The consequences of uncivil media*. Princeton University Press.

Nielsen, Jakob. 2016. "The Distribution of Users Computer Skills: Worse Than You Think." *Nielsen Norman Group* .

Nyhan, Brendan, and Jason Reifler. 2010. "When corrections fail: The persistence of political misperceptions." *Political Behavior* 32 (2): 303–330.

Nyhan, Brendan, Jason Reifler, Sean Richey, and Gary L Freed. 2014. "Effective messages in vaccine promotion: a randomized trial." *Pediatrics* 133 (4): e835–e842.

Phillips, Whitney. 2015. *This is why we can't have nice things: Mapping the relationship between online trolling and mainstream culture*. Mit Press.

Prior, Markus. 2007. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press.

Sanders, Sam. 2017. "Upworthy Was One Of The Hottest Sites Ever. You Won't Believe What Happened Next." *NPR* Juen 20, 2017.

Searles, Kathleen, and Johanna Dunaway. 2017. "News and Information Loss in the Mobile Era." *Working Paper* .

Searles, Kathleen, Mingxiao Sui, Paul Newly, and Johanna Dunaway. 2017. The Limits of Digital Citizenship: Constraints on News Consumption and Recall in the Mobile Setting. In *Unpublished Manuscript*.

Silverman, Craig. 2016. This Analysis Shows How Fake Election News Stories Outperformed Real News On Facebook. Technical report Buzzfeed.

Sydell, Laura. 2016. "We Tracked Down A Fake-News Creator In The Suburbs. Here's What We Learned." *NPR.com* .