

# Fake News and Advertising on Social Media: A Study of the Anti-Vaccination Movement

Lesley Chiou<sup>1</sup> and Catherine Tucker<sup>2</sup>

July 6, 2018

## Abstract

Online sources sometimes publish information that is false or intentionally misleading. We study the role of social networks and advertising on social networks in the dissemination of false news stories about childhood vaccines. We document that anti-vaccine Facebook groups disseminate false stories beyond the groups as well as serving as an "echo" chamber. We also find that after Facebook's ban on advertising by fake news sites, the sharing of fake news articles on Facebook fell by 75% on Facebook compared to Twitter.

---

<sup>1</sup>Economics Department, Occidental College, CA

<sup>2</sup>MIT Sloan School of Management, MIT, Cambridge, MA.

<sup>3</sup>We thank Yinbo Gao, Tamara Kawash, Gyan Prayaga, and Andrea Tuemmler for excellent research assistance.

# 1 Introduction

The Internet has significantly changed the type of news that consumers receive. In the past consumers relied on traditional media, such as radio and television, which involved relatively fewer and more established sources of news. Nowadays consumers are exposed to online sources of information, through, for example, social networking sites, which allow any individual to share content without "fact-checking or editorial judgment" (Allcott and Gentzkow, 2017a). Many worry that online sources may publish false information, but present it as facts or "real" news. We document how anti-vaccine groups on Facebook disseminate false information to users, and we also study whether Facebook's ban on the advertising of "fake" news prevents the spread of false or misleading news stories on childhood vaccines.

While nearly 60% of adults in the US have searched for health information online in the past year (Fox and Duggan, 2013), online information for consumer health is "often unreliable" and difficult for consumers to assess (Fu et al., 2016; Fahy et al., 2014). Studies demonstrate how consumers do not accurately determine the reliability of health content on the Internet (Allam et al., 2014; Knapp et al., 2011; Kutner et al., 2006). In particular, individuals do not take into account the credibility of the content when presented with online information that is critical of vaccination (Nan and Madden, 2012; Betsch et al., 2010, 2013; Allam et al., 2014).

We focus on childhood vaccines for several reasons. First significant confusion and misleading information on the Internet surrounds the adverse effects of vaccinations on children. For example online articles allege that the vaccine for measles, mumps, and rubella causes autism even though academic studies in the medical literature have since debunked these myths. Second although the Centers for Disease Control and Prevention (CDC) recommend that individuals receive their first vaccinations during childhood, parents report concerns

about safety as among primary reasons why their children are not vaccinated (Smith et al., 2016). Finally vaccines represent an important health concern for the general public because when children receive vaccinations, they also protect the community through herd immunity by preventing further spread of the disease to those individuals unable to be vaccinated.

We explore the role of Facebook groups in spreading false information. We collect data on the content and types of posts shared by Facebook groups that promote the discussion of anti-vaccine beliefs. We find that a handful of authors account for a disproportionately large number of posts and that the posts focus on promoting articles from fake news sites and other online social media. Our results suggest that anti-vaccine groups on Facebook serve as an alternative channel of information for users| both as an "echo" chamber (when users "like" anti-vaccine posts by other users) and as a means of disseminating false stories (when users share a post with others in their social network).

We then study the role of advertising in propagating fake news. In response to criticism over the potential influence of fake news on political outcomes, Facebook banned fake news ads from their advertising networks on November 14, 2016 (Dillet, 2016; Seetharaman, 2016; Wingfeld et al., 2016). The intervention marks a major shift in policy from one of the largest social networking sites in the US and occurred when scrutiny heightened over the role that online misinformation may have played in the outcome of the 2016 US presidential election. Since the ban, Facebook does not display ads that link to websites with misleading or illegal content. Because this ban is unrelated to health news, this provides us with the opportunity

other hand, advertising may convince users to share an article that they would not otherwise.

To circumvent challenges in measuring the effects of advertising (Gordon et al., 2017; Lewis and Reiley, 2014; Lewis et al., 2011), we exploit a difference-in-differences framework. We study how Facebook's ban on the advertising of fake news affects shares of fake news on Facebook, and we use another prominent social media platform, Twitter, that did not experience any policy change during this period as a control group. We compare the number of shares on Facebook with Twitter for news stories about childhood vaccines before and after Facebook's advertising ban on fake news. Our results suggest that the advertising ban


online advertising exists (Chiou and Tucker, 2016; Goldfarb and Tucker, 2011, 2015).

## **2 Fake News and Health Information on Social Media**

### **2.1 Facebook and Twitter**

Facebook and Twitter rank as the two largest social media platforms in the US. Users rely

Figure 1: Screenshot of Advertising of Fake News on Facebook



*Note: Source is gofundme. com.*

## 2.2 Misleading Health Information on the Internet

Vaccines protect the health of individuals as well as members of the community by preventing further spread of the disease. In fact, some vulnerable populations (people with allergies or weakened immune systems due to cancer, HIV/AIDS, and certain diseases) are unable to receive vaccinations and thus rely on community protection from the disease. The level of vaccination required to achieve this type of community protection or "herd" immunity ranges between 83 to 95 percent (of Health and Services, 2018).

Stories with false or misleading information surround health topics on the Internet. Public health officials voice concerns about the influence of fake news because one-third of US consumers use social media for health care information (Miller, 2017), and more than 40% of consumers say that "information found via social media affects the way they deal with their health."

In particular, false news stories surround the safety of vaccinations. Esposito et al. (2014) finds that the "dissemination of misinformation and anecdotal reports of alleged vaccine reactions by the media, the Internet and anti-vaccination groups leads parents to question the need for immunization." For instance, the vaccine for measles, mumps, and rubella is among the most "frequently omitted of the recommended vaccines, usually because of concerns about the vaccine safety." Fake news articles allege that the vaccine may cause autism even though the medical literature has since debunked such claims.

A growing body of evidence demonstrates that consumers struggle to evaluate the credibility and accuracy of online content. Experimental studies find that exposure to online information that is critical of vaccination leads to stronger anti-vaccine beliefs, since individuals do not take into account the credibility of the content (Nan and Madden, 2012; Betsch et al., 2010, 2013; Allam et al., 2014). Survey evidence also shows that only half of low-income parents of children with special healthcare needs felt "comfortable determining the quality of health websites" (Knapp et al., 2011). Since only 12% of US adults are proficient in health literacy with 36% at basic or below basic levels (Kutner et al., 2006), Fu et al. (2016) state that the influence of "low-quality antivaccine web pages that promote compelling but unsubstantiated messages."

Public health officials across the world express concerns about how fake news may influence parents' decision to vaccinate their children. The president of the Irish Medical Association states that the uptake rates for the HPV vaccine are declining to a "worrying

fake news jeopardizes the future health of young women (Power, 2017). Health Minister Beatrice Lorenzin of Italy indicates that the current measles epidemic in Italy and the corresponding declining vaccination rates presents an "emergency generated by fake news" (Press, 2017).

### 3 Social Sharing of Fake News within Facebook Groups

We first document how users disseminate false or misleading information in Facebook groups. Facebook groups serve as a place of communication for people to share common interests and to express their opinions. In Facebook groups, users may organize around a common cause, issue, or activity, and they may express objectives, discuss issues, post photos, and share related content. These groups are similar to discussion forums where groups of users may share photos, links, and updates (Singh, 2014).

We identify anti-vaccine Facebook groups by performing a keyword search on "anti vaccine" in Facebook and filtering the search results by Facebook groups. Using a script, we collect data from these groups for all posts between May and October 2017.<sup>1</sup>

For each post, we observe the author, message, date of the post, and the type of post. The types of posts include those that display an album, event, link, note, photo, status, or video. A note is a longer message that users access on a separate page; it may be edited and formatted as well as set to different privacies. A status is a short message that users post at the top of the page of the group. We also observe the cumulative number of likes, comments, and reactions for the post as of October 2017. On Facebook, users can respond to a post by "liking" it, commenting on it, or designating a reaction (e.g., "wow," "love," "haha", "sad," "angry") that captures their emotional response to the post.

Figure 2 shows a typical post from an anti-vaccine group.<sup>2</sup> Here an author posts a false

---

<sup>1</sup>We focus on Facebook groups with at least 100 users. Note that we only observe Public groups on








Figure 2: Screenshot of post on Anti-vaccine Facebook groups




Figure 3: Types of posts on Anti-vaccine Facebook groups

Table 1: Top 20 domains shared in Facebook groups

	Pct
facebook.com	0.36
youtube.com	0.057
naturalnews.com	0.022
youtu.be	0.018
bolenreport.com	0.016
bit.ly	0.015
ncbi.nlm.nih.gov	0.012
vaccineimpact.com	0.012
vaxxter.com	0.011
greenmedinfo.com	0.0095
ourfamilymagazine.com	0.0095
thevaccinereaction.org	0.0091
yournewswire.com	0.0085
inmymindtoday.com	0.0081
go.thetruthaboutvaccines.com	0.0080
healthimpactnews.com	0.0075
m.youtube.com	0.0059
ow.ly	0.0048
vactruth.com	0.0048
t.co	0.0044

Table 2: Summary statistics for Facebook groups

	Mean	Std Dev	Min	Max
Number of comments	1.69	7.80	0	578
Number of shares	2.35	17.4	0	2101
Number of reactions	7.63	23.3	0	2075
Observations	24025			

Notes: Each observation represents a post on the top 20 anti-vaccine Facebook groups.

network, or indicating an emotional reaction (i.e., likes, loves, wows, hahas, sads, angrys, special). A typical post receives 7 emotional reactions and 2 shares. To a lesser extent, users post comments; a post receives about 1.7 comments on average. Figure 4 reveals how the vast majority of reactions are "likes" that indicate users respond positively to posts on the Facebook groups.

1	2	3	4	5	6
7	8	9	10	11	12
13	14	15	16	17	18
19	20	21	22	23	24
25	26	27	28	29	30
31	32	33	34	35	36
37	38	39	40	41	42
43	44	45	46	47	48
49	50	51	52	53	54
55	56	57	58	59	60
61	62	63	64	65	66
67	68	69	70	71	72
73	74	75	76	77	78
79	80	81	82	83	84
85	86	87	88	89	90
91	92	93	94	95	96
97	98	99	100	101	102
103	104	105	106	107	108
109	110	111	112	113	114
115	116	117	118	119	120
121	122	123	124	125	126
127	128	129	130	131	132
133	134	135	136	137	138
139	140	141	142	143	144
145	146	147	148	149	150
151	152	153	154	155	156
157	158	159	160	161	162
163	164	165	166	167	168
169	170	171	172	173	174
175	176	177	178	179	180
181	182	183	184	185	186
187	188	189	190	191	192
193	194	195	196	197	198
199	200	201	202	203	204
205	206	207	208	209	210
211	212	213	214	215	216
217	218	219	220	221	222
223	224	225	226	227	228
229	230	231	232	233	234
235	236	237	238	239	240
241	242	243	244	245	246
247	248	249	250	251	252
253	254	255	256	257	258
259	260	261	262	263	264
265	266	267	268	269	270
271	272	273	274	275	276
277	278	279	280	281	282
283	284	285	286	287	288
289	290	291	292	293	294
295	296	297	298	299	300
301	302	303	304	305	306
307	308	309	310	311	312
313	314	315	316	317	318
319	320	321	322	323	324
325	326	327	328	329	330
331	332	333	334	335	336
337	338	339	340	341	342
343	344	345	346	347	348
349	350	351	352	353	354
355	356	357	358	359	360
361	362	363	364	365	366
367	368	369	370	371	372
373	374	375	376	377	378
379	380	381	382	383	384
385	386	387	388	389	390
391	392	393	394	395	396
397	398	399	400	401	402
403	404	405	406	407	408
409	410	411	412	413	414
415	416	417	418	419	420
421	422	423	424	425	426
427	428	429	430	431	432
433	434	435	436	437	438
439	440	441	442	443	444
445	446	447	448	449	450
451	452	453	454	455	456
457	458	459	460	461	462
463	464	465	466	467	468
469	470	471	472	473	474
475	476	477	478	479	480
481	482	483	484	485	486
487	488	489	490	491	492
493	494	495	496	497	498
499	500	501	502	503	504
505	506	507	508	509	510
511	512	513	514	515	516
517	518	519	520	521	522
523	524	525	526	527	528
529	530	531	532	533	534
535	536	537	538	539	540
541	542	543	544	545	546
547	548	549	550	551	552
553	554	555	556	557	558
559	560	561	562	563	564
565	566	567	568	569	570
571	572	573	574	575	576
577	578	579	580	581	582
583	584	585	586	587	588
589	590	591	592	593	594
595	596	597	598	599	600
601	602	603	604	605	606
607	608	609	610	611	612
613	614	615	616	617	618
619	620	621	622	623	624
625	626	627	628	629	630
631	632	633	634	635	636
637	638	639	640	641	642
643	644	645	646	647	648
649	650	651	652	653	654
655	656	657	658	659	660
661	662	663	664	665	666
667	668	669	670	671	672
673	674	675	676	677	678
679	680	681	682	683	684
685	686	687	688	689	690
691	692	693	694	695	696
697	698	699	700	701	702
703	704	705	706	707	708
709	710	711	712	713	714
715	716	717	718	719	720
721	722	723	724	725	726
727	728	729	730	731	732
733	734	735	736	737	738
739	740	741	742	743	744
745	746	747	748	749	750
751	752	753	754	755	756
757	758	759	760	761	762
763	764	765	766	767	768
769	770	771	772	773	774
775	776	777	778	779	780
781	782	783	784	785	786
787	788	789	790	791	792
793	794	795	796	797	798
799	800	801	802	803	804
805	806	807	808	809	810
811	812	813	814	815	816
817	818	819	820	821	822
823	824	825	826	827	828
829	830	831	832	833	834
835	836	837	838	839	840
841	842	843	844	845	846
847	848	849	850	851	852
853	854	855	856	857	858
859	860	861	862	863	864
865	866	867	868	869	870
871	872	873	874	875	876
877	878	879	880	881	882
883	884	885	886	887	888
889	890	891	892	893	894
895	896	897	898	899	900
901	902	903	904	905	906
907	908	909	910	911	912
913	914	915	916	917	918
919	920	921	922	923	924
925	926	927	928	929	930
931	932	933	934	935	936
937	938	939	940	941	942
943	944	945	946	947	948
949	950	951	952	953	954
955	956	957	958	959	960
961	962	963	964	965	966
967	968	969	970	971	972
973	974	975	976	977	978
979	980	981	982	983	984
985	986	987	988	989	990
991	992	993	994	995	996
997	998	999	1000	1001	1002

## 4 Does Advertising Spread Fake News?

### 4.1 Data on the Popularity of Fake News on Facebook and Twitter

We investigate the effect of advertising of fake news articles on Facebook and Twitter for

Table 3: Summary statistics

	Mean	Std Dev	Min	Max
shares	2409.5	11342.1	1	168397
Postban	0.64	0.48	0	1
Facebook	0.50	0.50	0	1
days	391.3	39.0	341	460
Observations	354			

Notes: Each observation represents a website and keyword combination.

indicator variable *Facebook* equals one for shares on Facebook and equals zero for shares on Twitter. The variable *days* measures the number of days that the article has been published online.

In our sample, an article receives on average 2000 shares on either Facebook or Twitter. Nearly 20% of the articles in our sample were published after the advertising ban. We have a matched sample with exactly half of observations for shares on Facebook and the other half for shares on Twitter. The articles have been in publication for approximately seven months on average.

Table A-2 in the Appendix describes the demographics of users of Facebook and Twitter. Reassuringly the demographics are relatively similar with the exception that Facebook has a lower proportion of users that are male. Facebook and Twitter have a similar age profile of their users, and both platforms have a sizable fraction (35 to 45 percent) of users in the highest income bracket. Overall the table suggests that users of Twitter provide a plausible control for users of Facebook.

## 4.2 Estimating the Effect of Advertising on Sharing

As a preliminary analysis, we compare the number of shares of articles with fake news published before and after the adoration- [(of)e83321754(127lysis,sn)-346ivy58S ook an39(of)-477(trol)-3Su




Figure 5: Logarithm of shares on Facebook drop relative to Twitter comparing the month before and after Facebook's ban on advertising by fake news sites

advertising ban on Facebook, shares for Facebook decline sharply after the ban.

For a formal regression analysis, we estimate the following equation. For every search term  $j$  on platform  $k$ , we regress the logarithm of the number of shares received by news article  $i$ :

$$\begin{aligned} \log(\text{shares}_{ijk}) = & \beta_0 + \beta_1 \text{Facebook}_k \text{ PostBan}_i + \beta_2 \text{Facebook}_k \\ & + \beta_3 \text{days}_k + \alpha_i + \alpha_j + \alpha_{ijkt} \end{aligned} \quad (1)$$

where  $\text{Facebook}$  is an indicator variable equal to 1 if the platform is Facebook and 0 if Twitter;  $\text{PostBan}$  is an indicator variable that equals one if the news article was published after Facebook's ban of advertising by fake news sites. The variable  $\text{days}$  is the number of days that the article has been in publication. The control  $\alpha_i$  is a fixed effect for the article's month of publication, and  $\alpha_j$  is a fixed effect for keyword. We cluster our standard errors by



platform and keyword.

Our identification strategy compares shares on Facebook and Twitter for articles published before and after Facebook's advertising ban. We exploit the advertising ban as an exogenous shifter of advertising for fake news for vaccines because the ban was unrelated to health topics and enacted due to concerns about political misinformation. The number of shares on Facebook for articles published after the advertising ban on Facebook reflect popularity in the absence of advertising while the number of shares on Facebook for articles published prior to the advertising ban include the influence of advertising. We control for seasonal differences in popularity of articles by using shares on Twitter over this same time period as a control group. Because we observe shares in Facebook and Twitter for the same article and control for the number of days in publication, we can isolate the effect of advertising.

We interpret our estimated coefficients of the semi-log specification as the "ratio-of-ratios" (Mullahy, 1999). For instance, to determine the effect of the advertising ban on shares, we compute the corresponding ratio-of-ratios:

$$\frac{\frac{E[\text{shares}_{Facebook=1; PostBan=1}]}{E[\text{shares}_{Facebook=1; PostBan=0}]}}{\frac{E[\text{shares}_{Facebook=0; PostBan=1}]}{E[\text{shares}_{Facebook=0; PostBan=0}]}} = \exp(\beta_1) \quad (2)$$

In Equation (2) above, the fraction in the numerator (proportionately) compares the expected number of shares on Facebook before and after the advertising ban. The fraction in the denominator compares the expected number of shares to the control Twitter before and after the advertising ban. The formula avoids the "retransformation bias" for estimating the number of shares from the semi-log regression and offers a natural interpretation for the estimated coefficients directly (Mullahy, 1999).

In other words,  $\exp(\beta_1)$  captures the extent to which shares on Facebook fall proportionately relative to shares on Twitter after the advertising ban. If the expression is less than

Table 4: Facebook shares drop relative to Twitter after Facebook's ban on advertising of fake news

Overall, a direct ban on ads of fake news dramatically decreases the number of shares of fake news by 75%. We explore the magnitude in two ways. First, we perform a back-of-the-envelope calculation of the total decline in shares from the advertising ban. Since our sample includes articles that were shared a total of 1.6 million times on Facebook before any of the advertising bans, a decrease of 75% equates to a decline in total shares of 1.12 million for the fake news sites in our sample.

Second, we calculate a benchmark of how referrals from Facebook to fake news sites change after the advertising ban. We collect additional data from comScore that tracks the incoming traffic to each of the fake news sites in our sample from Hoaxy.<sup>7</sup> Table 5 compares the average percentage of incoming traffic to the fake news sites on Hoaxy that originate from Facebook and the top three search engines. We consider the top three search engines, since they represent platforms that consumers use for information and that also feature fake news sites and stories. Facebook accounts for a large fraction (13 percent) of incoming referrals while the other three search engines each account for significantly fewer (less than 10% of referrals). If incoming traffic declines by the same proportion as the number of shares,

Table 5: Incoming traffic for fake news sites

	Pct
Facebook	13.3
Google	6.51
Bing	2.46
Yahoo	1.63

## 5 Robustness and Falsification Checks

### 5.1 Controlling for Differences between Articles

We perform several robustness and falsification checks in this section. First, we consider whether underlying differences in popularity between articles drive our results; did the decline in shares of fake news on Facebook occur because articles published before and after the ban differ in underlying popularity? As a robustness check, we include fixed effects for each article in Column (2) of Table 4. This identification strategy compares the number of shares

Table 6: No negative pre-trend for Facebook shares prior to its ban on advertising by fake news sites

		(1)	(2)
FakePost	Facebook	0.464 (0.519)	0.464 (0.390)
	Facebook	3.275 (0.348)	3.275 (0.356)
	days	0.0175 (0.0257)	2.874 (2.613)
	Month Fixed Effects	Yes	Yes
	Keyword Fixed Effects	Yes	Yes
	Observations	286	286
	R-Squared	0.412	0.772

Notes: Robust standard errors. \* $p < 0.1$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . The dependent variable is the logarithm of the number of shares of a news article from a platform| either Facebook or Twitter.

Table 4, the negative coefficient on the interaction of *PostBan* *Facebook* *Early* supports our hypothesis that the ban on fake news has a larger effect on articles about vaccines administered earlier in childhood.

### 5.3 Checking for a Pre-Trend

We investigate whether the decline in Facebook shares relative to Twitter after the

Table 7: Summary statistics for placebo group of health conditions

	Mean	Std Dev	Min	Max
Number of comments	1.69	7.80	0	578
Number of shares	2.35	17.4	0	2101
Number of reactions	7.63	23.3	0	2075
Observations	24025			

Notes: Each observation represents a post on the top 20 anti-vaccine Facebook groups.

controversial health topic of vaccines. We expect that the advertising ban would have a smaller (if any) effect on the sharing of fake news for medical conditions than for vaccines.

We identify health conditions listed under the "Most Viewed Health Topics" on the Dr. Sears website (<http://www.askdrsears.com/>). Dr. Sears is a prominent pediatrician and his website "Ask Dr. Sears" serves as an online resource for parents about children's health. In the Appendix, Table A-3 lists the keywords for these medical conditions.

Then we perform keyword searches for these health conditions on Hoaxy and collect data for the top 20 articles of each keyword in January 2017. We run a regression similar to equation 1 using this dataset. Table 7 lists the summary statistics. Our results in table 8 suggest that the advertising ban did not have an effect on the sharing of fake news articles for medical conditions because the estimated coefficient is smaller in magnitude than the results for vaccine keywords and is also not statistically significant.

## 6 Conclusion

This paper examines how false information spreads on social networkingci5[16

net

Table 8: No effect on placebo group of health conditions

	(1)	(2)
PostBan Facebook	-0.592 (0.445)	-0.592 (0.414)
Facebook	1.987 (0.375)	1.987 (0.430)
days	0.00199 (0.0204)	1.321 (0.826)
Month Fixed Effects	Yes	Yes
Keyword Fixed Effects	Yes	Yes
Article Fixed Effects	No	Yes
Observations	186	186
R-Squared	0.401	0.694

Notes: Robust standard errors. \* $p < 0.1$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . The dependent variable is the logarithm of the number of shares of a news article from a platform | either Facebook or Twitter.

We also study the effectiveness of advertising in spreading fake news about vaccines on social media sites. We examine a policy experiment where Facebook banned ads containing links to fake news sites in response to criticism over the influence of fake news on the US presidential election. Our results indicate that this ban on advertising led to a dramatic decline of 75% in the number of shares on Facebook relative to Twitter, which had no change in its advertising policy during this time.

Our results suggest potential ways of curbing the influence of fake news on social networking sites. We explore how false information spreads within Facebook groups dedicated to promoting false information. In addition, we illustrate how advertising regulations may effectively curtail the popularity of articles with fake news.

Our study also suggests another potential mechanism for counteracting fake news: creating Facebook pages of real news and using this advertising to disseminate accurate information. Future work can focus on whether positive advertising can counteract the effects of negative advertising.

In the context of our study on childhood vaccines, fake news may potentially harm consumers, as public health officials fear that fake news' misleading claims about the safety

of vaccines lowers vaccination rates. The community as a whole depends upon vaccinations by individuals, which makes this is a pressing public concern.

Consumers find health information to be extremely difficult to evaluate and to determine validity, so consumers benefit when firms and policymakers take direct action by preventing the spread of fake news articles.



## References

- Allam, A., P. Schulz, and K. Nakamoto (2014). The impact of search engine selection and sorting criteria on vaccination beliefs and attitudes: Two experiments manipulating Google output. *Journal of Medical Internet Research* 16, e100.
- Allcott, H. and M. Gentzkow (2017a). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives* 31, 211{236.
- Allcott, H. and M. Gentzkow (2017b). Social Media and Fake News in the 2016 Election. *Stanford University, working paper*.
- Baron, D. (2004). Persistent Media Bias. *Journal of Public Economics* 90, 1{36.
- Beech, H. (2018). As Malaysia Moves to Ban 'Fake News,' Worries About Who Decides the Truth. *New York Times*.
- Besley, T. and A. Prat (2004). Handcuffs for the Grabbing Hand? Media Capture and Government Accountability. *London School of Economics*.

- De Los Santos, B., A. Hortacsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review* 102(6), 2955{80.
- Dillet, R. (2016). Google and Facebook Ban Fake News Sites from Their Advertising Networks. *TechCrunch*.
- Esposito, S., N. Principi, and G. Cornaglia (2014). Barriers to the vaccination of children and adolescents and possible solutions. *Clinical Microbiology and Infection* 20, 25{31.
- Fahy, E., R. Hardikar, A. Fox, and S. Mackay (2014). Quality of Patient Health Information on the Internet: Reviewing a complex and Evolving Landscape. *Australasian Medical Journal* 7, 24{28.
- Fox, S. and M. Duggan (2013). Pew Internet and American Life Project.
- Friggeri, A., L. A. Adamic, D. Eckles, and J. Cheng (2014). Rumor cascades. In *ICWSM*.
- Fu, L., K. Zook, Z. Spoehr-Labutta, P. Hu, and J. Joseph (2016). Search Engine Ranking, Quality, and Content of Web Pages That Are Critical Versus Noncritical of Human Papillomavirus Vaccine. *Journal of Adolescent Health* 58, 33{39.
- Gentzkow, M. and J. Shapiro (2006). Media Bias and Reputation. *Journal of Political Economy* 114(2), 280{316.
- Gentzkow, M. and J. M. Shapiro (2011). Ideological segregation online and offline. *The Quarterly Journal of Economics* 126, 1799{1839.
- Goldfarb, A. and C. Tucker (2011). Advertising Bans and the Substitutability of Online and Offline Advertising. *Journal of Marketing Research* 48(2), 207{227.

- Goldfarb, A. and C. Tucker (2015). Standardization, standards and online advertising. *Management Science* 61, 2707{2719.
- Gordon, B., F. Zettelmeyer, N. Bhargava, and D. Chapsky (2017). A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook. *working paper*.
- Hoaxy (2016). Frequently Asked Questions About Hoaxy.
- Knapp, C., V. Madden, and H. Wang (2011). Internet use and eHealth literacy of low-income parents whose children have special health care needs. *Journal of Medical Internet Research* 13, e75.
- Kutner, M., E. Greenberg, Y. Jin, and C. Paulsen (2006). The health literacy of America's adults: Results from the 2003 National Assessment of Adult Literacy. *US Department of Education, National Center for Education Statistics* 13.
- Lewis, R., J. Rao, and D. Reiley (2011). Here, There, and Everywhere: Correlated Online Behaviors Can Lead to Overestimates of the Effects of Advertising. *Proceedings of the 20th International Conference on the World Wide Web*, 157{166.
- Lewis, R. and D. Reiley (2014). Online ads and online sales: measuring the effect of retail advertising via a controlled experiment on Yahoo! *Quantitative Marketing and Economics* 12, 235{266.
- Miller, M. (2017). 33 percent of us consumers use social media for health care info. *Search Engine Watch*.
- Montgomery, A. L., S. Li, K. Srinivasan, and J. C. Liechty (2004). Modeling online browsing and path analysis using clickstream data. *Marketing Science* 23(4), 579{595.



Singh, K. (2014). What's really the difference between facebook pages, groups and profiles?  
*Mashable*.

Smith, L., R. Amlot, J. Weinman, J. Yiend, and G. Rubin (2016). A systematic review of

# A Appendix

## A-1 Facebook policy on fake news sites (excerpts from April 6, 2017)

### Working to Stop Misinformation and False News

By Adam Mosseri, VP, News Feed

We know people want to see accurate information on Facebook and so do we.

False news is harmful to our community, it makes the world less informed, and it erodes trust. Its not a new phenomenon, and all of us tech companies, media companies, newsrooms, teachers have a responsibility to do our part in addressing it. At Facebook, were working to ght the spread of false news in three key areas:

- disrupting economic incentives because most false news is nancially motivated;

- building new products to curb the spread of false news; and

- helping people make more informed decisions when they encounter false news.

### Disrupting Economic Incentives

When it comes to ghting false news, one of the most e ective approaches is removing the economic incentives for tra ckers of misinformation. Weve found that a lot of fake news is nancially motivated. These spammers make money by masquerading as legitimate news publishers and posting hoaxes that get people to visit their sites, which are often mostly ads.

Some of the steps were taking include:

- Better identifying false news through our community and third-party fact-checking organizations so that we can limit its spread, which, in turn, makes it uneconomical.

- Making it as di cult as possible for people posting false news to buy ads on our platform through strict enforcement of our policies.

- Applying machine learning to assist our response teams in detecting fraud and enforcing our policies against inauthentic spam accounts.

- Updating our detection of fake accounts on Facebook, which makes spamming at scale much harder.

- Update on May 10, 2017: Weve made updates so people see fewer posts and ads in News Feed that link to low-quality web page experiences.

- Update on August 9, 2017: Weve made updates to address cloaking so that what people see after clicking an ad or post matches their expectations.

- Update on August 28, 2017: Weve made an update in which repeat o enders that repeatedly share stories marked as false will no longer be allowed to advertise on Facebook.

## **Building New Products**

Were building, testing and iterating on new products to identify and limit the spread of false news. We cannot

News Integrity Initiative: Weve joined a group of over 25 funders and participants including tech industry leaders, academic institutions, non-pro ts and third party organizations to launch the News Integrity Initiative, a global consortium focused on helping people make informed judgments about the news they read and share online. Founding funders of this \$14-million fund include Facebook, the Craig Newmark Philanthropic Fund, the Ford Foundation, the Democracy Fund, the John S. and James L. Knight Foundation, the Tow Foundation, AppNexus, Mozilla and Betaworks. The initiatives mission is to advance news literacy, to increase trust in journalism around the world and to better inform the public conversation. The initiative, which is administered by the CUNY Graduate School of Journalism, will fund applied research and projects, and convene meetings with industry experts.

Update on August 3, 2017: We made an update where if an article has been reviewed by fact checkers, we may show the fact checking stories below the original post in Related Articles.

We need to work across industries to help solve this problem: technology companies, media companies, educa-



Table A-1: Keywords for vaccines

<b>keyword</b>
chicken pox
diphtheria
dtp
u
haemophilus influenzae
hepatitis a
hepatitis b
hib
hpv
human papillomavirus
measles
meningitis
meningococcal
mmr
mumps
pertussis
pneumococcal
pneumonia
polio
rotavirus
rubella
shingles
smallpox
tetanus
varicella
whooping cough

Table A-2: Demographic description of users

Measure	Facebook	Twitter
Male	47.2	55.6
Age 18-24	12.1	14.9
Age 25-34	17.2	19.1
Age 35-44	16.4	15.6
Age 45-54	18	17.2
Age 55+	29.7	25.2
Income <25k	10.9	7.5
Income 25-60k	26.1	20.6
Income 60-100k	28.0	27.0
Income >100k	35	44.9

*Source: comScore*

*Note:* This table reports the fraction of users within each demographic category. Statistics are reported for users of Facebook and Twitter.

Table A-3: Keywords for placebo group of health conditions

**Health Keywords**

---

---

allergies  
asthma  
breastfeeding  
coughs and colds  
dehydration  
diarrhea  
ear infection  
fever  
head injuries  
lice  
rashes  
ringworm  
sore throat  
vomiting

---

---

Notes: This table lists the health keywords for the "Most Viewed Health Topics" on the Dr. Sears website <http://www.askdrsears.com/>.