The Impact of Fake News:

Evidence from the Anti-Vaccination Movement in the US

Meradee Tangvatcharapong *

October 2019

Abstract

The increasing amount of fake news has generated significant debate about the proper role of government and social media platforms in combating it. However, little is known about whether fake news can actually change behavior. This paper addresses this question by examining how vaccination rates responded to the unexpected surge in media coverage in 2007 of the verifiably false claim that the MMR vaccine caused autism. Specifically, I use a difference-in-difference approach to compare the MMR vaccination rates of children whose parents were most and least likely to be affected by the news over time. I determine parents' susceptibility using three predetermined characteristics: whether their child is a firstborn, the child's gender, and the parents' age. Results show that susceptible parents were 3.3 percentage points less likely to vaccinate their children with an MMR shot by the recommended age of 15 months and 4.1 percentage points less likely to do so by 29 months. This indicates that at a minimum, fake news caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children.

*Department of Economics, Texas A&M University

1 Introduction

Recent advancements in technology have enabled information to travel faster and reach far more people than before. Unfortunately, this also means that it has become easy to deliberately spread misinformation or 'fake news'. Additionally, while inconsequential fake stories such as the flat-earth conspiracy have always existed, many of the current fake stories are more likely to affect important outcomes. For example, Allcott and Gentzkow (2017) report that during the 2016 presidential election cycle fake news stories regarding presidential candidates were shared 37.6 million times on Facebook. In addition, they also estimate that average American adults likely saw several fake news stories in the months before the election. The increasing amount of fake news and its potential consequences has thus generated significant debate about the government's role in regulating fake news and social media platforms' responsibility to fight it. However, little is known about the actual impact of fake news (Lazer et al., 2018). In theory, fake news stories can be seen as distorted signals uncorrelated with the truth. These distorted signals could then lead consumers to make different decisions than they otherwise would have (Allcott and Gentzkow, 2017). Nevertheless, to my knowledge, there has not been a study that provides empirical evidence to confirm this theory. The purpose of this paper is to ask whether the dissemination of completely false news leads to meaningful changes in behavior.

I estimate the effect of fake news by studying how vaccination rates responded to the unexpected surge in media coverage in 2007 of the claim, which was shown to be false in the early 2000s, that the MMR (Measles-Mumps-Rubella) vaccine causes autism. This exogenous shock in misinformation, along with the fact that some parents are ex-ante more likely to be sensitive to this misinformation than others, allows me to identify the effects of false news about vaccine safety on parents' vaccination decisions.

There are several reasons why the surge in media coverage on the alleged link between vaccines and autism is the ideal setting in which to study the impact of fake news. First, the claim that the MMR vaccine, or any vaccine, causes autism is false and could be easily verified by 2007. The claim that the MMR vaccine causes autism stems from a now-retracted paper by Wakefield et al., which was published in 1998 in The Lancet, a major British medical journal. However, major medical and scientific bodies have since conducted further studies and refuted the claim as false; the Institute of Medicine (IOM) in May 2004, the Food and Drug Administration (FDA) in September 2006, and the Centers for Disease Control and Prevention (CDC) in July 2007. Therefore, whenever the media covered the stories or gave the platform to anti-vaccination activists to propel the false claim without explicitly refuting the claim, especially after 2007, they were broadcasting false information. Second, in contrast to some fake news topics, false information about vaccine safety can affect important health outcomes. Parents who do not vaccinate their children not only expose their own children to the risk of serious diseases but also makes it harder for the community to retain herd immunity as well. [1] Third, the surge in misinformation about vaccine safety by the media in 2007 was unexpected to parents, because it was largely driven by high-profile court cases alleging that vaccines cause autism and celebrities' decisions to speak out on the issue. The most notable instance of this was Jenny McCarthy making multiple appearances on talk shows, including The Oprah Winfrey Show.

I start my analysis by looking at the media coverage on the alleged link between vaccines and autism to confirm that there is a surge in false news. Specifically, I collect news transcripts from six major television networks in the US (ABS, CBS, NBC, CNN, MSNBC, and Fox News) from 2001 to 2012 via LexisNexis. I use coverage on major television networks as a proxy for media coverage because although many people get their news through other sources, 44% of Americans still prefer television as the platform they most preferred for news (Mitchell, 2018). I classify each news story as either only reporting the alleged link between vaccines and autism, or reporting that as well as actively refuting the claim. Coverage that only reported on the alleged link between vaccines and autism without refuting it with the correct information is what I then count as fake news. I found that the number of fake news stories about vaccine safety reported on these six networks rose dramatically from an average of 10 stories per year from 2001 to 2006 to 44 stories in 2007 and then 105 stories in 2008.

To identify the effects of fake news, I exploit this shock in fake news along with its differential impact on parents. Specifically, I expect fake news should have larger effects on parents who are ex-ante more likely to be sensitive and receptive to the news. Therefore, I identify the effects by comparing the vaccination rates of children whose parents are ex-ante most sensitive and least

[1] Herd immunity is defined as the resistance to the spread of contagious disease within a population that results if a sufficiently high proportion of individuals are immune to the disease, especially through vaccination. For example, the vaccination rate required to achieve herd immunity is 83-94% for measles (Fine, 1993).

sensitive to the news over time. While this approach will likely result in an underestimation of effects given all parents were likely somewhat affected by the fake news, it enables me to use a difference-in-differences approach to distinguish effects from other time-varying factors. Specifically, I do so using individual-level vaccination data obtained from healthcare providers of 19-35-month-old American children surveyed in the 2002-2012 National Immunization Surveys (NIS). I determine parents' susceptibility to the news using three predetermined characteristics: whether the child is a firstborn, a boy, and the mother is over 30 years old. Parents are classified as most sensitive if they have all these three characteristics present and least sensitive if they have none. I use these three characteristics to determine parents' sensitivity for the following reasons. First, experienced parents are more likely to have already formed an opinion on the issue based on past experience. As a result, the fake news stories after 2007 likely only accounted for a small fraction of their information. Second, the child's gender and parental age are predictors of parents' sensitivity to stories involving autism risks because boys and children of older parents are known to be at a much higher risk of autism than their counterparts.[2] Importantly, the identifying assumption behind this approach is that the least sensitive parents and the most sensitive parents would have changed their vaccination behavior in the same direction in the absence of the surge in fake news.

Results indicate that the surge in fake news caused susceptible parents to become 3.3 percentage points less likely to vaccinate their children with an MMR shot by 15 months old. Importantly, this is the maximum age at which the CDC recommends the first MMR shot be administered. To assess whether parents were delaying the MMR shot or completely forgoing it, I examine the effects on take-up at 29 months old, which is the oldest age at which vaccination rates are consistently recorded in the survey. Results indicate that the resistance to the vaccine persisted. I estimate a 4.13 percentage point reduction in MMR shot take-up at 29 months old. This indicates that at a minimum, fake news caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children. These results are robust to including time-varying controls, state-specific linear time trends, and allowing family and state characteristics to have different effects on the MMR vaccine take-up rates in each year. In addition, I also find that the results are qualitatively similar when using more loosely defined treatment and control groups.

[2] Autism is four times more common among boys than girls (CDC,2007). Children of older parents could be as much as five times as likely to be on the autism spectrum than children of younger parents (Reichenberg et al., 2006; Durkin et al., 2008).

Indeed, the estimated reduction in vaccine take-up of 3 to 4 percentage points (4 to 4.5 percent), which is likely an underestimate given the approach, is economically meaningful. Lo and Hotez (2017) estimate that a 5 percent decline in the MMR vaccine coverage would result in a three-fold increase in annual measles outbreaks.

In providing the first evidence that an increase in fake news can lead to meaningful changes in behavior, this paper contributes to two bodies of literature. First, it complements the literature studying vaccine controversies. Smith, Ellenberg, Bell, and Rubin (2008), Anderberg, Chevalier, and Wadsworth (2011), and Chang (2018) study the impact of the vaccine controversy in 1998 when the MMR vaccine was first linked to autism and found that the MMR vaccine take-up rate decreased after 1998. My study differs from these studies in that while the claim in 1998 was believed to be true given it was published in a prestigious medical journal, by 2007 this claim had been clearly refuted. In this way, while these studies estimated the effect of new information expected to be reliable, my study identifies the effect of blatant misinformation. Combined with previous findings, this paper shows that misinformation about vaccines reported by the media can affect people's decisions as much or even more than perceived reliable information.

Second, this paper complements research on misinformation and media bias. Consistent with the theoretical framework of fake news provided by Allcott and Gentzkow (2017), results here suggest that fake news can have important consequences. Importantly, these reductions in immunizations affect not only people's own welfare but also the welfare of those around them. Furthermore, my results also indicate that the general population does not easily detect misinformation, especially when it is reported by major media outlets. This is in line with the finding that consumers do not accurately determine the reliability of health content on the internet documented in Allam, Schulz, and Nakamoto (2014), Knapp, Madden, Wang, Sloyer, and Shenkman (2011), and Kutner, Greenburg, Jin, and Paulsen (2006). Lastly, this paper also speaks to related literature on the effects of media bias (DellaVigna and Kaplan, 2007; Gentzkow and Shapiro, 2006, 2010; Gerber, Karlan, and Bergan, 2009; Chiang and Knight, 2011; Enikolopov, Petrova, and Zhuravskaya, 2011; Prat, 2017; Martin and Yurukoglu, 2017). These studies built theoretical frameworks and provided empirical evidence that media slant can change individual beliefs and behavior. My results show that, in addition to media slant, completely false information reported by the media can also change behavior, even when it is easy for both the media and consumers to verify the information is wrong.

2 Background: Media Coverage of the Anti-Vaccination Movement in the US

Although vaccines are regarded as one of the most successful medical interventions of the 20th century (CDC, 1999), some opposition to vaccines has always existed (Hussain et al., 2018). In 1998, however, the claim that vaccines are dangerous was propelled into the mainstream by the media when an article by Wakefield et al. (1998) suggested a causal link between the MMR vaccine and autism. The article was published in the Lancet, a major British medical journal. Anderberg, Chevalier, and Wadsworth (2011) studied the effects of this 1998 vaccine controversy and found that the MMR vaccine take-up rate declined sharply in the immediate years following the controversy. While the controversy did not garner as much media attention in the US as in the UK, Smith, Ellenberg, Bell, and Rubin (2008) and Chang (2018) also observed that the MMR take-up rates in children 19-35 months old in the US dropped by approximately 1-2 percentage points immediately following the Wakefield publication, but returned to pre-controversy levels by 2003. Importantly, the article was eventually retracted by the Lancet in 2010 after several subsequent studies disproved its results.

In the US, the topic of vaccine safety gained popularity again in 2007 when the media coverage on vaccine safety increased dramatically. This rise in the coverage was due in part to several vaccine court hearings of a case alleging that vaccines cause autism, and in part to the increasing number of celebrities publicly claiming that vaccines cause autism. Notably, Jenny McCarthy, an actress and TV host, famously went on talk shows including the Oprah Winfrey Show to talk about her belief that the MMR vaccine causes autism and how her son got diagnosed with autism after the MMR shot. For example, during the interview with Winfrey, McCarthy talked about her experience:

"Right before his MMR shot, I said to the doctor, I have a very bad feeling about this shot. This is the autism shot, isn't it? And he said, 'No, that is ridiculous. It is a mother's desperate attempt to blame something on autism.' And he swore at me.... And not soon thereafter, I noticed that change in the pictures: Boom! Soul, gone from his eyes."

Mnookin (2011) estimated McCarthy's message to have reached at least 15-20 million viewers based on her appearance on The Oprah Winfrey Show, Larry King Live, and Good Morning America alone.

Figure 1 shows the number of news coverage on six major television networks (ABC, CBS, NBC, CNN, MSNBC, and FNC) on the topic of vaccine safety from 2001 to 2012. As stated earlier, the coverage was few and far in between from 2001 to 2006 before rising dramatically in 2007.

A critical aspect of the surge in media coverage on vaccine safety in 2007 is that at that point prominent medical bodies had already refuted the claim of any link between vaccines and autism. This includes the Institute of Medicine (IOM) in May 2004, the Food and Drug Administration (FDA) in September 2006, and the Centers for Disease Control and Prevention (CDC) in July 2007. In addition, as alluded earlier, the paper that had initially proposed the link had been disproved by multiple papers. Despite all that, Figure 1 shows that the majority of television coverage from 2007 to 2009 only reported on the alleged claim of vaccines' link to autism without refuting it. This means that although the alleged link between vaccines and autism had been thoroughly debunked, the public was dramatically and increasingly exposed to the false claim in 2007 by television networks. I leverage this unanticipated increase in misinformation to estimate the causal impact of misinformation by the media or fake news.

3 Data

To analyze the exposure to fake news, I look at the number of television news stories on the alleged link between vaccines and autism over time. I use coverage on major television networks as a proxy for media coverage because although many people get their news through other sources, 44% of Americans still report television as the platform the most preferred for news (Mitchell, 2018). I obtained the news transcripts of six major television networks in the US from January of 2001 to December of 2012 from Lexis Nexis. The six networks were ABC, CBS, NBC, CNN, MSNBC, and FNC. To get at the number of news reporting on vaccines' alleged link to autism, I first identify the news that mentions vaccines (or vaccination) and autism in the same section and hired two research assistants to read these news transcripts. The research assistants then individually read all the relevant news transcripts and classified the news into three categories: 1. news stories reporting a link between vaccines and autism or the possibility of the link, 2. news stories explicitly refuting the claim that vaccines cause autism, and 3. news stories that coincidentally mentioned vaccines and autism in the same section but do not actually report on vaccine safety.[3] For each news transcript, if both of the research assistants classified the story as reporting the link between vaccines and autism, I counted it as one fake news story. If both classified it as refuting the false claim, I count it as not a fake news story. If only one of them classified it as reporting on the link between vaccines and autism, I count it as half of a fake news story.

To identify the impact of fake news regarding vaccines on individual behavior, I look at parents' decisions regarding vaccination. In particular, since the MMR vaccine is the vaccine at the center of the vaccine-autism claim, I look at the MMR vaccine take-up rate as my main outcome. Individual-level data on vaccination decisions used in this paper come from the 2002-2012 National Immunization Survey (NIS), which is conducted yearly by the Centers for Disease Control and Prevention (CDC). For each survey, the CDC surveys parents of 19-35 month-old children about their children's vaccination history. In addition, the CDC also asks for consent to obtain the vaccination records from their medical providers. Approximately 70% of the parents consent to the CDC acquiring vaccination records from their healthcare providers. Since healthcare provider records offer much more accurate information than parents' memory or a shot card, I only include children whose provider data is available in my analysis. For the analysis in this paper, I only include the data starting from 2002 to avoid the confounding effects from the first MMR vaccine controversy in 1998 when the Wakefield et al. paper first published. I only include the data up until 2012 because I only have media data up until 2012. I show in the Appendix section that the results are robust to alternative starting and ending years.

The National Immunization Surveys classify children into three age groups: 19-23 month olds, 24-29 month olds, and 30-35 month olds. I use the vaccination information of children from all age groups, i.e. all 19-35 month olds whose provider data is available, to look at the MMR take-up rate at 15 months old. Since the CDC recommends that the first MMR shot is given to a child at 12-15 months old, looking at the MMR take-up rate at 15 months old allows me to see if parents follow the CDC's recommendation. In addition, it is also important to see if parents only delay vaccinating their children or refusing to vaccinate altogether. To address this question, I examine the MMR take-up rate of older children. The oldest children in my data set are 30-35 months old. This means

[3] The research assistants were instructed to sort and read the news transcripts in random order, rather than chronologically. This is to avoid any bias that could occur if they associate a certain time period with news of certain types.

that I have complete vaccination information up to when these children were 29 months old. I thus use the vaccination information of children 30-35 months old to look at the MMR take-up rate at 29 months old to see if parents have caught up to the vaccination schedule.

Table 1 provides summary statistics of children included in my analysis. Panel 1 reports on all children in the 2002- 2012 National Immunization Surveys whose provider data is available, i.e. all 19-35 month olds, while Panel 2 reports the statistics of only 30-35 month-old children. Overall, 78% of children are vaccinated with an MMR shot by 15 months old and 93% are vaccinated by 29 months old. This suggests that approximately 15% of parents do not strictly follow the CDC's recommendation, but eventually vaccinate their children. In addition, the vaccination rates at both ages are in general higher among the children most likely to be affected by fake news (boy/firstborn/mom/$\geq$30) than those least likely to be affected (girl/not firstborn/mom$<$30).

4 Empirical Method

4.1 Measuring Fake News Exposure and Identifying the Post Period

I begin my analysis by identifying first which cohorts of children were affected by the increase in fake news. I do so by looking at the number of fake news stories parents are exposed to. I first define the period when parents are most likely to pay attention to information about vaccine recommendations and vaccine safety as the 'news exposure period'. For each child, I consider the news exposure period to start in the month that the child was born and end in the month that I measure the child's MMR vaccine take-up. If I had information on each child's birthdate, I would identify each child's news exposure period and then count the number of fake news stories reported on television in this exposure period and use this number as a measure of parents' fake news exposure. However, although the National Immunization Survey (NIS) data is rich in many ways, it does not provide information on the date of birth, the date of the interview, or age at the time of the interview. Therefore, I cannot directly back out the birth month and calculate parents' fake news exposure for each child in my dataset individually. The NIS data does, however, provide information on which age group the child falls into at the time of the interview (19-23, 24-29, 30-35 months old). I thus calculate for the average news exposure for children in each age group in each interview year using this age group information along with two hypotheses. First, I assume that children of all ages are as equally likely to appear in the survey. Second, I assume that the probability of getting interviewed in each month is uniformly distributed throughout the year.

Figures 2 and 3 show the average fake news exposure of parents interviewed in each survey year. Figure 2 shows the average fake news exposure up until when the child was 15 months old. Panel A shows that for parents whose child was 19-23 months old at the time of the interview, the first cohort that experienced the surge in fake news was those interviewed in 2008. Panels B and C show the average fake news exposure of parents whose child was 24-29 months old and 30-35 months old at the time of the interview, respectively. Both panels show that for both groups of parents, the first cohort that experienced the surge in fake news was those interviewed in 2009. Figure 3 shows the average fake news exposure up until when the child was 29 months old. I only look at the average fake news exposure for parents whose child was 30-35 months old at the time of the interview here, because they are the only group with relevant information of child at 29 months old. We can see the average fake news exposure rose dramatically for the cohort interviewed in 2008.

4.2 Classifying Treatment and Control Groups

To identify the effects of fake news, ideally, we would compare a group that was randomly exposed to fake news to a group that was not exposed to fake news. However, this is difficult for several reasons. First, people usually choose what they watch on television. It could be the case that people who are less likely to vaccinate are the ones more likely to watch fake news about vaccines on television. Second, more than 95% of US homes have television service and therefore almost everyone was exposed to television and subsequently fake news to some degree. This makes it hard to identify a control group. In this paper, I overcome these issues by using a difference-in-differences approach that compares the groups that are ex-ante most and least sensitive to the fake news over time. Using this approach, I view the least sensitive group as my control group. The advantage of this approach is that I am able to distinguish the effect of fake news exposure from other common time-varying factors, as well as group-specific factors. The disadvantage is because all parents are to some extent treated, this approach will underestimate the effect of fake news on immunizations.

To identify which group of parents is the most sensitive and which group is the least sensitive to the fake news stories about vaccines, it is important to consider which factors would make some parents more sensitive to the news than others. Here, I propose that parents' sensitivity to fake news stories about vaccines is based on both their parenting experience and their child's risk of being on the autism spectrum. This is because when the surge in fake news happened, experienced parents would have already learned about vaccine safety from past experience. Therefore, by comparison to new parents who only started learning about vaccine safety after the surge in fake news, the fake news stories after 2007 would only account for a small percentage of information that experienced parents have and therefore would not be as impactful. Additionally, experienced parents are also more likely to have already formed their opinion on the issue from past experience and therefore less likely to be receptive to the new information than new parents. Hence, among parents in the same exposure period, experienced parents would likely be less sensitive to new information and thereby less affected by the increase in fake news.

Furthermore, since the fake news stories link vaccines to autism risk, parents whose child is at higher risk of being on the autism spectrum would likely be more sensitive to the news. In terms of autism risk, two characteristics—parental age and gender—have been consistently reported by both the CDC and media outlets to be associated with higher autism risk. For example, the CDC put out a press release in February of 2007 stating that the autism spectrum disorder is 3-5 times more common among boys than girls (CDC, 2007). And several news networks reported on a study by Reichenberg et al. (2006) that found that children of men over 40 years old were 5.75 times more likely to have autism spectrum disorder compared with children of men under 30 years old.[4] A large study by Durkin et al. (2008) also found that firstborn children of 2 older parents were 3 times more likely to develop autism than were third- or later-born offspring of mothers aged $20 - 34$ years and fathers aged $<$40 years.

I, therefore, determine parents' sensitivity to the news using three predetermined characteristics: whether the child is a firstborn, a boy, and the mother is over 30 years old. Mother's age is used as a proxy for parental age as that is the only consistent information about parental age available from the survey and the majority of couples are not more than 5 years apart in age.[5] Parents are classified as most sensitive to the fake news if they have all these three characteristics present and

[4]McNamara, M. (2006) 'Men's Biological Clocks Are Ticking, Too', CBS, 15 November (https://www.cbsnews.com/news/mens-biological-clocks-are-ticking-too/)
Robin, R. (2007) 'It Seems the Fertility Clock Ticks for Men, Too', The New York Times, 27 Feb (https://www.nytimes.com/2007/02/27/health/27sper.html)

[5]Based on the 2013 Current Population Survey, for 76.7% heterosexual married couples, the husband and wife are less than 5 years apart in age.

least sensitive if they have none. As a result, within my sample, I define the group that is the most sensitive to the fake news as boys who are a firstborn and whose mom is over 30 years old and the group that is the least sensitive as girls who are not a firstborn and whose mom is younger than 30 years old.

Now that I have defined the most treated and the least treated group, I will use them as my treatment and control groups respectively with a generalized difference-in-differences approach to identify the impact of fake news. Specifically, I compare the MMR vaccine take-up rates of boys who are a firstborn and whose mom is over 30 years old to the take-up rates of girls who are not a firstborn and whose mom is younger than 30 years old before and after the surge in misinformation about vaccines. Formally, I estimate the impact of the dramatic increase in fake news on parents' decision to vaccinate their child using the following model:

$$MMR_{it} = \alpha_t + \theta MostSensitive_i + \beta_x X_{it} + \beta MostSensitiveXPost_{it} + u_{it} \tag{1}$$

Where the outcome, MMR_{it}, is a binary variable indicating whether child i whose parent was interviewed in year t has been given at least one shot of MMR vaccine. In this paper, I focus on looking at this outcome at two points in time: when child i was 15 months old and 29 months old. I look at whether child i has been given any MMR shot at 15 months old because the CDC recommends that parents vaccinate their children with a dose of MMR vaccine at 12-15 months old, and therefore this will show whether parents stop following the CDC's recommendation. Additionally, it is also important to see whether fake news has long-term effects on vaccination take-up, because if fake news only results in parents delaying the vaccination, perhaps it is not as harmful. This is why I use children in the oldest age group, 30-35 months old at the time of the interview, to look at MMR take-up rate at 29 months old.

α_t is survey year fixed effects. $MostSensitive_i$ is an indicator variable for whether child i's parents are classified as the most sensitive, i.e. whether child i is a boy, a firstborn, and has a mother who is 30 over years old. X_i is a matrix containing child i's characteristics including state fixed effects, race, poverty status, mother's education, mother's marital status, child's age at the time of the interview, whether they live in a state they were born in, and whether their state allows personal belief exemption from vaccination. $MostSensitiveXPost_{it}$ is an indicator variable for whether child i is in the most sensitive group in the post period. The post-period starts in the year when we first see the dramatic increase in fake news exposure as discussed in Section 4.1. Importantly, the coefficient of interest here is β which measures the effects of fake news on parents' decision to vaccinate. Specifically, it measures whether parents most sensitive to the surge in fake news vaccinate their children differently than parents who are the least treated.

In all specifications, survey weights are used and robust standard errors and their corresponding p-values are reported. In addition, accounting for within-cluster dependence in estimating standard errors of regression estimates is important. Ideally, we want to cluster at the level of treatment or higher. However, since I only have two clusters, I follow the wild bootstrap method proposed in Cameron, Gelbach, and Miller (2008) which clusters at the year level. These wild-bootstrap p-values are reported for all specifications.

As with any difference-in-differences design, the underlying assumption for this approach is that MMR vaccine take-up rates of children in the control group and treatment group would have changed similarly over time in the absence of the increase in fake news. I provide support for this assumption by first showing the visual representation of the raw data that shows the MMR take-up rates for control and treatment groups track each other prior to the post period. Second, I also formally test for the divergence in outcomes between the treatment and control groups in the pre periods using a dynamic difference-in-differences approach.

One potential concern with my approach is that since I am only using the two extremes as my treatment and control groups, it is possible that my results are dependent on how I define treatment and control groups. To provide further support for my identification strategy, I also perform multiple analyses using more loosely defined treatment and control groups. Specifically, I do this in 3 different ways. First, I include more children in my control group. Namely, instead of excluding children in the middle between the two extremes, I include them in my control group. Second, I include more children in my treatment group, i.e. instead of excluding children in the middle between the two extremes, I include them in my treatment group. And lastly, I use two instead of three predetermined characteristics to determine treatment and control groups. With more loosely defined treatment and control groups, we would expect the effects to be weaker, but not completely disappear.

5 Results

5.1 Main Results

I begin by looking at the raw data of the MMR vaccine take-up rates over the year. Figures 4 and 5 show the MMR vaccine take-up rates at 15 months old and 29 months old, respectively. Time is recentered so that year=+1 is the first year parents experienced the surge in fake news exposure. For both the MMR vaccine take-up rates at 15 ad 29 months old, Figures 4 and 5 show that prior to the surge in fake news exposure, the take-up rates among children in the treatment group (boys who are a firstborn and whose mom is over 30 years old) and control group (girls who are not a firstborn and whose mom is younger than 30 years old) track each other well over the years. This is important since the validity of a difference-in-differences approach hinges on the parallel trend assumption. Additionally, the figures also show that before the increase in fake news exposure, children in the treatment group are consistently more likely to be vaccinated than children in the control group both at 15 months old and 29 months old. However, after the increase in fake news, the gap in vaccination rates between the two groups closes. The gap closes by about half for the MMR vaccination rate at 15 months old and closes completely for the MMR vaccination rate at 29 months old.

To check the parallel trends assumption more rigorously, I estimate a dynamic difference-in-differences model, controlling for year fixed effects and children characteristics, to check if the treatment group diverges from the control group in any years before the increase in fake news exposure. Figures 6 and 7 plot the dynamic difference-in-differences estimates for MMR vaccine take-up at 15 months old and 29 months old respectively. Both figures reaffirm that for both outcomes, there is no evidence of divergence in trends before the increase in fake news exposure. In addition, both figures also show that after the increase in fake news exposure, both the MMR take-up rates at 15 months old and 29 months old of children in the treatment group fall. This suggests that increased exposure to fake news about vaccine safety does not only lead parents to deviate from the CDC's recommended schedule, but also delays vaccination by a minimum of a year, and possibly much longer.

Next, I formally estimate the average treatment effects of the increase in exposure to fake news and report the results in Table 2. Column 1 shows the average treatment effect of fake news on the MMR vaccination rates using the simplest difference-in-differences model, without any controls. Based on this specification, it appears that the rise in fake news causes the MMR take-up rates at 15 months old to drop by 4.57 percentage points and the MMR take-up rates at 29 months old to drop by 4.53 percentage points. Column 2 reports the estimates from the preferred specification, shown in equation 1, which also includes controls for observable characteristics, state fixed effects, and state vaccination exemption law. If my results are driven by the change in the characteristics of children in my control or treatment groups and not by the increased exposure to fake news, then these controls should absorb my treatment effects. The estimates from this specification are only slightly smaller than those reported in column one but are still in the same direction and statistically significant. Based on these estimates, the increased exposure in fake news causes the MMR take-up rates at 15 and 29 months old to decreases by 3.27 percentage points and 4.13 percentage points, respectively. Finally, families with different characteristics, such as income, parents' education level, and race, may respond differently to year-to-year shock. For example, richer parents might have better access to vaccines in the year where there is a vaccine shortage. Since my treatment and control groups are considerably different in terms of family income and mother's education, in Column 3, I allow observable characteristics to affect the MMR take-up rate differently each year. The estimate reported from this specification for the MMR take-up rate at 15 months old is no longer statistically significant at the conventional level but the magnitude still remains at the similar level of -2.31 percentage points. The estimate for the effect on the MMR take-up rate at 29 months old is robust and remains stable at a statistically-significant 4.16 percentage points reduction. This shows that the effects were not driven by the differences in characteristics between the two groups. In this table, wild-bootstrap p-values, which allow the correlation between take-up rates within the same year, are also reported alongside with the robust p-values. As shown in the table, wild-bootstrap p-values and robust p-values are very similar and using the wild-bootstrap approach does not change my results.

Overall, these results suggest that misinformation about vaccines' link to autism causes both the MMR vaccination rates at 15 months old and 29 months old to drop by at least 3-4 percentage points. This indicates that at a minimum, fake news caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children.

5.2 Robustness Checks

5.2.1 Effects on MMR Take-up at Other Ages

In addition to the main results discussed in Section 5.1, I also look at the effects of fake news on the MMR take-up rates at other ages besides 15 and 29 months old. I estimate the average treatment effects on the MMR vaccine take-up rates at each age from 15-29 months old using the preferred specification shown in equation 1. For the estimate at each age, I only include children who at the time of the interview are older than the age at which I measure the MMR vaccine take-up. This is because we only have information on the vaccination history of each child up until the time of the interview. For example, when the outcome is the MMR vaccine take-up at 20 months old, I only include children who were older than 20 months old at the time of the interview in the analysis. Since there are three age groups of children in my dataset: 19-23 months, 24-29 months, 30-35 months, this means that only children in age groups 24-29 and 30-35 months old are included in the analysis of the MMR vaccine take-up rate at 20 months old. In addition, since the age at which I measure the MMR vaccine take-up changes the exposure period, specifically, the exposure period relevant for when I look at MMR take-up at z months old as an outcome would be from when the child was born until when the child was z months old and not after, I also reexamining the news exposure period, the fake news exposure and revise the first post year for each estimation.

The results, reported in Table 3 and visually in Figure 8, show that the estimates are relatively similar across ages. They are all negative and range from -1.3 to -4.6 percentage points with 80% of them being statistically significant from zero at the 10% level. This indicates that the negative effects of fake news observed in the earlier section are not driven by the selection of the 15 month and 29 month ages.

5.2.2 Using More Loosely Defined Control and Treatment Groups

In the main analysis, I compare children who are most and least likely to be affected by the treatment. I classify children into these two groups using three characteristics that are associated with susceptible parents; whether the child is a firstborn, the child's gender, and the parents' age. Children with all of these three characteristics presents are classified as most likely to be affected whereas children with none of these characteristics are classified as least likely to be affected and these two groups are then used as my treatment and control groups. In this section, I perform multiple analyses using more loosely defined treatment and control groups to test the robustness of my findings to alternative classifications. As explained in the earlier section, I redefine my control and treatment groups in three major ways: 1. including more observations in the control group, 2. including more observations in the treatment groups, and 3. defining treatment and control groups using only two risk factors. Using the more loosely defined treatment and control groups, we would expect to see the treatment effects become smaller in magnitude, but not completely disappear.

The results of this exercise for the MMR take-up rate at 15 months old are reported in Table 4 and the same results for the MMR take-up at 29 months old are reported in Table 5. Column 1 shows the results of the main identification strategy. Columns 2-3 show the estimates when I add more children into my control groups by including children with only one or two of the three characteristics associated with susceptible parents in the control group as well. Columns 4-5 show the estimates when I increase my treatment group by including children with only one or two of the three characteristics associated with susceptible parents in my treatment group. Columns 6-8 show the estimates when I only use 2 characteristics in defining my control and treatment groups. For any 2 characteristics I use, my treatment group is the children with both 2 characteristics present and the control group is the children with neither of the 2 characteristics present. All the estimates reported are, as expected, smaller in magnitude than the estimates from the main identification. And although some estimates lost its significance at the conventional level, all of them are still negative, which is in the same direction as the effects found using the main identification strategy, and all but one of them still report a relatively low p-value. In particular, the estimates for the MMR take-up at 29 months old are very robust to alternative definitions of treatment and control groups.

5.2.3 Other Robustness Checks

In addition, since the dependent variable is binary, I also use logistic regression to estimate my main results. The results are shown in Table A1 in the Appendix section. Similar to the linear regression results, the logistic regression results show reductions in the MMR vaccine take-ups. Furthermore, currently, I am using data from 2002-2012. I provide further evidence in Tables 2A-6A in the Appendix showing that changing the start or end years does not change the results. The estimates become a little smaller in terms of magnitude when I extend the analysis to include years up to 2014. This could be because parents in the later years did not get as much exposure to the fake news about vaccine safety. As observed in Figures 1 and 2, the exposure to fake news about vaccine safety has been decreasing since 2010-2011.

6 Discussion and Conclusion

This paper studies the effect of fake news on individuals using the unanticipated rise in television coverage of the alleged link between vaccines and autism in 2007 as an exogenous shock in misinformation to parents. Using vaccination data obtained from healthcare providers of 19-35-month-old children surveyed in the 2002-2012 National Immunization Surveys (NIS), I find that fake news about vaccine safety resulted in a drop of at least 3.3 percentage points (4.2 percent) in the MMR vaccine take-up rate at 15 months old which is the CDC's recommended age. In addition, fake news also led to a drop of at least 4.1 percentage points (4.4 percent) in the MMR vaccine take-up rate at 29 months old. This indicates that at a minimum, fake news caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children.

The estimates here are economically meaningful, especially considering that my identification strategy of comparing the most and least sensitive groups likely results in the underestimation of effects. Lo and Hotez (2017) estimate that a similar-sized reduction in the MMR vaccination rate would result in a three-fold increase in annual measles outbreaks. Importantly, results here suggest that people can change behavior in important ways that not only affect their own welfare but also the welfare of those around them. In addition, my estimates are comparable or even bigger than the reported effects of new and reliable information found in prior literature. For example, Smith, Ellenberg, Bell, and Rubin (2008) estimate that the number of American children who received all childhood immunizations except for the MMR vaccine rose from 0.8 percent to 2.1 percent after the publication of Wakefield et al. (1998) which first suggested a link between the MMR vaccine and autism. Chang (2018) also examines the effects of the 1998 vaccine controversy in the US and estimate that the overall MMR vaccine take-up declined by 1.1 to 1.5 percentage points in the immediate year following the Wakefield et al. (1998) publication. Additionally, Chang (2018) also finds that an increase of 10 news stories about the vaccine controversy led college-educated mothers to be 0.4 percent less likely to vaccinate their children with an MMR shot. Combined with these findings, my results thus suggest that blatant misinformation, when reported by the media, can change individual behavior as much as perceived reliable information and that the general public is not able to discern false information even when it is easy to verify.

These results also have clear relevance for public policy regarding fake news. Much of the debate over the responsibility of social media companies and the government in combating fake news depends on whether the news actually matters. Results presented here provide clear evidence that fake news can change a behavior that not only affects those individuals but also potentially imposes negative externalities on those around them. This suggests that there are potentially large social benefits from preventing the dissemination of fake news.

References

Allam, A., P. J. 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*(4), e100.

Allcott, H. and M. Gentzkow (2017). Social media and fake news in the 2016 election. *Journal of economic perspectives 31*(2), 211–36.

Anderberg, D., A. Chevalier, and J. Wadsworth (2011). Anatomy of a health scare: education, income and the mmr controversy in the uk. *Journal of Health Economics 30*(3), 515–530.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics 90*(3), 414–427.

Chang, L. V. (2018). Information, education, and health behaviors: Evidence from the mmr vaccine autism controversy. *Health economics 27*(7), 1043–1062.

Chiang, C.-F. and B. Knight (2011). Media bias and influence: Evidence from newspaper endorsements. *The Review of Economic Studies 78*(3), 795–820.

DellaVigna, S. and E. Kaplan (2007). The fox news effect: Media bias and voting. *The Quarterly Journal of Economics 122*(3), 1187–1234.

Durkin, M. S., M. J. Maenner, C. J. Newschaffer, L.-C. Lee, C. M. Cunniff, J. L. Daniels, R. S. Kirby, L. Leavitt, L. Miller, W. Zahorodny, et al. (2008). Advanced parental age and the risk of autism spectrum disorder. *American journal of epidemiology 168*(11), 1268–1276.

Enikolopov, R., M. Petrova, and E. Zhuravskaya (2011). Media and political persuasion: Evidence from russia. *American Economic Review 101*(7), 3253–85.

Fine, P. E. (1993). Herd immunity: history, theory, practice. *Epidemiologic reviews 15*(2), 265–302.

for Disease Control, C., P. (CDC, et al. (1999). Ten great public health achievements–united states, 1900-1999. *MMWR. Morbidity and mortality weekly report 48*(12), 241.

for Disease Control, C., Prevention, et al. (2007). Prevalence of the autism spectrum disorders (asds) in multiple areas of the united states, 2000 and 2002. *Community Report from the Autism Developmental Disabilities Monitoring (ADDM) Network*.

Gentzkow, M. and J. M. Shapiro (2006). Media bias and reputation. *Journal of political Economy 114*(2), 280–316.

Gentzkow, M. and J. M. Shapiro (2010). What drives media slant? evidence from us daily newspapers. *Econometrica 78*(1), 35–71.

Gerber, A. S., D. Karlan, and D. Bergan (2009). Does the media matter? a field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics 1*(2), 35–52.

Hussain, A., S. Ali, M. Ahmed, and S. Hussain (2018). The anti-vaccination movement: A regression in modern medicine. *Cureus 10*(7).

Knapp, C., V. Madden, H. Wang, P. Sloyer, and E. Shenkman (2011). Internet use and ehealth literacy of low-income parents whose children have special health care needs. *Journal of medical Internet research 13*(3), e75.

Kutner, M., E. Greenburg, Y. Jin, and C. Paulsen (2006). The health literacy of america's adults: Results from the 2003 national assessment of adult literacy. nces 2006-483. *National Center for Education Statistics*.

Lazer, D. M., M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, et al. (2018). The science of fake news. *Science 359*(6380), 1094–1096.

Lo, N. C. and P. J. Hotez (2017). Public health and economic consequences of vaccine hesitancy for measles in the united states. *JAMA pediatrics 171*(9), 887–892.

Martin, G. J. and A. Yurukoglu (2017). Bias in cable news: Persuasion and polarization. *American Economic Review 107*(9), 2565–99.

Mitchell, A. (2018). Americans still prefer watching to reading the news-and mostly still through television. *Pew Research Center, Washington, DC*.

Mnookin, S. (2011). *The panic virus: a true story of medicine, science, and fear*. Simon and Schuster.

Prat, A. (2017). Media power. forthcoming. *Journal of Political Economy*.

Reichenberg, A., R. Gross, M. Weiser, M. Bresnahan, J. Silverman, S. Harlap, J. Rabinowitz, C. Shulman, D. Malaspina, G. Lubin, et al. (2006). Advancing paternal age and autism. *Archives of general psychiatry 63*(9), 1026–1032.

Smith, M. J., S. S. Ellenberg, L. M. Bell, and D. M. Rubin (2008). Media coverage of the measles-mumps-rubella vaccine and autism controversy and its relationship to mmr immunization rates in the united states. *Pediatrics 121*(4), e836–e843.

7 Figures

Figure 1: Number of television coverage on the topic of vaccines and its link to autism

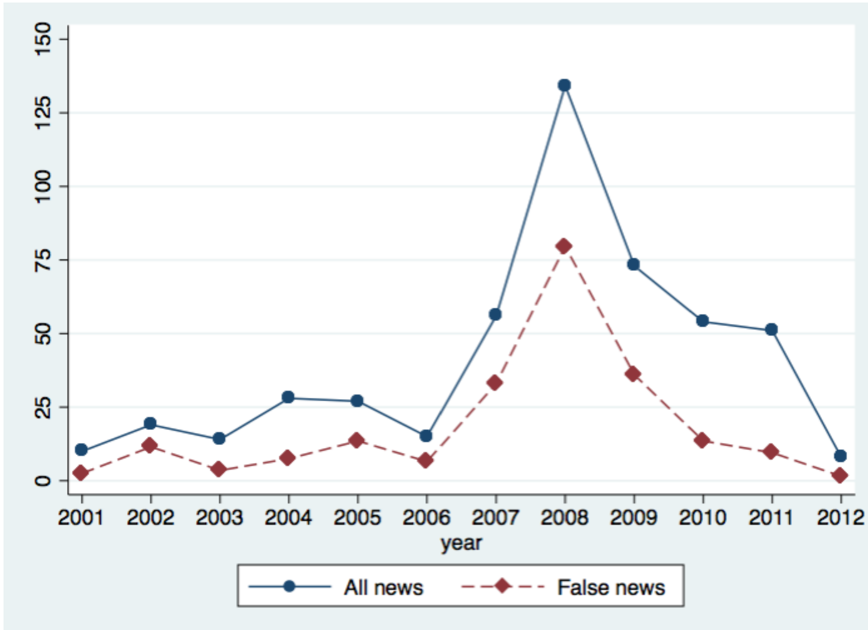

Source: LexisNexis
These numbers are based on coverage on 6 television networks: ABC, CBS, NBC, CNN, MSNBC, Fox News.
False news refers to the news that does not explicitly refute the claim that vaccines cause autism.

Figure 2: Fake news exposure from when child was born to 15 months old

(A) 19-23 months old at time of interview

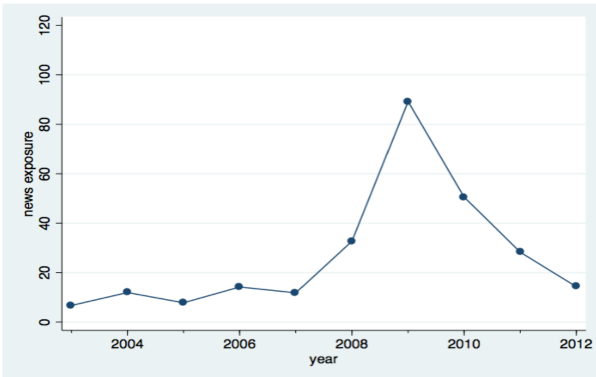

(B) 24-29 months old at time of interview

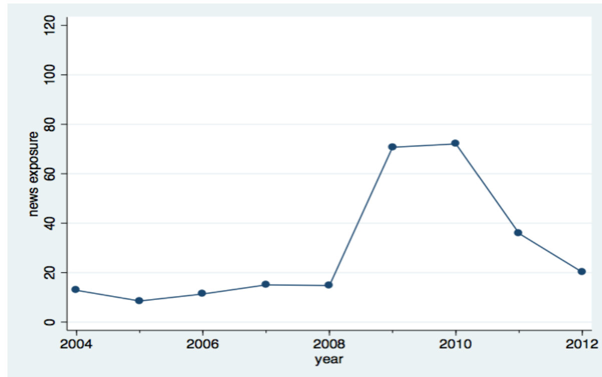

(C) 30-35 months old at time of interview

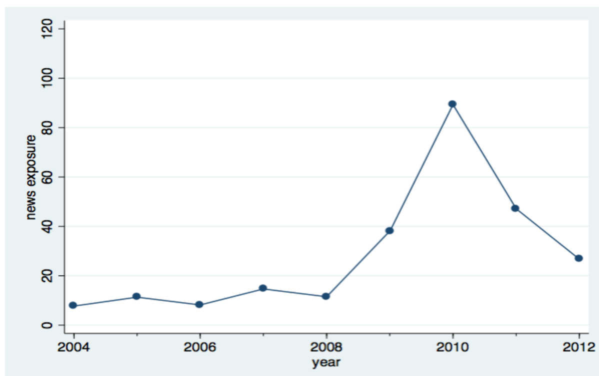

Figure 3: Fake news exposure from when child was born to 29 months old

Child was 30-35 months old at time of interview

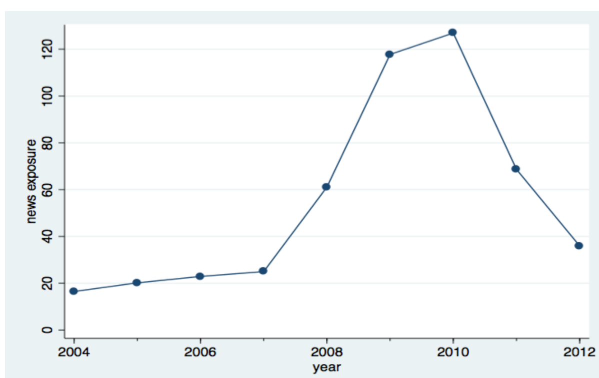

Figure 4: MMR take-up rate at 15 months old

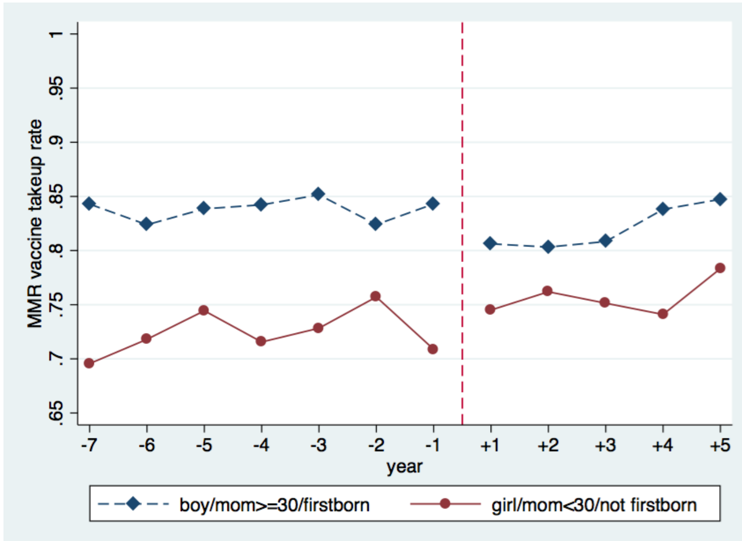

Figure 5: MMR take-up rate at 29 months old

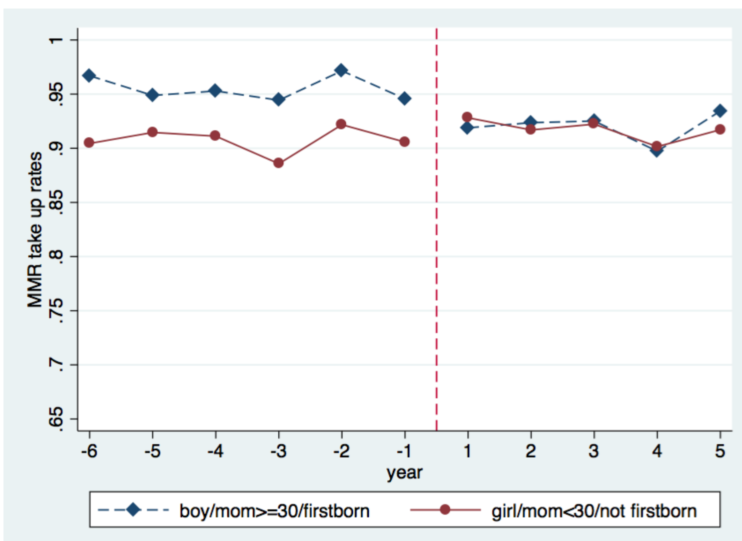

Figure 6: Dynamic difference-in-differences estimates for MMR take-up rate at 15 months old

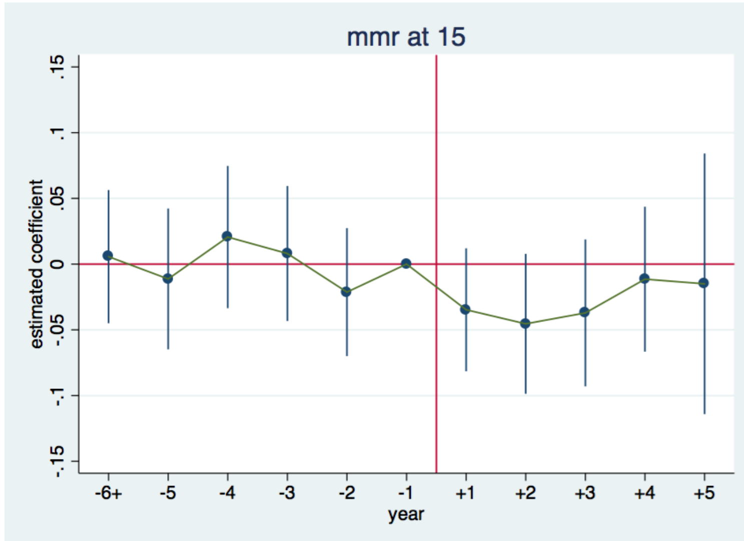

Figure 7: Dynamic difference-in-differences estimates for MMR take-up rate at 29 months old

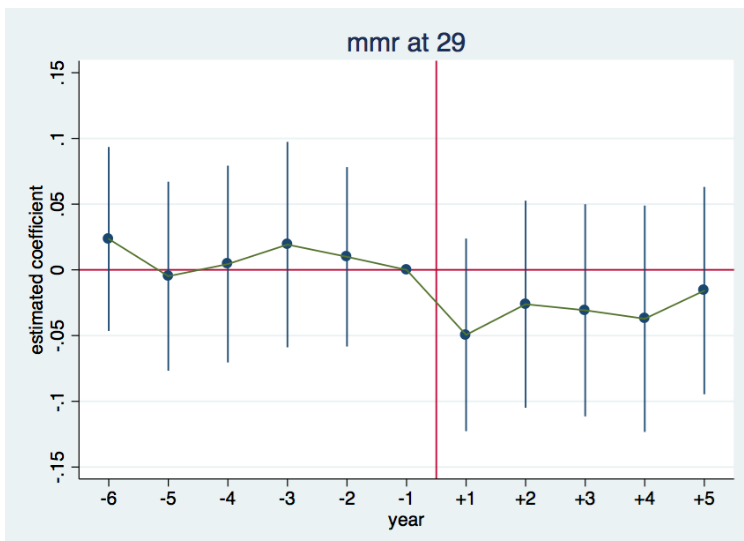

Figure 8: Estimated effects of fake news on MMR vaccine take-up rate at age 15-29, using the main specification (3 risk factors present vs. 0 risk factor present)

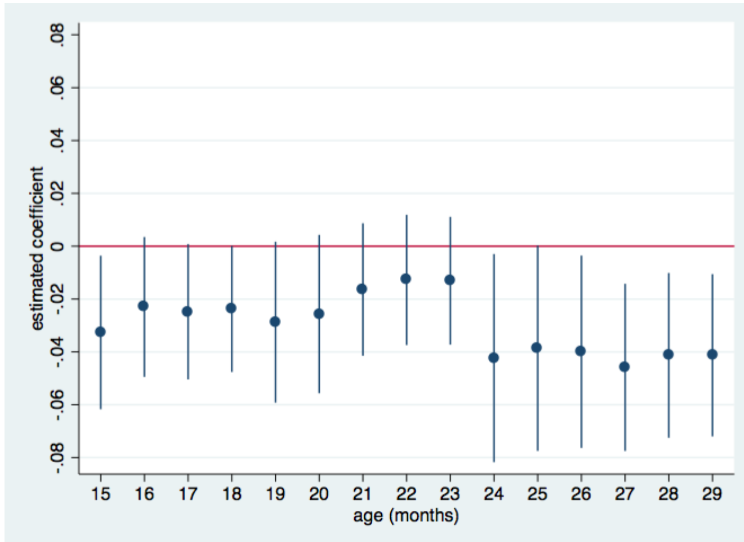

8 Table

Table 1: Summary Statistics

Panel 1: children 19-35 months old

	All		**Least Sensitive to Fake News**		**Most Sensitive to Fake News**	
	mean	sd	mean	sd	mean	sd
MMR shot at 15 months	0.78	(0.41)	0.74	(0.44)	0.83	(0.38)
Male	0.51	(0.50)	0.00	(0.00)	1.00	(0.00)
Firstborn	0.43	(0.49)	0.00	(0.00)	1.00	(0.00)
Mother $\geq$ 30	0.56	(0.50)	0.00	(0.00)	1.00	(0.00)
White	0.73	(0.44)	0.68	(0.47)	0.76	(0.43)
Black	0.15	(0.36)	0.20	(0.40)	0.11	(0.31)
In poverty	0.31	(0.46)	0.50	(0.50)	0.13	(0.34)
Mother with college degree	0.31	(0.46)	0.09	(0.28)	0.56	(0.50)
Mother is married	0.68	(0.46)	0.52	(0.50)	0.83	(0.37)
19-23 months old	0.30	(0.46)	0.32	(0.47)	0.29	(0.46)
24-29 months old	0.34	(0.47)	0.34	(0.47)	0.34	(0.47)
30-35 months old	0.36	(0.48)	0.34	(0.47)	0.36	(0.48)
Moved state after birth	0.08	(0.27)	0.08	(0.27)	0.08	(0.28)
Observations	196684		16987		22239	

Panel 2: children 30-35 months old

	All		**Least Sensitive to Fake News**		**Most Sensitive to Fake News**	
	mean	sd	mean	sd	mean	sd
MMR shot at 29 months	0.93	(0.26)	0.91	(0.28)	0.94	(0.24)
Male	0.51	(0.50)	0.00	(0.00)	1.00	(0.00)
Firstborn	0.42	(0.49)	0.00	(0.00)	1.00	(0.00)
Mother $\geq$ 30	0.58	(0.49)	0.00	(0.00)	1.00	(0.00)
White	0.73	(0.44)	0.68	(0.47)	0.76	(0.43)
Black	0.15	(0.36)	0.22	(0.41)	0.10	(0.30)
In poverty	0.30	(0.46)	0.49	(0.50)	0.13	(0.34)
Mother with college degree	0.31	(0.46)	0.09	(0.28)	0.54	(0.50)
Mother is married	0.69	(0.46)	0.51	(0.50)	0.82	(0.38)
Moved state after birth	0.09	(0.29)	0.09	(0.29)	0.09	(0.29)
Observations	70702		5655		8196	

Source: 2002-2012National Immunization Surveys
All estimates obtained using sampling weights provided by the National Immunization Survey.
The 'least sensitive to fake news' group refers to girls who are not a firstborn and whose mother is <30 years old. The 'most sensitive to fake news' group refers to boys who are a firstborn and whose mother is $\geq$ 30 years old.

Table 2: Effects of fake news on MMR take-up rates

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0457***	-0.0327**	-0.0231
	(0.0144)	(0.0148)	(0.0183)
P-value	0.0015	0.0276	0.2054
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	39226	39226	39226

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0453***	-0.0413***	-0.0416**
	(0.0154)	(0.0157)	(0.0187)
P-value	0.0033	0.0084	0.0264
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	13851	13851	13851

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008).
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table 3: Effects of fake news on MMR take-up rates at 15-29 months old

	15 months	16 months	17 months	18 months	19 months
MostSensitive X Post	-0.0327**	-0.0230*	-0.0248*	-0.0238*	-0.0288*
	(0.0148)	(0.0135)	(0.0131)	(0.0122)	(0.0155)
P-value	0.0276	0.0886	0.0575	0.0510	0.0641
Wild bootstrap p-value	0.0260	0.0200	0.0661	0.0320	0.1572
Outcome mean	0.78	0.82	0.84	0.87	0.88
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
N	39226	39226	39226	39226	27504

	20 months	21 months	22 months	23 months	24 months
MostSensitive X Post	-0.0257*	-0.0164	-0.0128	-0.0131	-0.0424**
	(0.0153)	(0.0128)	(0.0126)	(0.0123)	(0.0201)
P-value	0.0928	0.1996	0.3096	0.2888	0.0350
Wild bootstrap p-value	0.1882	0.3423	0.4715	0.3994	0.0130
Outcome mean	0.89	0.89	0.90	0.90	0.91
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
N	27504	27504	27504	27504	13851

	25 months	26 months	27 months	28 months	29 months
MostSensitive X Post	-0.0386*	-0.0400**	-0.0459***	-0.0413***	-0.0413***
	(0.0198)	(0.0186)	(0.0161)	(0.0159)	(0.0157)
P-value	0.0516	0.0314	0.0045	0.0094	0.0084
Wild bootstrap p-value	0.0220	0.0080	0.0080	0.0080	0.0020
Outcome mean	0.91	0.92	0.92	0.92	0.93
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
N	13851	13851	13851	13851	13851

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008).
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.
All estimates are obtained using the main specification, i.e. difference-in-difference with year fixed effects, state fixed effects, observable controls.

Table 4: Effects of fake news on MMR take-up rates at 15 months old with more loosely defined treatment and control groups

	Baseline	Increase control group		Increase treated group				
	3 characteristics vs. 0 characteristic	3 characteristics vs. 0/1 characteristic	3 characteristics vs. 0/1/2 characteristics	3/2 characteristics vs. 0 characteristic	3/2/1 characteristics vs. 0 characteristic	boy & mother≥30 vs. girl & mom<30	boy & firstborn vs. girl & not firstborn	mother≥30 & firstborn vs. mother<30 & not firstborn
MostSensitive x Post	-0.0327**	-0.0128	-0.0127	-0.0271***	-0.0250***	-0.0154	-0.0039	-0.0284***
	(0.0148)	(0.0110)	(0.0103)	(0.0102)	(0.0089)	(0.0098)	(0.0093)	(0.0106)
P-value	0.0276	0.2455	0.2186	0.0078	0.0050	0.1178	0.6777	0.0071
Wild bootstrap p-value	0.0260	0.1291	0.1291	0.0110	0.0050	0.2322	0.6587	0.0220
Outcome Mean	0.78	0.78	0.78	0.78	0.78	0.78	0.78	0.78
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	39226	114537	196684	121373	196684	98830	98013	78293

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008).
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table 5: Effects of fake news on MMR take-up rates at 29 months old with more loosely defined treatment and control groups

	Baseline	Increase control group		Increase treated group		Only using 2 characteristics to define treatment group		
	3 characteristics vs. 0 characteristic	3 characteristics vs. 0/1 characteristic	3 characteristics vs. 0/1/2 characteristics	3/2 characteristics vs. 0 characteristic	3/2/1 characteristics vs. 0 characteristic	boy & mother≥30 vs. girl & mom<30	boy & firstborn vs. girl & not firstborn	mother≥30 & firstborn vs. mother<30 & not firstborn
MostSensitive X Post	-0.0413***	-0.0275**	-0.0238**	-0.0253**	-0.0223*	-0.0226**	-0.0153	-0.0277**
	(0.0157)	(0.0114)	(0.0106)	(0.0128)	(0.0124)	(0.0106)	(0.0100)	(0.0109)
P-value	0.0084	0.0158	0.0249	0.0481	0.0726	0.0335	0.1281	0.0114
Wild bootstrap p-value	0.0030	0.0360	0.0220	0.0100	0.0240	0.1071	0.0500	0.0060
Outcome mean	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	13851	40628	70702	43925	70702	35649	35096	27659

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008).
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

9 Appendix

Table A1: Effects of fake news on MMR take-up rates using logistic regressions

	MMR at 15 months	MMR at 15 months	MMR at 15 months
outcome			
MostSensitive X Post	-0.2814***	-0.2112**	-0.1371
	(0.0877)	(0.0910)	(0.1122)
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	39226	39226	39226

	MMR at 29 months	MMR at 29 months	MMR at 29 months
outcome			
MostSensitive X Post	-0.7405***	-0.7010***	-0.6761**
	(0.2243)	(0.2285)	(0.2693)
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	13851	13851	13851

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table A2: Effects of fake news on MMR take-up rates: Data from 2001-2012

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0499***	-0.0375***	-0.0301*
	(0.0141)	(0.0144)	(0.0178)
P-value	0.0004	0.0093	0.0915
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	43548	43548	43548

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0465***	-0.0433***	-0.0438**
	(0.0151)	(0.0153)	(0.0184)
P-value	0.0020	0.0047	0.0170
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	15285	15285	15285

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table A3: Effects of fake news on MMR take-up rates: Data from 2003-2012

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0406***	-0.0279*	-0.0198
	(0.0148)	(0.0153)	(0.0188)
P-value	0.0062	0.0682	0.2930
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	35243	35243	35243

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0425***	-0.0401**	-0.0366*
	(0.0162)	(0.0165)	(0.0195)
P-value	0.0087	0.0154	0.0611
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	12492	12492	12492

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table A4: Effects of fake news on MMR take-up rates: Data from 2004-2012

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0442***	-0.0339**	-0.0258
	(0.0154)	(0.0159)	(0.0195)
P-value	0.0042	0.0332	0.1869
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	31241	31241	31241

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0451***	-0.0443**	-0.0370*
	(0.0173)	(0.0175)	(0.0206)
P-value	0.0092	0.0115	0.0721
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	11081	11081	11081

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table A5: Effects of fake news on MMR take-up rates: Data from 2002-2013

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0320**	-0.0237	-0.0216
	(0.0141)	(0.0145)	(0.0177)
P-value	0.0230	0.1027	0.2206
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	41823	41823	41823

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0381**	-0.0346**	-0.0378**
	(0.0152)	(0.0152)	(0.0175)
P-value	0.0119	0.0226	0.0306
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	14890	14890	14890

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table A6: Effects of fake news on MMR take-up rates: Data from 2002-2014

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0315**	-0.0237*	-0.0232
	(0.0134)	(0.0139)	(0.0167)
P-value	0.0191	0.0884	0.1646
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	44542	44542	44542

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0372**	-0.0351**	-0.0413**
	(0.0145)	(0.0147)	(0.0169)
P-value	0.0104	0.0167	0.0145
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	15994	15994	15994

*p<0.10, **p<0.05, ***p<0.010
Robust standard errors in parentheses.
Controls include state fixed effects, mother's age, firstborn status, race, poverty status, mother's education, mover status, state's personal exemption law.
All regressions are estimated using the sampling weights provided by the National Immunization Surveys.