Evaluating the fake news problem at the scale of the information ecosystem

Jennifer Allen[1], Baird Howland[2], Markus Mobius[3], David Rothschild[4], and Duncan J. Watts[5,*]

[1]MIT Sloan School of Management, 100 Main St., Cambridge, MA 02142
[2]Harmony Labs, 311 W 43rd St, New York, NY 10036
[3]Microsoft Research New England, 1 Memorial Dr, Cambridge, MA 02142
[4]Microsoft Research New York, 641 Avenue of the Americas, 7th Floor, New York, NY, 10011
[5]University of Pennsylvania, 3401 Walnut St, 459C, Philadelphia, PA 19104
[*]Correspondence to: djwatts@seas.upenn.edu

One sentence summary: Fake news is vastly outweighed by mainstream news, mostly on television, and news itself is a small fraction of overall US media consumption.

Abstract "Fake news," broadly defined as deliberately false or misleading information masquerading as legitimate news, is frequently asserted to be pervasive on the web, and on social media in particular, with serious consequences for public opinion, political polarization, and ultimately democracy. Using a unique multimode data set that comprises a nationally representative sample of mobile, desktop, and television consumption across all categories of media content, we refute this conventional wisdom on three levels. First, news consumption of any sort is heavily outweighed by other forms of media consumption, comprising at most 14.2% of Americans' daily media diets. Second, to the extent that Americans do consume news, it is overwhelmingly from television, which accounts for roughly five times as much as news consumption as online, while a supermajority of Americans consume little or no news online at all. Third, fake news comprises only about 1% of overall news consumption and 0.15% of Americans' daily media diet. Although consumption data alone cannot determine that online misinformation in any dose is not dangerous to democracy, our results suggest that the origins of public mis-informedness and polarization are more likely to lie in the content of ordinary news--especially on television--or alternatively in the avoidance of news altogether as they are in overt fakery.

Introduction

Since the 2016 US presidential election, the deliberate spread of online misinformation, in particular on social media platforms such as Twitter and Facebook, has generated extraordinary interest across several disciplines (*1–10*). In large part this interest reflects a deeper concern that the prevalence of "fake news" has increased political polarization, decreased trust in public institutions, and undermined democracy (*11–14*). Recently, a handful of papers have attempted to measure the prevalence of fake news on social media (*1*, *8*, *9*), finding that exposure is rare compared with other types of news content and generally concentrated among older, political conservative Americans. In spite of these findings, many researchers and other observers continue to advocate that deliberately engineered misinformation disseminated on social media is sufficiently prevalent to constitute an urgent crisis (*15*, *16*). Disagreements over the prevalence and importance of misinformation are difficult to evaluate empirically for three reasons. First, Americans consume news online via desktop computers and increasingly mobile devices as well as on television; yet no single source of data covers all three modes. As a result, researchers select data sources on the basis of their availability, which may not correspond with either representativeness or comprehensiveness. For example, many studies rely exclusively on Twitter, whose users are highly unrepresentative of the general population (*17*), while even studies that rely on representative online panels omit television consumption (*18*). Second, analyses of fake news often fail to account for how much of it is consumed relative to other types of news or non-news-related content. Because the volume of online content is so vast, even a very large numerator may constitute only a tiny fraction of the total (*19*). Third, even if its prevalence is low relative to other types of content, fake news could be important either because it is disproportionately impactful, or because it is concentrated on small subpopulations. While comprehensive measures of prevalence are intrinsically interesting and can indicate how much relative impact different types of content would have to have in order to dominate, they cannot on their own resolve questions about influence.

In this paper we address the first two of these three challenges, leaving the third for future research. We assembled a unique dataset that drew on three different sources to capture consumption across the two principal modes of news production, television and online, where we integrate total consumption across the modes by demographic bucket (see Materials and Methods and Supplementary Methods for a more detailed description of the datasets, definitions of key terms, and estimation methods). Content is defined by the mode on which it is consumed not produced; thus, for example, video consumed on desktop or a mobile device is categorized as online consumption even when it is produced by mainstream television stations.

Materials and Methods

First, we measure national television programming using Nielsen's nationally representative television panel (N $\approx$ 100,000). In addition, we measure local programming using a subset of the national panel (N $\approx$ 50,000) sampled from the 25 largest local markets. Television news consumption is defined as time devoted to watching any of the roughly 400 programs that are classified by Nielsen as "news"--a category that includes "hard news" (e.g. evening cable and network news), magazine news (e.g. Inside Edition, Dateline), morning shows (e.g. Good Morning America, Today Show), and entertainment news (e.g. TMZ, Access Hollywood)--as well as late night comedy shows (e.g. The Daily Show with Trevor Noah, the Late Show with Stephen Colbert), which are frequently viewed as a source of news-related information, especially for younger viewers (*20*).

Second, we measure desktop and mobile media consumption (including media consumed through apps on mobile) using Comscore's nationally representative desktop and mobile panel which breaks out total time spent on different types of media sites including news, search and social media by demographic bucket. Online (mobile and desktop) news consumption is defined as time spent on any article published on one of more than 800 websites, adapted from (*21*) , that primarily cover "hard" news topics like politics, business, and US and international affairs. Correspondingly, fake news consumption is the time spent on one of 98 websites previously identified by researchers (8), professional fact checkers, and journalists as sources of fake, deceptive, low quality, or hyperpartisan news. Thus, in accordance with the previous literature, with the notable exception of YouTube, fake news is defined at the publisher or URL-level. We further categorize online non-news consumption for the top 2000 domains, ranked by mobile and desktop traffic, into one of 28 ComScore categories (e.g. entertainment, gaming, health, social media, sports, etc.)

Third, we use Nielsen's nationally representative desktop-only web panel (90,000 in 2016 decreasing to 60,000 in 2018; see SM for details), which records individual visits to URLs as well as the referral URL, to impute passive news consumption (e.g news snippets, images, headlines, and summaries that appear on a newsfeed or search results page but which the user does not click on) on the top four social media sites (Facebook, YouTube, Twitter, and Reddit), as well as on the top three search engines (Google, Bing, and Yahoo). For every site except YouTube, we estimate this fraction as the fraction of URLs that are referred to from the platform in question that we classify as news and fake news respectively. For YouTube, which hosts all of its own content, we compute the fraction of a random sample of 360,000 videos (10,000 per month, weighted by viewing time) that are classified as "news and politics" in YouTube's internal classification scheme. We further count as online news consumption all time spent on the three major portals: MSN, Yahoo, and AOL. Finally, we use a subset of the Nielsen web panel ($N \approx 15{,}000$) who also appear in the television panel to estimate the relation between desktop and television news consumption.

Results

Fig. 1 shows the breakdown of Americans' daily desktop, mobile, and television media consumption, measured in minutes per person, over the course of three years spanning January 2016 through December 2018. Fig. 1A shows this pattern in aggregate, while Figs. 1B and 1C show the same pattern for the youngest (18-24 yo) and oldest (55+) age brackets respectively (see Fig. S1 for remaining age categories). On average, Americans devote over seven and a half hours (460 mins) per day to media consumption, including television, streaming video or music, gaming, engaging with social media, or browsing the web either from desktop or mobile devices (Fig. 1A). This total is relatively stable over the 36 month period of our data, showing seasonal declines during the summers, and peaks coinciding with the 2016 presidential election and the presidential inauguration in January 2017 (because the shares devoted to different types of content remain generally stable over time, in subsequent figures we aggregate over time; however the full over-time results are available in the SM). As expected, younger Americans spend more time on mobile devices, and less time watching television than average (Fig. 1B), whereas the pattern is reversed for older Americans (Fig. 1C); however, the former watch so much less television than the latter that their total media consumption is about 30% less in spite of their higher mobile usage. Fig. 1 also reveals three results that directly undercut the conventional wisdom about the prevalence of fake news online, and more broadly question the importance of online news relative to television news and other types of media consumption.

First, the bulk of daily media consumption is not news-related. As expected, young adults (Fig. 1B) spend less time consuming news (colored green) than average and far less time than the oldest group (Fig. 1C), but in all age groups news consumption is heavily outweighed by non-news consumption (colored blue). Of the 460 minutes per person per day of total media consumption, approximately 400 minutes (86%) is not related to news of any kind (see Table S6 for exact figures). Fig. 2 shows a more detailed breakdown of news and non-news categories of media consumption online (Fig. 2A) and on television (Fig. 2B). For online consumption, which includes mobile and desktop, news is dominated by several other categories such as entertainment, social media, and search. Even including passive exposure to news content on social media sites (Facebook, Twitter, Reddit, and YouTube), search engines (Google, Bing, and Yahoo!), and portals (Yahoo!, MSN, and AOL) news accounts for only 4.2% of total online consumption. Television news is more prominent, comprising the largest single category of television consumption and 23% of the total. In aggregate, however, television news is still heavily outweighed by non-news programming such as dramas, documentaries, movies, and sports (Fig. 2B). To the extent that Americans are uninformed about politics, economics and other issues relevant to democracy the reason may be simply that they are choosing not to inform themselves (*22*).

Second, to the extent that Americans do consume news they do so overwhelmingly by watching television. Overall, the ratio of television to online news--including both desktop and mobile devices--is more than five to one (54 minutes vs. 9.7 minutes), varying from a minimum of almost two to one for 18-24 year olds (9 mins vs. 5 mins) to a maximum of more than seven to one for 55 and older (94 mins vs. 13 mins). Online news (including both mobile and desktop activity) was more prominent in the vicinity of the 2016 election; however, the ratio of television to online remained similar (the minimum ratio in our 36 month time period is 4.5:1 during Dec 2016). Drawing on our sample of roughly 15K individuals who are members of both the Nielsen web and television panels, Fig. 3 shows that while essentially everyone is exposed to a substantial amount of daily television news, 44% of the sample is exposed to no online news at all and almost three quarters spends less than 30 seconds per day reading news online (see Fig. S2 for results broken down by age group, and Tables S8 and S9 for exact values). Because the Nielsen panel records only desktop activity these figures understate the true consumption of online news (i.e. including mobile). In light of our earlier result that average mobile news consumption is slightly less than desktop news consumption, however, and assuming that the distribution of news consumption is not dramatically different on mobile vs. desktop devices, then it follows that a majority of Americans spend less than a minute per day reading news online.

Third, fake news consumption (Fig. 1, colored red) is a negligible fraction of Americans' daily information diet. We emphasize here that both our definition of news and our definition of fake news are extremely broad. In the case of news we include, for example, morning shows and portals, while our definition of fake news includes highly biased and and hyper-partisan news sites such as Breitbart.com (i.e. corresponding to the "red" and "orange" categories defined in (*8*)) as well as outright fraudulent sites (i.e. the "black" category). Our estimates of the prevalence of news and fake news therefore likely overstate the true prevalence, although we also find that adopting stricter definitions makes no discernable difference to our main conclusions (see Fig. S3 for comparison of upper and lower bound estimates of news and fake news consumption respectively, and Table S10 for exact values). Fig. 4 shows a more detailed breakdown of news consumption online (Fig. 4a) and on television (Fig. 4b), also broken out by age group (see Table S11 for numerical values).

Referring first to online consumption, Fig. 4A shows that fake news stories were more likely to be encountered on social media (dark vs. light red), and that older viewers were heavier consumers than younger ones, consistent with previous findings (*6*, *8*, *9*). No age group, however, spent more than an average of a minute per day engaging with fake news, nor did it occupy more than 1% of their overall news consumption (i.e. including television), or more than 0.2% of their overall media consumption. Of potential concern, a very small fraction of desktop panelists (1.97%) did consume more fake news than mainstream news; however this number drops to 0.7% when restricting to people who consumed at least one minute of fake news per day. When restricting to just "black" and "red" fake news sites (i.e. excluding hyperpartisan sites), these numbers drop to 0.97% and 0.32% respectively. In other words, while majority-fake news consumers do exist they are extremely rare and most of them consume very little online news of any kind.

Turning to television, there are no objectively fake news stations of the sort that exist online; i.e. that are exclusively or near-exclusively devoted to disseminating deliberate falsehoods while masquerading as legitimate news organizations. Including TV news consumption in the previous calculation would therefore reduce the population of majority-fake news consumers even further. Nonetheless, misinformation construed more broadly can also manifest itself in regular news programming in the form of selective attention, framing, "spin," false equivalence and other forms of bias. Although a detailed analysis of false or misleading content contained in conventional news programming is beyond the scope of this paper, it is nonetheless interesting to examine how much collective attention is paid to different categories of news. Fig. 4b provides this breakdown, showing first that television news consumption greatly exceeds online news, and is sharply increasing with age, ranging from less than ten minutes per day (18-24 year olds) to over 90 minutes per day (55+). Local news is the dominant form of news consumption for all age groups except the oldest for whom national cable news (ranked second overall) is slightly more popular. In turn, the relative dominance of cable news in the 55+ category is driven by a small minority of voracious news consumers (roughly the top 10% by consumption). Hard network news (e.g. evening news shows) is ranked third for all age groups, while morning shows are ranked fourth for all age groups but the youngest, which slightly prefer late-night comedy shows. Given the large differences in total news consumption across age groups, the consistency of ranking of different types of news is striking. Also striking given its perceived importance for younger viewers, is the limited presence of late-night comedy (less than 5% overall, less than 7% for 18-24 year olds).

Discussion

Summarizing, we note that according to Google Scholar, 2,210 English language publications with “fake news” in the title have appeared since Jan 2017, compared with just 73 in all the years leading up to and including 2016. Not only has interest in fake news clearly exploded in the past two years, it has also far outstripped attention to television news: a comparable count yielded just 329 articles published since 2017 containing either "television news" or “TV news” in their titles, while 708 articles contained "online news," 394 contained “Twitter” or Facebook” and “news,” and 556 contained “social media” and “news.” Restricting further to studies that explicitly connect misinformation to a particular platform, Google Scholar yielded 99 results containing both “misinformation” and one of “online” or “social media” or “web” in the title since 2017, but just 1 result for “misinformation” and “television” or “tv”--an article about the unrealistic survival rates of cardiopulmonary resuscitation on tv shows. This evident focus of the recent research literature on online sources of fake news and misinformation is directionally and proportionately inconsistent with our results in three ways. First, whereas the research treats news consumption as the issue of primary importance we find that most media consumption, whether online or on television, is not news related. Second, whereas research on online news--and even more specifically news on social media platforms--dramatically outweighs research on television news, we find that television news consumption dominates online by a ratio of 5:1 (where the ratio is even more extreme for social media sites). Third, whereas the topic of fake news outstrips all other news-related research, we find that fake news itself is only 1% of overall news consumption, substantially lower for Twitter alone (*8*). Instead news consumption is heavily dominated by mainstream news sources both online and on television.

We emphasize that our results do not imply that fake news is not a problem worthy of attention. Arguably the deliberate circulation of false information with the objective of creating confusion and discord is intolerable in principle and should be combatted at any prevalence greater than zero. Moreover, it is possible that news consumed online could have more impact per minute of exposure than news consumed on TV, or that fake news could have an outsized impact compared with regular news, or that it could have large impacts on certain subpopulations. Finally, we note that our definitions of news and fake news are--with the exception of YouTube--dependent on site or program level classifications. News-relevant content on social media that is not tied to a particular URL, or false or misleading information that is promulgated by generally reliable news sources, would therefore be misclassified by our scheme. We hope that future work will address all of these areas of uncertainty. We note, however, that our methodology was designed to be consistent with previous work, which also has used list-based classification, and which has also relied on prevalence (i.e. not impact) to assess importance. On those terms, our finding that fake news is extremely rare, comprising only about one tenth of one percent of Americans' overall daily media diet, suggests that concerns regarding possible threats to democracy should be much broader in scope than deliberately engineered falsehoods circulating on social media. In particular, public ignorance or misunderstanding of important political matters could also arise out of a combination of (a) ordinary bias and agenda setting in the mainstream media (*23–25*)), and (b) the overall low exposure of many Americans to news content in general, especially in written form. We conclude that future work on misinformation and its potentially corrosive effects on democracy should consider all potential sources of problematic content, as well as the absence of relevant content, not simply the type that is most easily identified and least associated with conventional media interests (*19*).

1. H. Allcott, M. Gentzkow, Social Media and Fake News in the 2016 Election. *J. Econ. Perspect.* **31**, 211–236 (2017).

2. K. Shu, A. Sliva, S. Wang, J. Tang, H. Liu, Fake News Detection on Social Media: A Data Mining Perspective. *SIGKDD Explor. Newsl.* **19**, 22–36 (2017).

3. D. M. J. Lazer, M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, The science of fake news. *Science*. **359**, 1094–1096 (2018).

4. S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online. *Science*. **359**, 1146–1151 (2018).

5. C. Shao, P.-M. Hui, L. Wang, X. Jiang, A. Flammini, F. Menczer, G. L. Ciampaglia, Anatomy of an online misinformation network. *PLoS One*. **13**, e0196087 (2018).

6. A. Guess, B. Nyhan, J. Reifler, Selective Exposure to Misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign. *European Research Council* (2018).

7. M. Stella, E. Ferrara, M. De Domenico, Bots increase exposure to negative and

inflammatory content in online social systems. *Proc. Natl. Acad. Sci. U. S. A.* **115**, 12435–12440 (2018).

8. N. Grinberg, K. Joseph, L. Friedland, B. Swire-Thompson, D. Lazer, Fake news on Twitter during the 2016 U.S. presidential election. *Science*. **363**, 374–378 (2019).

9. A. Guess, J. Nagler, J. Tucker, Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Sci Adv*. **5**, eaau4586 (2019).

10. G. Pennycook, D. G. Rand, Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 2521–2526 (2019).

11. N. Persily, Can democracy survive the Internet? *Journal of democracy*. **28**, 63–76 (2017).

12. J. Tucker, A. Guess, P. Barberá, C. Vaccari, A. Siegel, Social media, political polarization, and political disinformation: A review of the scientific literature (2018) (available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3144139).

13. G. L. Ciampaglia, Fighting fake news: a role for computational social science in the fight against digital misinformation. *Journal of Computational Social Science*. **1**, 147–153 (2018).

14. K. Clayton, S. Blair, J. A. Busam, S. Forstner, J. Glance, G. Green, A. Kawata, A. Kovvuri, J. Martin, E. Morgan, M. Sandhu, R. Sang, R. Scholz-Bright, A. T. Welch, A. G. Wolff, A. Zhou, B. Nyhan, Real Solutions for Fake News? Measuring the Effectiveness of General Warnings and Fact-Check Tags in Reducing Belief in False Stories on Social Media. *Political Behavior* (2019), doi:10.1007/s11109-019-09533-0.

15. S. Aral, D. Eckles, Protecting elections from social media manipulation. *Science*. **365** (2019), pp. 858–861.

16. J. Roozenbeek, S. van der Linden, The fake news game: actively inoculating against the risk of misinformation. *J. Risk Res.* **22**, 570–580 (2019).

17. Wojcik Stefan And, Sizing up Twitter Users. *Pew Research Center* (2018) (available at https://www.pewinternet.org/2019/04/24/sizing-up-twitter-users/).

18. A. Mitchell, Americans Still Prefer Watching to Reading the News – and Mostly Still Through Television. *Pew Research Center* (2018) (available at https://www.journalism.org/2018/12/03/americans-still-prefer-watching-to-reading-the-news-and-mostly-still-through-television/).

19. D. J. Watts, D. Rothschild, Don’t blame the election on fake news. Blame it on the media. *Columbia J. Rev.* (2017) (available at https://www.cjr.org/analysis/fake-news-media-election-trump.php).

20. L. Feldman, The news about comedy: Young audiences, The Daily Show, and evolving notions of journalism. *Journalism*. **8**, 406–427 (2007).

21. S. Athey, M. M. Mobius, J. Pál, The Impact of Aggregators on Internet News Consumption (2017), (available at https://papers.ssrn.com/abstract=2897960).

22. S. Edgerly, E. K. Vraga, L. Bode, K. Thorson, E. Thorson, New Media, New Relationship to Participation? A Closer Look at Youth News Repertoires and Political Participation. *Journal. Mass Commun. Q.* **95**, 192–212 (2018).

23. M. E. McCombs, D. L. Shaw, The agenda-setting function of mass media. *Public Opin. Q.* **36**, 176–187 (1972).

24. R. Puglisi, J. M. Snyder, in *Handbook of Media Economics*, S. P. Anderson, J. Waldfogel, D. Strömberg, Eds. (North-Holland, 2015), vol. 1, pp. 647–667.

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

26. W. E. Deming, F. F. Stephan, On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. Ann. Math. Stat. 11, 427–444 (1940).

27. C. Silverman, In Spite Of Its Efforts, Facebook Is Still The Home Of Hugely Viral Fake News. BuzzFeed News (2018), (available at https://www.buzzfeednews.com/article/craigsilverman/facebook-fake-news-hits-2018).

28. C. Silverman, This analysis shows how viral fake election news stories outperformed real news on Facebook. BuzzFeed News. 16 (2016).

Acknowledgments: The authors are grateful to Harmony Labs for engineering and financial support, to Nielsen for access to their television and web panels, and to the Nathan Cummings Foundation for funding support. Author Contributions: All authors contributed to the design of the research. J.A. and C.H. analyzed the data and drafted the supplementary materials, and J.A., C.H., and D.R. produced the figures. D.W. drafted the paper, and all authors contributed to revisions. Competing Interests: All authors declare no competing interests. Data and Materials: All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Additional data related to this paper may be found at https://osf.io/cygta/ or requested from the authors.

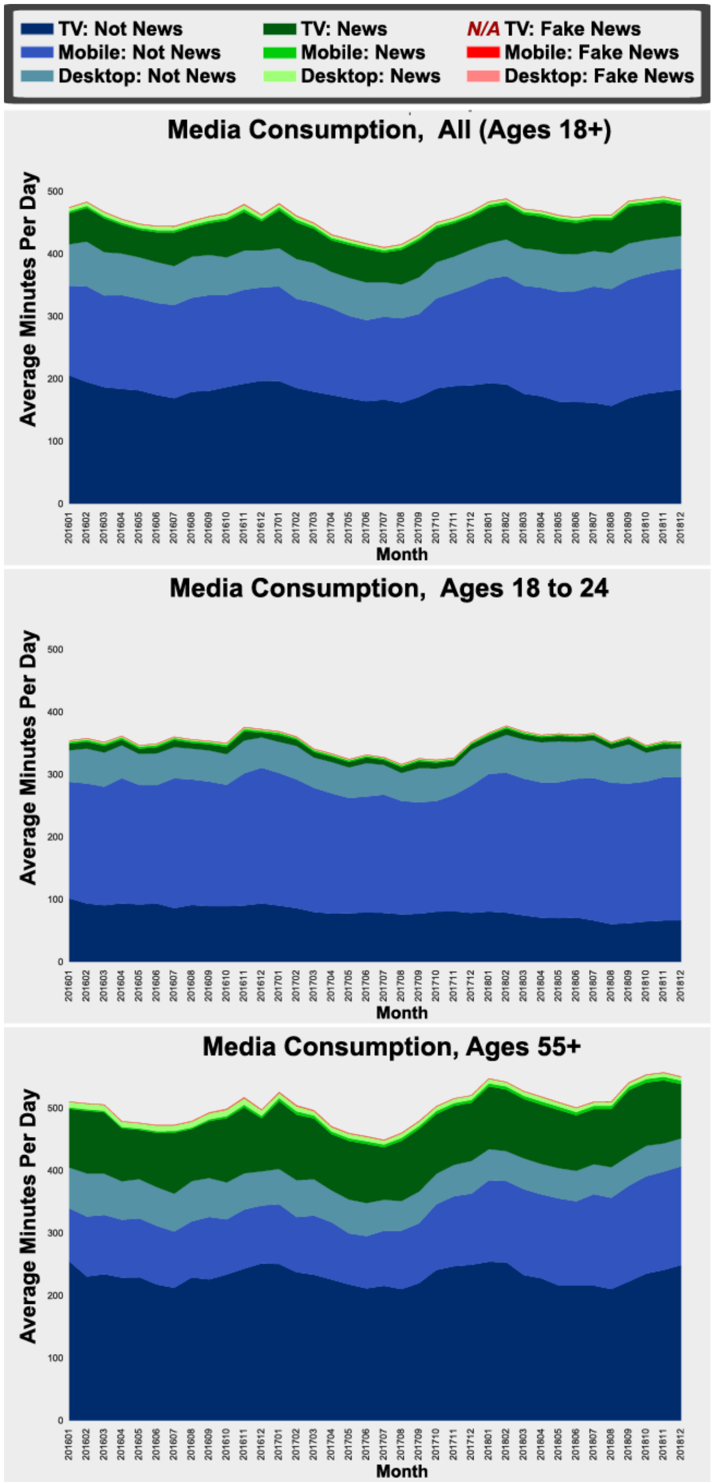

Figure 1. Overall information consumption by category and time, from Jan 2016 - Dec 2018 for (A) entire adult sample, 18 years and older, (B) 18-24 year olds, and (C) 55 years and older. See Table S6 for numerical values.

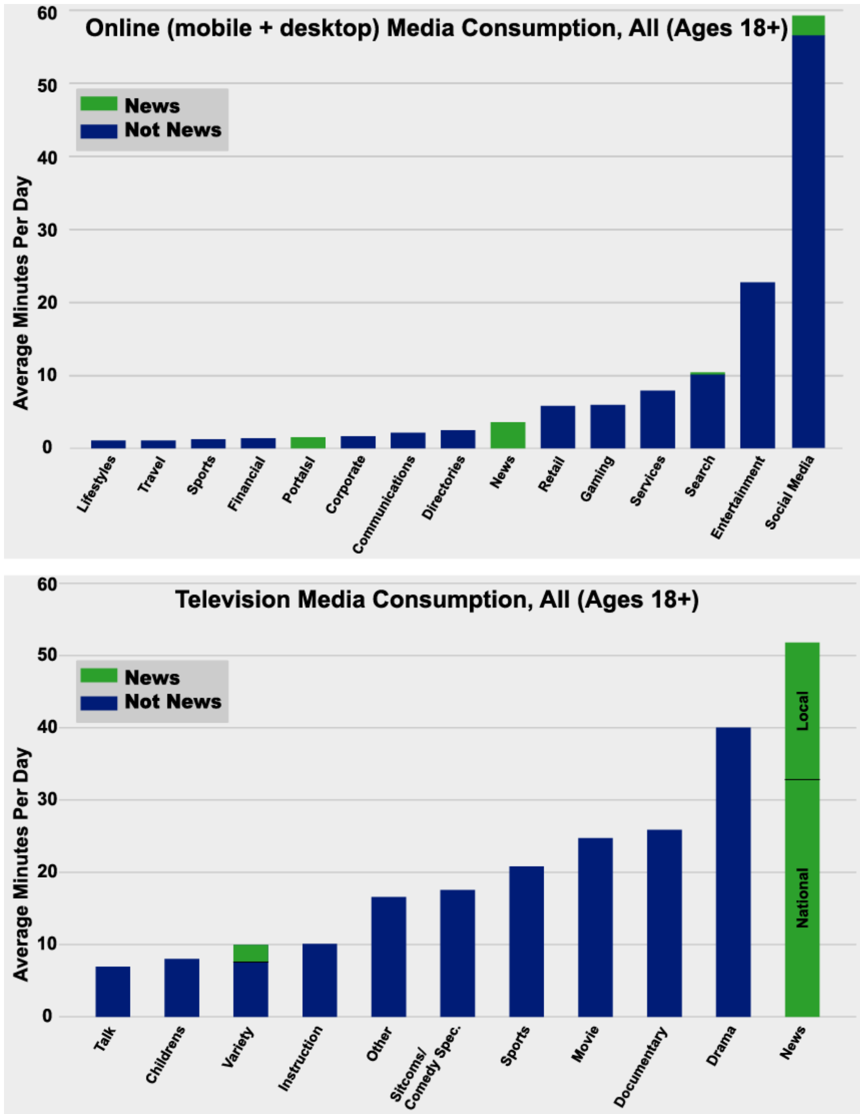

Figure 2. Detailed breakdown of overall media consumption (a) Online (including mobile and desktop) for the top 2000 sites/apps on ComScore, and (b) Television. Total online consumption is 227 mins per person per day, of which 58% is accounted for by the top 2000 sites, while total television consumption is 232 mins per person per day. To compute news consumption in Search and Social Media, excluding YouTube, we use the share of referrals from the site in question that redirect to news articles as a proxy for the share of time a user is exposed to news-related content on the platform. For YouTube, which does not redirect users to external sites, we randomly sampled 10,000 videos per month (weighted by viewing time) and computed the percentage that were classified "news and politics" Because portals such as MSN, Yahoo, and AOL almost always display some news-related stories on their landing pages, we count 100% of time spent on portals as news consumption. Finally, news consumption in the "Variety" category of television viewing is computed as 100% of time attributed to late night comedy programs, such as *The Daily Show with Trevor Noah*, which are known to contain commentary on politics and current events. For clarity 2A shows only the top 15 of 28 categories (see Table S7 for numerical values).

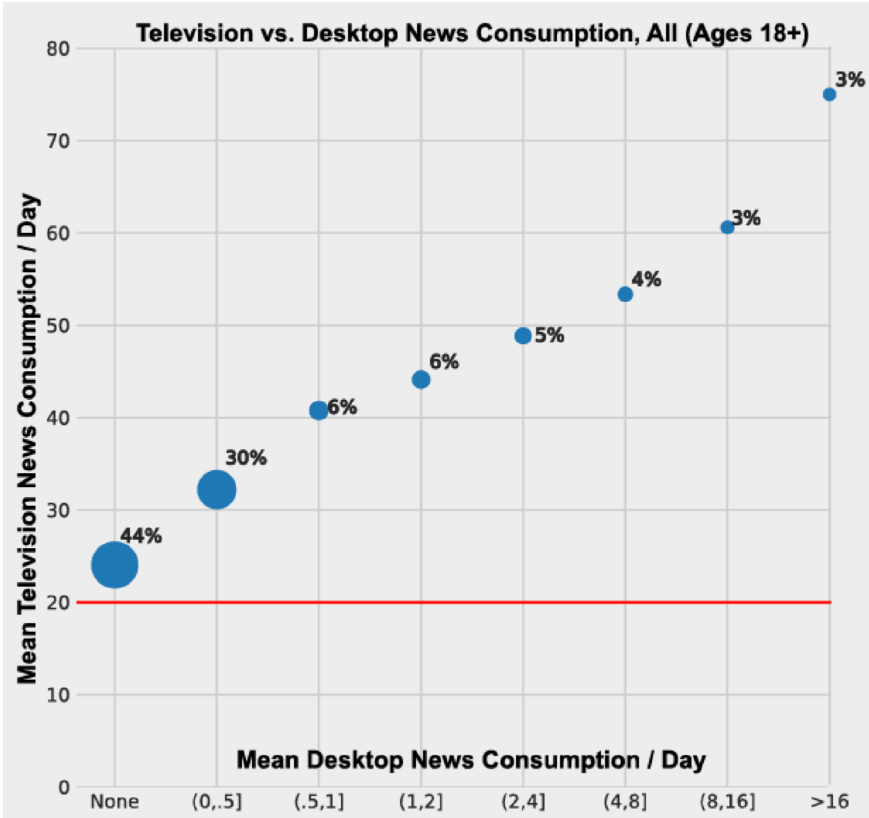

Figure 3. Television vs. desktop news consumption aggregated over all age categories 18-55+. For each month, the overlap panelists are separated into groups corresponding to different ranges of web news consumption. For each group, the mean television news consumption and group size as a percentage of all panelists are computed. Over-time averages for the mean television news consumption and size of each group are calculated by computing the mean television news mean and mean group size over all 36 months. Error bars are standard errors obtained via bootstrapping for group size and group television news consumption respectively, and are smaller than the symbols. See Fig. S2 for all results broken down by age group and Tables S8 and S9 for numerical values.

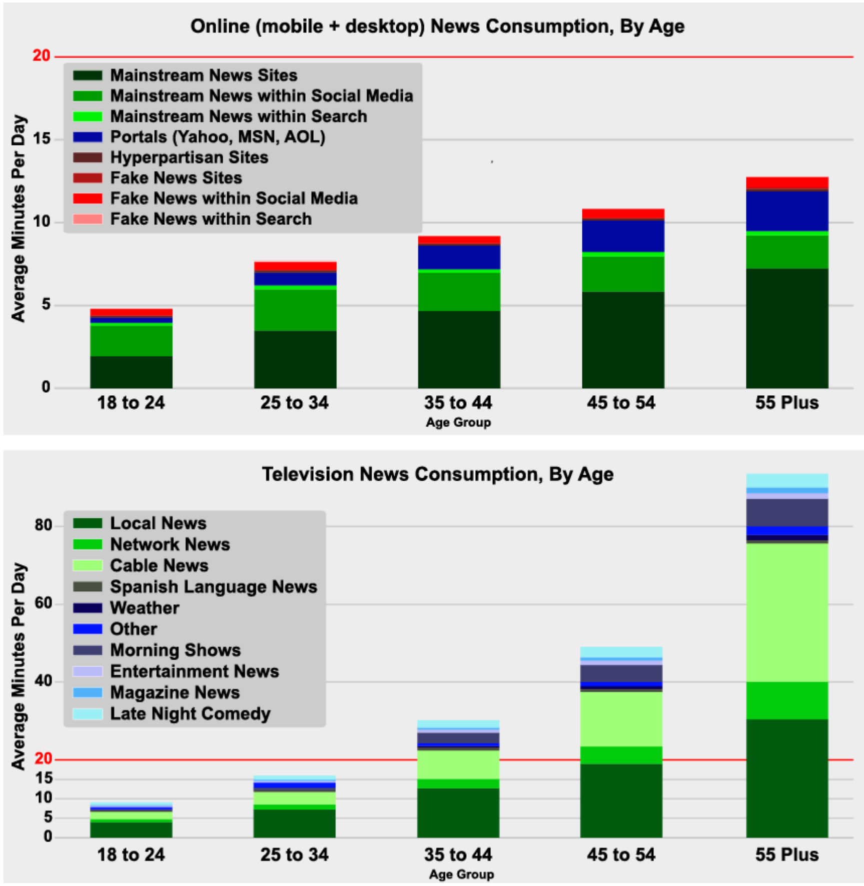

Figure 4. Detailed breakdown of news-only consumption by age group for (a) online (including mobile and desktop) and (b) television. See Figs. S4A and S4B for results plotted over time from Jan 2016-Dec 2018. See Table S11 for numerical values.

Evaluating the fake news problem at the scale of the information ecosystem

Supplementary Materials

I. Datasets

To measure Americans' consumption of media online and on television, we use data spanning 2016-2018, from two sources: Nielsen's individual-level television and desktop web panels and Comscore's aggregated app and browser (desktop and mobile) traffic data.

Overview

The Nielsen web panel was created in the early 2000s and tracks only desktop web traffic. It is constructed using a proprietary methodology which combines the industry-standard "Random Digital Dial" technique--wherein participating households are chosen for the panel in a minimally biased manner by randomly generating phone numbers--with higher-yield online recruitment (methodology guide).

The Nielsen television panels, which measure national television and local television viewing respectively, have existed since the 1950s. As with the web panel, the television panels are compiled using a proprietary methodology. We note, however, that television consumption metrics derived from the television panels are the media industry standard: Nielsen's ratings determine the price of advertising slots and are generally relied upon to gauge the commercial success of programs and networks (see https://en.wikipedia.org/wiki/Nielsen_ratings for more details). Panel participation is limited to two years and turnover happens on a continuous basis.

In addition to the Nielsen panels, we also use ComScore's aggregated digital traffic data. Like Nielsen for television, ComScore is a leading provider of digital traffic data, particularly of media and publisher data. Their competitive advantage is their proprietary "Unified Digital Measurement" method, which combines both census-based site analytics and panel-based audience measurement data to more accurately estimate digital consumption (described here). In addition to their desktop and mobile user panel, participating websites place tags on all their content – web pages, videos, apps and ads, that record calls by ComScore servers every time content is accessed. This combined approach allows them to validate their panel data with the census data, and vice versa, to more accurately estimate consumption.

Because ComScore data includes both desktop and mobile traffic we use ComScore for all of our main results. We note, however, that Comscore data is aggregated both over users and also over subdomain URLS, whereas the Nielsen web panel data tracks individual-level visits to unique URLs. For this reason, we use the Nielsen data to compute a number of intermediate results (see Table S1), which we then use to generate our main results.

Table S1. Datasets and their uses in this paper

Datasets	Uses in Paper
Nielsen National Television Panel Dataset	1. The set of national television program titles which count as news of any kind (Section II, Defining News Categories on Television) 2. Estimate of the average amount of time spent watching national news on television (Section III, Estimating News Consumption from national television and Section III, Estimating News Consumption from local television) 3. The set of all-content categories on national television (e.g. Feature Film, Situation Comedy) and their constituent program titles (Section IV, All-Content Category Consumption, Television) 4. Estimate of the total amount of time spent watching each all-content category of national television (Section IV, All-Content Category Consumption, Television)
Nielsen Local Television Panel Dataset	1. Estimate of the aggregate portion of all news that is local news (Section III, Estimating News Consumption from local television)
Nielsen Desktop Panel Dataset	1. Data for top 1000 news websites by traffic to track on ComScore (Section III, Estimating News Consumption via news websites) 2. Determining the percent referrals from social media and search sites to news websites (Section III, Estimating News Consumption on Platforms) 3. Identifying the set of YouTube videos watched by users of each age group that we sample based on time spent (Section III, Youtube) 4. Robustness checks on the ComScore data (Section III, Robustness)
Nielsen Overlap Television/Desktop Panel Dataset	1. Estimate of the correlation between individuals' time spent consuming news online (desktop) and time spent consuming news on television (national news) (Section V, Television News Consumption Conditional on Web News Consumption)
ComScore Aggregated Online Dataset	1. Data for aggregate time spent online (mobile and desktop) for the top news sites and fake news sites for estimating news consumption via news websites (Section III, Estimating News Consumption via news websites) 2. Data for time spent online on social media and search platforms for estimating news consumption on Platforms (Section III, Estimating News Consumption on Platforms) 3. Data for total online time spent and total digital population (mobile and desktop) (Section III, Estimating News

	Consumption) 4. List of top 2000 domains / entities that they track (Section IV: All Content Category Consumption, Online) 5. Set of 28 content categories and a list of top domains for each category (Section IV: All Content Category Consumption, Online) 6. Data for time spent online for the top 2000 domains / entities (Section IV: All Content Category Consumption, Online)

In the remainder of this section, we describe the Nielsen and ComScore panels in more detail, and also outline our weighting scheme.

Television panels: detail

- *Nielsen National Panel:* Nielsen's National Television panel consists of ~100,000 people in ~40,000 households. For each person, the dataset contains a log of all programs viewed, viewing session duration, time, and program channel. In each household, a "Nielsen box" is installed to each television that tracks, on a minute-by-minute basis, the program and station is being watched, including digitally recorded content, during national broadcasts. Information is tracked passively, except in multi-person households panelists must manually mark who is watching.
- *Nielsen Local Panel*: The Local Television panel consists of the subset of national panelists who live in the largest 25 television markets. For this group, consumption of local broadcasts (i.e., broadcasts from the local news stations that are affiliates of the four major networks) as well as national broadcasts is recorded; otherwise the measurement is the same as the national panel. The local panel contains ~50,000 people in ~20,000 households. We only have data from 2018 from this panel (see the section "Estimating News Consumption - Local News" for details on how local news consumption is imputed over the whole time period).

Online panels: detail

We use two different data sources to compute online consumption: the Nielsen user panel and ComScore's per-site metrics. Each method has its own advantages and disadvantages, but taken together they can give a unique and accurate picture of online information consumption.

- *Nielsen (Desktop Only)*: Nielsen's Desktop Web panel ranges in size from ~90,000 people in the beginning of 2016 to ~65,000 people by the end of 2018. The decrease in active users is approximately constant over time and corresponds to a systemic shift in the United States away from Desktop usage. Importantly, Nielsen consistently updates its weighting schema so that national-level projections remain accurate over time. In each panelist household, software is installed onto each personal computer that tracks, on a second-by-second basis, what URL is being actively visited. Information is tracked passively, except in multi-person households panelists must manually mark who is browsing. The dataset we have that is derived from this panel does not include all streaming data (Netflix, Hulu, etc.), which is tracked separately.
- *Comscore (Desktop and Mobile):* Comscore releases aggregated metrics for many websites and apps for both mobile and desktop. Most relevantly for our purposes, they track time spent and total viewers per month for most websites and apps for both mobile

and desktop. They also have the same data broken down by age group. For mobile, their estimates include both browser and app use, including that of news aggregator apps like Apple News and most apps from large publishers like the NYTimes app. However, unlike the Nielsen data, they do not share the individual level browsing data, nor do they have data for many of the smaller websites that account for very little traffic.

Television and Online: details

There is an overlap between the Nielsen Television and Web panels comprising 10,000-15,000 people, depending on the month. For this subset of people, we can observe both desktop web consumption behavior. This group is not crafted to be representative of the US as a whole.

Weighting:

The consumption data from Nielsen (Television & Desktop Web) is individual-level and the individuals are a subset of the US population. For aggregate consumption statistics derived from this data, we weight each individual panelist's contribution according to whether their demographics are over or under represented in the panel with respect to the distribution of demographics in the USA.

The weights used in our analysis are calibrated using a technique called iterative proportional fitting (IPF) *(26).* The aim of IPF is to adjust the values of a contingency table so that the marginal distribution over each variable in the adjusted table matches a specified target distribution. Cell values are adjusted one variable at a time (i.e. first each row, then each column, etc.), and the process is repeated until some specified convergence or maximum iterations is reached. In our case, the contingency table represents the panel in terms of the individual-level attributes and the target distributions are based on the Census values of the populations. Finally, the weight for a given panelist is defined as the ratio of the his or her corresponding cell's value after adjustment to its value and before adjustment.

Nielsen uses IPF to provide each panelist a weight targeting accurate demographic and behavioral metrics (Nielsen methodology guide provides a high level summary). Panelists in the overlap group have separate weights for the television and desktop web panels. We also create a third set of panelist weights for the overlap panel, for which Nielsen does not ensure representativeness and does not provide weights. We target correct distributions over age, sex, race, education level, and a binary *hispanic* variable, according to the 2017 US Census micro-level data.

Using these weights makes a negligible difference for the aggregated metrics we report.

II. Defining News Categories

Television news

We assign all television programs to one of three categories: *hard news*, *soft news,* and *not news*. *Soft news* generally mixes news and entertainment and includes programs such as *Good Morning America*, which devote some but not all of their time to news as well as programs like *The Daily*

Show with Trevor Noah, that are not "news" programs per se, but which are nonetheless popular with some viewers on account of their commentary on the news. We consider hard news only as a lower-bound of mainstream news consumption and hard news plus soft news as an upper-bound.

To classify each program, we take as a starting point the set of nearly 4,000 programs which were labeled NEWS by Nielsen over the three years of our analysis. We note, however, that Nielsen's naming convention often assigns many *program titles* to what is in reality the same program (e.g. the program Good Morning, America appears as Good Morning, America, GMA, and Good Morning, America-SUN, among others); thus the real number of programs is lower. To estimate the number of unique programs considered NEWS by Nielsen, we sampled one thousand programs and ordered them alphabetically to reveal at least some of the program title duplicates (e.g. *Good Morning, America* and *Good Morning, America--SUN* are adjacent after ordering). We found that 885 of the 1000 sampled titles were duplicates, leading us estimate that there are roughly 400 unique news programs in the dataset. Moreover, many of these are one-off news specials (e.g. the inauguration of a president), thus an even smaller number of regular programs accounts for the vast majority of viewing time.

Soft News

To distinguish *soft news* programs, we manually looked at a large subset of the programs marked as NEWS (the subset accounts for >90% of total consumption of the NEWS) and identified morning news, magazine news, and entertainment news as the three broad categories of soft news. We then referenced the lineups of the major networks from 2016-2018 and identified which programs fell into these categories--see table below for details. Finally, we developed string-matching heuristics to identify the various program names which resolve to the determined *soft news* programs. We also count as *soft news* programs from the weather channel and several late night shows which are marked by Nielsen as "General Variety", but are often considered a significant source of news for people (*20).*

Hard News

We define as *hard news* all programs marked as NEWS by Nielsen that we do not explicitly mark as *soft news*.
We define as *not news* all programs not marked as *soft news* or *hard news*. Table S2 describes the full set of categories of news in terms of which our analysis of television news (Figure 4B) is presented.

Table S2. Classification scheme for television news

Category	**Constituent Programs**	***Soft* or *Hard* News**
Magazine News	*Inside Edition, Right this*	Soft

	Minute, *Dateline*, *EXTRA*, *Insider*	
Soft Morning Shows	*Good Morning*, *America*, *Today Show*, *America this Morning*	Soft
Entertainment News	*TMZ*, *Access*, *Access Hollywood*, *Entertainment Tonight*, *Made in Hollywood*, and all programs designated as NEWS on the network E!	Soft
Weather	All weather programs on the network The Weather Channel	Soft
Late Night Shows	*Jimmy Kimmel Live*, *The Late Show with Stephen Colbert*, *The Late Late Show with James Corden*, *The Tonight Show with Jimmy Fallon*, *Late Night with Seth Meyers*, *Last Call with Carson Daly*, *Saturday Night Live*, *The Daily Show with Trevor Noah*	Soft
Hard Network News	All programs designated as NEWS on the networks ABC, CBS, NBC, and FOX *not* categorized as *soft news*	Hard
Spanish Language	All programs designated as NEWS on the networks Galavision, UniMas, Univision, Telemundo, Azteca America, Estrella, CNN en Espanol, and WAPA America.	Hard
Cable	All programs designated as NEWS on the networks MSNBC, Fox News, CNN, Fox Business, and CNBC.	Hard
Local News	All programs designated by Nielsen as Local News*	Hard

Other	All programs designated as NEWS that don't fall into any category. For example, news programs on the network PBS and presidential inauguration specials on traditionally non-news networks.	Hard

Online news

We categorize online news at the domain (i.e. publisher) level along two different dimensions: hard vs. soft news and fake vs. mainstream news. The distinction between hard and soft news depends on the category of news; hard news publishers primarily cover topics like politics, business, and U.S. and international affairs while soft news publishers cover topics like entertainment, sports, and lifestyle news. Fake vs. mainstream news refers to the legitimacy of the publisher and the extent to which the news publisher upholds journalistics standards of objective, fact-based reporting. Example stories and sources are given table S3.

Table S3. Online news classification

	Mainstream	**Fake**
Hard	*Headline*: "Trump Delays Auto Tariffs in Press for Deal With Japan, Europe" *Sources*: nytimes.com, wsj.com	*Headline:* "FBI Says No One Killed at Sandy Hook" *Sources:* infowars.com, dailywire.com
Soft	*Headline*: "'The Big Bang Theory' showrunner answers burning questions about the series finale" *Sources*: ew.com, espn.com	*Headline*: "R. Kelly Visits Bill Cosby In Prison for Advice" *Sources:* huzlers.com, viralhatch.com

For the purposes of this paper, we restrict attention to hard mainstream news; however, we include both hard and soft fake news on the grounds that (a) it is more difficult in practice to differentiate between purveyors of hard and soft fake news than for mainstream sources, and (b) previous research on fake news also fails to differentiate between them. For example, Buzzfeed's 2018 investigation of fake news sites *(27)* identifies the *fake soft* sites huzlers.com and worldnewsdailyreport.com as the top disseminators of fake news on Facebook, accounting for a combined 18 / 50 top fake news stories. Fortunately the inclusion of soft news in our fake category only increases our estimates of the prevalence of fake news, and hence strengthens our conclusions.

Mainstream News

As noted above, we classify as mainstream online news websites that cover primarily "hard" news topics like politics, business, and U.S. and international affairs as mainstream news. First, we take a list of mainstream news domains adapted from Athey, Mobius, and Pal (2017), which used online raters to classify whether a domain was a news site and if so, what category of news. We then add a list of the most prominent European and Asian news sites taken from Wikipedia. We filter out "soft news" sites that were classified as "sports", "entertainment", "technology" or "specialty" (e.g. espn.com, ew.com). The resulting list of 9798 domains includes the websites for print newspapers and magazines (e.g. nytimes.com, time.com, theguardian.com), digital-only news sites (e.g. Vox.com), local news sites (fox5ny.com), as well as the news landing pages of several large portal sites (e.g. yahoo.com/news). We do not exclude the "soft news" subpages of websites that primarily cover hard news (e.g. nytimes.com/section/sports) since in most cases we only have data for time spent on the parent domains. However, since we are primarily interested in determining the upper bound of news consumption online this inclusion of possible non-news content is not a major concern.

We use this full list of 9798 mainstream news sites in order to estimate the amount of news on platforms like social media and search (see the section "Estimating News Consumption on Platforms"). Due to technical limitations with the ComScore dataset, we use a truncated list of 806 of the top mainstream news sites, accounting for over 90% of traffic, to estimate the amount of news consumed directly on news websites. See the section "Estimating News Consumption on News Sites" for more details.

The intermediate and final lists of news sites can be found on the OSF project website (https://osf.io/cygta/).

Fake news

Our definition of fake news builds on three different sources:

1. *Academic Literature (Grinberg et al, 2018)* First, we take the three lists-- black, red, and orange-- laid out in Grinberg et al:
    a. The "black" list refers to domains that published entirely fabricated stories. The list itself is comprised by combining lists of fake news domains that have appeared on fact-checking websites and other academic publications including Politifact, FactCheck.org, Guess et al.
    b. The "red" and "orange" lists were created by identifying websites that appear frequently in the Snopes archives and then manually rating them based on the sites "propensity to elicit fake news and information." Sites were rated "red" if they had "little regard for the truth" and "orange" if they were "negligent or deceptive"
2. *NewsGuard:* We also use site ratings from the organization NewsGuard, an organization started by veteran journalists to rate sites based on journalistic principles of credibility and transparency. NewsGuard hires trained raters with experience in journalism to rate news and information sites based on nine criteria of credibility and transparency, each of which have a different point value based on the importance of the category. The total point value of all the categories is 100. If a site meets enough of the criteria to earn 60 points, it is given a "green" rating by NewsGuard. If the site merits a score of less than 60 points, it is given a "red" rating. We take all sites rated by NewsGuard as "red", which signals that the site does not adhere to journalistic standards of credibility and

transparency. A more detailed description of NewsGuard's rating system can be found here on its website.

3. *Buzzfeed:* Finally, we add the domains identified by the Buzzfeed News annual investigation of the most popular fake news sites on Facebook. While not an academic or traditional fact-checking organization, Buzzfeed's work *(27, 28)* is updated regularly, frequently cited by academics studying fake news, and commands the attention of the popular press which is why we have chosen to include them.

Using the three sources described above, we differentiate between (a) "truly" *fake news,* defined as sites (e.g. breaking-cnn.com, infowars.com) that are mostly or wholly fabricated and deliberately masquerading as legitimate news sources and (b) *hyperpartisan/low-quality news*, which refers mostly to hyperpartisan sites (breitbart.com, dailywire.com) that publish some amount of highly misleading or outright false material. The *fake news* category has 549 domains total and is composed of domains in the "red" and "black" lists from Grinberg et al and the domains from Buzzfeed's list. In addition, we add the "red" sites from NewsGuard that fail to meet the criteria "Doesn't repeatedly publish false content," the most stringent requirement. The *hyperpartisan/low-quality news* has 93 domains total and is composed of domains in the "orange" list from Grinberg et al as well as all "red" sites from NewsGuard that do not fall into the strict *fake news* category, but nonetheless lack other standards for journalistic integrity, like "Gathers and presents information responsibly" and "Avoids deceptive headlines."

There has been disagreement over whether or not these low-quality or hyperpartisan sites should be included in definitions of "fake news", since many of them do not publish outright fabricated stories or might only report falsehoods as a small percentage of their overall content. However, they do not uphold typical standards of journalistic integrity and can lead readers to incorrect conclusions nonetheless. Reflecting this disagreement, we compute upper and lower bound estimates of fake news prevalence, where the upper bound includes "hyperpartisan/low-quality" news consumption and the lower bound includes only the "truly" fake news websites. All results in the main text reflect our upper bound definition, while Fig. S3 and Table S10 report both upper and lower bounds.

Similar to our process for mainstream news, we use this full list of 642 fake news domains in order to estimate the amount of news on social media and search platforms. However, because ComScore does not track all fake news outlets due to their low overall traffic, we use a truncated list of 98 of the most-trafficked fake news sites accounting for over 80% of traffic to estimate fake news consumption via fake news websites. See the section "Estimating News Consumption on News Sites" for more details.

The final list of fake sites and their categorizations (fake, hyperpartisan/low-quality) can be found on the OSF project website (https://osf.io/cygta/).

Table S4. Classification scheme for Online News

Category	**Constituent Websites**
Mainstream News	Example domains:

	Cnn.com, foxnews.com, msn.com/en-us/news, nytimes.com, washingtonpost.com
Fake News	Example domains: conservativetribune.com, wnd.com, bipartisanreport.com, thegatewaypundit.com, beforeitsnews.com, thepoliticalinsider.com, rt.com
Hyperpartisan or Low Quality News	Example Domains: Dailykos.com, breitbart.com, worldstarhiphop.com, theblaze.com, dailycaller.com, rushlimbaugh.com

III. Estimating News Consumption

Television:

Because the local television panel contains a subset of the people in the national television panel, we estimate the aggregate time spent watching local and national news separately and then sum estimates to estimate the total.

National News

For all nationally broadcast programs, we measure the time spent watching news on television by simply tabulating the average time spent by the national panel watching *hard news* and *soft news* program.

Local News

The local television panel consists of the subset of national panelists who live in the twenty-five largest television markets. Instead of calculating the average time spent by this subset watching local news and assuming it holds for the entire country, we calculate the portion of the local panelists' total television news diet made up of local news, and assume this percentage holds for the entire country.

Specifically, first, we calculate the portion of news consumption that is local news for the local panelists.

$$P_L = \frac{local_L}{local_L + national_L}$$

Where $local_L$ is the average time spent watching local news by the local panel and $national_L$ is the average time spent watching national news--hard or soft--by the local panel. Second, we solve the same equation for the national panel assuming $P_L = P_N$:

$$local_N = P_L \times \frac{national_N}{1 - P_L}$$

We calculate the estimated $local_N$ conditional on each age group. In other words, a given age group's local news estimate is a function of that age group's calculated portion of news consumption that is local news P_L, not the whole panel's. Since we have local news data for 2018 only, we compute a single P_L based on the entire year and assume this is constant from 2016-2018.

Lower and Upper Bounds

The lower and upper bounds of television news consumption follow directly from the classification of each program as *hard news, soft news,* or *not news*. We define the lower bound as the time spent watching programs marked as *hard news* (i.e. programs categorized as Local News, Cable, Hard Network News, Spanish Language News, and Other). We define the upper bound as the time spent watching programs marked as *hard news* or *soft news* (i.e. programs categorized as Magazine News, Soft Morning Shows, Weather, Entertainment News, and Late Night Shows). As in the discussion of fake news, all results in the main text use the upper-bound estimate of television news consumption; however, we have also checked that our conclusions--including the finding that every age group consumes more news on television than online--hold when using the lower bound definition of television news.

Online:

The most common way of measuring online news consumption is by measuring the time spent on *news sites* like nytimes.com or aol.com/news. However, other websites that are not specifically news-related often feature news content, like a news article that appears on someone's Twitter feed. We call this phenomenon news consumption on *platforms*, and we attempt to measure it along with the more straightforward news consumption on news websites. However, we only include time spent on reading news platforms in the "Upper Bound" of news and not the "Lower Bound", since time spent exposed to news on platforms differs from news consumption in the traditional sense.

Estimating News Consumption on News sites

Since the ComScore panel includes both mobile and desktop data, we use it as our primary source for measuring online news consumption. However, ComScore does not have complete web traffic for every domain in our dataset, especially the less-trafficked ones. Moreover, ComScore limits the number of domains that one user can track at a given time.

Since we can only track a limited number of domains via ComScore, we truncate our mainstream media list to only the top 1000 trafficked mainstream media domains. We do this by calculating the average time spent on each mainstream news site by panelists in our Nielsen desktop web panel, ranking each site by time spent, and taking the top 1000, which accounts for 95.1% of total mainstream news traffic according to Nielsen. We then attempt to find a match for these domains within ComScore. While we only get a match for 806 / 1000 domains, since ComScore does not track domains that are very low-trafficked, we are able to find a match for domains that account for 91.5% of Nielsen desktop traffic.

For fake and hyperpartisan/low-quality sites, since we have only 642 sites total, we do not need to truncate our list to the top domains. We simply attempt to find a match for all sites on ComScore and successfully find a match for 98 / 642 sites, accounting for 81.0% of fake news desktop traffic according to Nielsen. While this share is lower than that of mainstream news, since the majority of fake news consumption that we estimate takes place on platforms like social media (described in the next section "Estimating News Consumption on Platforms"), it does not make a meaningful difference in our final upper bound estimate of fake news (see Robustness section below).

Then, we use the ComScore interface to get the total monthly minutes spent on each site for U.S. users 18+ on both mobile and desktop, respectively, for each month from Jan 2016 - Dec 2018. Additionally, for users 18+, we get the total monthly minutes spent on *all* sites as well as the total number of users (both desktop web and mobile) for each month. We repeat this process for users in the following age buckets: 18 - 24, 25 - 34, 35 - 44, 45 - 54, 55+.

We then use this data to calculate the daily average minutes spent per person per month for each site in our ComScore list.

Estimating News Consumption on Platforms

In most cases, we are able to measure news consumption by adding up the time spent on news-specific sites accessed via news websites, like nytimes.com or yahoo.com/news. However, we also know that people consume a significant amount of news online on sites that are not news specific, including social media (e.g. Facebook), search (e.g. Google), and portals (e.g. Yahoo Homepage). Unlike regular websites, where each article is accessed via a unique URL, all activity internal to these platforms occurs on the same domain-level URL (e.g. facebook.com). As a result, the Nielsen data only allows us to observe news exposure on these platforms when a user clicks on a story and exits the platform to read the story at the original URL. If a user simply reads the snippet in their newsfeed or if the article is hosted directly on the site (as with Facebook's instant articles feature) we would see only that the user spent time on facebook.com, not what portion of that time was devoted to news. Since we cannot measure news consumption on these portals, social media, and search directly, we estimate the upper bound of news content on these sites using the below methods, respectively.

Portals

AOL, Yahoo, and MSN homepages are classified as "portals" by virtue of their historical roles as the entry point to the web. While people use them for a variety of reasons like email or search, external news stories do appear on homepages that could expose readers to news headlines. Since we cannot determine what percent of time on the homepage is spent paying attention to news, we take the most conservative possible estimate and decide to consider all time spent on portal homepages as "mainstream news" as an upper bound on news consumption.

Social Media

Facebook, Twitter, and Reddit. We estimate on-platform exposure to mainstream and fake news (including hyperpartisan/low-quality) as follows.

1. For each over 18 user in the Nielsen web panel, we identify all visits the user made to a social media site and note their age bracket (18 - 24, 25 - 34, 35 - 44, 45 - 54, 55+)
2. For each social media site visit, we look at the site that the user visited immediately after (within a 1 minute window). We exclude all visits to a homepage (e.g. nytimes.com rather than a link to a particular nytimes article), search client, or email client, since they are more likely to be new browsing sessions rather than referrals from a social media site.
3. If a user goes from a social media site to a mainstream news site or fake news site (including hyperpartisan sites) we count that as a mainstream news referral and fake news referral, respectively. If the user goes to any other non-news site, we count that as an non-news referral.
4. We then group the referrals by age bracket and sum the number of mainstream news and fake news referrals over all users in that bracket and divide by the sum of all referrals to calculate the mainstream news and fake news referral rate, respectively, for each social media site and age bracket
5. We use these referral rates to calculate two numbers, "Social Mainstream News" and "Social Fake News", by multiplying the time spent on each social media site by the proportion of mainstream and fake news referrals, respectively. The equation is given below, where *d* is a domain in our set of social media domains, $ts_{d,a}$ is the average time spent per user of age bracket *a* on that domain *d*, and $r_{fake,d,a}$ and $r_{mainstream,d,a}$ are the referral rates for fake news and mainstream news sites, respectively, for domain *d* and age bracket *a.*

$$Social\ Mainstream\ News = \sum \quad ts_{d,a} * r_{mainstream,d,a}\ where\ d\ \in Social\ Media$$

$$Social\ Fake\ News\ = \sum \quad ts_{d,a} * r_{fake,d,a}\ where\ d\ \in Social\ Media$$

We note that this method likely overestimates on-platform exposure to real and fake news, for four reasons. First, the trends in news consumption we see in our referral data closely match the trends seen in other papers of news on social media. Second, although we know that only a portion of content on social media contains a link of any kind, we use the percentage of referrals from social media to news sites as a direct estimate of the percentage of news consumption on social media, i.e. if 10% of the referrals from Facebook are to news sites, we estimate that 10% of the total time spent on Facebook is consuming news. Thus, while it is certainly possible that we are overestimating news consumption on Facebook, it is highly unlikely that we are underestimating it. Third, the clickbait nature of fake news suggests that it would be overrepresented in clicks versus time spent, further adding to the unlikeliness of our undercounting fake news. Recent research has suggested that conditional on appearing in feed, users are more likely to engage with fake news rather than real news (Vosoughi, Roy, & Aral, 2018). Thus, using clicks on fake news, which is designed in most cases to attract quick clicks rather than deep engagement, as a proxy for time spent biases the data toward fake news over mainstream and non-news. Finally, we note that even if in spite of these precautions we are underestimating--rather than overestimating--news consumption on Facebook, our error would have to be very large in order to alter our qualitative findings. Given that all of our other robustness checks (e.g. changing the restrictiveness of the category for fake news) yield very little appreciable difference in consumption time, we find it extremely unlikely that we have somehow missed a large and substantively meaningful amount of news consumption on Facebook.

YouTube. YouTube hosts all videos on its own platform, but each video has a unique URL that we can query details from, so we do not need the referral method: we can directly compute mainstream and fake news consumption:

1. Using the Nielsen panel data for each month Jan 2016 - Dec 2018, we identify all visits by panelists to youtube URLs, the view time of that visit, and the age bracket of the panelist (18 - 24, 25 - 34, 35 - 44, 45 - 54, 55+ -- we exclude views by panelists below 18)
2. We then group YouTube visits by age bracket and calculate the overall view time across all panelists of that age bracket for each individual YouTube URL. This gives the list of all YouTube videos watched by viewers of each age bracket in the form *age_bracket, URL1, total view time*
3. Next, for each age bracket separately, we randomly select a sample of 10,000 videos URLs weighted by view time, without replacement
4. Using the YouTube public API, we pull metadata for each video including video title, the channel name and video category. The video category is selected by the video uploader from the following categories: "Film & Animation", "Autos & Vehicles", "Music", "Pets & Animals", "Sports", "Short Movies", "Travel & Events", "Gaming", "Videoblogging", "People & Blogs", "Comedy", "Entertainment", "News & Politics", "Howto & Style", "Education", "Science & Technology", "Nonprofits & Activism", "Movies", "Anime/Animation", "Action/Adventure", "Classics", "Comedy", "Documentary", "Drama", "Family", "Foreign", "Horror", "Sci-Fi/Fantasy", "Thriller", "Shorts", "Shows", "Trailer".
5. For each month, we calculate the percent of videos for each age bracket that are designated in the “News & Politics” category. We consider this percentage as the overall percentage of news, mainstream and fake, on YouTube.

Age Bracket	**Percent News (by watchtime)**
18_24	2.1%
25_34	2.8%
35_44	2.7%
45_54	4.1%
55plus	5.3%

6. In order to calculate the percentage of these news videos that were fake and mainstream, we then take a time-weighted sample of 100 news videos from the month of Dec 2018 (which are the most current and thus most likely to not have been taken down). Researchers labeled the videos into the following categories:
    a. Mainstream News: channels of legitimate news outlet (CNN, ABC) or channel that only reposts mainstream news
    b. Junk News: channels of existing fake / hyperpartisan news outlets (RT, Ben Shapiro) or clearly fake news or conspiracy theories (e.g. Hillary Clinton gets indicted at George H.W. Bush's funeral)
    c. Youtuber: YouTube bloggers that put out commentary on news, current events, culture, politics etc. (e.g. Philip DeFranco, Secular Talk)
    d. Other YouTube News: YouTube-only news channels that produce clips for YouTube only. Focus mostly on human interest and non-political news, e.g. Barcroft television
    e. Citizen News: Video footage of news events taken by ordinary citizens (e.g. video of a hurricane)
    f. Non-News: Videos that were miscategorized as news by the uploader (e.g. a music video)

citizen news	1
junk news	14
Mainstream news	44
not news	13
Other youtube news	9
youtuber	19

7. We group "citizen news", "mainstream news" and "other youtube news" into the overall "mainstream news" category and "junk news" and "youtuber" into the "fake news" category. We then calculate the proportion of mainstream news on YouTube as 54%, fake news as 33%, and not news as 13%.
8. For each month and age category, we then calculated the proportion of mainstream and fake news on YouTube by the following method:
    a. Mainstream News Proportion = % of monthly overall news Youtube consumption * 0.55 (the percentage of overall news that was mainstream news)
    b. Fake News Proportion = % of monthly overall news Youtube consumption * 0.33 (the percentage of overall news that was fake news)

origin	**Average Proportion of Overall**	**Average Proportion of Overall Content,**

	Content, Mainstream News	Fake and Hyperpartisan/Low-Quality News
facebook.com	6.94%	1.40%
reddit.com	4.72%	0.30%
twitter.com	9.54%	0.85%
youtube.com	1.67%	1.02%

Search

Similar to social media news, we estimate an upper bound of news and fake news consumption on social media from the Nielsen web panel data by calculating the proportion of referrals from Google, Bing and Yahoo to mainstream and fake news (including hyperpartisan/low-quality) sites for each age group (18 - 24, 25-34, 35-44, 45-53, 55+), respectively using the Nielsen data. We compute these referral statistics by the following method:

1. For each over 18 user in the Nielsen web panel, we identify all visits they made to a search site and note their age bracket (18 - 24, 25 - 34, 35 - 44, 45 - 54, 55+). Then, for search site visit, we look at the site that the user visited immediately after (within a 1 minute window). We exclude all visits to a homepage (e.g. nytimes.com rather than a link to a particular nytimes article), since those are considered "navigational searches" and do not consistently highlight a set of news stories in the same way as other news searches.
2. We then calculate the mainstream news and fake news referral rate in the identical manner as for social media referrals.
3. We again use these referral rates to calculate two numbers, "Search Mainstream News" and "Search Fake News", by multiplying the time spent on each search site by the proportion of mainstream and fake news referrals, respectively, for each age bracket. Note that we only consider at search traffic for Google, Yahoo, and Bing and not all traffic to the parent websites. Time spent on yahoo.com homepage would NOT be counted as search time spent (and would in fact be counted in the "portals" category), but time spent on yahoo.com/search would be counted. The equation is given below, where d is a domain in our set of search, $ts_{d,a}$ is the average time spent per user of age bracket a on that domain d, and $r_{fake,d,a}$ and $r_{mainstream,d,a}$ are the referral rates for fake news and mainstream news sites, respectively for domain d.

$$Search\ Mainstream\ News = \sum ts_{d,a} * r_{mainstream,d,a}\ where\ d \in Search$$

$$Search\ Fake\ News = \sum ts_{d,a} * r_{fake,d,a}\ where\ d \in Search$$

origin	Average Referral Rate, Mainstream News	Average Referral Rate, Fake and Hyperpartisan/Low-Quality News
google.com/search	4.44%	0.20%
yahoo.com/search	3.02%	0.19%

bing.com/search	4.08%	0.17%

Upper and Lower Bounds, Mainstream News

Mainstream News Lower Bound: The time spent on sites in the mainstream news category accessed via news websites

Mainstream News Upper Bound: The time spent on mainstream news accessed via news websites as well as the time spent on mainstream news on platforms including portals, search, and social media, estimated via the method described above.

Upper and Lower Bounds, Fake News

Fake News Lower Bound: The time spent on sites in the “truly” fake news category (see section *Online news-Fake news* for definition) accessed via news websites

Fake News Upper Bound: The time spent on sites in the fake news and hyperpartisan news categories accessed via news websites as well as the time spent on fake and hyperpartisan news on search and social media platforms, estimated via the method described above.

We present this information of online news consumption as part of figures 1 and 4. In figure 1, we show the upper bounds of both news and fake news. The graphs calculated with lower bounds can be found in the appendix in figure S3 , but do not change substantively. In figure 4, we show a breakdown of time spent on sites in each of our five categories, bucketed by age.

Robustness

As a robustness check, we use the Nielsen panel to calculate the time spent on sites in the full list of news domains, fake and mainstream, and add the time spent news on social media and search estimated using the methods described above. We then do the same process, but using the truncated list of domains for which we were able to find ComScore matches. Using this method, we find that using only the domains for which we have ComScore matches accounts for 94.9% of fake news (upper bound) and 97.6% of mainstream news (upper bound) desktop traffic in the Nielsen panel.

IV. All-Content Category Consumption

Television:

In addition to classifying news content on television, we also examine the non-news content. Nielsen maps each program to a general content category, including the NEWS category we use in our definition of news on television. We create our own categories, which are simply the combination of one or more of Nielsen’s categories, and tabulate the total time spent watching each category. Table S5 shows how we create our all-content categories from Nielsen's categories.

Table S5. Classification scheme for non-news television content

All-Content Category	**Constituent Nielsen Categories**
Childrens	Child Multi-Weekly, Children's News, Child - Day Animation, Child - Live,
Comedy	Comedy Variety, Situation Comedy
Documentary	News Documentary, General Documentary
Drama	Daytime Drama, General Drama, Western Drama, Private Detective
Instruction	Instruction, Advice
Movie	Feature Film
News	News, Political
Other	Official Police, Award Ceremonies, Unclassified, Devotional, Concert Music, Evening Animation, Quiz Panel, Quiz Give Away, Science Fiction, Format Varies, Adventure, Popular Music, Suspense/Mystery, Audience Participation
Sports	Sports Commentary, Sports Anthology, Sports Event, Sports News
Talk	Conversations, Colloquies
Variety	General Variety*, Participation Variety

*We count several programs marked as Variety as news in our estimation of news consumption on television

Online:

In addition to classifying news content online, we also examine the non-news content. ComScore provides lists of domains in a variety of categories: *automotive, business-to-business, career, corporate, directories, education, entertainment, family, finance, financial, gambling, gaming, government, health, info, isp, lifestyles, news, portals, real estate, retail, search, services, social media, sports, technology, telecommunications, travel.* In addition, ComScore provides a list of the top 2000 domains. We use this information to classify the top 2000 domains in the following way.

1. If the domain is in the “mainstream”, “fake”, or “hyperpartisan / low-quality” categories (as identified in the “Defining News Categories” section), we classify the domain as “news”
2. If the domain is in the “social media” or “portal” category (as identified in the “Defining News Categories” section), we classify it as its respective category
3. If the domain is found on a category list provided by ComScore, we classify it as that category. Domains that are in ComScore’s “news-info” category that are not on our list are reclassified as “info” after confirming that they do not contain news.
4. If the domain is not found on a list provided by ComScore (66/2000), the researchers agree on a category for that domain.

The final list of the top domains and their categories can be found here: https://osf.io/cygta

V. Television News Consumption Conditional on Web News Consumption

To investigate how online news consumption is related to television news consumption on an individual level, we measure both modes of news consumption for each of the people that are in Nielsen’s web and television panels (the overlap group is between 10k and 15k people depending on the month). We can not use the Comscore data or imputations of aggregate metrics for this, so we modify the measurements of news consumption as follows:

Television News Consumption:
Television news consumption for each individual panelist in the overlap group is calculated the same as is described in the section *Estimating News Consumption-Television* for the *upper bound*, except we exclude that which we don’t have individual-level data for: local news consumption.

Online News Consumption:
Online news consumption for each individual panelist in the overlap group is defined as the average time per day spent on News Sites (mainstream, hyperpartisan, and fake), Portals, Social Media, and Search combined. We estimate these nearly as described in *Estimating News Consumption-Online*, with three differences corresponding to the need to estimate individual-level consumption. First, as noted above, we use the Nielsen online panel instead of comscore, so only online consumption via desktop is estimated. Second, we adapt the measure of *Social Media* news and *Search* news to apply to individuals: an individual's imputed time spent consuming news on the social media platforms and search is the product of their time spent on the social media or search site and the proportion of their referrals from that site to mainstream and fake news sites. Third, the estimate of news on Youtube is entirely excluded.

To measure the relationship between television and online news consumption, for each month, the overlap panelists are separated into groups corresponding to different ranges of web news consumption. For each group, the mean television news consumption and group size as a percentage of all panelists are computed. Over-time averages for the mean television news consumption and size of each group are calculated by computing the mean television news mean and mean group size over all 36 months.

Standard errors for group size and group television news means are obtained via bootstrapping, i.e., we create many replicates of the data by sampling an equally sized dataset with replacement from the original set of 36 monthly overlap panels and repeat the described process for each replicate; standard errors are then the standard deviation of the resulting distribution. As with all aggregated consumption metrics derived from the Nielsen data, we use panelist demographic weights, which in this case of bootstrapped sample replicates are slightly adjusted so that the mean weight is one for each sample replicate.

VI. Supplementary Figures

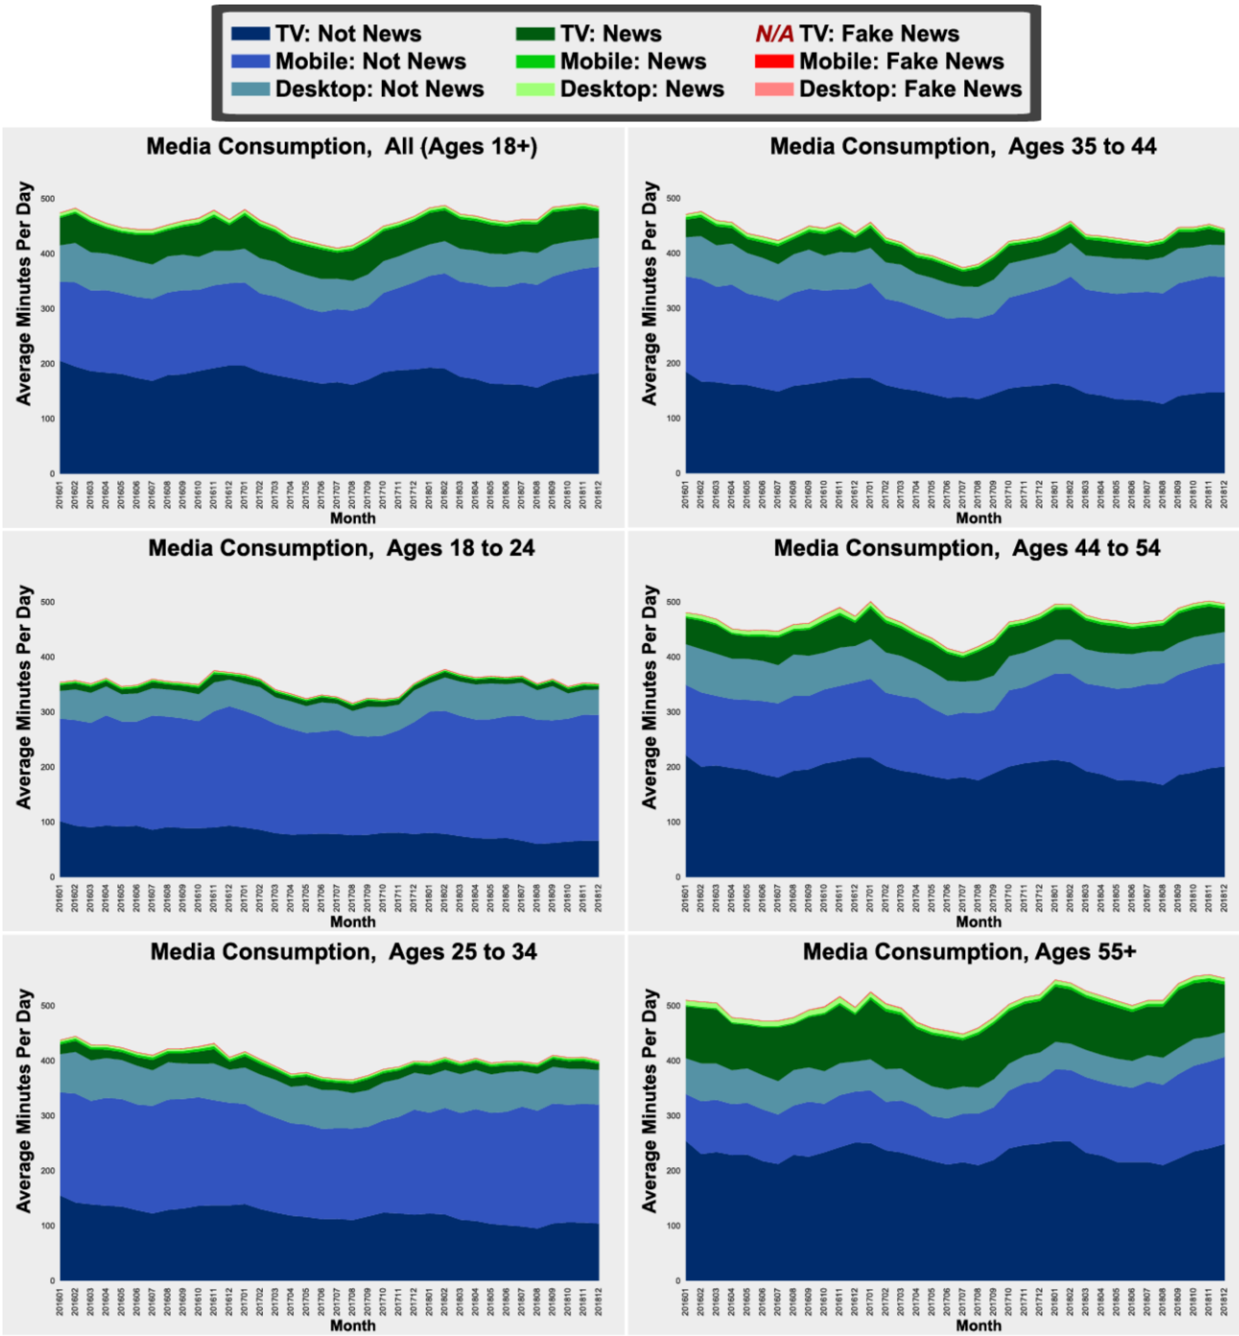

Figure S1 Expanded version of Figure 1 in the main text, showing all age groups including three -- 25 - 34, 35 - 44, 45 - 54-- not shown in the original figure. The general trends seen in figure 1 hold across age groups: overall television consumption and news consumption increase with age, while mobile consumption as a proportion of total consumption decreases with age.

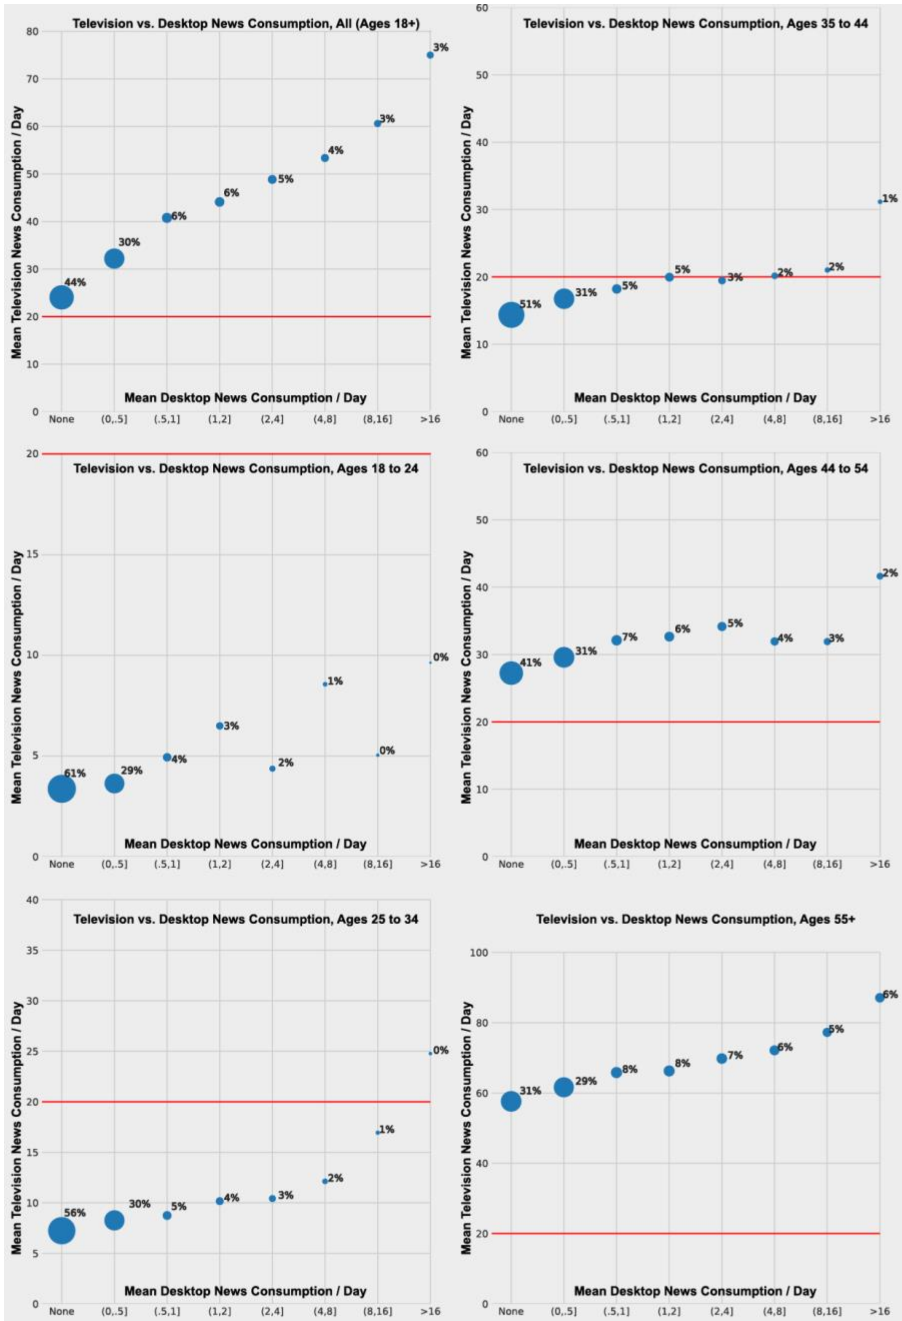

Figure S2. Same as Fig. 3 in main text but broken down by age group (top left panel is identical to Fig 3). Much of the positive trend between television news and desktop web news depicted in Figure 3 is driven by the oldest age group, which contains many people who consume a significant amount of both online and television news. However, there is no age group with a negative trend i.e. there is no age group that appears to, on average, substitute desktop web news for television news.

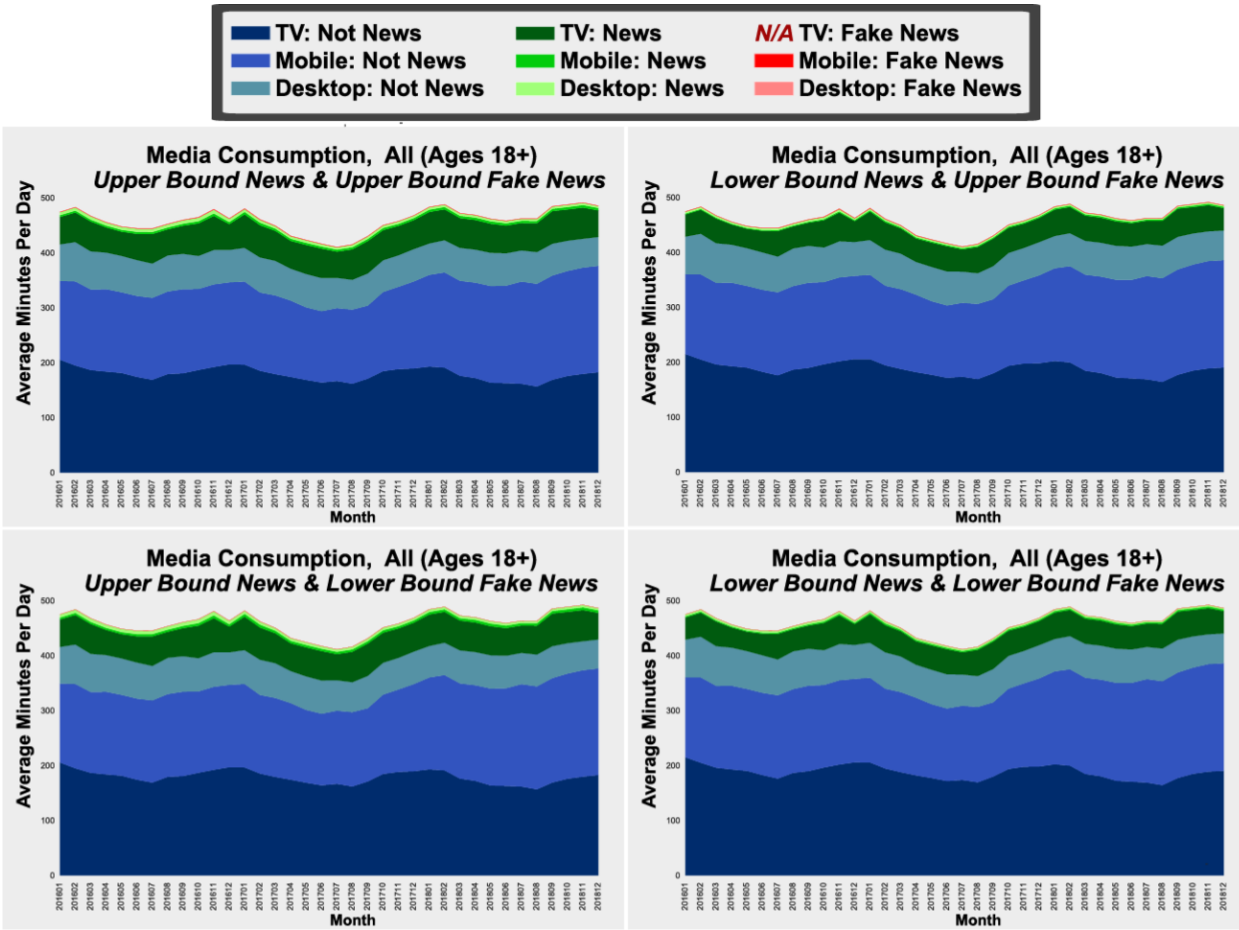

Figure S3. Information consumption under lower and upper bound definitions of news and fake news, respectively. Upper left panel is identical to Fig. 1A in the main text. See Supplementary Methods for descriptions of upper and lower bounds. In all panels, consumption of non-news content vastly outweighs news, television news vastly outweighs online news, and online news vastly outweighs fake news.

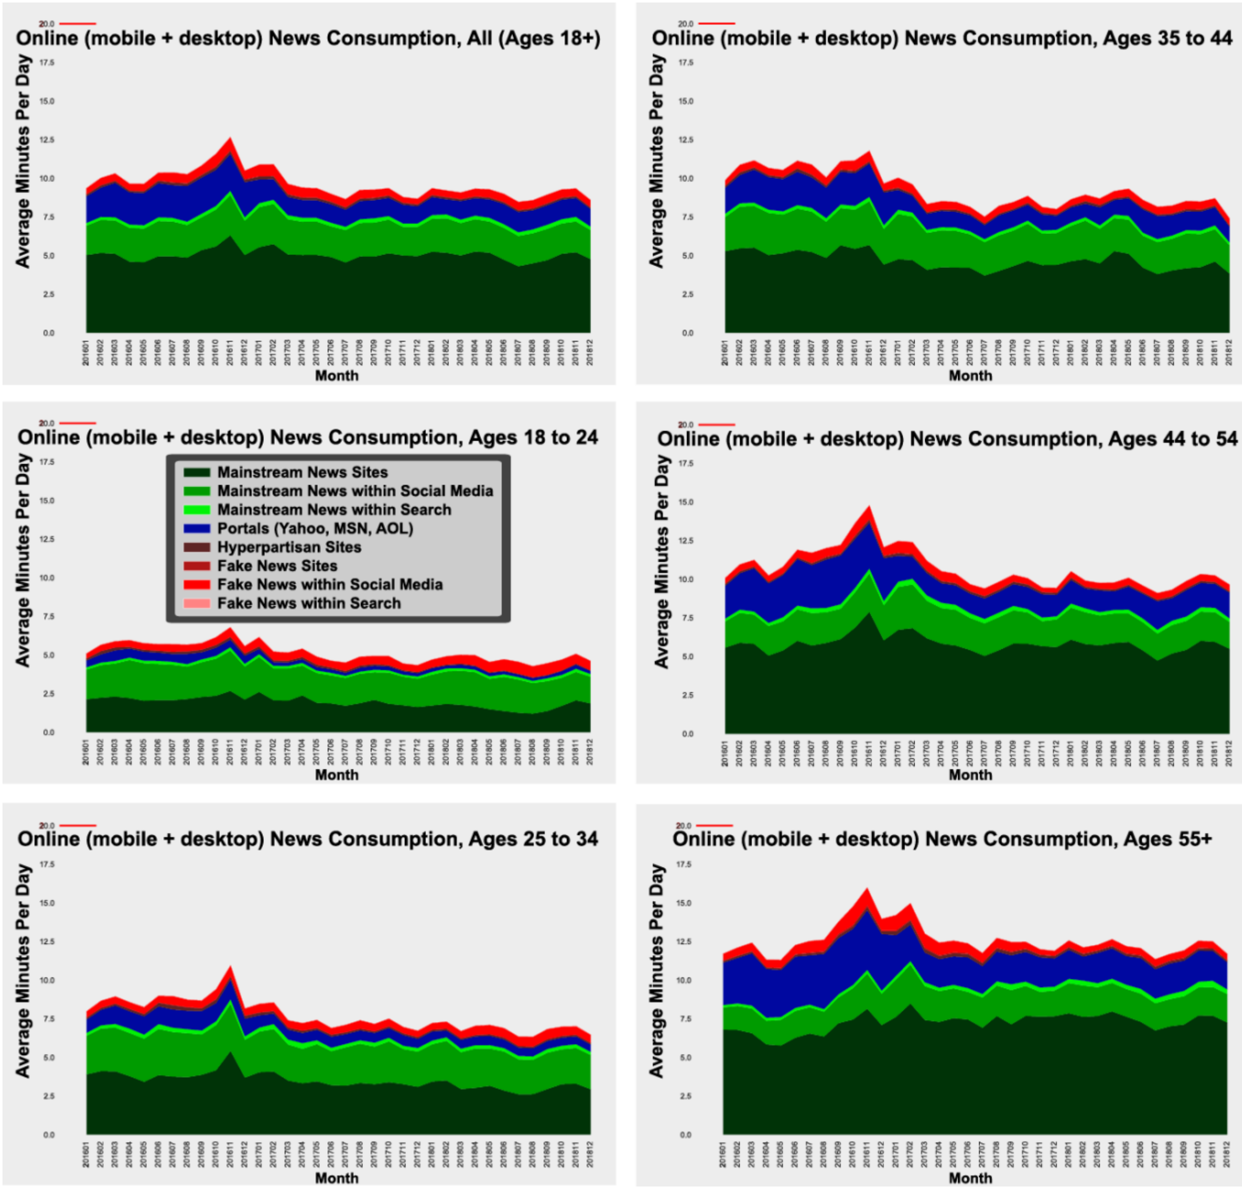

Figure S4A Same information as in Fig. 4A in main text-- online news consumption broken down by age groups--but here also plotted over time (i.e. each bar of Fig 3A is an over-time average of the corresponding panel here). The trends identified in figure 3A hold fairly constant throughout the three year period, but there are some additional secondary findings that emerge when looking at the data over time. In particular, news consumption is highest in 2016, with a clear spike around November 2016, the month of the U.S. Presidential election. Consumption of fake news, in particular fake news on social media denoted in red, also is highest during the months surrounding the 2016 election but then decreases over time. Additionally, as more and more users shift consumption to mobile, time spent on portals -- a primarily desktop phenomenon-- decreases. Data is available on OSF project website (https://osf.io/cygta/).

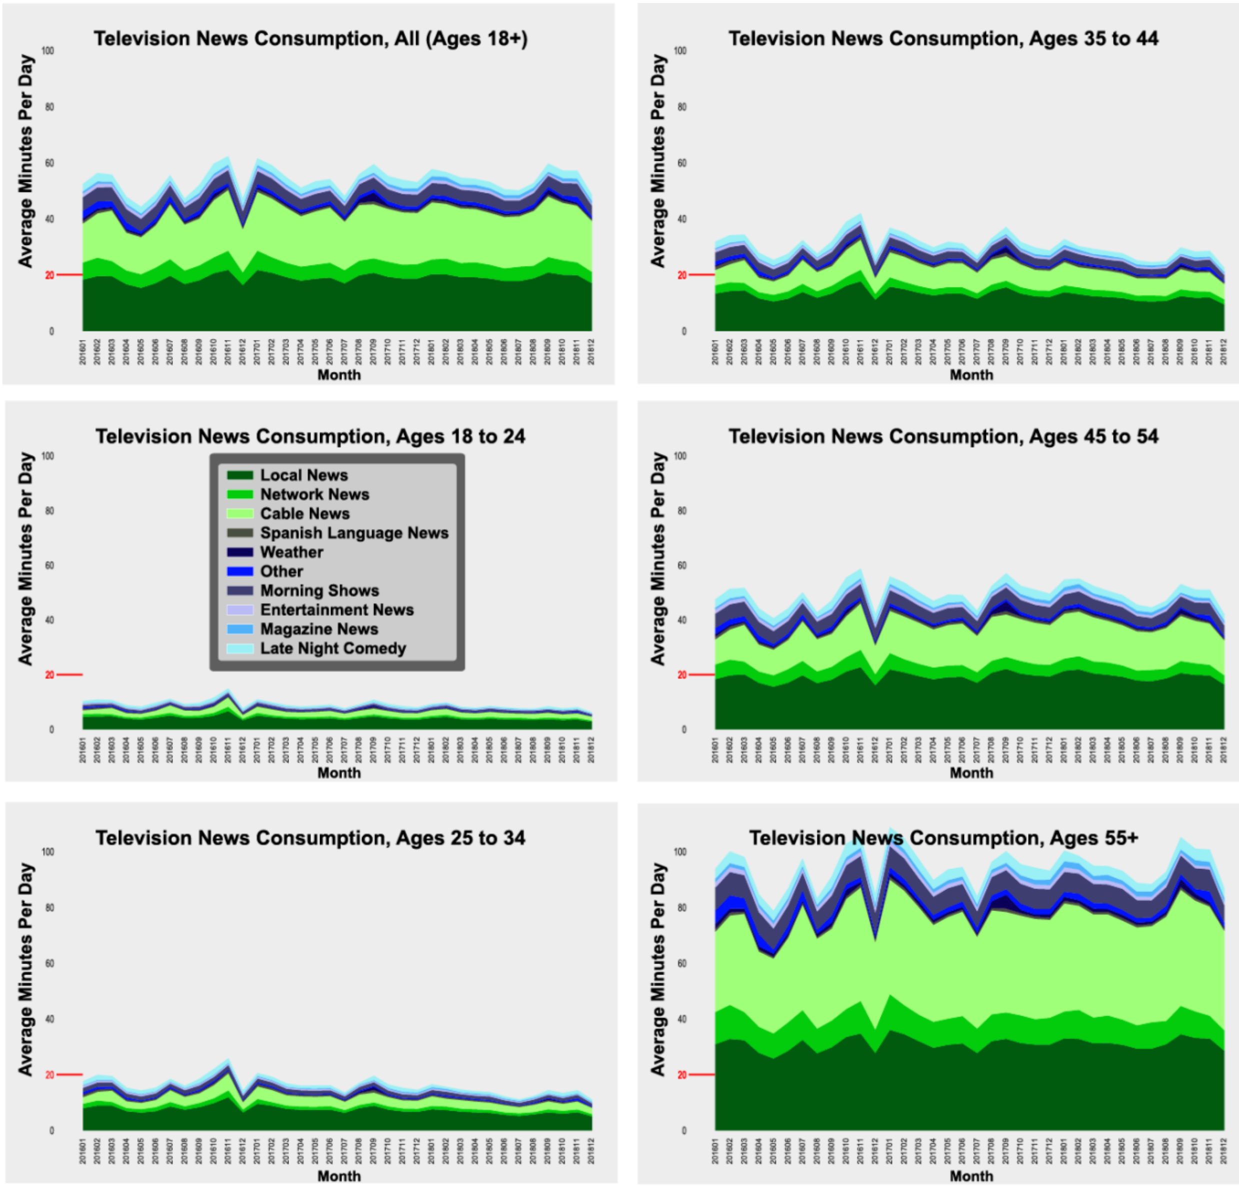

Figure S4B Same information as in Fig. 4B in main text -- television news consumption broken down by age groups--but here also plotted over time (i.e. each bar of Fig 4B is an over-time average of the corresponding panel here). Television news consumption is fairly stable across categories over time. There are slight and opposite trends in the amount of cable and hard network news consumed over 2016-2018. Both in terms of average minutes per person and portion of total television news diet, each year of our analysis cable news increases (32%, 35%, and 36%, for 2016, 2017, and 2018 respectively) and hard network news decreases (11%, 10%, 9%). Data is available on OSF project website (https://osf.io/cygta/).

VII. Supplementary Tables

Data corresponding the the figures is shown below. All data is aggregated over the 36 month period from Jan 2016 - Dec 2018 for ease of understanding. Monthly level data is available at https://osf.io/cygta/

Table S6. Numerical values corresponding to Figs. 1 and S1, averaged over 36 months.

Overall Information Consumption, Minutes per Day

	18_24	25_34	35_44	45_54	55plus	all
news - desktop	1.54	2.96	4.2	5.88	7.56	4.9
other - desktop	52.82	68.1	64.81	65.82	54.32	60.57
fake - desktop	0.16	0.23	0.2	0.29	0.41	0.28
news - mobile	2.74	4.03	4.41	4.26	4.34	4.05
other - mobile	204.31	190.77	174.84	145.38	110.56	157.2
fake - mobile	0.36	0.41	0.4	0.41	0.49	0.43
news - television	9.15	16.15	30.47	49.48	94.51	53.90
other - television	80.3	120.96	152.62	194.46	230.96	178.36
total	351.39	403.62	431.95	465.97	503.14	459.68

Table S7. Numerical values corresponding to Fig 2 - Television and Online by Detailed Category

	News, Avg Mins	Non-news, Avg Mins
automotive	0.00	0.20
business-to-business	0.00	0.11
career	0.00	0.35
corporate	0.00	1.65
directories	0.00	2.48
education	0.00	0.43
entertainment	0.00	22.72
family	0.00	0.58
finance	0.00	0.05
financial	0.00	1.37
gambling	0.00	0.50
gaming	0.00	5.92
government	0.00	0.22
health	0.00	0.69
info	0.00	0.40
isp	0.00	0.07
lifestyles	0.00	1.06
news	3.55	0.00
portals	1.52	0.00
real-estate	0.00	0.76
retail	0.00	5.80
search	0.25	10.14
services	0.00	7.88
social-media	2.67	56.52
sports	0.00	1.22
technology	0.00	0.44
telecommunications	0.00	2.12
travel	0.00	1.06

Television Consumption By Category, Minutes per Day

Category	News, Avg Mins	Non-News, Avg Mins
CHILDRENS	0.00	8.01
COMEDY	0.00	17.55
DOCUMENTARY	0.00	25.86
DRAMA	0.00	40.00
INSTRUCTION	0.00	10.10
MOVIE	0.00	24.72
NEWS	51.51	0.00
OTHER	0.00	16.56
SPORTS	0.00	20.79
TALK	0.00	6.93
VARIETY	2.38	7.56

Table S8. Numerical values corresponding to Fig 3 - Television vs. Desktop news

	Mean television News	SE of Mean television News	% of U.S.	SE of % of U.S.
0 Avg Web News	24.03	0.02	44.28	0.01
(0, .50]	32.19	0.02	29.87	0.01
(.50, 1.00]	40.75	0.05	6.33	<0.01
(1.00, 2.00]	44.10	0.06	5.62	<0.01
(2.00, 4.00]	48.83	0.07	4.68	<0.01
(4.00, 8.00]	53.34	0.09	3.67	<0.01
(8.00, 16.00]	60.60	0.13	2.77	<0.01
>16.00	74.99	0.13	2.67	<0.01

Table S9. Numerical values corresponding to Fig S2 - Television vs. Desktop news (separate tables for each age group)

Age Group: 55+

	Mean television News	SE of Mean television News	% of U.S.	SE of % of U.S.
None	57.62	0.04	31.15	0.01
(0,.5]	61.61	0.05	29.32	0.01
(.5,1]	65.82	0.1	8	0.01
(1,2]	66.27	0.09	7.75	0.01
(2,4]	69.79	0.11	7.09	0.01
(4,8]	72.13	0.12	6.11	0.01
(8,16]	77.25	0.15	5.07	0.01
>16	87.1	0.17	5.5	0.01

Age Group: 45-54

	Mean television News	SE of Mean television News	% of U.S.	SE of % of U.S.
None	27.24	0.04	41.42	0.02
(0,.5]	29.59	0.04	31.44	0.02
(.5,1]	32.11	0.1	6.86	0.01
(1,2]	32.66	0.1	5.97	0.01
(2,4]	34.16	0.11	5.03	0.01
(4,8]	31.95	0.17	3.98	0.01
(8,16]	31.92	0.11	2.86	0.01
>16	41.63	0.2	2.42	0.01

Age Group: 35-44

	Mean television News	SE of Mean television News	% of U.S.	SE of % of U.S.

None	14.37	0.02	51.08	0.02
(0,.5]	16.76	0.03	30.63	0.02
(.5,1]	18.21	0.09	5.46	0.01
(1,2]	19.96	0.11	4.64	0.01
(2,4]	19.46	0.12	3.29	0.01
(4,8]	20.17	0.14	2.31	0.01
(8,16]	21.03	0.2	1.51	0.01
>16	31.15	0.4	0.98	0

Age Group: 25-34

	Mean television News	SE of Mean television News	% of U.S.	SE of % of U.S.
None	7.24	0.01	56.39	0.02
(0,.5]	8.26	0.02	29.76	0.02
(.5,1]	8.74	0.05	4.64	0.01
(1,2]	10.17	0.07	3.58	0.01
(2,4]	10.42	0.08	2.64	0.01
(4,8]	12.12	0.13	1.53	0.01
(8,16]	16.94	0.27	0.78	0
>16	24.77	0.38	0.42	0

Age Group: 18-24

	Mean television News	SE of Mean television News	% of U.S.	SE of % of U.S.
None	3.37	0.01	60.52	0.03
(0,.5]	3.63	0.01	28.53	0.03

(.5,1]	4.93	0.05	4.25	0.01
(1,2]	6.5	0.08	3.02	0.01
(2,4]	4.37	0.06	1.86	0.01
(4,8]	8.56	0.3	0.98	0.01
(8,16]	5.04	0.11	0.39	0
>16	9.64	0.27	0.11	0

Table S10. Numerical values corresponding to Fig S3 - Upper and Lower Bounds

Information Consumption with Bounds, Average Minutes Per Day

	Fake, Lower - News, Lower	**Fake, Upper - News Lower**	**Fake Lower - News Upper**	**Fake, Upper, News Upper**
fake - desktop	0.02	0.28	0.02	0.28
fake - mobile	0.02	0.43	0.02	0.43
news - desktop	3.18	2.92	5.16	4.9
news - mobile	2.53	2.12	4.46	4.05
news - television	44.51	44.51	53.04	53.04
other - desktop	62.81	62.55	60.83	60.57
other - mobile	159.54	159.13	157.61	157.2
other - television	187.44	187.44	178.91	178.91

Table S11. Numerical values corresponding to Fig 4 - News by Platform

Online News Consumption, Average Minutes Per Day

	18_24	**25_34**	**35_44**	**45_54**	**55plus**	**all**
Fake News, Search	0.01	0.01	0.01	0.01	0.01	0.01
Fake News, Social Media	0.39	0.48	0.44	0.53	0.67	0.53
Fake Sites	0.01	0.02	0.03	0.04	0.07	0.04
Hyperpartisan Sites	0.1	0.13	0.11	0.12	0.15	0.13

Mainstream News Sites	1.94	3.49	4.67	5.83	7.25	5.05
Mainstream News, Search	0.17	0.25	0.22	0.27	0.26	0.24
Mainstream News, Social Media	1.85	2.48	2.3	2.13	1.99	2.14
Portals (Yahoo, MSN, AOL)	0.33	0.77	1.41	1.9	2.4	1.52

Television News Consumption, Average Minutes Per Day

	18_24	**25_34**	**35_44**	**45_54**	**55plus**	**all**
Local News	4.10	7.40	12.91	19.28	31.17	18.95
Hard Network News	0.75	1.24	2.45	4.50	9.70	5.27
Cable News	1.88	3.26	7.33	14.07	35.71	18.41
Spanish Language News	0.35	0.52	0.77	0.79	0.84	0.71
Weather	0.12	0.19	0.38	0.70	1.32	0.75
Other (uncategorized hard news)	0.22	0.41	0.73	1.09	2.26	1.23
Soft Morning Shows	0.59	1.27	2.69	4.36	7.08	4.25
Entertainment News	0.32	0.49	0.78	1.12	1.48	1.01
Magazine News	0.18	0.31	0.52	0.90	1.48	0.89
Late Night Comedy	0.64	1.06	1.91	2.66	3.49	2.38