

# Deceptive Claims Using Fake News Advertising: The Impact on Consumers

Anita Rao\*

Booth School of Business

University of Chicago

## Abstract

Fake news advertising can be harmful if it misleads consumers to take actions they otherwise would not have taken (e.g., purchase an inferior product). However, little is known whether fake news ads bring in new customers or are merely viewed by people already in the market for the advertised products. Exploiting an FTC enabled shutdown of fake news advertisements for various products such as acai cleanses and teeth whiteners but where the product sites continued to remain operational, I identify the extent of consumer interest, i.e., purchase propensity, in the presence and absence of fake news advertising. The findings indicate interest wanes after the shutdown of fake news, but there is some substitution to other channels such as regular advertisements.

---

\*Thanks to seminar participants at Boston University, University of Chicago Law and Economics Workshop, University of Florida and to participants at the 2019 Marketing Dynamics Conference, 2019 Marketing Science Conference, 2019 Workshop on Economics of Advertising and Marketing, 2019 Competition Policy Enforcement Conference, 2019 Midwest IO Fest and the 2019 Frank M. Bass UT Dallas FORMS Conference for their valuable suggestions. Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Email: [anita.rao@chicagobooth.edu](mailto:anita.rao@chicagobooth.edu)

# 1 Introduction

The use of fake news advertising has been prevalent for a long time. Advertising that mimics the format of its surrounding news content in both digital and print media, promotional segments aired on news programs without disclosure to the public,<sup>1</sup> advertisement sites posing as legitimate news channels are all examples of firms using fake news stories to deceive consumers into buying their products. Such deceptive practices are so prevalent that the FTC's guidelines explicitly states that ads are deceptive if they convey that they are independent, impartial or from a source other than the sponsoring advertiser. However, little is known on whether such fake news advertising affects consumers. Do such fake news stories change consumers' purchase propensity, or are they viewed by a select set of consumers already in the market for such products? The former, if true, is more harmful, because it implies a direct treatment effect of fake news advertising on consumers.

In this paper, using a set of fake news sites rendered inoperational by the FTC, I measure the impact of such campaigns on consumers' purchase propensity. Moreover, because these product domains were using both regular advertising and fake news advertising, I can disentangle the impact of fake news from that of regular advertising. In April 2011, the FTC brought to a halt the operations of 10 fake news website operating companies. These 10 companies operated over 150 fake news websites with names such as `onlinenews6.com` and `consumerdigestweekly.com`. The product domains to which these fake news sites referred consumers typically sold purported weight-loss products and colon-cleansing products. To propagate their products, these companies used both direct advertising through sites such as google and fake news "ads". The fake news sites falsely reported, in a journalistic manner, the positive impact of using the products. The FTC's actions were directed toward the fake news operating companies, and not the advertising companies, which were under no restraint and could continue to operate. The difference in product-domain visits before and after the FTC action identifies the causal impact of the visits due to the fake news.

If fake news advertising were the only means of reaching these product sites the answer to "what happens to product site visits after the fake news shutdown?" is trivial, i.e. products would face a 100% decline in visits after the FTC shutdown. However, consumers can reach these sites through organic links (e.g. google search) as well as regular non-fraudulent ads (e.g. legitimate sponsored ads on google). If, after the shutdown, consumers continue finding these sites at the same rate, then fake news advertising did not play an important role in product discovery. If this rate decreases substantially it implies the fake news ads were an important driver of product discovery without which consumers are no longer able to find

---

<sup>1</sup>FTC in re Synchronal Corp 1993, FCC fines local TV stations, WSJ 2011

these products. Therefore, apart from quantifying the total effect, this paper aims to identify the extent of substitution to other pathways (if any).

Using detailed browsing data from comScore, which are well-suited for this purpose because it also tracks referral domains, I can identify whether consumers visited the product's website directly, were referred from a fake news ad site, or were referred via regular ads (e.g., sponsored sites on google). Such data help identify whether consumers, after the shutdown of the fake news sites, were still able to reach the product website, and if so how. If consumers no longer visit the product websites, the fake news sites likely had a direct treatment effect, without which consumers were no longer able to find the sites. If consumers still visit these sites, I can identify if they reach them directly or are now being referred there by regular ads. To the extent that consumers reach these sites directly (e.g., through organic search results), advertisements are not influential in this setting; to the extent that consumers use regular ads to reach these sites, fake news ads and regular ads are likely to be substitutes. In either case, the implication is that the fake news sites did not have a treatment effect and that consumers are still able to reach these product sites. Quantifying the degree to which fake news sites had a direct treatment effect is crucial in understanding the extent of harm such campaigns cause, and therefore the extent of regulatory oversight required.

The findings indicate that the probability of merchants receiving a visit in a given month after the FTC shutdown decreases by 22%-41%. Total domain visits decrease by 12.5-31 visits per month, a 50% decline relative to the same period the previous year. To understand whether the intensity of the decline is different across various pathways, I break down the total visits into organic visits, fake news referrals and regular ad referrals. I find that referrals via fake news ads drop (mechanistically) and direct visits drop. However, referrals via regular ads exhibit heterogeneity across merchants with some merchants seeing a reduction in visits and others seeing an increase in visits. For those merchants that see an increase in visits, the average increase is 25 visits per month and is statistically significant, suggesting fake news and ads are substitutes at least in these product domains.

Quantifying the effect of the drop in visits coming from each pathway, I find that without regular ad referrals the drop in visits would be 15 visits per domain-month and with these referrals the drop is smaller at 12.5 visits. Therefore, regular ads are able to bring in at least 2.5 new visitors, an increase of 17%. Taken together, these findings imply the regulatory action had a treatment effect on reducing the number of visitors arriving at these merchant sites. However, some new users continue to find these product domains indicating they are already in the market for such products and the degree of harm of the fake news campaigns, for this subset of users, is less severe.

I supplement this analysis with data from consumer-complaint boards, which are in-

dicative of actual purchases and amount spent on these product websites. I find that the probability of a merchant receiving a complaint declines by 8% following the FTC shutdown, further confirming the impact of the regulatory action. Using two other proxies for purchase from the browsing data: duration spent at the merchant site and visits to potential order management sites, I find a significant decline in such metrics. Taken together with visits and the complaints data, these provide evidence on the impact of the shutdown on product purchases.

I also analyze supply-side responses in terms of changes in ad spend and the introduction of (potentially) new fake news campaigns after the FTC-enabled shutdown. I do not find evidence supporting a change in advertising nor do I find evidence that merchants introduced new fake news campaigns. Moreover, the demand-side results are robust to controlling for advertising intensity. I also verify that negative publicity (arising from FTC's press release or news PR effects) alone does not explain the decline in visits. I do so by exploring heterogeneity among consumers as defined by their level of news consumption, finding the biggest impact is on those who consume the least news. If news effects were driving the decline, we would expect the group consuming more news to see the largest decline. This finding is consistent with the hypothesis that those with low news consumption are likely to belong to the susceptible population. Finally, the treatment effect could include changes search engines such as Google might have made to their algorithms in response to the shutdown. Therefore, this measure is likely to be an upper bound on the true impact of the removal of the fake news advertising campaigns.

The specific setting of the paper involves fraudulent products that were using fake news advertising as one means of marketing. However, that the products were fraudulent was discovered (by consumers and regulators<sup>2</sup>) only later on. Therefore, these findings would likely apply to legitimate products using such a fake news advertising strategy. Future work in other settings would help establish the generalizability of these findings.

The rest of the paper is organized as follows. The next section briefly describes the related literature in this domain. In the third section I present a framework illustrating the various pathways a consumer can use to reach the merchant. I then describe the data and the institutional setting. In the fifth section, I present empirical analysis on browsing data. In the sixth section, using data from complaint boards, I analyze purchase behavior after the shutdown. In the seventh section, I analyze possible firm-side responses. The final section concludes.

---

<sup>2</sup>Subsequent investigations by the FTC found that some of the merchants had engaged in deceptive "trial" offers. These merchants stated that consumers need only pay a nominal shipping fee to obtain a trial product, but in fact charged consumers for the product. Two of the merchants – LeanSpa and NutriSlim – were shutdown in December 2011, eight months after the fake news advertising channel was shutdown.

## 2 Literature Review

Advertisements designed to look like editorial content have been prevalent since newspapers began national circulation circa 1879 (Petty 2013). However, this form of deceptive advertising has received little empirical attention until recently. Chiou and Tucker (2018) study the impact of Facebook's ban on fake news ads on sharing of fake news articles. This paper contributes to this limited literature by studying the effect of a regulatory ban on fake news ads on consumers' purchase propensity. Moreover, by observing the path a user takes to reach a product domain, I can analyze whether consumers nevertheless find these product domains without fake news ads either directly or through regular online ads.

This paper is also closely related to the literature on deceptive practices that firms undertake to win consumers' walletshare. This literature has received attention from both the behavioral and, more recently, the empirical literature. Early work in marketing (Shimp and Preston 1981, Olson and Dover 1978) shows that evaluative (non-factual) claims in advertising are more likely to deceive consumers by influencing consumers' beliefs. Recent empirical work on deceptive practices include review fraud (Luca and Zervas 2016, Mayzlin et al 2014, Anderson and Simester 2014), false claims (Rao and Wang 2017, Zinman and Zitzewitz 2013, Chiou and Tucker 2016) and use of deceptive ad formats (Aribarg and Schwartz 2020, Sahni and Nair 2016, Edelman and Gilchrist 2012). This paper contributes to this growing literature by studying ads designed to appear like news, that are deceptive in both format and content.

The impact of fake news has received a lot of attention in the political literature (e.g., Guess, Nyhan, and Reifler 2018; Guess, Nagler, and Tucker 2019). Fake news has been studied by researchers aiming to understand the factors that influence the perceived accuracy of such news (Pennycook et al 2018, Pennycook et al 2020), whether such beliefs can be corrected (Porter et al 2018) and the reasons for the spread of such news (Vosoughi et al 2018). Fake news sites have recently garnered a lot of attention because of their potential impact on the 2016 elections. Sites such as Facebook and Google are under fire for their role as distributors of fake news. However, the empirical impact of fake news on actual outcomes is unknown. Alcott and Gentzkow (2017), in the context of the 2016 elections, point out fake news may have merely strengthened voters' predetermined beliefs, but did not change their voting behavior. For example, people likely to vote for Trump are the ones who believe pro-Trump fake news. Measuring the true impact of the fake news is empirically challenging because we do not observe voting behavior prior to the exposure to fake news. This paper extends this literature by studying the impact of fake news in the domain of consumer goods, with the added benefit this particular setting provides, which allows me to observe outcomes

before and after the fake news campaigns.

Work in the political setting has also shown that being persuaded in general is hard (Berelson, Lazarsfeld and McPhee, 1954; Kalla and Broockman 2018) and that being persuaded by misinformation might be harder (e.g. Little 2018) either because receivers discount such information or because such misinformation forms a small part of all other messages a receiver sees. This paper aims to understand the extent of such persuasion by misinformation in the context of marketing. Specifically, do fake news ads influence visits to product sites, or are such fake ads viewed by consumers already in the market for such products.

That consumers might choose to consume content geared towards their preferences has been shown in several settings: Gentzkow and Shapiro (2010) and Simonov (2018) in the context of news consumption, and Blake et al. (2015) and Chiou and Tucker (2016) in the context of paid search ads. This paper aims to understand such selection in the context of deceptive fake news ads in a market where consumers are likely to be more susceptible, where understanding whether such campaigns cause harm and thus need increased regulatory surveillance is important.

This paper studies advertisements geared to appear like news, bringing together the literature that studies news consumption and the literature that studies firms' deceptive practices and their impact on consumers' purchase decisions. Because the timing of such fake news ad campaigns are unlikely to be exogenous to demand and because one cannot create such exogenous variation using field experiments in this setting (federal law prohibits false or misleading advertising), this paper uses an event-study approach using the timing of the FTC shutdown as an exogenous event.

### **3 Framework**

Consumers can reach a product domain in three ways (1) directly, (2) referrals through the fake news sites, or (3) referrals through other regular advertisements. Direct visits occur when the consumer either directly types the product domain's address into the search bar or reaches the product site through an organic search. Referrals through fake news sites occur when the consumer clicks on a domain such as `Channel16HealthBeat.com` and then clicks on a link to the product domain present on that fake news site. Referrals through regular advertisements occur when a consumer clicks on an ad such as `LeanSpaAcai.com` and then reaches the product domain. Examples of regular ads are provided in Web Appendix W1. This paper asks what happens when the fake news path is shutdown: Do consumers no longer find the product domains, or do they merely substitute to finding these sites through regular advertisements? The former, if true, implies a treatment effect of the fake news sites;

that is, consumers find the product domains because of the fake news sites. The latter, if true, implies a selection effect; that is, consumers interested in these weight-loss sites will reach them through other means, and the fake news sites were just one way of doing so.

To understand the two effects of interest to the paper, i.e. the treatment and selection effects, I present a framework in this section. Let  $u_{ad}$  represent the utility<sup>3</sup> a consumer gets from viewing a regular advertisement. Let the utility she gets from viewing a fake news ad be  $u_{fn,ad}$ .

**Treatment effect** Fake news advertising has a direct treatment effect when a consumer reaches the merchant site only when fake news ads are present. In the absence of fake news ads, she never visits the merchant and remains unconvinced even if she views regular ads. Such an effect would occur when

$$u_{fn,ad} > 0 \text{ and } u_{ad} < 0 \tag{1}$$

**Selection effect** In this scenario, consumers visit the merchant sites in both the presence and absence of fake news ads. If fake news ads are more persuasive than regular ads, they might visit the merchant after viewing a fake news ad. Nevertheless, in the absence of fake news ads, users still reach the merchants because the regular ad is still persuasive (albeit not as persuasive as fake news ads), i.e., when

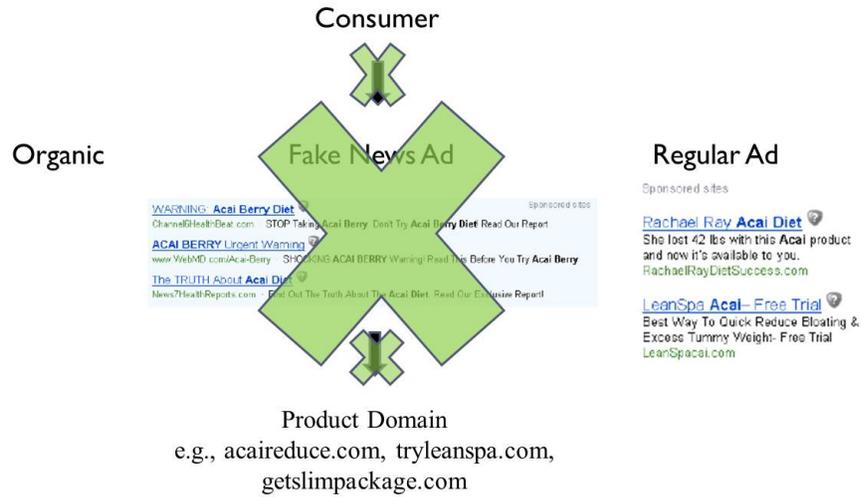
$$u_{fn,ad} > u_{ad} > 0 \tag{2}$$

Both types of consumers can exist in the market, and this paper aims to understand whether this specific setting has more consumers for whom regular ads are a reasonable substitute (equation 2) or whether regular ads do not suffice (equation 1). Understanding this split has implications for how harmful fake news ads are in this domain: for example, if all consumers just substitute to regular ads and continue reaching the site, fake news might have been more persuasive but did not alter consumer behavior drastically. On the other hand, if consumers completely stop visiting the merchant site, it implies that fake news ads were drastically altering consumer behavior.

---

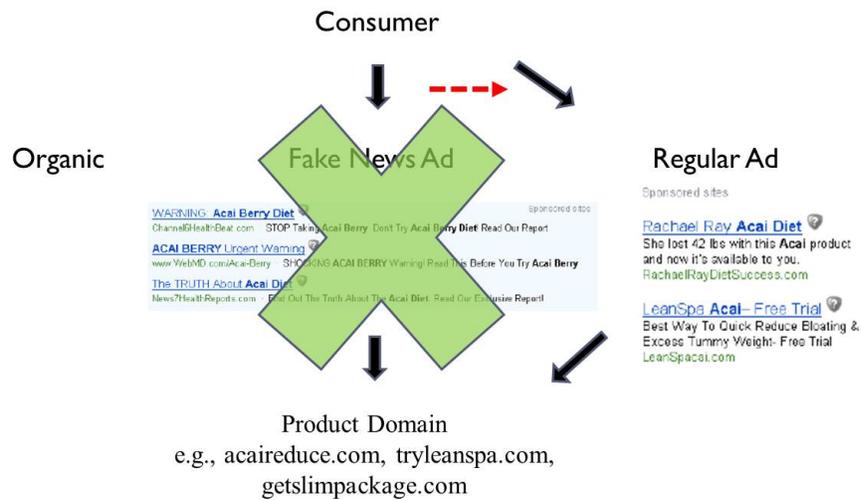
<sup>3</sup>This utility could be driven by 1) information (e.g. availability, price) revealed in the advertisement, 2) an enhanced consumption experience of the product because of the ad, or 3) persuasion. The advertising literature has broadly categorized advertisements as having an informative, persuasive or complimentary role. All of these create shifts in utility but in different manners (persuasion changes consumers preferences, information alters uncertainty in attributes, complementarity creates an interaction with consumption. For the purposes of this framework, a net shift in utility created by regular versus fake advertising (be it due to persuasion or information) suffices to capture the main research question.

Figures 1 and 2 illustrate these two possible outcomes after the shutdown of the fake news advertisements.



Note: Before the shutdown, consumers could reach the product domains using one of three pathways: (1) Directly/Organic searches (2) Fake News ads (3) Regular ads

Figure 1: Treatment: After shutdown of fake news ads, consumers no longer find the product sites



Note: Before the shutdown, consumers could reach the product domains using one of three pathways: (1) Directly/Organic searches (2) Fake News ads (3) Regular ads

Figure 2: Selection: After shutdown of fake news ads, consumers continue to reach the product sites

## 4 Data and Descriptives

In April 2011, the FTC identified 10 companies as operating fake news websites.<sup>4</sup> Table 1 lists a subset of the fake news websites operated by the 10 companies as recovered from the FTC complaint files accessed using the Bloomberg Law database. Figure 3 shows an excerpt of such a fake news advertisement posing as a news article, extracted from one of the FTC court dockets (FTC v Coulomb Media Inc, Declaration of Loretta Kraus). In the figure, I highlight four features common to these campaigns: the presentation of facts by a journalist and the presence of legitimate news-channel names such as CNN and abc. However, neither the channel name, “Health News”, in this example, nor the website, `consumerhealthwarning.com`, are legitimate news sites. In short, in these fake articles, the “journalists” voice their skepticism of the products, claim to try it for themselves, and report their “fake” findings of weight loss. Web Appendix W2 shows two other examples of such articles.

I combine this information on the identity of the fake news site with consumer browsing data in 2010 and 2011 from comScore. comScore tracks detailed browsing and buying behavior for 50,000 internet users across the United States. The panel is based on a random sample from a cross-section of more than 2 million global Internet users who have given comScore explicit permission to confidentially capture their web-wide activity (source: comScore).

Table 1 also lists the number of yearly site visits to these fake news domains, using the comScore data. Following the FTC order in April 2011, these sites were inoperational. To show that after the FTC order, the fake news websites indeed discontinued operation, Figure 4 records the site visits across all fake news websites at the monthly level. This figure provides evidence that the fake news sites were inoperational after April 2011.

---

<sup>4</sup><https://www.ftc.gov/news-events/press-releases/2011/04/ftc-seeks-halt-10-operators-fake-news-sites-making-deceptive>

Not legit news channel



Legit news-site name

### Acai Berry Diet Recipe Exposed: Miracle Diet Recipe or Scam?

As part of a new series: "Diet Trends: A look at America's Top Diets" we examine consumer tips for dieting during a recession



Journalist

Julia investigates the Acai Berry diet to find out for herself if this super diet works.

Everywhere you look there's an internet blog, an ad, or a doctor calling Acai Berry the new miracle fruit—guaranteeing it will stimulate weight loss. But can this thick-skinned little berry live up to all the hype? Or have we embraced the latest Hollywood trend without real merit or scientific evidence that acai really does help you cinch those inches?

To get started, I volunteered to be the guinea pig. I applied for a bottle of the Acai Berry & Advance Colon. While there are tons of Acai Berry ads online, this particular Acai Berry combined with Advance Colon comes from one of the most credible and trustworthy suppliers on the market. It included a free\* bottle of the product with coupon and it did not try to fool me into agreeing to additional hidden offers. Another reason why I chose Acai Berry combined with Advance Colon is because it is the most concentrated and purest acai products on the market. This would give me the most accurate results for my test.



Health and Diet writer, Julia Miller of the News 7 team recently put the Acai Diet to the test. She spent four weeks testing the effects of America's Newest Superfood combined with a Colon Cleanse to see for ourselves what this diet was all about. And, the results were surprising.

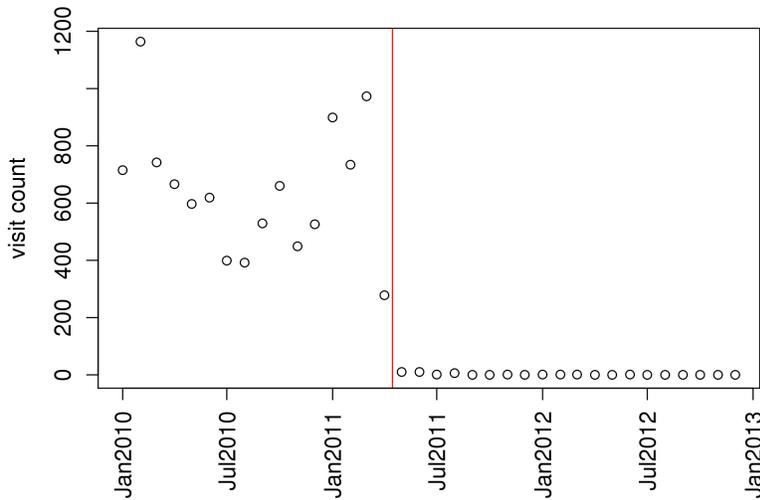
She lost 25lbs in 4 weeks.

[http://consumerhealthwarning.com/diet/\[12/10/2010 11:36:07 AM\]](http://consumerhealthwarning.com/diet/[12/10/2010 11:36:07 AM])

Presentation of facts

Source: FTC Document, Declaration of Loretta Kraus, Coulomb Media, Inc., et al., <https://www.ftc.gov/sites/default/files/documents/cases/2011/04/110419coulomb-kraus-pt1.pdf>

Figure 3: Example of a fake news advertisement from consumerhealthwarning.com



Note: Red line indicates the date of the FTC shutdown order

Figure 4: Fake news domains' visit counts

Table 1: A subset of Fake news domains and their visit count

Company	Fake News Domain names	2010	2011	2012
Ambervine	usahealthnews.org	8	0	0
	usahealthnewsreports.com	8	0	0
CircaDirect	channel8health.com	313	0	1
	clkrd.com	428	0	0
	online6reports.com	961	1	0
	onlinenews6.com	826	4	0
	online6health.com	191	437	0
	online8report.com	252	148	0
Beony	channel6reports.com	423	15	0
	healthnews10.com	43	0	0
	consumertipsdaily6.com	90	37	0
	rachelrayblogs.com	50	0	0
	consumersdigestweekly.com	580	0	0
	consumerstipweekly.com	85	0	0
	internetjobnews.net	60	0	0
	kljsdf.com	47	11	0
Coulomb Media	news6reports.com	982	0	0
	fdx8news.com	90	0	0
	new6reports.com	86	127	0
DLXM	health8news.com	51	43	0
	dlxm.net	84	36	0
	channel9newstoday.com	41	1	0
	channel5healthnews.com	70	66	1
Charles Dunlevy	acai-berry-trial-offers.com	206	0	0
	colon-cleanse-trial-offers.com	44	0	0
Garrett Vaughn	channel9newstoday.com	41	1	0
	health8news.com	51	43	0
	channel6reports.com	423	15	0
	news6report.com	259	1	0
IMM Interactive	channel9healthbeat.com	131	3	0
	nbsnewsat6.com	51	0	0
	consumerproductsdaily.com	137	0	0
	fdxnewsat8.com	217	9	0
	consumeracaicleanse.com	47	47	0
Ricardo Jose Labra	consumerstricksweekly.com	169	0	0
	consumerworkjournal.com	49	1	0

Consumers can reach the fake news site either organically or be referred there by another website/advertisement. In the data, more than 50% of traffic came in organically, whereas the rest were referred by advertising platforms such as `adshuffle`, `facebook`, and `rubicon-`

project. Because some of the organic searches might have come from users who discovered the site earlier through a referral, I keep only the first visits per user domain and find the percentages to be fairly identical; that is, over 50% of visits are organic. Figure 5 shows examples of each of these paths. Users do not spend much time at these fake news websites: the median number of visits per user is 1 website, and the median duration spent is 1 minute. The 90th percentile fake news site received 256 visits in 2010 and served as a referral to a product merchant 64 times while the median fake news site received 25 visits in 2010 and served as a referral 6 times. These numbers suggest that a few fake news sites appear to contribute the most.



Organic links

Sponsored Ads

Investigator’s search for commonalities, in this case “this stuff truley”, resulted in many fake news sites appearing organically. Source: FTC v Circa Direct, Exhibit 5

Investigator’s search for “acai berry” resulted in sponsored ads for fake news sites. Source: FTC v Ambervine, McKenney Attachment B

Figure 5: Organic links and Sponsored Ads linking to fake news sites

#### 4.1 Product Domains/Merchants

Of direct relevance are the domains to which these fake news sites referred users. The com-Score data not only records the domain visited, but also the referrer domain that led the user to a given domain. We can therefore observe the set of product domains that were using the fake news sites as referrers. Extracting all domains the fake news sites referred users to gives us a set of 191 product domains, including sites such as [fibradetox.com](http://fibradetox.com), [getslimpackage.com](http://getslimpackage.com), and [tryacaiberrypure.com](http://tryacaiberrypure.com). Some of these referred-to sites are normal domains,

such as [accuweather.com](http://accuweather.com), [live.com](http://live.com), and [newsvine.com](http://newsvine.com) and a few legitimate ad networks, such as [crwdcntrl.net](http://crwdcntrl.net). I conduct the analysis excluding these normal domains, resulting in 174 product domains that likely used affiliate marketers to advertise their products via fake news. I use the term “product domains” or “merchants” to refer to the websites to which these fake news sites referred consumers.

I assume these product domains were operational after the FTC order in April 2011. This assumption is valid because the FTC order only required the affiliate marketers to shut down their news website operations. The merchants were not ordered to shut down (two of the merchants, LeanSpa and NutriSlim, were subsequently investigated and were required to shut down much later, starting in December 2011<sup>5</sup>). However, to be precise, I also restrict analysis to the set of merchants mentioned in the FTC documents resulting in a smaller set of 48 product domains. To ensure the product domains were operational post FTC order, I further restrict attention to those that received at least one consumer visit, in the data, post-April 2011, resulting in a much smaller subset of 17 merchants. Note this measure is extremely conservative, because comScore keeps track of only a sample of consumers, and all 48 sites are extremely likely to have had visits that might just not be visible in our sample.

The most frequently sold category relates to weight-loss products (including products related to acai berry and colon-cleansing), followed by merchants that engage in earn-money schemes, skin care and teeth whitening products. Finally, there are sites that appear to be ad-trackers which could either be jump links to the main site or a tracking site, as well as what appear to be order-placing portals. The remaining are either categorized into Other or Not classified. Other consists of smaller categories like quit smoking, cure cancer, and health sites. Not classified sites have no information about them on the web or on the Wayback machine internet archive. Table 2 provides these statistics across the three samples studied in the paper. The nature of the products/services was inferred from historical data available through the Wayback machine and/or the name of the domain. Web Appendix W3 provides examples of sites in each category from screenshots taken from Wayback Machine. Web Appendix W4 provides the classification and number of visits for each domain in the sample “Cited by FTC”.<sup>6</sup> Based on the difficulty in finding parent-company related information for the merchant domains, it appears that many of these are short-term businesses aiming to make profits in the short-run by luring susceptible consumers with their claims.

---

<sup>5</sup>These merchants were selling products such as LeanSpa, LeanSpa Cleanse, NutraSlim, QuickDetox, and SlimFuel through sites [LeanSpa.com](http://LeanSpa.com), [TryLeanSpa.com](http://TryLeanSpa.com), [GetLeanSpa.com](http://GetLeanSpa.com), [LeanSpaCleanse.com](http://LeanSpaCleanse.com), [TryNutraSlim.com](http://TryNutraSlim.com), [GetNutraSlim.com](http://GetNutraSlim.com), [TryQuickdetox.com](http://TryQuickdetox.com), [GetQuickdetox.com](http://GetQuickdetox.com), and [SlimFuel.com](http://SlimFuel.com). These merchants not only used fake news campaigns but also made unauthorized recurring charges to consumers’ credit cards of the order of \$79.99 a month.

<sup>6</sup>Because of the large number of websites in the entire sample, I showcase the list for the smaller sample “Cited by FTC” consisting of 48 websites.

Table 2: Classification of merchants by type of service

	(1) <b>All</b>	(2) <b>Cited by FTC</b>	(3) <b>Active</b>	<b>Examples</b>
<b>Total</b>	174	48	17	
Weight loss	31	17	7	getleanspa.com
Earning money online	35	7	4	workfromhomeheaven.com
Skin care	20	7	3	dermitage.com
Teeth whitening	13	3	0	mysuperstarsmile.com
Offers/Savings/Coupons	10	3	0	dropdowndeals.com
Ad malware/services/jump links	22	2	1	dmcpmtrack.com
Order management services	10	2	1	securesiteorders.net
Other	11	1	1	ecigsbrand.com
Not classified	22	6	0	akjhf.com

## 4.2 Pathways to the Merchants

Using data at the individual user level, Table 3 shows the percentage of visits coming from each of the three different paths for all domains (All), as well as restricting attention to those sites cited by the FTC in the court documents (Cited by FTC) and those that are definitely active after the FTC shutdown order (Active). This table shows that fake news sites form a small portion, about 0.8%-4.2% of the way consumers reach these product domains. Because attribution is difficult—subsequent visits could have been informed by previous visits—I keep only the first visit per individual per domain. The last three columns of Table 3 show that fake news sites still represent a small portion of how consumers reach the domains. Regular advertisements followed by organic visits seem to be the most common pathway by which consumers reach these domains. The magnitude of this statistic is in line with the extant literature; for example, Allcott and Gentzkow (2017) estimate the average US adult saw about one fake news story in the months before the election; Guess et al. (2018) estimate that about one in four Americans visited a fake news website from October - November 2016 and that fake news consumption is concentrated in a small subset of people; Guess et al.(2019) show that 91.5% Facebook users do not share even one fake news link contrasting this to regular article shares where a large majority (61.3%) share 100 to 1000 links.

Aggregating the individual-level data, using each individual’s first visit per domain, to the domain-month level, Table 4 provides the statistics for the various pathways across the main dataset (All 174 domains). Statistics for the other two subsets are similar and not reported for brevity.

Table 3: How do consumers reach product domains

	All visits			First visits		
	All	Cited by FTC	Active	All	Cited by FTC	Active
Number of unique websites	174	48	17	174	48	17
Total visits per individual-domain-year-month	2.39	1.41	1.47	1	1	1
Direct	38.4%	34.2%	26.6%	44.9%	39.6%	31.9%
Referred by fake news	0.8%	3.0%	1.9%	2.4%	4.2%	2.7%
Referred by Other	60.8%	63.2%	71.6%	52.7%	56.2%	65.5%
N obs	149,103	29,333	23,228	110,134	26,243	20,471
N individual	50,277	19,035	16,784	50,277	19,035	16,784

Note: (All) refers to all domains the fake news sites referred people to, (Cited by FTC) restricts attention to those sites cited in the FTC court documents, and (Active) refers to those that are definitely active after the FTC shutdown order in the comScore data.

Table 4: Summary Statistics for the Dependent Variables Used in Analyses

Dependent Variable	Mean	SD	Min	Max	N obs	N domains
(1) Browsing Behavior						
Number of visits, Total	30.8	170.4	0	5,843	3,306	174
Number of visits, Direct	14.2	68.3	0	1,748	3,306	174
Number of visits, Fake News Referral	0.8	5.8	0	175	3,306	174
Number of visits, Regular Ad Referral	15.9	129.8	0	5,746	3,306	174
(2) Purchase Data						
Complaint (Yes/No)	0.1	0.3	0	1	966	46
Number of complaints	0.1	0.9	0	17	966	46
Total Dollar	16.1	135.4	0	2,976	966	46
(3) Proxies for Purchase						
Duration Spent, Total	62.0	392.2	0	16,667	3,306	174
Number of visits, tracking	87.6	267.3	0	2,505	380	20
Number of visits, tracking/ordering	67.9	220.2	0	2,505	608	32

Notes: Table presents summary statistics for all dependant variables used in the paper. (1) Browsing behavior presents for All 174 domains, the average monthly visits via each pathway per domain. Only individuals' first visit per domain used. (2) Purchase data obtained from complaint boards presents for all domains where there was a complaint issued, whether there was a complaint in a given month, the total number of complaints and the dollar amount cited in that complaint. (3) Proxies for purchase include duration spent and number of visits at sites with names containing trk and trac, and names containing secure, clk, click, order, trk and trac. Each observation is at the domain-year-month level.

### 4.3 Purchase Data

The comScore data, which also track transactions, show that no one in the sample actually purchased the purported weight-loss products. However, per the FTC documents, consumers

were subject to considerable harm and injury, likely in the order of \$4.68m, which was the judgment issued across all 10 defendants investigated. This judgment implies consumers were deceived into purchasing the products. To quantify the effect of the fake news articles on actual purchases, we need a source of purchase data.

Because the comScore data cannot speak to purchases, I turn to a secondary source of data: consumer-complaint boards. Using two main sources, Ripoff Report and Complaints Board, I collect, for each of the 174 product websites, the date and content of each complaint. An excerpt<sup>7</sup> of a complaint for `theadvancedcleanse.com`, written in April 2010, is “I ordered the Acai Berry Pure and Advanced Cleanse trial products based off the internet report from Newschannel 6 that states the offer is legitimate. The site states you will receive the product in a few days and if not interested you must cancel within 14 days. [...] Needless to say I never received the Acai Berry but I did receive a deduction of \$149.95 from my bank account. [...]”. 46 out of 174 websites had at least one complaint that was filed against them in these two complaint boards between the years 2010 and 2012; the others had no complaints filed against them in this timespan.

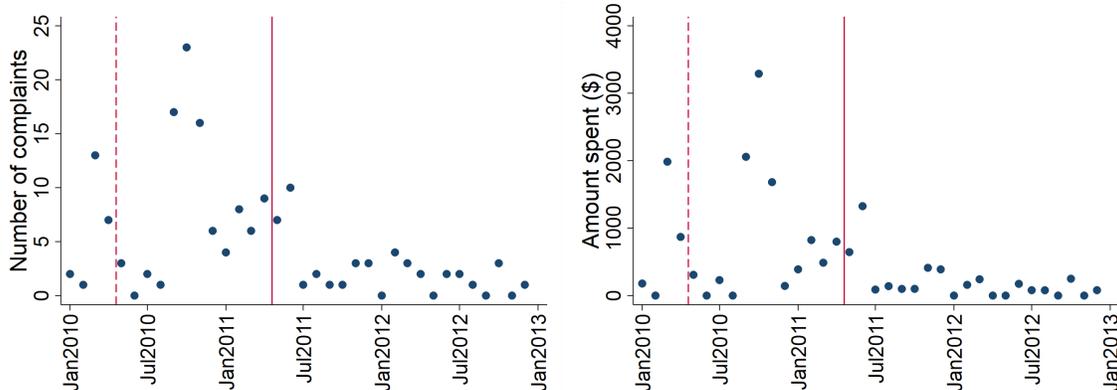
Figure 6 plots the total number of complaints on these sites from 2010-2012, as well as the amount consumers reported having been charged wrongfully. The solid red line indicates the date of the FTC order. A decline in purchases appears to exist, but nearly two months after the FTC order. Such a lag in decline is due to consumers often not immediately reporting a scam; consumers usually take a few weeks to realize they have been scammed, and subsequently post a complaint on a board (e.g., the credit card gets charged one month after a trial subscription). To empirically verify this, I extract the product purchase date from the text of the complaints data and compare it to the date of the complaint. Across all instances where such data is available, the average difference between complaint and purchase date is 2.18 months and the median is 1.50 months suggesting such a lag is indeed a feature of the data.

I also use two additional proxies for purchases. First, I use duration spent at each site as a proxy for purchases. Second, I use visits to what could be order management sites indicating that the user moved to an ordering stage: 1) visits to websites that have `trac` or `trk` in their name and 2) visits to sites that have `click` or `clk`, `trac` or `trk`, `order` or `secure` in their name. These sites include those classified as “order management services” such as `secuesslcenter.com`, `secure81.com`, `securesiteorders.net`, `securesiteoffers.com`, `orderwave.com`, `cpatrac.com`, `trkcpa.com`, `xyztrk.com`, `securedorderweb.com` and `imatrack.com`.

Table 4 provides the statistics for the purchase data from complaints’ boards as well as the statistics for duration spent and visits to order management sites.

---

<sup>7</sup>[https://www.ripoffreport.com/reports/specific\\_search/advancedcolonmax.com](https://www.ripoffreport.com/reports/specific_search/advancedcolonmax.com)



Note: Red line indicates the date of the FTC shutdown order. Dotted red line indicates the placebo date.

Figure 6: Number of complaints and total amount spent on product websites

## 5 Empirical Analysis

I first present results using a simple before-after analysis, relative to a placebo (previous) year. I then investigate, in the next subsection, the pathways through which site visits change.

**Dependent variable:** In all regressions, I use an individual’s first visit to a domain to construct the dependent variable. I do so for two reasons. First, the number of first visits to a domain is a direct measure of the number of new users a site is able to attract. Knowing whether new users are able to reach these domains even in the absence of the fake news ads is a measure of the effectiveness of fake news marketing. If users are unable to reach the merchant sites, it implies the fake news ads had a large treatment effect without which users are unable to find the merchants. However, if users are still able to reach the merchants, it implies a lesser degree of harm. Therefore, knowing whether the first visits drop after the fake news ads shutdown is crucial from a regulator’s perspective. Second, the first visit is a measure of a user’s exposure to the product domain, after which she might have stored it in her browsing history, bookmarked it, or reached it from recall. For example, consider an individual who was referred to a merchant by a fake news website and subsequently visited the site organically. The second visit should not be counted if the fake news site was the reason she discovered this site in the first place. Using only the first visit helps correct for any wrongful attribution.

## 5.1 Difference-in-differences Analysis

Using the FTC order date as an event, I run a before-after regression on all product domains to which the fake news sites referred people. In this regression, I also compare this estimate with a placebo date<sup>8</sup>. To understand the impact at the domain-level, I conduct the analysis on data aggregated to the domain-year-month level. To account for the excess zeros in the data, I conduct the analysis using a zero-inflated Poisson model (Lambert 1992).

I first illustrate the basic differences-in-differences approach using a linear specification, and then use this setup to specify the zero-inflated poisson model. The basic differences-in-differences setup is specified by equation 3 as:

$$y_{jt} = \beta \times Post_t + \beta^{FTC} \times Post_t \times FTC_Y + \alpha_m + \alpha_{m1} \times FTC_Y + \alpha_{m2} \times t + \varepsilon_{jt} \quad (3)$$

Here,  $y_{jt}$  represents the number of visits received by domain  $j$  in month  $t$ .  $Post_t = 1$  if  $t$  is between May and July—the three months after the FTC shutdown month in either the treated or the untreated years;  $FTC_Y$  is an indicator for the actual year when the shutdown occurred, 2011.  $\beta$  captures the average visits attributable to the months May - July relative to the rest of the year, in any year.  $\beta^{FTC}$  is the treatment effect, and captures the additional change attributable to the FTC shutdown. Domain  $j$  is classified into merchant-type  $m$  as specified in the Data section.  $\alpha_m$  is the merchant-type fixed-effect that allows some merchant-types to have more visits than others,  $\alpha_{m1}$  is the merchant-type year fixed-effect specific to the year 2011, and  $\alpha_{m2}$  is the merchant-type specific monthly time trend.

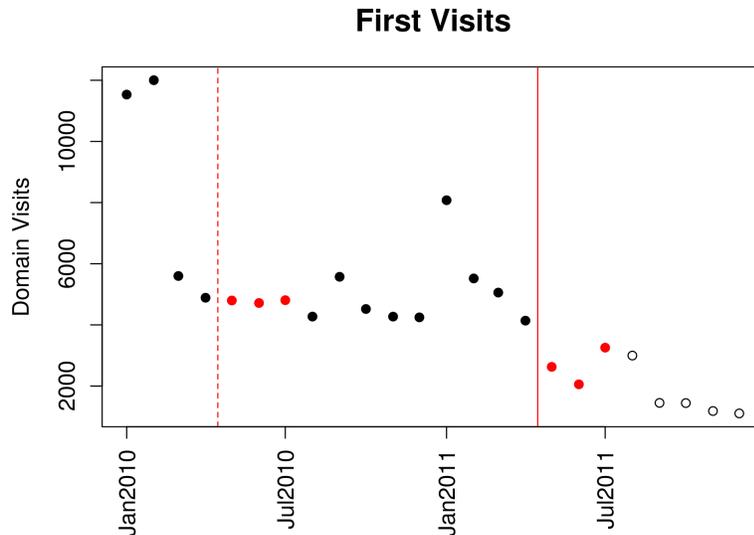
Figure 7 illustrates the idea behind this specification. The three (red) filled dots after the FTC shutdown are compared not only against the (black) filled dots before the shutdown (with appropriate year and month controls), but also against the three (red) filled dots after the placebo shutdown. Although this figure shows a decline in new visits to merchant sites, the estimation equation helps control for merchant-type fixed effects and merchant-type specific time trends. Because the object of interest is the short-term treatment effect, the unfilled dots after the treatment window are not included in the analysis.

Because the data consist of a large number of zero-visits (nearly 60% of observations at the domain-month level), I estimate a zero-inflated Poisson model specified by Equation 4.

$$y_{jt} \sim \begin{cases} 0 & \text{with probability } p_{jt} \\ Poisson(\lambda_{jt}) & \text{with probability } 1 - p_{jt} \end{cases} \quad (4)$$

---

<sup>8</sup>In Appendix A I also present results of a difference-in-difference-in-differences analysis using control sites.



Note: Figure plots total number of first-time visits per month across all impacted merchants. Red line indicates the date of the FTC shutdown order. Dotted Red line indicates the placebo date

Figure 7: First Visits: Measuring the change after the FTC Fake News Ads Shutdown

Here,  $y_{jt}$  represents the number of visits received by domain  $j$  in month  $t$  and  $p_{jt}$  is the probability that domain  $j$  receives a zero-visit in month  $t$ . Conditional on a domain getting a visit, the distribution is assumed to follow a Poisson process with mean  $\lambda_{jt}$ .

The probability of a zero visit,  $p_{jt}$ , is modeled using a logit model, such that  $p_{jt} = \frac{\exp(\gamma X)}{1 + \exp(\gamma X)}$  where  $\gamma X$  follows the same specification in the right-hand side (RHS) of equation 3. The mean,  $\lambda_{jt}$ , of the Poisson distribution (for non-zero visits) is specified such that  $\ln(\lambda_{jt}) = \beta X$  also follows the same specification as the RHS of equation 3.

Data from all months in 2010 and 2011 are included with the exception of the months after the treatment window which are deliberately excluded because they are treated<sup>9</sup>. Standard errors are clustered at the domain level.

I run this regression separately for the three subsets of the data described above: (1) All, excluding the normal domains, comprising 174 unique product domains; (2) Cited by the FTC, those websites that are mentioned in the FTC court documents, comprising 48 unique domains; and (3) Active, those websites that are definitely operational post-April 2011, comprising 17 domains.

An observation in this regression is at the domain-year-month level, created by aggregating up data at the individual browsing level. The comScore dataset records only sessions of active browsing. To allow for no visits, I expand the dataset so that every domain has

<sup>9</sup>Alternatively, one can conduct a long-term analysis which includes all months in 2011 making the treatment window longer (8 months as opposed to the 3 months). Results are robust to analyses conducted using the long-term window.

an observation for every month of the data, accounting for zero visit counts when no visits to domains occurred. This expansion allows search behavior to change; that is, domains might experience a decline in visits after the shutdown but domains might also experience a complete stop to its visits resulting in zero-visits. Not accounting for zero visits would result in an incorrect estimate.

Table 5 presents the results of these regressions. Because of the non-linear nature of the specification, I also present the marginal effects of each component (zero visits and number of events) along with the total marginal effect. All marginal effects are computed at  $t = 5$ , i.e., one month after the shutdown (May). First, consider the probability of a zero-visit. Across all specifications, the interaction term Post X FTC is positive and significant indicating the probability of a zero-visit increases after the FTC shutdown. The marginal effect for the increase in probability of a zero-visit occurring after the FTC shutdown ranges from 22%-41% across the three subsets of the data. Next, conditional on a visit, domains do not experience a change in the total number of visits. The marginal effect is not statistically different from zero. The total marginal effect, which combines both the zero-visit probability and the number of events conditional on a visit, exhibits a decline albeit statistically insignificant. The biggest effect of the shutdown therefore is on the probability a domain receives a visit (and not on the number of visits). I next examine the visits via the various pathways as outlined in the Framework section.

Table 5: Difference in Domains' Visits after Fake News Sites Shutdown

	(1)		(2)		(3)	
	All		Cited by FTC		Active	
N unique treated websites	174		48		17	
	coeff	t-stat	coeff	t-stat	coeff	t-stat
Pr(zero visits): Logit						
Post	0.151	1.56	0.266	1.16	-0.015	-0.06
<b>Post X FTC</b>	<b>1.000</b>	<b>5.87</b>	<b>1.722</b>	<b>3.95</b>	<b>1.934</b>	<b>2.60</b>
Marginal Effect	22%	6.04	33%	4.28	41%	2.7
Number of Events: Poisson						
Post	-0.191	-1.44	-0.153	-0.43	-0.250	-0.57
<b>Post X FTC</b>	<b>0.136</b>	<b>0.54</b>	<b>0.253</b>	<b>0.41</b>	<b>0.360</b>	<b>0.37</b>
Marginal Effect	11.19	0.54	17.99	0.39	37.61	0.35
Total Marginal Effect	-12.46	-1.47	-17.78	-1.16	-31.37	-0.51
N obs	3,306		912		323	
N	174		48		17	
Controls	merchant type				year	
	merchant type x year				month	
	merchant type x month					
Cluster	domain					

Note: Table presents, using first visits to each domain across individuals, results of a before-after zero-inflated poisson regression, relative to a placebo year. The dependent variable is the count of visits per month. The 3-month window after the FTC shutdown is the treatment period. The coefficients corresponding to Post X FTC are the relevant treatment effects. (All) refers to all domains that the fake news sites referred to, excluding normal domains, (Cited by FTC) restricts attention to those sites cited in the FTC court documents, and (Active) refers to those that are definitely active after the FTC shutdown order in the comScore data. Due to insufficient observations, merchant-type specific effects cannot be estimated in the subset (Active). Data aggregated to the domain-year-month level. Standard errors clustered at the domain level.

## 5.2 Pathways to the Merchants

The above results show that new visits to product-domain sites dropped after the FTC's efforts to halt these fake news operations. Although visits to the domains from the fake news sites have to go to zero (by definition, because the fake news sites do not exist anymore), whether direct visits to these sites go up or down post-shutdown is not clear. Similarly, we do not know whether referrals coming in through other ads/sources such as facebook go up or down. Therefore understanding the pathways by which these site visits change is useful. Knowing the total impact is crucial, because if consumers find their way to these product

domains through other means, the impact of the policy (i.e., the shutdown) is unclear.

To this end, I classify visits as “Referred via Fake News”, “Referred via Regular Ads”, and “Direct”. The total visits to a product domain consist of one of these three forms of visits. Table 6 presents the result of the zero-inflated poisson regression with this classification for the main dataset<sup>10</sup> (All 174 domains).

First, consider the probability of a zero visit. Across all pathways, the coefficient for the term Post X FTC is positive and significant indicating the probability of a zero-visit increases after the FTC shutdown. Next, conditional on a visit, the number of visits increases only for “referred by regular ads” after the FTC shutdown: a statistically significant increase of 25.07 visits. This pattern for regular ad referrals suggests an underlying heterogeneity across merchants: some merchants see an increase in zero-visits after the shutdown and other merchants experience an increase in visits. This increase in visits via regular ad referrals for some merchants implies consumers are reaching these merchants through this new path now that the fake news sites are shutdown. In other words, the absence of the fake news advertisements does not entirely prevent consumers from finding these merchants, regular ads serve as a good substitute.

I next quantify the impact attributable to each pathway to understand whether consumers substitute to any particular pathway after the shutdown. Note because of the non-linearity of the zero-inflated poisson model one cannot simply add up the total marginal effect from each additional pathway.

---

<sup>10</sup>The large number of coefficients in the zero-inflated poisson model makes estimation infeasible across all pathways for the smaller subsets “Cited by FTC” and “Active”.

Table 6: Difference in Domains' Visits Across Various Pathways

	Total		Direct		Referred Fake News		Referred Regular Ads	
	coeff	t-stat	coeff	t-stat	coeff	t-stat	coeff	t-stat
<b>(1) All</b>								
Pr(zero visits): Logit								
Post	0.151	1.56	0.156	1.47	0.096	0.41	0.100	0.98
<b>Post X FTC</b>	<b>1.000</b>	<b>5.87</b>	<b>0.837</b>	<b>4.68</b>	<b>3.036</b>	<b>4.75</b>	<b>0.734</b>	<b>4.23</b>
Marginal Effect	22%	6.04	17%	4.77	32%	5.14	15%	4.26
Number of Events: Poisson								
Post	-0.191	-1.44	0.004	0.03	-0.434	-1.16	-0.232	-1.63
<b>Post X FTC</b>	<b>0.136</b>	<b>0.54</b>	<b>-0.339</b>	<b>-1.11</b>	<b>-1.240</b>	<b>-1.71</b>	<b>0.501</b>	<b>2.38</b>
Marginal Effect	11.19	0.54	-17.24	-1.05	-8.10	-1.81	25.07	1.98
Total Marginal Effect	-12.46	-1.47	-12.66	-2.17	-3.03	-4.35	1.05	0.31
N obs	3,306							
N	174							
Controls	merchant type merchant type x year merchant type x month							
Cluster	domain							

Note: Table presents, for only the first visit to each domain across individuals, for each possible path to a domain, results of a before-after zero-inflated poisson regression, relative to a placebo year. Results under column (Total Visits) reproduce results from Table 5. Subsequent columns break these Total Visits into Direct, Referred by Fake News Ad, and Referred by Regular Ads. The dependent variable is the number of site visits per month. The coefficients corresponding to Post X FTC are the relevant treatment effects. (All) refers to all domains that the fake news sites referred to, excluding normal domains. Data aggregated to the domain-year-month level. Standard errors clustered at the domain level.

### 5.2.1 Quantifying the effect

To quantify the effect from each pathway, I evaluate the cumulative effect of each additional individual pathway. Starting with visits coming from fake news ad referrals, which mechanistically have to drop, I evaluate the drop in direct visits *and* fake news referrals, and finally the drop in direct visits, fake news *and* regular ad referrals taken together. To do so, I construct cumulative visits to the merchant starting with fake news referrals and subsequently adding direct visits, and finally regular ad referrals. In other words, I run the same zero-inflated poisson regression on three dependent variables constructed as:

$$\sum_p Visits_p \text{ where } p \in \begin{cases} \text{Fake News} \\ \text{Fake News, Direct} \\ \text{Fake News, Direct, Regular} \end{cases}$$

Table 7 presents the results of this analysis. The first column “Fake News Ads Only” presents the change in visits occurring via only fake news referrals and is identical to the column “Referred Fake News” in Table 6. The second column “Fake News + Direct Visits” presents the change in visits occurring via both these paths. The total marginal effect of the FTC shutdown shows that total visits drop by 15.02 visits when considering fake news referrals and direct visits. However, when regular ad referrals are included (column (3) of Table 7), this drop is smaller at 12.46 visits, implying that the regular ad referrals are able to bring in 2.55 new visitors per month. Note that the drop in direct visits could be cushioned by regular ads as well. For example, if the mere exposure to regular ads make users visit the merchant directly, then the absence of regular ads would imply a larger than 15 visits drop. Therefore, regular ads are able to bring in at least 2.55 new visitors, an increase of 17%.

Table 7: Cumulative effect of each additional pathway

	(1) Fake News Only		(2) Fake News + Direct Visits		(3) Fake News + Direct Visits + Regular Ads	
	coeff	t-stat	coeff	t-stat	coeff	t-stat
<b>(1) All</b>						
Pr(zero visits): Logit						
Post	0.096	0.41	0.170	1.5	0.151	1.56
<b>Post X FTC</b>	<b>3.036</b>	<b>4.75</b>	<b>1.051</b>	<b>5.6</b>	<b>1.000</b>	<b>5.87</b>
Marginal Effect	32%	5.14	22%	5.83	22%	6.04
Number of Events: Poisson						
Post	-0.434	-1.16	-0.017	-0.12	-0.191	-1.44
<b>Post X FTC</b>	<b>-1.240</b>	<b>-1.71</b>	<b>-0.348</b>	<b>-1.08</b>	<b>0.136</b>	<b>0.54</b>
Marginal Effect	-8.10	-1.81	-16.54	-1.03	11.19	0.54
Total Marginal Effect	-3.03	-4.35	-15.02	-2.36	-12.46	-1.47
N obs	3,306					
N	174					
Controls	merchant type merchant type x year merchant type x month					
Cluster	domain					

**Discussion** Alternative factors other than the shutdown could have contributed to the estimated decline. In Appendix B, I provide evidence that news-driven PR effects are unlikely to contribute to the decline. I also show that the declining popularity of acai during this time period does not contribute to the decline following the shutdown. Moreover, any other macro-trends facing the industry are controlled for using the monthly time-trend. I also point out that merchants were not likely to have increased fake news advertising, but could have altered their levels of regular ad spending. I control for such changes to advertising spend, and find the effect continues to hold.

I show that an information effect, where the mere presence of the fake news ad might be sufficient for consumers to visit the merchant site (clicking on the ad but not using it as a referral, not clicking on the Ad but absorbing the information in the Ad) might contribute to the decline. I also show that the shutdown of the fake news advertising can have spillovers to other merchants an individual would have otherwise visited. These two effects are implicitly effects of the shutdown. Finally, changes made by search engines like google could have had an impact on the ranking of these merchants and hence their visits.

Click-throughs are a relevant metric used by the industry as well as in various academic studies, e.g. Aribarg and Schwartz (2020), Chatterjee, Hoffman and Novak (2003) and Blake, Nosko and Tadelis (2015). Visits to the site have also been used as proxies (e.g. Ilfeld and Winer 2002, Sherman and Deighton 2001). However, because actual conversion to purchases might be a small fraction of click-throughs (see for example Manchanda et al 2006), I examine the impact of the shutdown on stricter proxies for purchase such as complaints arising from those who purchased, duration spent at the merchant sites and visits to potential ordering sites in the next section.

## 6 Purchases

Using data collected from complaint boards, I run a logit<sup>11</sup> regression to examine the change in propensity of receiving a complaint after the FTC shutdown. The probability of merchant  $j$  receiving a complaint in month  $t$  is specified by equation 5 as:

$$Pr(\text{complaint}_{jt}) = \frac{\exp(u_{jt})}{1 + \exp(u_{jt})} \quad (5)$$

where  $u_{jt}$  follows the same specification in the RHS of equation 3. The only difference here is that I take into account the 2-month lag in complaints (because complaints occur 2

---

<sup>11</sup>The complaints data consist of a large portion of 0s and 1s (95% 0s, 4% 1s and the remainder 1% having more than 1 complaint at the domain-month level) making a logit model an appropriate choice to model the complaints data.

months after the actual purchase) described in the Data section, i.e. the treatment period under consideration is 2 months after the FTC shutdown, July - September. Table 8 presents the results of the analysis showing that the probability of receiving a complaint dropped significantly after the FTC shutdown. The coefficient for the interaction term Post X FTC is negative and statistically significant and the marginal effect indicates a significant decline of 8%. In Appendix C, I present results from two additional proxies for purchases, duration spent and visits to possible order management sites, which further confirms the decline in purchases.

Taken together with visits (which have been used as a proxy for purchases in the extant literature), these provide further evidence on the impact of the shutdown on product purchases.

Table 8: Change in Probability of a Complaint after Fake News Sites Shutdown

	coeff	t-stat
<hr/>		
Base: No Complaint		
<hr/>		
Complaint (Yes)		
Post	-0.656	-1.32
<b>Post X FTC</b>	-1.605	-2.04
Marginal Effect	-8%	-1.97
<hr/>		
N obs	957	
N	46	
<hr/>		
Controls	merchant type	
	merchant type X year	
	merchant type X month	
Cluster	domain	
<hr/>		

Note: Table presents results of a logit (0: No complaint, 1: Complaint) regression on the period after the FTC shutdown relative to a placebo year (2010). Data aggregated to the domain-year-month level and is expanded to include months of no complaints. The coefficient corresponding to Post X FTC is the treatment effect.

## 7 Supply-Side Responses

### 7.1 Advertisements

The impacted merchants might have increased or decreased their ad expenditure on regular advertisements after the FTC shutdown. To measure a change in ad spend, I use the Nielsen Ad Spend data, and extract any ads relevant to the merchants. To do so, I search for various

combinations of keywords pertaining to the merchant’s website name in the Nielsen Brands database. For example, the merchant site `fibradetox.com` corresponds to the brand Fibra Detox Nutritional Supplement in the Media data. I run a difference-in-differences analysis, specified in Equation 6, on total ad expenditure, frequency, and impression at the brand-month level.

$$Ad_{jt} = \gamma \times Post_t + \gamma^{FTC} \times Post_t \times FTC_Y + \delta_m + \delta_{m1} \times FTC_Y + \delta_{m2} \times t + \varepsilon_{jt} \quad (6)$$

where  $Ad_{jt}$  is the corresponding measure of advertising by merchant  $j$  in month  $t$ .  $Post_t = 1$  if  $t$  is between May and July– the three months after the FTC shutdown month in either the treated or the untreated years.  $FTC_Y$  is an indicator for the actual year when the shutdown occurred, 2011.  $\delta_m$  is the merchant-type specific fixed-effect that allows some domains to have more ads than others.  $\delta_{m1}$  is the merchant-type year fixed effect specific to the year 2011 and  $\delta_{m2}$  is the merchant-type specific monthly time trend.

Table 9 presents the results of this analysis showing no significant change in advertising across all measures used. In a robustness check presented in Web Appendix W5, I control for advertising intensity and find the results are robust.

Table 9: Changes in Ad measures for the treated brands

	Frequency		Spend (\$)		Impressions	
	coeff	t-stat	coeff	t-stat	coeff	t-stat
Post	-5.63	-2.26	-1803.18	-0.4	-280425	-0.33
Post X FTC	4.49	1.04	1614.19	0.2	320188.4	0.22
N obs	589					
N brands	31					
Controls	merchant type merchant type x year merchant type x month					
Cluster	domain					

## 7.2 Launch of Other Fake News Sites

Firms may open and operate new fake news websites with different names, after the FTC shutdown of the 150 websites. To identify if firms engage in such conduct, I use two means of identifying potential fake news websites. First, I extract all websites that contain the terms `health`, `news`, `report`, `today`, `review`, `channel`, `daily`, `weekly`, `evening` or `journal`, resulting in over 163,000 sites. Most of the original fake news websites contained at least one of these terms.

In the first approach, I use the Wayback Machine to capture the start date of the websites<sup>12</sup>. Figure 8 provides an example of how such start dates are captured: for each domain, Wayback Machine tracks not only screenshots but also the start date of the site. In this example, below “27 captures” is the start date 1 Nov 2009 of the site weekly-bulletin.com. However, scraping the large number of sites is infeasible both because Wayback Machine does not allow such large-scale scraping and because it would take a long duration to collect all this information. A middle ground was to scrape a smaller sample of sites from Wayback Machine: I chose a random 10% sample and for this sample obtained the start date of each domain.

In the second approach, I consider all domains that have a number in their domain name, for example, `new6reporter.com`. This approach relies on the fact that many FTC shutdown sites had such names (e.g., `online8report.com` and `news6report.com`).

I then run a difference-in-differences analysis on the resulting count of unique domain names at the monthly level. If firms indeed started new alternative fake news sites, then we should see an increase in sites classified as potential fake news sites. Table 10 shows no such increase, under both approaches: (1) start dates and (2) contains number. As a reference, I present the regression for the fake news sites that were actually shut down, which shows a significant decline. However, if merchants come up with domain names very different from those used pre-shutdown then these two approaches do not capture the launch of new fake news domains. To the extent such fake news domains exist, they are currently captured under regular ad referrals.

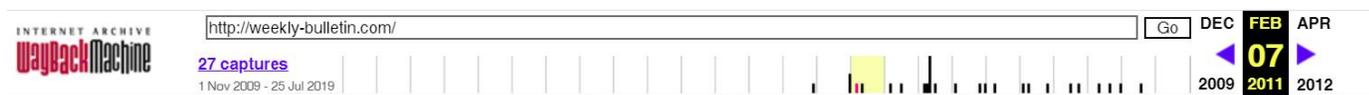


Figure 8: Wayback Machine provides start date of each website

<sup>12</sup>In another exercise, I use current domain status as an indicator of potential fake news sites. This strategy assumes fake news sites that opened in 2011 (following the FTC shutdown) are no longer operational today, while legitimate sites continue operation and should return an active status when pinged. I ping each of the 163,000 sites and record the resulting status code of the domain, for example, No Response, 302 Found, and 403 Forbidden. I count all websites that return a No Response or Timeout after three seconds as potential fake news sites. I also count, using data from only those individuals who visited the impacted merchants, unique domains comprising the terms `health`, `news`, `report` etc. The results are robust to these analyses.

Table 10: No evidence for increase in number of new potential fake news sites

(1) Start Dates	coeff	t-stat
Post	14.67	0.15
Post X Treatment	-197.00	-1.39
<hr/>		
(2) Contains number		
Post	-46.67	-1.14
Post X FTC	-155.00	-2.68
<hr/>		
(3) Fake news shutdown by FTC		
Post	7.33	5.11
Post X FTC	-27.33	-13.48
N Months	12	

Note: Table presents a difference-in-differences of the change in the total number of unique domains per month after the FTC shutdown, using the previous year as a placebo. (1) Start Dates counts the number of domains that started operation in a given month and contain terms such as "health". (2) Contains number counts those domains that contain terms such as "health" and also a numeral in their name. (3) Fake news shutdown by FTC are those sites the FTC actually shut down.

## 8 Conclusion

This paper examines the role of fake news as advertisements, and finds that fake news can cause increased interest in product domains. The FTC shutdown of the fake news campaigns appears to have a treatment effect, emphasizing the role of a regulator in a context where users are likely to be more susceptible. This paper also aims to study the path through which users arrive at the merchants employing such fake news advertisements. The absence of fake news results in a drop in users arriving not only through fake news referrals but also through direct visits, suggesting the two paths have positive spillovers. However, the shutdown of fake news seems to have diverted some traffic to regular advertisements, indicating the two are substitutes to some extent.

Some of the decline in direct visits could be due to consumers' prior exposure to fake news advertisements and spillovers to other merchants an individual would have otherwise visited in the presence of fake news referrals. These two effects are implicitly effects of the shutdown. Other possible factors (which I rule out) are the declining popularity of acai and the effects of negative publicity. Finally, actions taken by google and other search engines might have impacted organic rankings. I also show that merchants did not increase

fake news advertising after the shutdown, and were unlikely to do so in light of the FTC investigation. However, these sites might also have used fewer/more regular ads following the FTC shutdown. Although the ad-spend data suggest otherwise and the results are robust controlling for such advertising changes, to the extent there exist other potential supply-side responses (that are unobservable to the researcher), the effect measured includes such supply-side actions.

This paper uses both purchase propensity as measured by visits to the domains, duration spent at the sites and visits to potential order management sites, and actual purchases as measured by consumer complaints to understand the impact of the absence of the fake news ads. Ideally, both outcome measures would have come from the same data source. However, this paper and setting represent a step toward understanding the impact of fake news on consumers' purchase decisions.

## References

- [1] Allcott, A. and M. Gentzkow (2017), "Social Media and Fake News in the 2016 Election", *Journal of Economic Perspectives*, 31(2), 211-236.
- [2] Anderson, E. T., and D.I. Simester (2014), "Reviews without a Purchase: Low Ratings, Loyal Customers, and Deception", *Journal of Marketing Research*, 51(3), 249–269.
- [3] Aribarg, A., and E.M. Schwartz (2020), "Native Advertising in Online News: Trade-Offs Among Clicks, Brand Recognition, and Website Trustworthiness", *Journal of Marketing Research*, 57(1), 20–34.
- [4] Berelson, B. R., P. F. Lazarsfeld, and W. N. McPhee (1954), "Voting: A Study of Opinion Formation in a Presidential Campaign", University of Chicago Press.
- [5] Blake, T., C. Nosko and S. Tadelis (2015), "Consumer Heterogeneity And Paid Search Effectiveness: A Large-scale Field Experiment", *Econometrica* 83(1), 155-174.
- [6] Chatterjee, P., D.L. Hoffman, and T.P. Novak (2002), "Modeling the Clickstream: Implications for Web-Based Advertising Efforts", *Marketing Science*, 22 (4), 520–41.
- [7] Chiou, L. and C. Tucker (2016), "How Do Restrictions on Advertising Affect Consumer Search?", Available at SSRN: <http://ssrn.com/abstract=1542934>.
- [8] Chiou, L. and C. Tucker (2018), "Fake News and Advertising on Social Media: A Study of the Anti-Vaccination Movement", Available at SSRN: <https://ssrn.com/abstract=3209929>.

- [9] Edelman B.G., and D.S. Gilchrist (2012), “Advertising Disclosures: Measuring Labeling Alternatives in Internet Search Engines. *Information Economics and Policy*”, 24, 75-89.
- [10] Gentzkow, M. and J. M. Shapiro (2010), “What drives media slant? evidence from us daily newspapers”, *Econometrica* 78 (1), 35-71.
- [11] Guess, A., B. Nyhan, and J. Reifler (2018), “Selective Exposure to Misinformation: Evidence from the consumption of fake news during the 2016 U.S. presidential campaign”, working paper, Dartmouth College.
- [12] Guess, A., J. Nagler, and J. Tucker (2019), “Less than you think: Prevalence and predictors of fake news dissemination on Facebook”, *Science Advances*, 5(1):eaau4586
- [13] Ilfeld, J.S. and R.S. Winer (2002), “Generating Website Traffic”, *Journal of Advertising Research*, 42 (5), 49–61.
- [14] Kalla, J. and D. Broockman (2018), “The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments”, *American Political Science Review*, 112(1), 148-166.
- [15] Little, A.T. (2018), “Fake News, Propaganda, and Lies Can Be Pervasive Even If They Aren’t Persuasive”, *Comparative Politics Newsletter*, 28(2), 49–55.
- [16] Luca, M. and G. Zervas (2016), “Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud”, *Management Science*, 62(12), 3412-3427.
- [17] Manchanda, P., J.P. Dube, K.Y. Goh, and P. K. Chintagunta (2006), “The Effect of Banner Advertising on Internet Purchasing”, *Journal of Marketing Research*, 43 (1), 98–108.
- [18] Mayzlin, D., Y. Dover, and J. Chevalier (2014), “Promotional Reviews: An Empirical Investigation of Online Review Manipulation”, *American Economic Review*, 104 (8): 2421-55.
- [19] Olson J.C and P.A. Dover (1978), “Cognitive Effects of Deceptive Advertising”, *Journal of Marketing Research*, 15(1):29-38.
- [20] Pennycook, G., A. Bear, E. T. Collins and D. G. Rand (2020), “The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings”, *Management Science*, in press.

- [21] Pennycook, G., T. D. Cannon and D. G. Rand (2018), "Prior exposure increases perceived accuracy of fake news", *J. Exp. Psychol. Gen.*, 147(12), 1865–1880
- [22] Petty, R.D. (2013), "From puffery to penalties: a historical analysis of US masked marketing public policy concerns", *Journal of Historical Research in Marketing*, 5(1), 10-26.
- [23] Porter, E., T. J. Wood and D. Kirby (2018), "Sex trafficking, Russian infiltration, birth certificates, and pedophilia: A survey experiment correcting fake news", *J. Exp. Polit. Sci.* 5, 159–164.
- [24] Rao, A. and E. Wang (2017), "Demand for "Healthy" Products: False Claims and FTC Regulation", *Journal of Marketing Research*, 54(6), 968-989.
- [25] Sahni, N. and H. Nair (2016), "Native Advertising, Sponsorship Disclosure and Consumer Deception: Evidence from Mobile Search-Ad Experiments", *Marketing Science*, forthcoming.
- [26] Sherman, L. and J. Deighton (2001), "Banner Advertising: Measuring Effectiveness and Optimizing Placement", *Journal of Interactive Marketing*, 15 (2), 60–64.
- [27] Shimp T.A. and I.L. Preston IL (1981), "Deceptive and Nondeceptive Consequences of Evaluative Advertising", *Journal of Marketing*, 45(1):22-32.
- [28] Simonov, A. (2018), "What Drives Demand for Government-Controlled News in Russia?", working paper, Columbia Business School.
- [29] Vosoughi, S., D. Roy and S. Aral (2018), "The spread of true and false news online", *Science* 359, 1146–1151.
- [30] WSJ (2011), "FCC Fines Local TV Stations for Airing Fake News", Accessed via Factiva.
- [31] Zhu, Y. and A. Duke (2015), "Selective Reporting of Factual Content by Commercial Media", *Journal of Marketing Research*, 52, 56-76.
- [32] Zinman, J. and Zitzewitz, Z. (2013), "Wintertime for Deceptive Advertising", working paper, Dartmouth College

## A Difference-in-difference-in-differences Analysis

In this section, I include control sites to control for any industry trends beyond seasonality that might influence the results. The specification follows the same zero-inflated poisson distribution as specified in the main analysis, but the mean of the Poisson process,  $\lambda_{jt}$ , is now specified such that:

$$\begin{aligned} \ln(\lambda_{jt}) = & \beta \times Post_t + \beta^{FTC} \times Post_t \times FTC_Y \\ & + \beta^{Fake} \times Post_t \times Fake_j + \beta^{FTC,Fake} \times Post_t \times FTC_Y \times Fake_j \\ & + \alpha_m + \alpha_{m1} \times FTC_Y + \alpha_{m2} \times t + \varepsilon_{jt} \end{aligned} \quad (7)$$

Here,  $Post_t = 1$  if  $t$  is between May and July—the three months after the FTC shutdown month in either the treated or the untreated years;  $FTC_Y$  is an indicator for the actual year when the shutdown occurred, 2011.  $Fake_j$  is an indicator for those domains impacted by the FTC shutdown.  $\beta$  captures the average additional visits in the months May - July relative to the rest of the year, in any year.  $\beta^{FTC}$  is the treatment effect, and captures the additional change across all domains during the year of the FTC shutdown, 2011.  $\beta^{Fake}$  and  $\beta^{FTC,Fake}$  capture the additional change in these estimates specific to the domains impacted by the FTC shutdown.  $\alpha_m$  is the merchant-type fixed-effect that allows for the fact that some merchant-types have more visits than others,  $\alpha_{m1}$  is the merchant-type fixed effect specific to the year 2011, and  $\alpha_{m2}$  is the merchant-type specific monthly time trend. Control domains are included as a separate merchant-type. The portion that accounts for the excess zeros,  $p_{jt}$ , is modeled using a logit model, such that  $p_{jt} = \frac{\exp(\gamma X)}{1 + \exp(\gamma X)}$  where  $\gamma X$  is the same specification used in the RHS of equation 7.

Ideal control sites would be domains/merchants that used fake-news-style advertising but did not face an FTC shutdown during the period of analysis. Although such merchants exist in the post-shutdown period, they did not exist in the pre-shutdown period (January-March 2011), preventing a difference-in-difference-in-differences style analysis. These merchants were those whose affiliate networks faced shutdown orders in 2012 and 2014. Also, because these merchants were in the initial stages in the treated period (May - June 2011), the results of a simple difference analysis with the treated and control sites might not be representative. Instead, to control for time trends in this industry, I use sites related to regular weight-loss products such as `weightwatchers`, `mytrueweightloss`, and `weightlossforall`. In the comScore dataset, I search for all domains that include the word “weight” in them, and treat the resulting domains as control sites. In Appendix B, because acai products form a major portion of the sold products on the merchants’ sites, I also use sites that contain the term “acai” as controls.

Table 11 present the results of this regression. The results are qualitatively identical to the difference-in-differences analysis presented in the Empirical Analysis section of the paper. To ensure parallel trends in the pre-treatment period between the control sites and the treated sites, I compare data patterns of the control domains with the treated sites. A visual verification shows the pre-treatment trends look fairly similar across control and treated sites. I further add a pre-treatment term in the estimation equations for the control and treated groups and find that the treated groups are not significantly different validating the parallel-trends assumption.

Table 11: Difference in Domain First Visits after Fake News Sites Shutdown: Difference-in-Difference-in-Differences

	All	
N unique treated websites	174	
	coeff	t-stat
Pr(zero visits): Logit		
Post	-0.081	-2.44
Post X FTC	0.031	0.57
Post X Fake	0.240	2.38
<b>Post X FTC X Fake</b>	<b>0.950</b>	<b>5.43</b>
Marginal Effect	10%	5.42
Number of Events: Poisson		
Post	-0.243	-2.76
Post X FTC	-0.023	-0.12
Post X Fake	-0.075	-0.37
<b>Post X FTC X Fake</b>	<b>-0.435</b>	<b>-1.24</b>
Marginal Effect	-18.76	-1.16
Total Marginal Effect	-4.81	-2.18
N obs	72,390	
N	3,810	
Controls	merchant type merchant type x year merchant type x month	
Cluster	domain	

Note: Table presents, for only the first visit to each domain by an individual, results of a before-after zero-inflated poisson regression, relative to a placebo year and relative to control sites related to weight. The dependent variable is the number of site visits. The coefficients corresponding to Post X FTC X Fake are the treatment effects. (All) refers to all domains that the fake news sites referred to, excluding normal domains.

## B Robustness Checks

### News consumption and heterogeneous effects

One alternative explanation for the decline in direct visits is the effect of PR following the FTC shutdown, where bad press might have reduced consumers' propensity to visit these domains. However, I show below that the biggest impact of the shutdown is on those who consume the least news. If news effects were driving the decline, one would expect the group consuming more news to see the largest decline. This finding is consistent with the hypothesis that those with low news consumption are likely to belong to the susceptible population.

To verify if news consumption (FTC's efforts or news PR effects) explains the decline in visits, I explore heterogeneity among consumers as defined by their level of news consumption. I create an indicator variable,  $LowNews_i = 1$ , if individual  $i$ 's news consumption is below or equal to the median consumption across users. I then aggregate the data to the group-domain-month level where all individuals with  $LowNews_i = 1$  are grouped together into group  $Low$  and the remainder are grouped into another group. I use three measures that are indicators of consumers' news consumption: 1) visits to FTC's website, `ftc.gov` 2) visits to news sites that covered FTC's press release in April 2011 and 3) top news sites such as `cnn`. The median number of visits to (1) `ftc.gov` is 0, (2) the 17 news sites covering the release is 0, and (3) the top news sites is 4.

Equation 8 illustrates the basic differences-in-differences approach in a linear specification:

$$\begin{aligned}
 y_{gjt} = & \beta \times Post_t + \beta^{FTC} \times Post_t \times FTC_Y \\
 & + \beta^{Low} \times Post_t \times Low_g + \beta^{Low,FTC} \times Post_t \times FTC_Y \times Low_g \\
 & + \alpha_m + \alpha_{m1} \times FTC_Y + \alpha_{m2} \times t \\
 & + \alpha_m^{Low} \times Low_g + \alpha_{m1}^{Low} \times FTC_Y \times Low_g + \alpha_{m2}^{Low} \times t \times Low_g + \varepsilon_{gjt}
 \end{aligned} \tag{8}$$

Here,  $y_{gjt}$  represents the number of visits received by domain  $j$  in month  $t$  by group  $g$  (i.e., low-news consumers or high-news consumers). The rest of the terms are as specified in the main regression analysis of the paper. The only difference between equation 8 and the main regression specification is that each term is now interacted with the group.

Because of the excess zeros in the data, I then estimate a zero-inflated poisson regression, where

$$y_{gjt} \sim \begin{cases} 0 & \text{with probability } p_{gjt} \\ Poisson(\lambda_{gjt}) & \text{with probability } 1 - p_{gjt} \end{cases} \tag{9}$$

Here,  $y_{gjt}$  represents the number of visits received by domain  $j$  in month  $t$  by group  $g$

(i.e., low-news consumers or high-news consumers ) and  $p_{gjt}$  is the probability that domain  $j$  receives a zero-visit in month  $t$  by group  $g$ . Conditional on a domain getting a visit, the distribution is assumed to follow a Poisson process with mean  $\lambda_{gjt}$ .

The probability of a zero visit,  $p_{gjt}$ , is modeled using a logit model, such that  $p_{gjt} = \frac{\exp(\gamma X)}{1 + \exp(\gamma X)}$  where  $\gamma X$  follows the same specification in the right-hand side (RHS) of equation 8. The mean,  $\lambda_{gjt}$ , of the Poisson distribution (for non-zero visits) is specified such that  $\ln(\lambda_{gjt}) = \beta X$  also follows the same specification as the RHS of equation 8.

The results reported in Table 12 show those who visit fewer news sites – either *ftc.gov* or sites that cover the FTC announcements – drop their visits to the questionable product domains. Considering the probability of a zero visit, the coefficient for the interaction term Post X Low X FTC is positive and significant for sites that cover the FTC press release (column (2)) and marginally significant for visits to *ftc.gov* (column (1)) indicating that the probability of a zero-visit increases for those who consume less news. (There is no difference between users who consume more popular news versus less (column (3))). This result provides evidence that this drop is unlikely to be driven by the press effect (e.g., consumers read news and therefore decide to go less to the domains), because these users consume very little news by construction.

This test helps rule out the press-effect and also helps recover an important layer of heterogeneity: It suggests those who consume less news are likely more influenced by the fake-news and are the group of users for whom the regulation is impactful.

## Popularity of Acai

Another source is the popularity of acai (which form a major portion of the sold products on the merchants' sites) which might be declining in the period corresponding to the FTC release. To test if acai-driven popularity might be contributing to this effect, I use google search volume data on acai as a control, as well as regular “acai” sites as controls.

**Acai trend:** Using Google Trends, I capture the search volume per month for the search term “Acai”. If popularity of “acai” is following a downward trend, this might explain the drop in visits around the FTC press-release. Figure 9 plots this search volume for the years 2010 and 2011 exhibiting a decline over time. I therefore use the acai search volume as an additional control in the analysis. The monthly time-trend captures any other trends that might be occurring in the industry. Table 13 presents the results of this regression. The coefficients for the search volume are statistically insignificant in both the zero-probability and number of events distributions, suggesting the merchant-specific monthly time trends capture most of the decline. Considering the probability of a zero-visit, the main coeffi-

cient of interest corresponding to Post X FTC is still positive and significant indicating the probability of a zero-visit increases after the FTC shutdown.

**Acai sites as controls:** If acai products saw a downward trend during this period, all sites containing the term “Acai” in them should experience such a decline and would serve as a good control because it represents the industry trend. Moreover, a PR effect (if any) will likely effect normal acai sites if consumers cannot easily distinguish between a merchant using the fake news site and a regular merchant, and there is a negative spillover to the normal acai sites. I therefore extract all sites that contain the term acai in them (but are not part of the merchants associated with the fake news ads) including sites like [takeacainow.com](http://takeacainow.com), [acaimaxcleanse.com](http://acaimaxcleanse.com) and [acaiberryforsale.net](http://acaiberryforsale.net). Table 14 presents the results of the difference-in-difference-in-differences analysis using legitimate “acai” sites as controls. The results indicate only the treated sites see a decline in this period, providing an additional robustness check to the main finding.

## **Past Exposure to Fake News Ads**

An information effect, where the mere presence of the Ad is sufficient for consumers to visit the site (clicking on the Ad but not using it as a referral, not clicking on the Ad but absorbing the information in the Ad), might be responsible for the decline in site visits. One measure of consumers’ past exposure are observations where they clicked on the fake news ad but did not use it as a referral to visit the merchant site. Such observations form 0.83% of the data (compared to 2.4%), suggesting consumers read these Fake news articles and then visit the site subsequently (without a direct referral).

While this provides a direct measure of exposure, other means of exposure such as viewing the ad (without clicking and reading the fake news article) are unmeasurable but likely to contribute to consumers’ exposure as well.

## **Spillovers to Other Merchants**

On average, using the pre-treatment year of 2010, an individual visits 2 merchants per year. Restricting attention to the subset of users who have used fake news referrals once: the average number of visits to merchants by such individuals is much higher at 3.25 visits. This statistic implies that users who came in through fake news referrals are visiting more merchants, suggesting there can be spillovers of the shutdown to other domains visited by an individual. That is, an individual does not just stop visits to the merchant she became aware of through the fake news advertisement, but might stop visiting other such merchants.

## **Role of search engines**

Another alternative explanation is that google and other search engines might have stopped ranking these now-questionable merchants after the shutdown, causing a drop in direct visits because of a change in the search engine's algorithm, and not because of the absence of the fake news ads. This possible explanation would also lower the effect of regular ad referrals, implying the increase we observe is a conservative estimate.

## **Supply-side factors**

It is unlikely that firms increased any advertising of the fake news kind because the FTC will penalize such actions in the light of the investigation. However, firms could have increased or decreased other kinds of advertising. The analysis using the Ad expenditure data in the section on Supply-side Responses provides evidence that there was no significant ad spend increase. I also show that the results are robust after controlling for changes in advertising intensity. However, to the extent there exist other potential supply-side responses (that are unobservable to the researcher), the effect measured includes such supply-side actions.

Table 12: Analysis interacted with news consumption: zip model with Aggregate Data

	(1) ftc.gov		(2) news on press release		(3) top news	
	coeff	t-stat	coeff	t-stat	coeff	t-stat
Pr(zero visits): Logit						
Post	0.062	0.52	0.149	1.19	0.108	1.04
Post X FTC	0.697	3.28	0.724	3.56	0.950	5.44
Post X Low	0.090	0.88	-0.013	-0.13	0.047	0.79
<b>Post X Low X FTC</b>	<b>0.314</b>	<b>1.67</b>	<b>0.326</b>	<b>2.18</b>	<b>0.112</b>	<b>1.04</b>
Marginal Effect	6%	1.67	7%	2.17	2%	1.04
Number of Events: Poisson						
Post	-0.078	-0.54	-0.191	-1.19	-0.176	-1.25
Post X FTC	0.165	0.71	0.187	0.77	0.218	0.88
Post X Low	-0.123	-1.66	-0.016	-0.21	-0.029	-0.69
<b>Post X Low X FTC</b>	<b>-0.016</b>	<b>-0.1</b>	<b>0.004</b>	<b>0.03</b>	<b>-0.049</b>	<b>-0.6</b>
Marginal Effect	-1.32	-0.16	-0.20	-0.03	-2.23	-0.57
Total Marginal Effect	-2.90	-1.14	-2.83	-1.2	-1.78	-1.36
N obs	6,612					
N domains	174					
Controls	merchant type merchant type x year merchant type x month					
Cluster	domain					

Note: Low defined as those who visit fewer than the median number of visits of (1) ftc.gov (2) news sites that covered the FTC's press release and (3) top news sites such as cnn.com. Data at the domain-year-month level. Standard errors clustered at the domain level.

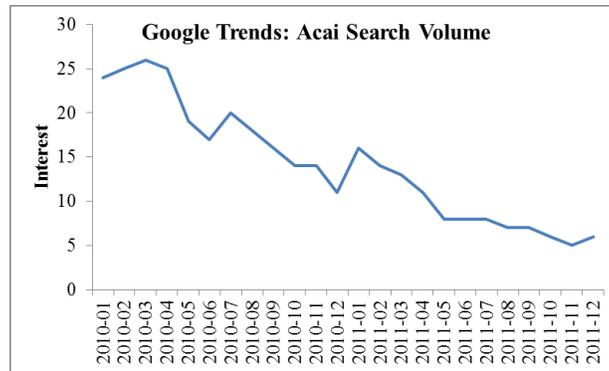


Figure 9: Google Trends: Acai Search Volume

Table 13: Difference-in-differences analysis controlling for search volume of acai

	coeff	t-stat
Pr(zero visits): Logit		
acai search volume	0.011	0.56
Post	0.166	1.5
<b>Post X FTC</b>	<b>0.994</b>	<b>5.78</b>
Marginal Effect	22%	5.96
Number of Events: Poisson		
acai search volume	-0.025	-0.93
Post	-0.229	-1.44
<b>Post X FTC</b>	<b>0.143</b>	<b>0.57</b>
Marginal Effect	12.39	0.56
Total Marginal Effect	-12.65	-1.43
N obs	3,306	
N	174	
Controls	merchant type	
	merchant type x year	
	merchant type x month	
Cluster	domain	

Notes: Search volume corresponds to the monthly search volume for the term "acai" captured from Google Trends data.

Table 14: Difference-in-difference-in-differences analysis using normal acai sites as controls

	All	
N unique treated websites	174	
	coeff	t-stat
Pr(zero visits): Logit		
Post	0.229	2.51
Post X FTC	-0.149	-0.68
Post X Treated	-0.078	-0.59
<b>Post X FTC X Treated</b>	<b>1.149</b>	<b>4.15</b>
Marginal Effect	14%	4.19
Number of Events: Poisson		
Post	-0.104	-0.58
Post X FTC	0.438	1.21
Post X Treated	-0.086	-0.39
<b>Post X FTC X Treated</b>	<b>-0.302</b>	<b>-0.68</b>
Marginal Effect	-16.32	-1.08
Total Marginal Effect	-5.41	-2.07
N obs	18,734	
N	986	
Controls	merchant type merchant type x year merchant type x month	
Cluster	domain	

Note: Table presents, for only the first visit to each domain by an individual, results of a before-after zero-inflated poisson regression, relative to a placebo year and relative to control sites related to acai. The dependent variable is the number of site visits. The coefficients corresponding to Post X FTC X Treated are the treatment effects. (All) refers to all domains that the fake news sites referred to, excluding normal domains.

## C Additional Proxies for Purchases

Here, I use two additional proxies for purchases, duration spent and visits to possible order management sites. Table 15 presents the results of difference-in-differences analyses on duration spent and visits to tracking sites (websites that have `trac` or `trk` in their name) and visits to possible ordering sites (sites that have `click` or `clk`, `trac` or `trk`, `order` or `secure` in their name). The results indicate a significant decline in purchases using both these metrics. Specifically, the probability a domain receives a zero visit after the shutdown increases by 15%-22% across these three proxies and this increase is statistically significant.

Table 15: Analyses using proxies for purchase: Duration spent, visits to possible ordering sites

	duration		tracking		ordering	
	coeff	t-stat	coeff	t-stat	coeff	t-stat
Pr(zero visits): Logit						
Post	0.159	1.62	0.181	1.08	0.150	1.03
<b>Post X FTC</b>	1.029	5.8	1.186	2.84	0.712	2.18
Marginal Effect	22%	5.92	27%	3.02	15%	2.2
Number of Events: Poisson						
Post	-0.200	-1.06	-0.003	-0.02	0.115	0.65
<b>Post X FTC</b>	-0.101	-0.34	0.295	1.4	0.055	0.23
Marginal Effect	-15.94	-0.35	48.10	1.39	6.73	0.23
Total Marginal Effect	-41.20	-2.25	-19.05	-0.89	-14.94	-0.9
N obs	3306		380		608	
N	174		20		32	
Controls	merchant type merchant type x year merchant type x month					
Cluster	domain					

# Web Appendix

## Appendix W1. Regular Ads and their Content

Because the focus of the FTC court exhibits was on the fake news ads, they are not a good source of the regular ads. However, in one search for the search term “acai” the investigator was served with 5 sponsored links one of which (highlighted in red) is provided in Figure 1 and appears to be a regular ad recommending a trial “Buy Acai, Evaporate Weight Loss”. The next one also highlights an ad that appears to be a trial offer for a product. Figure 2 shows some display ads, some of which were featured on the fake news articles themselves, that provide us a hint of the content.

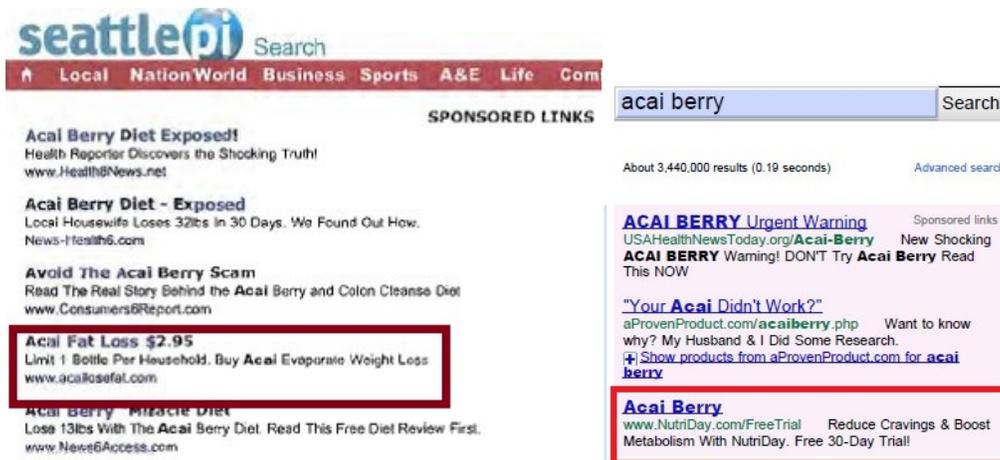


Figure 1: Regular Ad Examples: Sponsored Ad Links



Figure 2: Regular Ad Examples: Display Ads on Fake News Article Sites

# Appendix W2. Fake News Article Examples

Advertorial

SUBSCRIBE News Comments Email



**Health News**

**Jennifer Thouriau**

"Acai Berry Diet trend uncovers steps to weightloss success."





**HEALTH 6 BEAT**

**James Field**

"Acai Berry Diet, is it a scam or not?"



---

## Acai Berry Diet Exposed: Miracle Diet or Scam?

As part of a new series: "Diet Trends: A look at America's Top Diets" we examine consumer tips for dieting during a recession.

Friday, December 10, 2010

AS SEEN ON:     



Stacie investigates the Acai Berry diet to find out for herself if this super diet works.

Acai berries are the latest weight loss fad. These so called Super Foods that you take as a supplement to lose weight have been getting a lot of international attention. And like you have probably already seen; they are all over the internet in blogs and success stories of people who have apparently used the pills and lost a ton of weight. But we here at Channel 6 are a little skeptical and aren't sure that we've seen any real proof that these pills work for weight loss. So we decided to put these products to the test. What better way to find out the truth than to conduct our own study?

To get started, I volunteered to be the guinea pig. I applied for a bottle of the [Get Slim Acai](#). While there are tons of acai berry ads online, [Get Slim Acai](#) is one of the most credible and trustworthy suppliers on the market. It included the free trial of the product and it did not try to fool me into agreeing to additional hidden offers. Another reason why I chose [Get Slim Acai](#) is because it is the most concentrated and purest acai products on the market. This would give me the most accurate results for my test.

Here is what [Get Slim Acai](#) claimed on their website...

- Up To 4 Times More Weight Loss Than Standard Diets
- Boosts Energy without the 'jitters'
- Burn Calories
- Boost Immune System
- Rich in Antioxidants

Were pretty skeptical, but wanted to find out for ourselves if this product could actually do everything that it claimed. Most of the success stories talk about combining Acai berry with colon cleansing products to achieve maximum weight loss. I decided to do the same. The idea behind combining the products is that while the [Acai Berry](#) encourages weight loss and increases energy, the colon cleanse helps rid your body of toxins and allows your body to work and burn calories more efficiently. I chose [Get Slim Colon Cleanse](#) to test.

Here is what [Get Slim Colon Cleanse](#) claimed on their website...

- Helps Eliminate Bad Toxins that have Built Up Over the Years
- Removes 'Sludge' from the Walls of the Colon
- Helps Get Rid of Gas and Bloating
- Helps to Regulate the Metabolism

And the [Get Slim Colon Cleanse](#), like the [Acai Berry](#), had a free trial with a 100% satisfaction guarantee and had no hidden offers.

**Putting Acai to the Test**

Both the [Get Slim Acai](#) and [Get Slim Colon Cleanse](#) arrived within 4 days of having placed my order online for the free trials.



Stacie Sandler, our Health and Diet columnist, recently put the Acai Diet to the test. After four weeks of testing the effects of America's Newest Superfood combined with a Colon Cleanse, she has reached the conclusion to what this diet is all about, and the results were surprising.

**She lost 25 lbs in 4 weeks.**

The Acai Berry Diet benefits beat all of our initial skepticism. This diet was discovered to not only provide weight loss, but also boost energy levels and helped Stacie with a better night's sleep and waking up more rested.

**Step 1:**

First get the free trial of [Get Slim Acai](#)  
 Use our exclusive promo code "**TIPS6**" and pay just \$1.50 for s/h!

**Step 2:**

Then get the free trial of [Get Slim Colon Cleanse](#)  
 Use our exclusive promo code "**TIPS6**" and pay just \$1.50 for s/h!

This is key. Using both products is highly recommended for results like Stacie's.

**Note: Free Trials expiring on Saturday, December 11, 2010**

**» RELATED VIDEOS**

**Acai Berry Nutrition News**

Acai is the new superfood, loaded with antioxidants. My Fox reports on it's benefits.

**The Real Dangers of having a Toxic Colon**

Special CBS new report on the importance of colon health. Why it's important to remove toxins from your colon.

**» ADVERTISEMENTS**



**Acai Max CLEANSE**

**FLUSH POUNDS OF WASTE & TOXINS from your body**

**Acai Max CLEANSE**

Source: FTC Document, Declaration of Douglas McKenney  
 Accessed from: FTC v. Beony International LLC, Attachment C, Bloomberg Law

Figure 3: Example of a Fake News Article from consumertipsdaily6.com (Page 1)

The bottles I received held a month's worth of pills which worked out perfect as I was to follow the supplement routine for 4 weeks time and document my progress throughout

#### My Test 4 Week Acai Berry Diet: LeanSpa + Nature's Colon Rescue

- [Take one LeanSpa pill per day](#)
- [Take one Nature's Colon Rescue pill per day](#)

#### My Results

##### Week One

After one week on the diet using both products I was surprised at the dramatic results. My energy level was up, and I wasn't even hungry; an apparent side effect of the Acai Berry which curbs the appetite.

I honestly felt fantastic.

And I didn't even change anything about my daily routine. On day 7 I got on the scale and couldn't believe my eyes. I had lost 9 lbs. But I still wasn't convinced as they say you lose a lot of water at the beginning of any diet. I wanted to wait and see the results in the upcoming weeks. But it sure was looking up! I now weighed under 140 pounds for the first time in years!

##### Week Two

After two weeks of using both supplements, I started the week off with even more energy and was actually sleeping more soundly than before. I was no longer waking up during the night and tossing and turning because my body was actually able to relax (this is a result of getting rid of the toxins I think). Plus I still managed to lose another 7 pounds, putting me at an unbelievable 16 lbs of weight loss, in just 2 weeks.

I must admit that I'm starting to believe that this diet is more than just a gimmick.

##### Week Three

After 3 weeks all my doubts and skepticism had absolutely vanished! I am down, 2 full dress sizes, after losing another 6 pounds. And I still have a ton of energy. Quite often, around the third week of other diets, you tend to run out of steam. But with the LeanSpa and Nature's Colon Rescue diet my energy levels don't dip, but remain steady throughout the day. I no longer need that cat nap around 3pm in the afternoon! And I am even noticing that my stomach is digesting food so much better. No bloating or embarrassing gas after I eat!

##### Week Four

After the fourth week, my final results were shocking. I lost an unbelievable 25 lbs since starting the LeanSpa and Nature's Colon Rescue diet! Actually everyone at Channel 6 is kicking themselves for not having volunteered to be the guinea pig. Using the LeanSpa and Nature's Colon Rescue in week 4 I lost 3 more pounds. But to be honest I really didn't have much more than that left to lose. And I am definitely going to continue taking the LeanSpa afterwards because it has so many antioxidants and vitamins that makes my skin look unbelievable.

**I couldn't be any happier with the results.**

**I Lost 25lbs in 4 Weeks, No Special Diet, No Intense Exercise**



| Acai Berry |



#### » Healthy Weight Loss Tips

**1. Set a goal.** Identify your ideal weight and set up a plan to start reaching your goal.

**2. Don't be afraid to ask for and get help.** You're not going to lose weight alone. Tell your family. Get support.

**3. Vitamins are good for you.** The American diet lacks essential vitamins and minerals. Balance your health with the best supplements.

**4. Walk the Walk.** You burn calories when you walk, did you know that? Keep active and balance your diet with regular exercise.

**5. Sleep it off.** To be an efficient fat-burning machine, your body requires at least eight hours of sleep a night. If you think that you're doing yourself a favor by sleeping less, you're mistaken.

**Conclusion** Like us, here at Channel 6, you might be a little doubtful about the effects of this diet, but you need to try it for yourself; the results are real. After conducting our own personal study we are pleased to see that people really are finding success with it (myself included :). And you have nothing to lose. Follow the links to the free trials I have provided and know that you are getting a quality product

Source: FTC Document, Declaration of Douglas McKenney

Accessed from: FTC v. Beony International LLC, Attachment C, Bloomberg Law

Figure 4: Example of a Fake News Article from consumertipsdaily6.com (Page 2)

SUBSCRIBE News Comments Email Adverorial Thursday, September 30, 2010

**Channel 8 NEWS** Jennifer Theuriau "Acai Berry Diet trend uncovers steps to weightloss success." News 8 James Field "Acai Berry Diet, Is it a scam or not?"

HOME U.S. WORLD BUSINESS POLITICS ENTERTAINMENT LEISURE HEALTH SCITECH OPINION SPORTS ON AIR

## Acai Berry Diet Exposed: Miracle Diet or Scam?

As part of a new series: "Diet Trends: A look at America's Top Diets" we examine consumer tips for dieting during a recession

Thursday, September 30, 2010

» RELATED VIDEOS

**Acai Berry: Fox35 Special Report**



Julie investigates the Acai Berry diet to find out for herself if this super diet works.

(Chicago) - Acai berries are the latest weight loss fad. These so called Super Foods that you take as a supplement to lose weight have been getting a lot of international attention. And like you have probably already seen, they are all over the internet in blogs and success stories of people who have apparently used the pills and lost a ton of weight.

But we here at Channel 8 are a little skeptical and aren't sure that we've seen any real proof that these pills work for weight loss. So we decided to put these products to the test. What better way to find out the truth than to conduct our own study?

To get started, I volunteered to be the guinea pig. I applied for a free trial of the [LeanSpa Acai](#). While there are ton's of Acai berry ads online, [LeanSpa Acai](#) is one of the most credible and trustworthy suppliers on the market. It included the free trial of the product and it did not try to fool me into agreeing to additional hidden offers. Another reason why I chose [LeanSpa Acai](#) is because it is the most concentrated and purest acai products on the market. This would give me the most accurate results for my test.

### Here is what [LeanSpa Acai](#) claimed on their website...

- Accelerates Fat Loss
- 4 Times More Weight Loss Than Diet And Exercise
- Boosts Energy
- Rich in Antioxidants

We were pretty skeptical, but wanted to find out for ourselves if this product could actually do everything that it claimed. Most of the success stories talk about combining acai berry with colon cleansing products to achieve maximum weight loss. I decided to do the same. The idea behind combining the products is that while the [Acai Berry](#) encourages weight loss and increases energy, the colon cleanse helps rid your body of toxins and allows your body to work and burn calories more efficiently. I chose [Nature Detox](#) to test.

### Here is what [Nature Detox](#) claimed on their website...

- Helps Eliminate Bad Toxins that have Built Up Over the Years
- Removes 'Sludge' from the Walls of the Colon
- Helps Get Rid of Gas and Bloating
- Helps to Regulate the Metabolism

And the [Nature Detox](#), like the [Acai Berry](#), had a sample with a 100%



Julie Ayers, our Health and Diet columnist, recently put the Acai Diet to the test. After four weeks of testing the effects of America's Newest Superfood combined with a Colon Cleanse, she has reached the conclusion to what this diet is all about, and the results were surprising.

### She lost 25lbs in 4 weeks.

The benefits of the Acai berry diet beat all of our initial skepticism. We found the diet not only with weight loss, but it seemed to boost energy levels, and also helped Julie sleep better a night to wake up more rested.

#### Step 1:

First get [LeanSpa Acai](#)  
Use our exclusive promo "[LEAN100](#)" and get price reduced to **\$1.99!**

#### Step 2:

Then get [Nature Detox](#)  
Use our exclusive promo "[Nature](#)" and get price reduced to **\$3.87!**

**\*\*This is key. Use both for results like Julie**

Free Trials expire on Friday, October 01, 2010

#### Network Reviews:

ABC News Calls Acai Berry A Superfood! Many world-class athletes have started using Acai berry products as part of their personal training

Acai Health Benefits: America's #1 Superfood. Look younger and Live Longer and Healthier.

**The Real Dangers of having a Toxic Colon**

Special CBS new report on the importance of colon health. Why it's important to remove toxins from your colon.

» ADVERTISEMENTS

**Claim your RISK FREE Trial Offer**



[http://www.channel8health.com/HEALTH/Acai-Berry/index.php?202id=7628e202kw=\[9/30/2010 10:28:04 AM\]](http://www.channel8health.com/HEALTH/Acai-Berry/index.php?202id=7628e202kw=[9/30/2010 10:28:04 AM])

Source: FTC Document, Declaration of Douglas McKenney

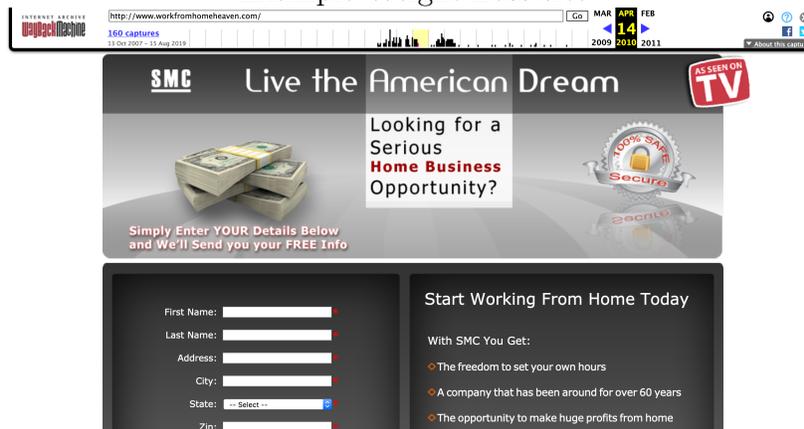
Accessed from: FTC v. Circa Direct LLC, Exhibit 2, Attachment E, Bloomberg Law

Figure 5: Example of a Fake News Article from channel8health.com (Page 1)

# Appnedix W3. Examples of Merchants in Various Categories



Example Weight Loss site



Example Earn Money Online site



Example Skin Care site

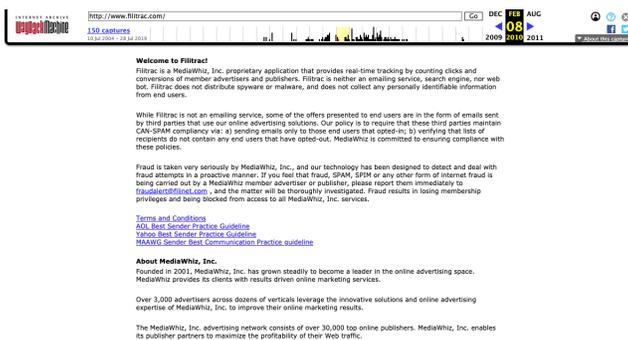
Figure 6: Screenshots of merchants – 1



Example Teeth Whitening site



Example Offers/Savings/Coupons site



Example Ad malware/service/jump links site

Figure 7: Screenshots of merchants – 2



### Example Order Management site



### Example Other site



### Example Not Classified site

Figure 8: Screenshots of merchants – 3

## Appendix W4. Merchants in Sample “Cited by FTC”

Table 1: Domains in “Cited By FTC”, their categorization and number of visits in 2010

<b>Domain</b>	<b>Category</b>	<b>Visits in 2010</b>
hostedcontent.net	Earning money online	18880
cpaclicks.com	Ad malware/services/jump links	5487
theonlinebusiness.com	Earning money online	1814
acaioptimum.com	Weight loss	1930
msicourse.com	Earning money online	1277
swipeauctions.com	Offers/Savings/Coupons	1627
dermitage.com	Skin care product	1135
tryhydroxatone.com	Skin care product	287
tryleanspa.com	Weight loss	258
myincomeconnection.com	Earning money online	396
mydietmaxcleanse.com	Weight loss	581
clkrd.com		428
trackleady.com		311
dermatalskincare.com	Skin care product	120
tryacaiberrypure.com	Weight loss	289
exclusivecashsecrets.com	Earning money online	263
tryadvancedcleanse.com	Skin care product	246
curecolon.com	Cancer cure	131
advancedcolonmax.com	Weight loss	178
getleanspa.com	Weight loss	43
getslimpackage.com	Weight loss	98
southbeachjava.com	Weight loss	16
dlxm.net		84
orderwave.com	Order Management Services	40
naturecleansing.com	Skin care product	97
healthierwaytogo.com	Weight loss	19
getbrightwhites.com	Teeth whitening	79
theinstantwealthbuilder.com	Earning money online	36
leanspacleanse.com	Weight loss	#N/A
trytotalhealth.com	Weight loss	#N/A
kljsdf.com		47
trytadalafil.com	Weight loss	4
maxcoloblast.com	Weight loss	46
xyztrk.com	Order Management Services	45
fibradetox.com	Weight loss	25
offers4all.info	Offers/Savings/Coupons	40
identprtct.com		3
trimsportweightloss.com	Weight loss	#N/A
dependive.com		6
dlxtrack.com	Ad malware/services/jump links	4
acaireduce.com	Weight loss	14
acaiberrydetoxmax.com	Weight loss	13
trypurewhitening.com	Teeth whitening	#N/A
top-web-offers.com	Offers/Savings/Coupons	8
purecleanse360.com	Skin care product	#N/A
getvitacleanse.com	Skin care product	4
applydebtnow.com	Earning money online	4
teethwhitenswab.com	Teeth whitening	4

## Appendix W5. Robustness check controlling for advertising

In this Appendix, I present the analysis of the main regressions on domain visits controlling for advertising. Doing so controls for changes firms might have made to their advertising in response to the FTC shutdown of the fake news campaigns. I also present the results on proxies for users' purchases, the propensity of receiving a complaint after the FTC shutdown controlling for advertising, in Table 3. The main effects continue to hold.

Table 2: Difference in Domain Visits Across Various Pathways Controlling for Advertising

	Total		Direct		Referred Fake News		Referred Regular Ads	
	coeff	t-stat	coeff	t-stat	coeff	t-stat	coeff	t-stat
<b>(1) All</b>								
Pr(zero visits): Logit								
Post	0.132	1.34	0.136	1.25	0.072	0.31	0.078	0.76
<b>Post X FTC</b>	1.025	5.88	0.856	4.68	3.064	4.74	0.760	4.25
Marginal Effect	22%	6.06	17%	4.78	32%	5.08	15%	4.3
Number of Events: Poisson								
Post	-0.178	-1.32	0.020	0.14	-0.457	-1.2	-0.230	-1.59
<b>Post X FTC</b>	0.130	0.52	-0.344	-1.14	-1.234	-1.66	0.498	2.39
Marginal Effect	10.69	0.52	-17.63	-1.08	-7.95	-1.77	24.97	1.97
Total Marginal Effect	-12.63	-1.49	-12.73	-2.16	-3.07	-4.31	0.91	0.27
N obs	3,306							
N	174							
Controls	merchant type merchant type x year merchant type x month advertising frequency							
Cluster	domain							

Note: Table presents, for only the first visit to each domain across individuals, for each possible path to a domain, results of a before-after zero-inflated poisson regression, relative to a placebo year. Results under column (Total Visits) are for visits irrespective of how user reached the domain. Subsequent columns break these Total Visits into Direct, Referred by Fake News Ad, and Referred by Regular Ads. The dependent variable is the number of site visits per month. The coefficients corresponding to Post X FTC are the relevant treatment effects. (All) refers to all domains that the fake news sites referred to, excluding normal domains. Data aggregated to the domain-year-month level. Standard errors clustered at the domain level.

Table 3: Change in Probability of a Complaint after Fake News Sites Shutdown controlling for Ads

	coeff	t-stat
Base: No Complaint		
Complaint (Yes)		
Post	-0.744	-1.4
<b>Post X FTC</b>	-1.501	-1.92
Marginal Effect	-8%	-1.86
N obs	957	
N	46	
Controls	merchant type merchant type X year merchant type X month advertising frequency	
Cluster	domain	

Note: Table presents results of a logit (0: No complaint, 1: Complaint) regression on the period after the FTC shutdown relative to a placebo year (2010). Data aggregated to the domain-year-month level and is expanded to include months of no complaints. The coefficient corresponding to Post X FTC is the treatment effect.