Forthcoming at
Research & Politics

Camouflaged Propaganda: A Survey Experiment on Political Native Advertising

Yaoyao Dai[1] and Luwei Luqiu[2]

Abstract

We examine a new form of propaganda, *political native advertising*, in which political actors, including foreign governments, buy space in independent media outlets to publish advertisements that are camouflaged as standard news stories. Those who engage in this form of propaganda hope to exploit the higher credibility of the hosting media site to enhance the persuasiveness of their message. Despite the obvious political implications and ethical issues at stake, political native advertising has received almost no scholarly attention. Our article begins to redress this imbalance. Using an online survey experiment with real political native advertisements in the *Washington Post* and *The Telegraph* bought by the Chinese government, we provide some of the first empirical evidence on basic but important features of political native advertising. We find, among other things, that respondents struggle to distinguish political advertisements from standard news stories regardless of their level of education and media literacy, that political advertisements are more convincing if they appear on and are perceived as news from an independent hosting media site than in a government-controlled news outlet, and that trust in the hosting media site declines if the political advertisement is detected.

Keywords

Political Native Advertising, Native Advertising, China Watch, Propaganda, Disinformation Campaign

Introduction

Since 2010, readers in the United States have been able to obtain news on China from a multi-page special section named *China Watch* in the *Washington Post*, the *Wall Street Journal*, and the *New York Times* (Cook 2017; Fallows 2011). Unfortunately, instead of being a special editorial column on China, the *China Watch* section is a paid supplement provided by *China Daily*, a Chinese government-controlled English-language newspaper (Fallows 2011). As of March 2018, *China Daily* had cooperated with, and provided *China Watch* content to, more than 40 legacy news media in over 20 countries with a circulation of 4 million people.[1] This is all part of *China Daily*'s strategy to use the platforms and reputations of partnership publishers to increase the worldwide audience for its news stories (China Daily 2018). China is not the only country that pays western legacy media outlets to publish news stories from government-controlled media. For example, *Russia Beyond*, a Russian government-controlled media outlet, has also paid to place news stories in the *Washington Post* under the title *Russia Now*.[2] Political parties in democracies have also engaged in similarly deceptive advertising activities. In the United States, for example, both Democratic and Republican candidates have paid to insert campaign advertisements that mimic news stories and other forms of standard editorial content in domestic media outlets (Iversen and Knudsen 2017; Dykhne 2017; Murtha and Gourarie 2016).

Communication and journalism scholars refer to paid content and advertisements camouflaged as standard editorial content as *native advertising* (Howe and Teufel 2014). Unlike conventional sponsored content or advertisements, *native advertisements* are camouflaged as standard editorial content coming from the hosting media outlet. As a result, people are often unaware that they are reading sponsored and paid content. To date, existing studies have focused on native advertising almost exclusively in the context of commercial products (Carlson 2015; Iversen and Knudsen 2017; Jamieson et al. 2000; Batsell 2017; Edmonds 2017; Einstein 2016; Mullin 2016). We use the term *political native advertising* to refer to situations where political actors engage in native advertising.

Despite the obvious political implications and ethical concerns, such as co-opted (foreign) political influence and the threat to media freedom, political native advertising has received almost no scholarly attention. While there is a large and growing literature dealing with political propaganda, it tends to focus on issues of media control and the types of hard and heavy-handed propaganda that are easily detected, such as pro-regime reporting in government-controlled news programs and government-sponsored nationalist commercials (Di Tella and Franceschelli 2011; Stockmann 2013; Gehlbach and Sonin 2014; Archer and Clinton 2018; Huang 2015, 2018; Little 2017). A common claim in the literature is that propaganda, perhaps because it is often so easily detected, is not designed to achieve persuasion but instead

[1]Social Science Division, New York University Abu Dhabi.
[2]School of Communication, Hong Kong Baptist University.

Corresponding author:
Yaoyao Dai, Social Science Division, New York University Abu Dhabi, Abu Dhabi, U.A.E..
Email: yaoyao.dai45@gmail.com

to identify loyalists or signal the coercive capacity of the state (Crabtree et al. 2018; Little 2018; Marquez 2018; Huang 2018, 2015). Relatively few studies have examined the type of soft propaganda that is harder to detect and is the hallmark of political native advertising. To our knowledge, there has been only one systematic study of political native advertising, which finds that clearly labeled hypothetical political native advertisements by political parties in Norway reduce citizen trust in political news in general (Iversen and Knudsen 2017).

In this paper, we begin to redress this imbalance in the literature by providing one of the first investigations of political native advertising, particularly as it relates to its use by foreign governments. We use an online survey experiment with real political advertisements in the *Washington Post* and *The Telegraph* that were paid for by the Chinese government to obtain some basic and important, but as yet unknown, information about political native advertising. To what extent can readers detect political native advertising? Is political native advertising considered persuasive? Do hosting media outlets suffer a cost if the political native advertising is detected? We find that people often struggle to distinguish paid political advertisements from standard news stories regardless of their level of education and media literacy. We also find that the content found in political native advertisements is perceived to be more persuasive than the same content on a government-controlled media source when the readers perceive the political advertisements as editorial content from the hosting media outlet. Finally, we also find that people significantly reduce their level of trust in the hosting media outlet if they detect the political native advertising.

An Experiment on Political Native Advertising: *China Watch*

Given the lack of basic empirical information about political native advertising, we use an online survey experiment to examine the extent to which readers are able to detect political native advertisements, whether political native advertisements are considered persuasive, and whether independent hosting media outlets suffer a cost in terms of trust if readers detect a political native advertisement. These three basic elements are important in understanding political native advertising: the persuasiveness of political native advertising speaks to the motivation of the political sponsors; the reputational cost of accepting political native advertising is likely to influence the decision of independent media outlets to host such advertisements; the likelihood of being detected is likely to influence the persuasiveness of the political native advertisement and the actual reputational costs to the hosting media outlets.

Our experiment focuses on *China Watch*, arguably the largest effort at political native advertising in the world. The survey experiment, which uses real *China Watch* articles as treatments, was conducted in March 2018. Respondents were recruited from Amazon's Mechanical Turk (MTurk) and were directed to an external Qualtrics survey where the experiment took place. A common concern with recruiting a convenience sample from a crowd-sourcing website such as MTurk is that the respondents may be unrepresentative of the population of interest, leading to low external validity (McDermott 2011). However, recent studies show that MTurk samples are more representative of the overall population than different types of in-person convenience samples (Paolacci et al. 2011; Berinsky et al. 2012). Importantly for external validity, several studies in political science, law, and psychology have also shown that the magnitude of average treatment effects estimated from MTurk samples is similar to the size of effects estimated from nationally representative samples (Berinsky et al. 2012; Clifford et al. 2015; Firth et al. 2017).

Experimental Design

To closely resemble the real experience of reading a *China Watch* article, respondents are provided with a link to the hosting media outlet's website to read *one* of two real articles produced by *China Daily*. Some respondents read an article that appeared on the *China Watch* page of the *Washington Post*, while some read an article that appeared on the *China Watch* page of *The Telegraph*. In effect, these particular respondents received a political native advertisement. A third set of respondents read one or the other of these two articles but as it appeared on *China Daily*'s own webpage. In effect, this third set of respondents did not receive a political native advertisement. All respondents received the same message before reading the treatment article: "In the next section, we will ask you to read a short news article. After reading the article, you will be asked to answer several related questions. Please click this link to read the article." We then display the actual URL to the treatment article.[3]

Although it would be ideal for the content of the two *China Watch* articles to be identical, this is not possible. *China Daily* deliberately avoids displaying the exact same articles in different hosting media outlets across countries. We therefore chose similar articles that addressed the same topic: China's plan to continue with market reforms. We chose to use articles that focused on economic news because this is the main type of news story found in the *China Watch* sections of both the *Washington Post* and *The Telegraph*. Importantly, respondent perceptions of the articles they received did not differ across the two *China Watch* articles (see Online Appendix A).

To examine whether increased awareness of political native advertising would increase the likelihood that respondents can detect the true source of the *China Watch* articles, we also include an *Education Priming* treatment related to political native advertising. Specifically, half of the respondents who receive *China Watch* articles are also randomly selected to receive information about the practice of political native advertising. This *Education Priming* treatment comes before the respondents receive their *China Watch* article. The exact wording of the *Education Priming* treatment is

> "Foreign countries, such as China and Russia, have started to launch international propaganda campaigns using a technique called native advertising. For example, they buy space on western mainstream media outlets to publish government-sponsored content produced

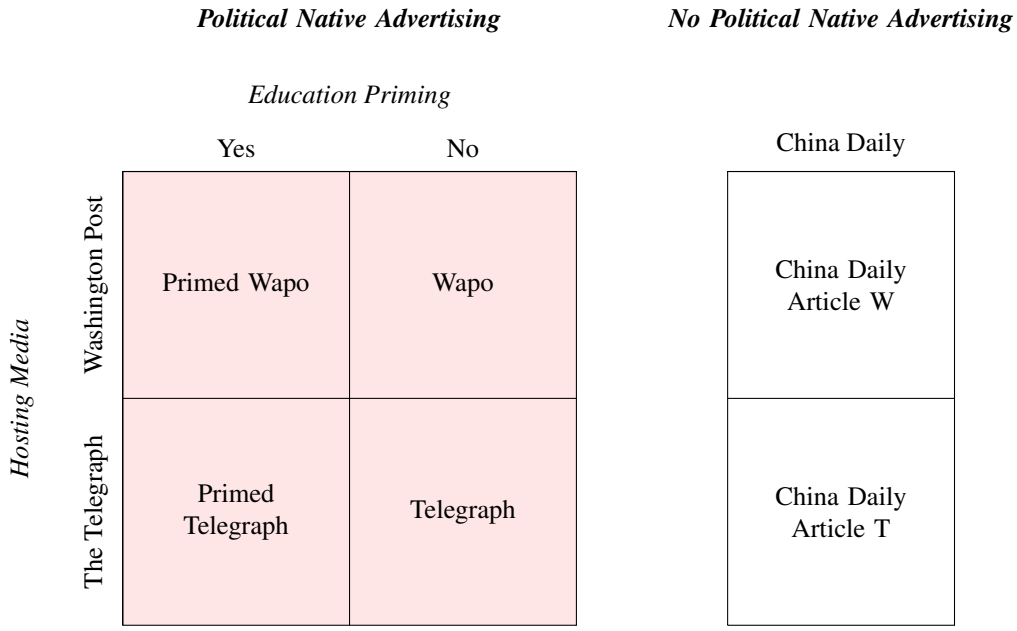

Figure 1. Full Experimental Design

Note: Respondents received one of two similar articles related to China's plan to continue with market reforms. These articles appeared in the *Washington Post* or *The Telegraph* (Political Native Advertising) or in *China Daily* (No Political Native Advertising). Half of the respondents who read a political native advertisement were randomly selected to receive an *Education Priming* treatment in which they were told about the practice of political native advertising.

> by government-owned media, such as *China Daily*."

A graphical overview of the full experimental design is shown in Figure 1. 660 respondents were randomly assigned among each of the six groups. Two groups of respondents received *Education Priming* before reading the *China Watch* article. Two groups of respondents read the same *China Watch* articles on China Daily.

Detecting Political Native Advertising To examine whether respondents are able to detect political native advertising as non-editorial content, respondents who receive one of the *China Watch* articles are asked to identify the source of the article immediately after reading it (left panel of Figure 1). They are given four options: *New York Times*, *China Daily*, *Washington Post*, and *The Telegraph*. Studies on commercial native advertising find readers are more likely to recognize native advertising when the disclosure is more explicit, higher in prominence, and positioned at the bottom (Amazeen and Wojdynski 2018; Amazeen and Muddiman 2018; Wojdynski and Evans 2016). Due to different disclosure rules, there is considerable variation in both the clarity and style of disclosure regarding the *China Watch* pages across the *Washington Post* and *The Telegraph*. Screenshots of the *China Watch* pages on the two hosting media outlets are shown in Figure 2. In *The Telegraph*, the *China Watch* page is located under *The Telegraph*'s World News section and closely resembles its own editorial content. Above the *China Watch* articles, there is a statement in small font that reads, "This content is produced and published by *China Daily*, People's Republic of China, which takes sole responsibility for its content." Significantly, there is no indication that the articles are paid supplements from a foreign government. The disclosure in the *Washington Post* is clearer and more prominent than that found on *The Telegraph*. More importantly, although the web address for the *China Watch* page clearly links it with the *Washington Post*, the material itself looks less like the standard editorial content produced by the *Washington Post*. The *Washington Post* also uses the term "advertisement" at the top of the page and additional information at the bottom of the page states that "This content is paid for and provided by an advertiser, and the site is managed by WP BrandStudio. The Washington Post newsroom and WP BrandStudio were not involved in the creation of this content." Given the variation in the clarity of disclosure, we expect that respondents who receive a *China Watch* article on the *Washington Post* will be more likely to detect the political native advertising than the respondents who receive a *China Watch* article on *The Telegraph*.

Those who receive the *Education Priming* treatment should be more likely to detect the political native advertisement irrespective of where they see the *China Watch* article. This is because they are primed to look for native advertising, especially as it relates to China. Whether this is, in fact, the case is important as it speaks to the possibility that information campaigns can immunize readers against the deceptive nature of political native advertising.

Reputational Costs of Political Native Advertising Given that respondents may be unable to detect political native advertising on their own, we inform the respondents that the *China Watch* story they have just read is a paid supplement from the Chinese government-controlled *China Daily* to examine whether there are costs to the hosting media's reputation for publishing political native advertisements. To measure possible reputational costs, we compare the level of respondent trust in the hosting media site *before* they receive their *China Watch* article with their level of trust

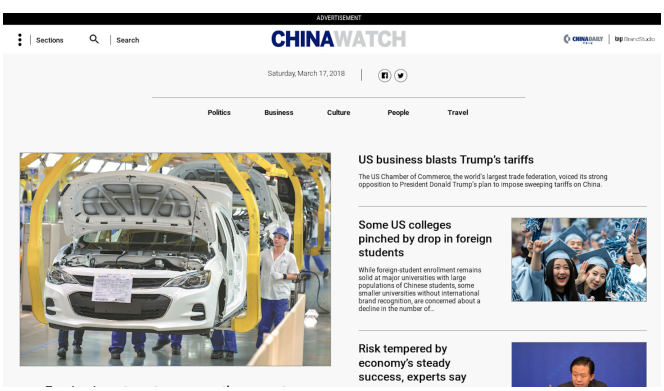

China Watch on the *Washington Post*

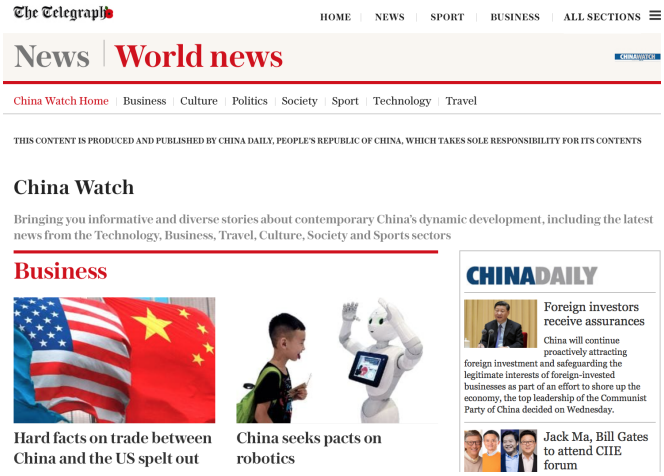

China Watch on the *The Telegraph*

Figure 2. *China Watch* on the *Washington Post* and *The Telegraph*

in the hosting media site *after* they are informed that the article was a paid supplement from a foreign government. A respondent's trust in the hosting media site is measured on a $1-6$ scale, where 1 indicates that the respondent has no trust in the hosting media site at all and 6 indicates that the respondent has a great deal of trust in the hosting media site.

While we expect that all respondents will lower their reported trust in the hosting media site when they are told that it is accepting political native advertisements from a foreign government, it is possible that this negative effect will be larger for those respondents who were unable to detect the political native advertising for themselves. These individuals learn that they have been deceived by the hosting media outlet and the foreign government advertiser into thinking that the political advertisement is editorial content by the hosting media. This additional sense of deception may cause their drop in trust in the hosting media outlet to be particularly marked.

Persuasiveness of Political Native Advertising To examine the persuasiveness of political native advertising, we compare how convincing the respondents find the two *China Watch* articles on the *Washington Post* and *The Telegraph* with how convincing they find the exact same articles when published on *China Daily*. Respondents are asked to rate how convincing their article is on a scale of $1-5$, where 1 indicates that the article is not at all convincing and 5 indicates that the article is very convincing. From the perspective of a political actor, the primary appeal of political native advertising is that her message is more convincing if it is camouflaged as standard editorial content on an independent, and thus more credible, hosting media site. Note, though, that the increased persuasiveness of political native advertising ultimately relies on successful deception. If respondents are able to identify the political native advertisement and hence the true source of the news story, then there is no reason to believe that they will find the *China Watch* articles any more (or less) convincing than the same articles in *China Daily*. Thus, political native advertising should only increase the persuasiveness of a news story among those respondents who fail to detect the political native advertisement.

Experiment Results

To make valid causal inferences, respondents must be assigned randomly to different treatment conditions. In other words, there should not be any factors that influence the treatment assignments (Trochim and Donnelly 2006). Although all of our treatments were randomly assigned, it is possible that the respondents in the various treatment groups could still differ demographically. To evaluate this, we conducted difference-in-means tests on reported demographics among the six treatment groups. We find that the six treatment groups are mostly balanced in terms of age, gender, education, and income. There is some evidence that the respondents in the Education Telegraph and Telegraph groups are slightly younger on average than the respondents in the Primed Wapo group.[4] As a result, we controlled for age in relevant analyses.

Detecting Political Native Advertising

To examine the factors that influence the probability that respondents are able to identify the true source of the *China Watch* story, we estimated a logit regression model in which the dependent variable, *Correct Source*, was coded 1 if the respondent identified the source of the news story as *China Daily* and 0 otherwise. The results from two slightly different model specifications are shown in Table 1.[5]

In Model 1, we find that the effect of the hosting media site, and implicitly the clarity of disclosure, is substantively large and statistically significant. The predicted probability that a respondent is able to identify the true source of the *China Watch* article increases by 0.569 $[0.503, 0.631]$ if a respondent sees the article on the *Washington Post* as opposed to *The Telegraph*.[6] Two-tailed 90% confidence intervals are shown in square brackets. In fact, 72% of respondents who read the *China Watch* article on the *Washington Post* were able to identify the true source of the article, while only 14% of respondents who read the

	Dependent Variable: Correct Source (0,1)	
	Model 1	Model 2
Washington Post	2.726***	2.854***
	(0.241)	(0.252)
Education Priming	0.413*	0.468*
	(0.238)	(0.246)
Media Experience		0.028
		(0.041)
Media Literacy		−0.022
		(0.035)
Media Trust		0.049
		(0.048)
Age	0.011	0.012
	(0.010)	(0.011)
Education		0.128
		(0.097)
Income		0.038
		(0.039)
Female		−0.110
		(0.252)
Constant	−2.271***	−3.754*
	(0.473)	(2.045)
Observations	444	439
Log Likelihood	−222.462	−213.998
Akaike Inf. Crit.	452.923	447.995

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$ (two-tailed)

Table 1. Identifying the Correct Source for the *China Watch* News Story

Note: The dependent variable, *Correct Source*, is a dichotomous variable that equals 1 if a respondent correctly identifies the true source of the *China Watch* article as *China Daily*, and 0 otherwise. Only those respondents who received a *China Watch* article are included. Estimates are based on logit regressions. Standard errors are shown in parentheses.

China Watch article on the *Telegraph* were able to do so. Although statistically significant, the substantive effect of the *Education Priming* treatment is smaller. To be precise, the predicted probability that a respondent is able to identify the true source of the *China Watch* article increases by 0.075 [0.003, 0.151] when a respondent receives the *Education Priming* treatment.[7]

It is possible that our respondents might be more familiar with the *Washington Post* than the *Telegraph* given that they are U.S. residents, although they indicate a similarly high level of trust towards the *Telegraph* in the survey. The potential extra experience with the *Washington Post* might help the respondents better recognize the stylistic difference between the *China Watch* page and the editorial page on the *Washington Post*. We might expect respondents with more media experience and higher media literacy to be more likely to recognize political native advertising. Therefore, in Model 2, we examine whether a respondent's general level of media experience, media literacy, media trust, and level of education affect the probability that she is able to identify the true source of the *China Watch* article. We employ standard media and education measures from the existing literature. *Media Experience* is based on five dimensions (news consumption, experience, expertise, familiarity, and access), each of which is measured on a 5-point scale. *Media Experience* is calculated as the sum of the scores for each of these five dimensions; its possible values run from 5 to 25, with higher numbers indicating more media experience (Flanagin and Metzger 2000). *Media Literacy* captures a respondent's understanding of media ownership, media and politics, media effects, news framing, agenda-setting, and the role of journalists. It draws on 14 survey items that are each measured on a 5-point scale. *Media Literacy* is calculated as the sum of the scores for each of these items; its possible values run from 14 to 70, with higher numbers indicating greater media literacy (Ashley et al. 2013). *Media Trust* uses seven survey items, each measured on a 5-point scale, to measure an individual's level of trust in the general media (Tsfati and Cappella 2003). *Media Trust* is simply the sum of the scores on the seven survey items; its possible values run from 7 to 35, with higher numbers indicating greater trust in the news media in general. *Education* is measured on a 1−8 scale, with larger numbers indicating higher levels of education. As the results in Model 2 indicate, none of these variables have any significant effect on the probability that a respondent is able to identify the true source of the *China Watch* articles. A likelihood ratio test indicates that the media and education variables are also jointly insignificant. These variables are also insignificant if we remove our two treatment variables, *Washington Post* and *Education Priming*. In sum, there is no evidence that more educated or more media-savvy individuals are any better at identifying political native advertising than other individuals.

Reputational Costs of Political Native Advertising

In Table 2, we compare the mean level of trust in each of the two hosting media sites before the respondents receive their *China Watch* article (Pre-Treatment) and after the respondents are informed that the article is a paid supplement from a foreign government (Post-Treatment). The third column, Paired Difference, indicates whether the mean of the paired differences in the pre-treatment and post-treatment levels of trust for each respondent are statistically different. [8] As expected, trust in both the *Washington Post* and *The Telegraph* declines in a statistically significant way when respondents are told that the *China Watch* article comes from the Chinese government-controlled *China Daily* and not the hosting media outlet. The magnitude of the decline in trust is almost identical for both hosting media outlets — the average decline in trust is 0.80 points for those respondents who received the *China Watch* article on the *Washington Post* and 0.77 points for those respondents who

	Trust in Hosting Media Outlet (1-6)			
	Pre-treatment	Post-treatment	Paired Difference (Post−Pre)	N
Washington Post	4.09 (1.36)	3.29 (1.33)	−0.80***	182
The Telegraph	3.42 (1.36)	2.65 (1.31)	−0.77***	121

*p<0.1; **p<0.05; ***p<0.01

Table 2. Comparing Pre- and Post- Treatment Trust in the Hosting Media Outlets

Note: The pre-treatment and post-treatment columns indicate the mean level of trust in each of the two hosting media sites before the respondents receive their *China Watch* article and after respondents are informed that the article is a paid supplement from a foreign government. Standard deviations are shown in parentheses. The paired difference column indicates the mean of the paired differences for each respondent between their pre- and post-treatment levels of trust. Only the respondents with complete answers to both pre- and post- treatment questions on the trust in hosting media outlets that are used for the paired t-tests are included to calculate the pre-treatment and post-treatment group means.

		Trust in Hosting Media Outlet (1-6)		
		Pre-treatment	Post-treatment	Paired Difference (Post−Pre)
Washington Post	Self-Detected	4.09 (1.39)	3.33 (1.36)	−0.76***
	Not Self-Detected	4.08 (1.25)	3.15 (1.25)	−0.93***
The Telegraph	Self-Detected	3.36 (1.47)	2.73 (1.16)	−0.64
	Not Self-Detected	3.43 (1.34)	2.64 (1.35)	−0.80***

*p<0.1; **p<0.05; ***p<0.01

Table 3. Change in Trust in the Hosting Media Outlets

Note: The pre-treatment and post-treatment columns indicate the mean level of trust in each of the two hosting media sites before the respondents receive their *China Watch* article and after respondents are informed that the article is a paid supplement from a foreign government. The two rows for each media outlet distinguish between those respondents who were able to detect the political native advertisement on their own and those respondents who were not.Standard deviations are shown in parentheses. The paired difference column indicates the mean of the paired differences for each respondent between their pre- and post-treatment levels of trust. Only the respondents with complete answers to both pre- and post- treatment questions on the trust in hosting media outlets that are used for the paired t-tests are included to calculate the pre-treatment and post-treatment group means.

received the article on *The Telegraph*. These differences are substantively meaningful as they equate to a 27.4% decline in trust for the *Washington Post* and a 29.2% decline in trust for *The Telegraph*. The fact that the decline in trust is similar across the two hosting media outlets is interesting as it suggests that the better disclosure provided by the *Washington Post* does not immunize it against potential reputational costs if the political native advertising is detected, which contradicts findings in commercial native advertising that readers have less unfavorable view towards the hosting media outlet when it is more transparent about native advertising (Amazeen and Wojdynski 2019). It is possible that news consumers are much less tolerant when it comes to political native advertising, especially when the advertisers are foreign governments, given the political implications that independent media outlets can be co-opted by foreign governments to deliver propaganda.

Recall that we expect respondents who are unable to detect political native advertising on their own to exhibit a particularly marked reduction in their trust towards the hosting media outlet after they are told that the *China Watch* article is a paid advertisement from the Chinese government due to the additional perceived deception. In line with our expectation, the results reported in Table 3 indicate that the decline in trust towards the hosting media outlet is in fact larger among the respondents who were unable to detect the political native advertising for themselves than among those respondents who were able to detect it. This is true regardless of whether the respondents saw the *China Watch* article on the *Washington Post* or *The Telegraph*. In robustness checks reported in Online Appendix A, we find that this effect of self-detection is also statistically significant after controlling for demographic and media related variables. Together the results show that hosting media outlets that accept political native advertising can expect to suffer a reputational cost if the political native advertising is revealed, especially if readers feel that they have been deceived.[9]

Persuasiveness of Political Native Advertising

To examine the persuasiveness of political native advertising, in Table 4 we compare how persuasive the respondents find the two *China Watch* articles based on whether they saw the article on one of the hosting media outlets (Political Native Advertising) or on *China Daily* (No Political Native Advertising) and whether they were able to detect the native advertising on their own. Note that we exclude the respondents who received *Education Priming* in this analysis. Consistent with our expectations, the articles that appear on the hosting media site are considered significantly more persuasive than the same article on *China Daily* when the respondents fail to identify the true source of the article. This is true both for those respondents who saw the article on the *Washington Post* and those who saw the article on

	Mean of Persuasiveness (1-5)			
	On Hosting Media			
	Identified True Source	Not Identified True Source	On China Daily	Difference
Washington Post China Watch Article	3.49 (0.83)		3.47 (0.89)	0.02
		3.77 (0.84)	3.47 (0.89)	0.31*
The Telegraph China Watch Article	3.27 (0.96)		3.43 (1.07)	−0.17
		3.87 (0.84)	3.43 (1.07)	0.44***

*p<0.1; **p<0.05; ***p<0.01

Table 4. Persuasiveness, Deception, and Political Native Advertising

Note: The first two columns indicate how persuasive the respondents found each of the two *China Watch* articles on average when they read the stories on the hosting media outlet. The first column focuses on those respondents who identified the true source of the *China Watch* article, while the second column focuses on those respondents who failed to identify the true source of the article. Standard deviations are included in parentheses. The third column indicates how persuasive the respondents found each of the two articles when they read the stories on *China Daily*. The final column indicates whether the mean level of persuasiveness was significantly higher when the articles were read on the hosting media outlet instead of *China Daily*. Cells shown in gray refer to those respondents who failed to detect the political native advertisement.

The Telegraph: the *China Watch* article that is perceived as from the *Washington Post* is, on average, 0.31 points more persuasive (9% increase) than the same article on the *China Daily*. The *China Watch* article that is perceived as from *The Telegraph* is, on average, 0.44 points more persuasive (13% increase) than the same article on the *China Daily*. This is exactly what we expected and highlights the fact that the effectiveness of political native advertising relies on successful deception.[10]

Conclusion

In this article, we study an overlooked yet important form of propaganda, political native advertising, in which political actors buy space on independent media sites to publish political advertisements that mimic the standard editorial content found on the hosting media site. As one of the first empirical studies on political native advertising, we examine some basic but important key features of this type of propaganda. Specifically, we examine the deceptiveness and persuasiveness of political native advertisements and the reputational costs associated with political native advertising to the hosting media outlets. Using an online survey experiment with real political native advertisements, we find that unless the political native advertising is clearly disclosed using visual separation as in the *Washington Post*, respondents are unlikely to detect the true source of political native advertising, irrespective of their level of education and experience as a news consumer. The deceptive nature of political native advertising makes it an effective tool for political actors to influence the citizenry. As we show in our results, the same message is perceived as much more persuasive when it comes in the form of native advertisement and is perceived as coming from the independent hosting media outlet. Note that the respondents are recruited from MTurk, which are previously shown to be better educated and more media savvy than the overall population in the U.S. (Berinsky et al. 2012), the effects found in the experiment are likely to be even stronger in the population: citizens, on average, are likely to have an even harder time detecting the political native advertising than the respondents in the experiment. The increased salience of foreign disinformation campaigns after the 2016 election in the U.S. is also likely to make citizens more alert of political native advertising and likely to punish the hosting media outlets more than citizens in other targeted countries. However, with the increased popular and scholarly concern over disinformation and misinformation campaigns around the world (e.g. Brady 2015, 2018; Tucker et al. 2017), we are likely to observe similar responses to political native advertising from citizens outside of the U.S.

Although our findings are troubling in that citizens are often unable to detect political native advertising unless it is clearly disclosed and are likely to be influenced by such propaganda, our analysis also suggests methods to deter such propaganda. Much of the effectiveness of political native advertising depends on the deceptiveness of the advertising. It suggests that decreasing the deceptiveness of the advertising and/or increasing the reader's ability to detect the political advertising are likely to reduce the incentives that political actors have to engage in these sorts of propaganda activities and deter independent media outlets from cooperating with political actors. However, there seem to be no significant reputational gains to the hosting media outlet to disclose political native advertising. As a result, disclosure regulations are needed to provide media outlets additional motivations to disclose political native advertising. In addition, although less effective than clear disclosure, informational campaigns that publicize the presence and use of political native advertising can also increase the reader's ability to recognize political native advertising.

Our results also have important regulatory implications for native advertising in general. Although studies on commercial native advertising find certain types of disclosure to be more effective in increasing the respondents' advertisement recognition, in general, only around 10% of the respondents were able to detect native advertising under

the better disclosure conditions (Amazeen and Wojdynski 2019; Wojdynski and Evans 2016). Our finding that respondents were much more likely to detect the true source of the *China Watch* article on *The Washington Post* suggests that increasing the degree of visual separation between the native advertisement and editorial content might be a better method than simply using text disclosures to alert readers of native advertising.

Acknowledgements

We thank Amber Boydstun, Matt Golder, Sona Golder, Arthur Lupia, Benjamin Radford, David Szakonyi, Rory Truex, and the anonymous reviewers, as well as audiences at the 2019 Annual Meeting of the Southern Political Science Association and the 2019 Annual Meeting of the American Political Science Association for their helpful comments on this paper.

Ethical Approval

The online survey experiment described in this paper received IRB approval (Study #00008806) at Pennsylvania State University.

Supplemental material

1. Additional analysis can be found in the Online Appendix A. The exact texts of the two treatment articles are included in the Online Appendix B and C. More information on the media variables can be found in the Online Appendix D.

2. The data, codebook, and R code necessary to replicate the results and figures in this analysis can be accessed through Research & Politics Dataverse at https://doi.org/10.7910/DVN/GBLKVP.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Notes

1. Legacy media refer to older and more traditional media outlets such as newspapers, television, and radio, in which the audience does not actively "interact" with the media content.
2. The hosted website and column, russianow.washingtonpost.com, disappeared in 2015. For information and reports on *Russia Now*, see Barton (2015) and the Washington Free Beacon (2014).
3. For example, the URL for the Washington Post's *China Watch* article is http://chinawatch.washingtonpost.com/2018/02/capital-market-to-open-further/. This URL no longer works, since the *Washington Post* removed the *China Watch* page in the summer of 2018, while continuing to deliver the print version of *China Watch* with the printed *Washington Post.*
4. Detailed results are included in Online Appendix A section 1.
5. Note that only the 444 respondents who are in the four treatment groups with *China Watch* articles on the hosting media outlets are relevant and hence included in this analysis.
6. This change in predicted probability is calculated for a respondent who is 41 years old (sample mean) and who did not receive the *Education Priming* treatment.
7. This change in predicted probability is calculated for a respondent who is 41 years old and who saw the *China Watch* article on the *Washington Post* site.
8. The numbers in the Paired Difference column reflect the average difference in the *paired* post-treatment and pre-treatment levels of trust from *each* respondent which is different from the difference in the two group means of the post-treatment levels of trust and pre-treatment levels of trust from *all* respondents. A paired difference-in-means test is appropriate for our within-subject before-and-after experimental design and is generally more powerful than an unpaired test because it reduces inter-subject variability (Larsen et al. 1986).
9. Additional robustness checks controlling for demographic and media related variables are included in Online Appendix A section 4.
10. Additional robustness checks controlling for demographic and media related variables are included in Online Appendix A section 5.

References

Amazeen MA and Muddiman AR (2018) Saving media or trading on trust? the effects of native advertising on audience perceptions of legacy and online news publishers. *Digital Journalism* 6(2): 176–195.

Amazeen MA and Wojdynski BW (2018) The effects of disclosure format on native advertising recognition and audience perceptions of legacy and online news publishers. *Journalism* : 1464884918754829.

Amazeen MA and Wojdynski BW (2019) Reducing native advertising deception: Revisiting the antecedents and consequences of persuasion knowledge in digital news contexts. *Mass Communication and Society* 22(2): 222–247.

Archer AM and Clinton J (2018) Changing owners, changing content: Does who owns the news matter for the news? *Political Communication* 35(3): 353–370.

Ashley S, Maksl A and Craft S (2013) Developing a news media literacy scale. *Journalism & Mass Communication Educator* 68(1): 7–21.

Barton E (2015) China, Russia pay *Washington Post* to publish their propaganda. *The Daily Caller*, October 29.

Batsell J (2017) 4 steps to bring ethical clarity to native advertising. Nieman Reports.

Berinsky AJ, Huber GA and Lenz GS (2012) Evaluating online labor markets for experimental research: Amazon.com's mechanical turk. *Political Analysis* 20(3): 351–368.

Brady AM (2015) China's foreign propaganda machine. *Journal of Democracy* 26(4): 51–59.

Brady AM (2018) China in xi's 'new era': New zealand and the ccp's 'magic weapons'. *Journal of Democracy* 29(2): 68–75.

Carlson M (2015) When news sites go native: Redefining the advertising–editorial divide in response to native advertising. *Journalism* 16(7): 849–865.

China Daily (2018) *china daily*'s global development.

Clifford S, Jewell RM and Waggoner PD (2015) Are samples drawn from mechanical turk valid for research on political ideology. *Research and Politics* 4(2).

Cook S (2017) Chinese government influence on the U.S. media landscape. Testimony before the U.S.-China Economic and

Security Review Commission Hearing on China's Information Controls, Global Media Influence, and Cyber Warfare Strategy, May 4.

Crabtree C, Kern HL and Siegel DA (2018) Cults of personality, preference falsification, and the dictator's dilemma. Unpublished manuscript, University of Michigan.

Di Tella R and Franceschelli I (2011) Government advertising and media coverage of corruption scandals. *American Economic Journal: Applied Economics* 3(4): 119–51.

Dykhne I (2017) Persuasive or deceptive-native advertising in political campaigns. *S. Cal. L. Rev.* 91: 339.

Edmonds R (2017) Native advertising grows up fast, shedding its rogue image. Poynter.

Einstein M (2016) *Black Ops Advertising: Native Ads, Content Marketing and the Covert World of the Digital Sell*. OR Books.

Fallows J (2011) Official Chinese propaganda: Now online from the WaPo! *The Atlantic*, February 3.

Firth K, Hoffman DA and Wilkinson-Ryan T (2017) Law and psychology grows up, goes online and replicates. Public Law and Legal Theory Research Paper Series # 17-38, University of Pennsylvania Law School.

Flanagin A and Metzger M (2000) Perceptions of internet information credibility. *Journalism & Mass Communication Quarterly* 77(3): 515–540.

Gehlbach S and Sonin K (2014) Government control of the media. *Journal of Public Economics* 118: 163–171.

Howe P and Teufel B (2014) Native advertising and digital natives: The effects of age and advertisement format on news website credibility judgments. *ISOJ Journal* 4(1): 78–90.

Huang H (2015) Propaganda as signaling. *Comparative Politics* 47(4): 419–444.

Huang H (2018) The pathology of hard propaganda. *Journal of Politics* 80(3): 1034–1038.

Iversen MH and Knudsen E (2017) When politicians go native: The consequences of political native advertising for citizens trust in news. *Journalism* : 1464884916688289.

Jamieson KH et al. (2000) *Everything you think you know about politics–and why you're wrong*. Basic Books New York.

Larsen RJ, Marx ML et al. (1986) *An introduction to mathematical statistics and its applications*, volume 2. Prentice-Hall Englewood Cliffs, NJ.

Little A (2018) Fake news, propaganda, and lies can be pervasive even if they aren't persuasive. In: Golder M and Golder SN (eds.) *Fake News and the Politics of Misinformation, Comparative Politics Newsletter*. 28(2), pp. 49–54.

Little AT (2017) Propaganda and credulity. *Games and Economic Behavior* 102: 224–232.

Marquez X (2018) Two models of political leader cults: Propaganda and ritual. *Politics, Religion, and Ideology* 19(3): 265–284.

McDermott R (2011) New directions for experimental work in international relations. *International Studies Quarterly* 55(2): 503–520.

Mullin B (2016) Vox Media pitches signature "Explainer" format to advertisers. The Wall Street Journal.

Murtha J and Gourarie C (2016) Do BuzzFeeds native political ads cross a line?

Paolacci G, Chandler J and Ipeirotis PG (2011) Running experiments on amazon mechanical turk. *Judgment and Decision Making* 5(5): 411–419.

Stockmann D (2013) *Media commercialization and authoritarian rule in China*. Cambridge University Press.

Trochim WM and Donnelly JP (2006) *Research Methods Knowledge Base*. Atomic Dog.

Tsfati Y and Cappella JN (2003) Do people watch what they do not trust? exploring the association between news media skepticism and exposure. *Communication Research* 30(5): 504–529.

Tucker JA, Theocharis Y, Roberts ME and Barberá P (2017) From liberation to turmoil: Social media and democracy. *Journal of democracy* 28(4): 46–59.

Washington Free Beacon (2014) WaPo Still Publishing Russian Propaganda. September 3.

Wojdynski BW and Evans NJ (2016) Going native: Effects of disclosure position and language on the recognition and evaluation of online native advertising. *Journal of Advertising* 45(2): 157–168.

Appendix A: Additional Analyses and Robustness Checks

1 Sample Information

	Group Means						
	Primed Wapo	Wapo	Primed Telegraph	Telegraph	China Daily Article W	China Daily Article T	Full Sample
Age	41.97	41.56	39.21*	39.27*	40.73	41.50	40.72
	(11.45)	(11.63)	(11.79)	(10.83)	(12.71)	(12.60)	(11.86)
Female	0.56	0.60	0.58	0.59	0.59	0.55	0.58
	(0.50)	(0.49)	(0.50)	(0.49)	(0.49)	(0.50)	(0.49)
Education	4.50	4.50	4.45	4.69	4.39	4.50	4.50
	(1.28)	(1.32)	(1.37)	(1.25)	(1.41)	(1.41)	(1.34)
Income	6.38	6.35	6.11	6.64	6.24	6.58	6.38
	(3.42)	(3.02)	(3.42)	(3.50)	(2.94)	(3.39)	(3.28)
N	116	109	112	107	105	111	660

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$ (two-tailed)

Table 1: Demographic Balance Across Treatment Groups

Note: Table 1 indicates the means for different demographic variables across the six treatment groups and the sample as a whole; standard deviations are shown in parentheses. The six treatment groups correspond to those shown in Figure 1. *Education* is measured on a $1-8$ scale, with larger numbers indicating higher levels of education. *Income* is measured on a $1-12$ scale, with larger numbers indicating higher income. The *Primed Wapo* treatment group is treated as the baseline group for conducting difference-in-means tests. Welch's t-test, which allows for unequal variances, was used for the difference-in-means tests. All demographic information was gathered prior to the experimental treatments.

In Table 1, we provide demographic information on the 660 respondents in our sample. The information is provided for each of the six treatment groups shown in Figure 1 in the main article as well as for the sample as a whole. Although all of our treatments were randomly assigned, it is possible that the respondents in the various treatment groups could still differ demographically. To evaluate this, we conducted difference-in-means tests in which the baseline group included those respondents who received the *Education Priming* treatment and received the *China Watch* article on the *Washington Post* (Primed Wapo). As Table 1 indicates, the six treatment groups are mostly balanced in terms of age, gender, education, and income. There is some

evidence that the respondents in the Education Telegraph and Telegraph groups are slightly younger on average than the respondents in the Primed Wapo group. As a result, we controlled for age in relevant analyses.

2 Similar Quality of Treatment Articles

In the experiment, we use one *China Watch* article on the *Washington Post* and one *China Watch* article on *The Telegraph* as political native advertising treatments, and the same two articles on their original media outlet, *China Daily*. Although it would be ideal for the treatment articles to have identical content to prevent any effects caused by the different content and quality of the treatment articles, *China Daily* intentionally avoids displaying identical articles in different news outlets and we are unable to identify identical *China Watch* articles on the two hosting news outlets. To minimize any effects associated with the different quality and content of different articles, we used two articles on the same topic making the same argument. In the main text, we claimed that respondents did not perceive any differences in the quality of the two articles. We now discuss the empirical evidence on which this claim is based.

To examine whether the two articles are perceived differently, we compare the two articles in terms of their perceived levels of accuracy, bias, and persuasiveness. An article's level of accuracy is measured on a scale of $1-5$, with larger numbers indicating greater accuracy. An article's level of bias is measured on a scale of $1-5$, with larger numbers meaning that the article was perceived as *less* biased. An article's level of

Table 2: Differences in the Quality of Two *China Watch* Articles

	China Daily Wapo	China Daily Telegraph	Difference in Means (t-statistics)
Accuracy	3.35 (0.76)	3.41 (0.79)	-0.06 (-0.59)
Bias	2.95 (0.965)	3.03 (1.040)	-0.08 (-0.54)
Persuasiveness	3.47 (0.89)	3.43 (1.07)	0.03 (0.26)

*p<0.1; **p<0.05; ***p<0.01

Note: The first column indicates the means for the perceived accuracy, bias and persuasiveness of the treatment article for the respondents who read the *Washington Post*'s *China Watch* article on the *China Daily* platform. The second column indicates the same information for those who read the *Telegraph*'s *China Watch* article on the *China Daily* platform. The third column indicates the differences in the means. Welch's t-test, which allows for unequal variances, was used for the difference-in-means tests.

persuasiveness is measured on a scale of $1-5$, with larger numbers meaning that the article was considered more convincing. We compare the quality of the two articles along these dimensions as they appeared on *China Daily*. The results in Table 2 indicate that the two articles are considered similar in terms of their accuracy, bias, and persuasiveness. The differences between the two articles are not only substantively small but also statistically insignificant.

3 Quality of Responses

The survey experiment relies on voluntary participation. To control the quality of the responses, we limit the participation to the top-ranking MTurk respondents. However, it is still possible that some of the respondents might simply give random answers to the survey and ignore the treatment. While we cannot monitor the online survey responses, there are two ways that can help us check the quality of the responses. First, we include a question on the source of the treatment article that the respondents read. For the respondents who receive *China Watch* articles, answers indicating a source of either the hosting media outlet or *China Daily* indicate some attention to the treatment. Since there are four options, there is a 50% chance that the respondents guess correctly without actually paying attention to the treatment. In our experiment, only 10 out of 444 respondents (2.25%) in the four China Watch treatment groups failed to answer the source as either the hosting media outlet or the *China Daily*. For the two groups of respondents who receive the China Daily article to read, only one option out of the four options is reasonable and there is a 25% chance they can guess correctly without paying attention to the treatment. Only 1 out of 215 (0.47%) respondents in the two *China Daily* treatment groups failed to identify the source as *China Daily*. The extremely low percentage of respondents who fail to identify either the hosting media outlets or the China Daily as the source indicate a high quality of responses of our survey. In other words, the majority of the respondents paid at least some attention to the treatment articles.

Second, we also record the duration of responses. On average, the respondents spend 4 minutes to complete the survey. 60 out of the 660 respondents completed the survey within 1 minute. Although it is a short and simple survey, completion within 1 minute is still alarming and possibly indicates a lack of proper attention to the questions. If the lack of attention does not correlate with the treatment assignment (missing treatment at random), there would be no threat to internal validity of the experiment. In other words, those

Table 3: Identifying Correct Source for the China Watch Story Among Attentive Respondents

	Dependent variable:	
	Correct Source	
	(1)	(2)
Washington Post	2.821***	2.963***
	(0.259)	(0.272)
Education Priming	0.590**	0.676**
	(0.255)	(0.263)
Age	0.012	0.013
	(0.011)	(0.012)
Education		0.128
		(0.103)
Media Experience		0.035
		(0.044)
Media Literacy		−0.023
		(0.037)
Media Trust		0.053
		(0.050)
Income		0.044
		(0.041)
Female		−0.207
		(0.269)
Constant	−2.448***	−4.082*
	(0.509)	(2.182)
Observations	404	401
Log Likelihood	−197.340	−190.458
Akaike Inf. Crit.	402.681	400.916

*p<0.1; **p<0.05; ***p<0.01

Note: Only the respondents who spend at least 1 minute to complete the survey and receive a China Watch article to read are included in the analysis. The dependent variable, *Correct Source*, is a dichotomous variable that equals 1 if a respondent correctly identifies the true source of the *China Watch* article as *China Daily*, and 0 otherwise. Estimates are based on logit regressions. Standard errors are shown in parentheses.

inattentive respondents would be random noise instead of systematic bias if the short attention does not correlate with the treatment assignment. In fact, we find that the inattentive respondents distribute evenly among the six treatment groups instead of clustering in certain groups, which indicates that the inattentive respondents are likely to be random noise.

In addition, we repeat all the analyses in the main paper on only those respondents with proper attention, the ones that spend at least 1 minute to complete the survey, as additional robustness checks. In Table 3, we report the likelihood of detecting the true source of a China Watch article among the respondents who spend at least 1 minute to complete the survey. Compared to the results in Table 2 in the main article in which all respondents are included, the results are in fact stronger, which again suggests that the inattentive respondents add random noise to the analysis rather than bias. The respondents who received a China Watch article to read are significantly more likely to detect the true source of the China Watch article as *China Daily*. Educational priming also significantly increases the likelihood of detection. Similarly, we report the difference in post- and pre-treatment trust in the hosting media outlet in Table 4 and the perceived persuasiveness of the article in different media outlets in Table 5. Compared to Table 3 and Table 5 in the main article where all respondents are included, the results are almost identical.

Table 4: Post- and Pre- Treatment Trust in the Hosting Media Outlets Among Attentive Respondents

	Trust in Hosting Media Outlet (1-6)		
	Pre-treatment	Post-treatment	Paired Difference (Post−Pre)
Washington Post	4.10 (1.33)	2.93 (1.41)	-0.83^{***}
The Telegraph	3.37 (1.33)	2.42 (1.28)	-0.68^{***}

*p<0.1; **p<0.05; ***p<0.01

Note: Only the respondents who spend at least 1 minute to complete the survey are included in the analysis. The pre-treatment and post-treatment columns indicate the mean level of trust in each of the two hosting media sites before the respondents receive their *China Watch* article and after respondents are informed that the article is a paid supplement from a foreign government. Standard deviations are shown in parentheses. The paired difference column indicates the mean of the paired differences for each respondent between their pre- and post-treatment levels of trust.

4 Additional Robustness Checks on the Reputational Costs

In the main paper, we argue that the respondent will decrease their trust in the hosting media outlets once they discover the political native advertising. Empirically, we expect a decrease in the reported trust in the

Table 5: Persuasiveness, Deception, and Political Native Advertising

	Mean of Persuasiveness (1-5)			
	On Hosting Media			
	Identified True Source	Not Identified True Source	On China Daily	Difference
Washington Post China Watch Article	3.47 (0.86)		3.40 (0.87)	0.07
		3.77 (0.84)	3.40 (0.87)	0.37**
The Telegraph China Watch Article	3.21 (0.97)		3.43 (1.09)	−0.22
		3.87 (0.80)	3.43 (1.09)	0.44***

*p<0.1; **p<0.05; ***p<0.01

Note: Only the respondents who spend at least 1 minute to complete the sutvey are included in the analysis of Table 5. The first two columns indicate how persuasive the respondents found each of the two *China Watch* articles on average when they read the stories on the hosting media outlet. The first column focuses on those respondents who identified the true source of the *China Watch* article, while the second column focuses on those respondents who failed to identify the true source of the article. Standard deviations are included in parentheses. The third column indicates how persuasive the respondents found each of the two articles when they read the stories on *China Daily*. The final column indicates whether the mean level of persuasiveness was significantly higher when the articles were read on the hosting media outlet instead of *China Daily*. Cells shown in gray refer to those respondents who failed to detect the political native advertisement.

hosting media outlets once we inform the respondents about political native advertising. We do not have predictions on other variables' effects on the magnitude of the decrease in trust. Therefore, simple t-tests are suitable for testing our hypothesis. In this section, we examine the other variables that might influence the post- and pre-treatment trust in hosting media outlets as additional robustness checks. In Table 6, we report results from three OLS models with slightly different specifications. The dependent variable is the difference between post-treatment trust and the pre-treatment trust in the hosting media outlet, which takes possible value from -5 to 5. A negative value in the DV means a decrease in trust, and a positive value in the DV means an increase in trust after being informed that the treatment article is a paid content from *China Daily*. In Model 1, we only include the pre-treatment level of trust and the dummy variable indicating whether the respondents identify the source of the China Watch article correctly. Model 2 includes additional variables measuring respondents' media literacy experience and trust. Model 3 includes demographic variables in addition to the media variables. Consistent with our expectation and the results in Table 4 of the main article, correctly identifying the true source of the article has a significant positive effect on the change in trust and this positive effect is robust and consistent in three model specifications. Pre-treatment level of trust also influences the change in trust significantly. Respondents with higher pre-treatment levels of trust in the

Table 6: Change in Trust in the Hosting Media Outlets

	Dependent variable:		
	Post-pre Treatment Difference in Trust Towards Hosting Media Outlets		
	(1)	(2)	(3)
Pre-treatment Trust	-0.625^{***}	-0.655^{***}	-0.661^{***}
	(0.052)	(0.053)	(0.053)
Media Trust		0.040	0.039
		(0.029)	(0.029)
Media Literacy		-0.023	-0.024
		(0.020)	(0.020)
Media Experience		0.068^{***}	0.066^{***}
		(0.025)	(0.025)
Identify Source Correctly	0.325^{**}	0.287^{**}	0.279^{*}
	(0.144)	(0.144)	(0.144)
Age		-0.004	-0.002
		(0.007)	(0.007)
Female			-0.031
			(0.148)
Education			0.098^{*}
			(0.056)
Income			-0.036
			(0.023)
Constant	1.428^{***}	0.644	0.528
	(0.215)	(1.119)	(1.159)
Observations	303	300	300
R^2	0.328	0.357	0.367
Adjusted R^2	0.324	0.344	0.347
Residual Std. Error	1.240 (df = 300)	1.224 (df = 293)	1.219 (df = 290)
F Statistic	73.367^{***} (df = 2; 300)	27.098^{***} (df = 6; 293)	18.658^{***} (df = 9; 290)

*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is created by subtracting the pre-treatment level of trust from the post-treatment level of trust in the hosting media outlet. The possible values range from -5 to 5. A negative value means a decrease in trust in the hosting media outlet after the treatment. Estimates are based on three liner regressions. Standard errors are shown in parentheses.

hosting media are likely to experience a greater drop in trust after learning about the practice of political native advertising by the hosting media outlet. Among the three media variables, only *Media Experience* has a significant effect on the change in trust. Respondents with more media experience are likely to be more tolerant of political native advertising and decrease their trust in the hosting media less.

Figure 1: Predicted Change in Post- and Pre- Treatment Trust in Hosting Media Outlets

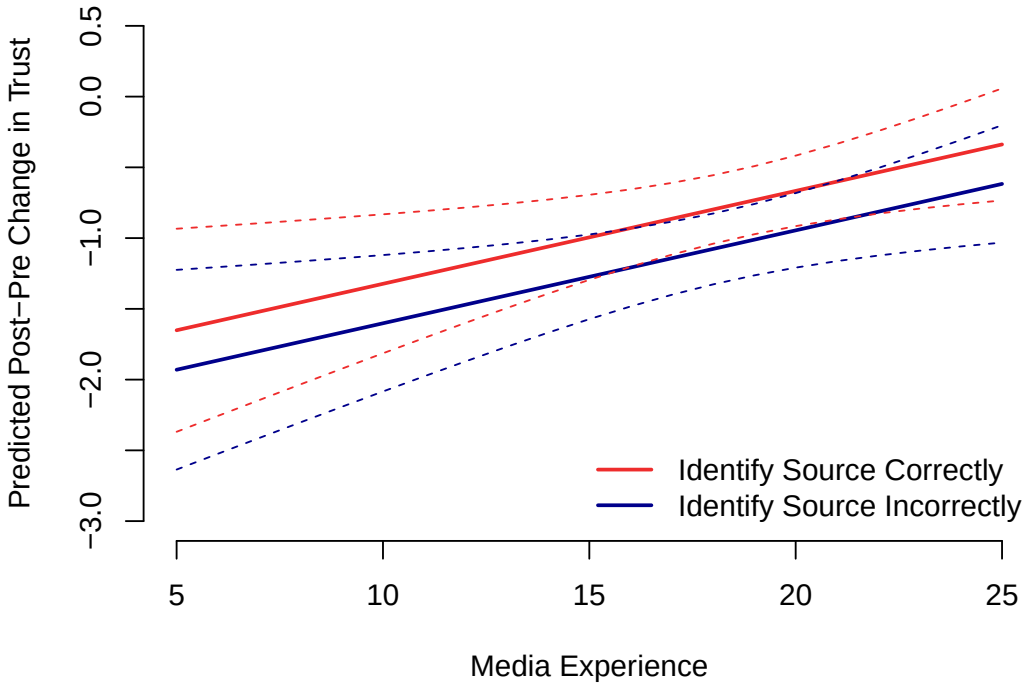

Note: Figure 1 shows the predicted change in trust based on Model 3 in Table 6. The x-axis indicates Media Experience which is an additive measure with possible values from 5 to 25. All other variables are set to be at their median values. The dash lines indicate 95% confidence intervals.

As mentioned earlier, although the regression analyses in Table 6 provide informative results on the variables that might affect the magnitude of the change in trust, they do not directly test our hypothesis that the respondents in general should lower their trust in the hosting media outlet after knowing that the media outlet has engaged in political native advertising. In other words, the post-treatment trust should be lower than the pre-treatment trust, and the difference in the post- and pre-treatment level of trust should be negative. Therefore, in Figure 1, we visualize the predicted change in trust for readers with different media experiences and for the cases where they identify the true source on their own and where they fail to identify the true source of the China Watch article. [1] The solid lines indicate predicted change in trust

[1]The predicted change is based on Model 3 in Table 6. All other control variables are set at their median values with pre-

and the dashed lines represent 95% confidence intervals. As can be seen from the figure, readers who fail to identify source correctly on their own drop their trust in the hosting media 0.3 points more than those who correctly identify the source after knowing the hosting outlet's involvement in political native advertising. As media experience increases, respondents are predicted to decrease their trust in the hosting media outlet less. However, the predicted change in trust always stays negative as media experience increases. The bottom line here is that the main result in the article that readers decrease their trust in the hosting media outlets after knowing about the hosting outlet's involvement in political native advertising is robust after controlling for demographic characteristics and media related variables. Although the magnitude of change in trust varies among respondents with different pre-treatment levels of trust and media experience, respondents in general decrease their trust in the hosting media significantly, especially among those who are not able to detect political native advertising on their own.

5 Additional Robustness Checks on the Persuasiveness

Using simple t-tests in the main article, we find that political native advertising is more persuasive when it is perceived to be from the hosting media outlet. In this section, we provide additional robustness checks by controlling demographic and media related variables. Given the ordered and discrete nature of the dependent variable, *Level of Persuasiveness*, we estimate four ordered logit regressions and report the results in Table 7. Model 1 and Model 3 include respondents who read the China Watch article on the *Washington Post* and the same article on the original *China Daily* site. Model 2 and Model 4 include respondents who read the China Watch article on *The Telegraph* and the same article on the *China Daily* site. All four models control the same demographic variables and the variables relate to media. Respondents who received the *Education Priming* prior to reading the China Watch article are excluded from the analysis. The variable *Washington Post China Watch* in Model 1 is a dichotomous variable with value 1 indicating the China Watch article is on the *Washington Post* and value 0 indicating the article is on the *China Daily*. Similarly, the variable *Telegraph China Watch* in Model 2 is also a dichotomous variable with 1 indicating the article is on *The Telegraph* and 0 indicating the article is on the *China Daily*. Results from Model 1 and Model 2 are consistent with our findings in the main article. The China Watch article on the *Washington Post* is more persuasive than the same article on *China Daily*. However, this positive effect is substantively small

treatment trust to the hosting media outlet as 4 (moderate amount of trust), age as 38, gender as female, level of education as

Table 7: Persuasiveness of Political Native Advertisement

	Dependent variable:			
	Persuasiveness of Treatment Article (1-5)			
	(1)	(2)	(3)	(4)
Washington Post China Watch	0.142			
	(0.266)			
Telegraph China Watch		0.649**		
		(0.258)		
Perceived as Washington Post			0.604	
			(0.422)	
Perceived as Telegraph				0.797***
				(0.270)
Perceived as China Daily			−0.020	−0.244
			(0.289)	(0.516)
Age	−0.003	−0.010	−0.003	−0.008
	(0.011)	(0.011)	(0.011)	(0.011)
Education	−0.164*	−0.069	−0.149	−0.072
	(0.097)	(0.101)	(0.098)	(0.101)
Media Experience	0.017	0.0003	0.022	0.002
	(0.048)	(0.046)	(0.048)	(0.046)
Media Literacy	−0.019	0.014	−0.019	0.008
	(0.040)	(0.036)	(0.040)	(0.036)
Media Trust	0.040	0.084*	0.042	0.084*
	(0.052)	(0.047)	(0.053)	(0.047)
Income	0.006	0.065	−0.001	0.074*
	(0.045)	(0.041)	(0.045)	(0.041)
Female	0.477*	−0.372	0.473*	−0.344
	(0.284)	(0.266)	(0.284)	(0.266)
Observations	208	217	208	217

*p<0.1; **p<0.05; ***p<0.01

Note: The estimates are based on ordered logit models. Standard errors are included in parentheses. Model 1 and Model 3 include respondents who read the China Watch article on the *Washington Post* and the same article on the original China Daily site. Model 2 and Model 4 include respondents who read the China Watch article on *The Telegraph* and the same article on the China Daily site. Respondents who received the *Education Priming* prior to reading the China Watch article are excluded from the analysis.

and statistically insignificant. The China Watch article on *the Telegraph* is much more persuasive than the same article on the *China Daily*, and this positive effect is statistically significant.

Next, we separate the respondents who successfully detect the true source of the China Watch article from the respondents who fail to detect the true source and perceive the article as from the hosting media outlet and compare them with respondents read the China Watch article on the *China Daily*. Based on Model 3, when the China Watch article is perceived as from the *Washington Post*, its hosting media outlet, it is likely to be considered as more persuasive than the same article on the *China Daily*. Based on Model 4, when the China Watch article is perceived as from *The Telegraph*, it is considered to be much more persuasive than the same article on *China Daily*. Both Model 3 and Model 4 suggest that When a China Watch article is perceived as from *China Daily*, it is not more persuasive than the same article on *China Daily* and might even be less persuasive than the same article on the *China Daily*. However, this negative effect is not statistically significant. Overall, the results from Table 7 are consistent with the findings in Table 5 of the main article.

bachelor's degree, income as between $50,000 to $59,000, trust in media as 25, and media literacy as 54.

Appendix B: Content of China Watch article on the *Telegraph*

By **Wang Yanfei and Zhong Nan**

China will continue to open up to foreign investment, especially in high-end manufacturing, services and green industries, to help transform its economy

China will become more accessible to foreign investors as it looks to achieve high-quality growth, with policy changes in areas such as finance, manufacturing and services. The government will reduce restrictions on market access in some areas, a Ministry of Commerce spokesman, Gao Feng, said last month.

The ministry will guide provinces and cities involved in economic pilot programmes as they overhaul regulations to introduce a unified national negative list due this year. A negative list determines sectors in which foreign participation is prohibited or limited.

Mr Gao said other priorities for the ministry include optimising the regional distribution of foreign investment and deepening institutional innovation in pilot free-trade zones. These efforts correspond with the goals of reform and opening-up set out by the Central Economic Work Conference that concluded on 20 December 2017.

"Granting more market access is expected to bring great business opportunities for foreign firms eyeing the evolving Chinese economy"

Under the new economic concept set out by Chinese President Xi Jinping, opening-up will be "further expanded in scope and level" and the "concept, structure and layout and system of opening up" will be further expanded, the statement said. The 19th CPC National Congress in October also emphasised that China will open up more.

To achieve its economic goals, the country must stick to Mr Xi's new economic thought, the statement said. "The new economic thought has gradually formed over the past five years," said Liu Shengjun, an economist at the Financial Reform Institute of China. "Its introduction means China will give priority to economic reforms over the coming five years."

China's pledge to further promote opening-up follows another round of increased foreign direct investment in the last quarter, as it continued to shift its focus to attract such investment in high-end manufacturing, services and green industries to transform its economy.

As China pushes for growth relying more on domestic consumption and less on exports, sectors that can help the country achieve high-quality growth are expected to benefit from efforts to open up further, experts said.

Feng Yaoxiang, a spokesman for the China Council for the Promotion of International Trade, said companies from developed markets have noticed the changing trend and have become more willing to invest in China-based research and development, as well as its science and technology and design sectors.

Granting more market access is expected to bring great business opportunities for foreign firms eyeing the evolving Chinese economy. Gordon French, HSBC's head of global banking and markets for the Asia-Pacific, said greater access helps the company to better deploy its strength.

Jing Shuiyu contributed to this story.

***Read more news and features at chinadaily.com.cn** (http://web.archive.org/web/20180131203414/http://europe.chinadaily.com.cn/)*

Appendix C: Content of China Watch article on the *Washington Post*

Capital market to open further
By WANG YU and LI XIANG | Published February 12, 2018

China will accelerate the opening of its capital market to foreign asset managers and the market will likely see a surge in new investment products offered by foreign funds in the first quarter of this year, according to a senior official of the Asset Management Association of China.

Recent stock market volatility and intensified regulation to curb financial risks will not derail China's effort to reform and open its capital markets, and foreign fund managers have expressed optimism regarding the Chinese market's long-term prospects, the AMAC official said.

"They will come with their mature investment strategies and business models, which will bring healthy competition and benefit the long-term development of the domestic sector," the official said.

The growing affluence of Chinese investors and their growing demand for increasingly sophisticated investment products have attracted many global investment firms to China's multi-trillion dollar asset management market.

So far, 10 foreign asset managers including Fidelity International, BlackRock, UBS, Man Group, and Schroders are qualified to sell onshore funds to Chinese clients, according to the AMAC.

Despite last week's 9.6 percent fall in the Chinese mainland's A-share market, Lynda Zhou, a portfolio manager at Fidelity International, said a greater opening of China's financial market was likely this year as foreign participation helps it become increasingly mature.

"My long-term view on China is not changed. … The recent market correction is a healthy one as the valuation of blue chips was very expensive. When value re-emerges, I believe foreign investors will come in again," Zhou said.

China opened its asset management industry to international players in 2016 by allowing them to set up wholly foreign-owned enterprises to raise onshore funds and invest in the mainland's securities market. Previously, they could only enter the market by owning a minority stake in a joint venture with a local Chinese partner.

Foreign fund managers have expressed great interest in raising onshore funds to invest in the Chinese interbank bond market, which is currently unavailable to them, as they see fixed-income products as an important asset class in their portfolios and a necessary tool to hedge against risks in the equities market for their clients.

Many of them have also been seeking regulatory approval to raise onshore funds and invest in overseas markets with a certain quota under the pilot program known as the Qualified Domestic Limited Partnership.

"The QDLP program has allowed us to further understand the risk preference of Chinese investors and the existing regulatory environment in China. We will continue to use our global advantage to create value for our Chinese clients," BlackRock said in a written response to China Daily.

China suspended the program in 2015 as the country faced mounting pressure from capital outflows. But the renewed foreign interest indicates that the authorities may lift the suspension in the near future.

The recent market correction is not expected to detract from China's effort to reform its capital market or boost the attractiveness of A-shares to domestic and foreign investors.

The Shenzhen Stock Exchange said on Friday in its 2018-20 development plan that it will reform the current listing requirements for its startup board and will make the exchange more inclusive for high-tech and innovative companies.

Contact the writer at lixiang@chinadaily.com.cn

Appendix D: Additional Information on the Media Variables

In the main text, we briefly described the 'media' variables we used to capture the level of respondent media experience, media trust, and media literacy. In this online appendix, we provide more detailed information about each of these variables.

Items for *Media Experience*

Following Flanagin and Metzger (2000), we use five survey items to calculate respondent media experience. Each item is measured on a 5 point scale. *Media Experience* is calculated as the sum of the scores for each of the five survey items. The exact wording of the survey items is:

1. How many hours per day on average did you read/watch/listen to news in the last seven days?
   - 0 hour
   - Less than 30 minutes
   - 30 -60 minutes
   - 1-2 hours
   - More than 2 hours
2. How would you rate your experience of consuming news?
   - No experience
   - Below average experience
   - Average experience
   - Above average experience
   - A great deal of experience
3. How would you describe your expertise of consuming news?
   - No expertise of consuming news
   - Low expertise of consuming news
   - Average expertise of consuming news
   - High expertise of consuming news
   - Complete expertise of consuming news
4. How would you evaluate your level of familiarity with the news?
   - No familiarity
   - Low familiarity with the news
   - Average familiarity with the news

- High familiarity with the news
- Complete familiarity with the news

5. How easy is it for you to access the news?
    - Extremely hard to access the news
    - Hard to access the news
    - Easy to access the news
    - Very easy to access the news
    - Extremely easy to access the news

Items for *Media Trust*

Following Tsfati and Cappella (2003), we use seven survey items to measure a respondent's general trust in media. Each item is measured on a 5-point scale. The variable *Media Trust* is calculated as the sum of the score for each of the seven survey items. The exact wording of the items is as follows:

1. How much of the time do you think you can trust media organizations to report the news fairly?
    - Never
    - Sometimes
    - About half the time
    - Most of the time
    - Always
2. The news media helps society to solve its problem.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
3. The news media should care about being the first to report a story.
    - Totally agree
    - Somewhat agree
    - Neither agree nor disagree

- Somewhat disagree
- Totally disagree

4. The news media should care about being accurate in reporting a story.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
5. The news media should care about public interests.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
6. The news media should care about the government's interests.
    - Totally agree
    - Somewhat agree
    - Neither agree nor disagree
    - Somewhat disagree
    - Totally disagree
7. The news media should care about the audience's interests.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree

Items for *Media Literacy*

Using the *Media Literacy* scale developed by Ashley, Maksl and Craft (2013), we evaluate a respondent's media literacy with 14 survey items that are each measured on a 5-point scale. The variable *Media Literacy* is calculated as the sum of the scores for each of these items. Its possible values run from 14 to 70, with higher numbers indicating greater media literacy. The exact wording of the items is as follows:

1. News media choose stories based on what will attract the largest audience.
    - Totally agree
    - Somewhat agree
    - Neither agree nor disagree
    - Somewhat disagree
    - Totally disagree
2. Individuals can find news sources that reflect their own political values.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
3. People pay more attention to news sources that fit with their beliefs than news sources that do not.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
4. People might interpret the same story differently.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree

5. People are influenced by news whether they realize it or not.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
6. News makes things seem more dramatic than they really are.
    - Totally agree
    - Somewhat agree
    - Neither agree nor disagree
    - Somewhat disagree
    - Totally disagree
7. News is designed to attract the audience's attention.
    - Totally agree
    - Somewhat agree
    - Neither agree nor disagree
    - Somewhat disagree
    - Totally disagree
8. A journalist's first obligation should be to the truth.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
9. Lighting is used to make certain people in the news look good or bad.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree

- Totally agree

10. Production techniques can be used to influence a viewer's perception.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
11. When taking pictures, photographers decide what is most important.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
12. A story about conflict is more likely to be featured prominently than other stories.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
13. News coverage of a politician candidate will influence people's opinions of him/her.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree
14. A news story that has good pictures is more likely to show up than other news stories.
    - Totally disagree
    - Somewhat disagree
    - Neither agree nor disagree
    - Somewhat agree
    - Totally agree

References

Ashley, Seth, Adam Maksl and Stephanie Craft. 2013. "Developing a news media literacy scale." *Journalism & Mass Communication Educator* 68(1):7–21.

Flanagin, Andrew and Miriam Metzger. 2000. "Perceptions of Internet information credibility." *Journalism & Mass Communication Quarterly* 77(3):515–540.

Tsfati, Yariv and Joseph N Cappella. 2003. "Do people watch what they do not trust? Exploring the association between news media skepticism and exposure." *Communication Research* 30(5):504–529.