Poison Parasite Counter: Turning Frequently-Encountered Duplicitous Mass Communications into Self-Negating Memory Retrieval Cues

Robert Cialdini[1], Jessica Lasky-Fink[2*], Linda J. Demaine[1], Daniel W. Barrett[3], Brad J. Sagarin[4], and Todd Rogers[5]

Abstract

We live in a time beset by communicators of false information. Asymmetrical reach in mass communications due to disparities in communicators' resources or power can lead to imbalances in who is heard, regardless of the soundness or veracity of their messages. The Poison Parasite Counter (PPC) involves inserting a strong (poisonous) counter-message, just once, into a close replica of a rival's communication. In parasitic fashion, the rival's communication then "hosts" the poisonous counter-message, which is recalled upon each new exposure to the actual rival communication. The strategy harnesses associative memory to turn the rival's communication into a retrieval cue for a poisonous counter-message embedded within the rival's message. In seven randomized tests (N = 3,678), we show that the PPC lastingly undermines a duplicitous rival's communication by influencing judgments of communicator honesty, as well as changing behavior in the form of real political donations to actual candidates.

Acknowledgments: The underlying idea for the PPC procedure is related to previous thinking by one coauthor (Cialdini) and several colleagues (L. J. Demaine, D. W. Barrett, B. J. Sagarin, P. Petrova, K. v. L. Rhoads, J., Maner, and N. J. Goldstein) who collected presently unpublished data on the topic. We thank Sam Madison, Tanya Kent, Steve Masnjak, Jennifer Murphy, Anna Valuev, and Alexa Weiss for assistance conducting this research. We thank Peter Clark, Maurice Schweitzer, Emma Levine, and Michael Norton for feedback on the research. **Funding:** We thank Harvard Kennedy School, the Foundations of Human Behavior initiative at Harvard, and Robert Cialdini for supporting this research. **Author contributions:** RBC conceived the initial idea along with LJD, DWB, BJS and PP. TR, RBC, and JLF developed idea further and designed the studies; JLF implemented and analyzed the studies; JLF, TR, and RBC wrote the paper. **Competing interests:** The authors declare no conflict of interest. **Data and materials availability:** The materials and data for all studies are available on OSF (https://osf.io/aps8h/).

[1]Arizona State University. [2]University of California, Berkeley. [3]Western Connecticut State University. [4]Northern Illinois University.[5]Harvard Kennedy School, Cambridge, MA, USA. ***Corresponding author. Email: jlaskyfink@berkeley.edu**

Two months before what many saw as the most consequential U.S. election of our time, Facebook announced it would ban any new political advertising in the week prior to Election Day, so as to reduce the corrosive impact of disinformation/misinformation on the electoral process (New steps to protect the US elections, 2020). Although various commentators concur that deceptive influence attempts are a societal scourge to be curbed (Lazer et al., 2018; Steinmetz, 2018; Truth vs. Lies; 2020) and that social media sites are of particular concern because false news is diffused more rapidly and extensively than truthful information online (Vosoughi, Roy, & Aral, 2018), critics wondered why Facebook's restriction applied only to new political ads. After all, existing political ads are likely to contain the same mischaracterizations and duplicities as new ones. Why, then, would existing ads but not new ones be permitted in the week before the election? In a follow-on post, Facebook's CEO Mark Zuckerberg explained: Only new political ads placed shortly ahead of Election Day would not allow rivals enough time to counter any false information with properly produced correctives (Zuckerberg, 2020).

Although the explanation highlights the crucial need for interventions to reduce the harmful consequences of misleading communications, most types of such correctives—disclaimers, disclosures, forewarnings, fact checks, retractions—have produced mixed success at best, including worrisome backfire effects (Argo & Main, 2004; Jang, Lee, & Shin, 2020; Johar & Simmons, 2000; Lewandowsky et al., 2020; Schwartz et al., 2007; Schwartz, Newman, & Leach, 2016; Skurnik et al., 2005; Wahlheim, Alexander, & Peske, 2020). One exception, however, involves the presentation of strong counterarguments, which have proven reliable, effective, and robust across a variety of content domains (Eagly et al., 2000; Maaravi, Ganzach, & Pazy, 2011; Petrova & Cialdini, 2011; Petty & Brinol, 2010; Siegel et al., 2008; Wood & Quinn, 2001). But, there is a problem with even sound counterarguments: their diminished salience and availability at the time of recipients' next exposure to the rival communication. If interposed information interferes with the memory of the counterargument or if the passage of time weakens that memory before a reintroduction of the rival message, its fitness as a corrective will fall (Koriat, Goldsmith, & Pansky, 2000).

For traditional forms of counterargument delivery, these conditions for effectiveness are onerous. They require that, optimally, counter-argumentation occur *immediately* after receipt of *every* rival message—so the challenging position will be as strong, salient, and accessible as the rival position each time the rival message is re-encountered. Such requirements pose a pair of daunting practical problems. First, as the Facebook ban recognized, communicators in competing political advertising campaigns are rarely able to produce and issue well-crafted refutational messages immediately. Second, within many large-scale persuasive undertakings such as commercial or political advertising, some communicators are better equipped to transmit their messages. When, often by virtue of superior budgetary resources, one adversary possesses the capacity to utilize the channels of communication more frequently, it is simply not possible to expose a desired audience to counterclaims in every instance (Koerth-Baker, 2018; OpenSecrets, 2018). This asymmetry in ability to expose audiences to one's arguments constitutes a serious difficulty, as recipients' belief in the truth of an argument increases with the number of exposures to it (Fazio & Sherry, 2020; Unkelbach et al., 2019).

The Poison Parasite Counter (PPC)

To address these challenges, we propose the Poison Parasite Counter technique that combines dual components. The first is parasitic, which operates by attaching one's

counterarguments to each presentation of a rival communication. Our examination of the literature of memory science uncovered a pair of promising candidates for achieving this end via mechanisms that prompt people to recall one set of previously encountered arguments while experiencing another: 1) Endel Tulving's (1983) "encoding specificity" mechanism, and 2) the use of retrieval cues as mnemonic devices. According to the encoding specificity principle, retrieval cues increase the probability a given memory will be recalled; and, the best retrieval cues are those stimuli that were present when the memory was formed (Rogers & Milkman, 2016 Tulving & Schacter, 1990; Tulving & Thomson, 1973). Much research demonstrates that reinstating the retrieval cues that were present during encoding greatly facilitates recall (see Petrova & Cialdini, 2011, and Roediger, Tekin, & Uner, 2017, for reviews).

There is an upshot of these findings for overcoming the problem of differential salience of competing arguments: Install retrieval cues in the rival message that call the counterarguments to consciousness during re-exposure to the rival message. This should cause one's points to be recalled (by virtue of the common retrieval cues) whenever audience members experience the opponent's message. Not only should this strategy solve the problem of the differential salience that naturally occurs when the competitor's arguments are presented at a later point, it should also solve the problem of differential access to audience attention when one's adversary can reach the audience more often. That is, if we have arranged for our own points to be raised to consciousness in audience members' minds each time an adversary raises his or her points, then the playing field of presentation opportunities will have been leveled. Our rival's advertising budget will have become shared by us.

The second functional component of the PPC is its poisonous character. Even if we have arranged for our counterarguments to live parasitically in the body of a rival message, they will prove ineffective unless they are strong (poisonous) enough to damage that message. Within the realm of misleading communications, certain counter-communications have proven deeply damaging. They are those that undermine the perceived honesty and rectitude, rather than the just knowledge or wisdom, of the rival communicator—as there appears to be an evolved human sensitivity and intense aversion to someone who appears deceptive (Cosmides & Tooby, 1992; Fehr & Gachter, 2002; Trafimow, 2001). Once a ruse is recognized or revealed, its targets strongly resist information associated with it and its perpetrator (Darke & Ritchie, 2007; Farrelly et al., 2005; Sagarin et al., 2002).

Thus, the combination of poisonous and parasitic elements of the PPC should be an effective counter to communications containing deceptive information, even when the source of the communications is advantaged asymmetrically in the ability to expose audiences to the dishonest messages. In seven randomized studies, we tested this possibility in the realm of political candidate advertising and commercial product advertising, two areas where asymmetric reach in mass communications is common and consequential (Lowery, 2014; Gerber, 1998; Gerber, 2004). Using both static and video ads, we showed that the PPC can be executed by inserting strong (poisonous) counter-messages economically—just once—into a close replica of a rival's communication, causing the counter-messages to then live (parasitically) in that communication. This parasitic element was designed to produce recurring, retrieval-cue-induced recall of the poisonous counter-messages whenever the rival's communication is later encountered, making the counter-messages resistant to normal memory degrading processes (Wixted, 2004). Additionally, we focused on the ability of the PPC to reduce the benefits of

asymmetric reach in its most troubling form, when the claims of a disproportionately advantaged communicator are deceitful.

Study 1: The Poison Parasite Counter Effect

Study 1 sought to test whether inserting strong counter-messages into a close replica of a deceptive rival political ad subsequently undermines the rival communication relative to presenting the same counter-messages in a traditional, visually independent fashion.[1]

Method

Participants. Participants were recruited through Amazon's Mechanical Turk (MTurk) to complete a 15-minute online survey for which they were paid $1.00 each. Only MTurk workers who were located in the United States, who had an MTurk approval rating of at least 95%, who had not participated in previous similar studies, and who passed an initial attention check were eligible to complete the survey. A total of 297 participants (mean age = 38 years, *SD* = 11.9; 50% female) completed the study. The necessary sample size was decided ex ante based on expected effect sizes from pilot testing.

Procedure. All workers who consented to participate and passed the attention check were randomly assigned by the survey platform to one of two experimental conditions: a Traditional Response condition or the Poison Parasite Counter (PPC) condition. Participants were not aware of their condition assignment.

In the approximately 15-minute long survey, participants saw a total of 10 political ads for five fictional candidates—one "pro" and one "response" ad for each candidate. For one of these candidates, Walter McKinley, we developed two response ads corresponding with each experimental condition. The first utilized the PPC procedure by overlaying counter-messages highlighting the duplicitous nature of the rival message on an exact visual replica of the rival ad (see Fig. 1B). The Traditional Response ad presented the same counter-messages as the PPC ad, but with a different visual aesthetic that provided no associative links to the rival's original ad (see Fig. 1C).

The fictional ads were interspersed with four excerpts from current, non-partisan news articles. Each ad or excerpt was shown on a separate page, and participants had to wait at least five seconds to advance from one page to the next. All participants saw the rival Walter McKinley ad on the second screen, and either the PPC or Traditional Response ad on the eighth. Filler questions were asked after every excerpt and ad, designed to distract participants from the focus of the study and interfere with memory processes.

[1] The original conceptualization for the PPC and, in particular, for this study was first developed by a team that included four authors of this article (R.B.C, L.J.D, D.W.B, and B.J.S.) as well as several others who collected presently unpublished data on the topic (Cialdini, Demaine, Barrett, Sagarin, Petrova, Rhoads, & Maner, 2012).

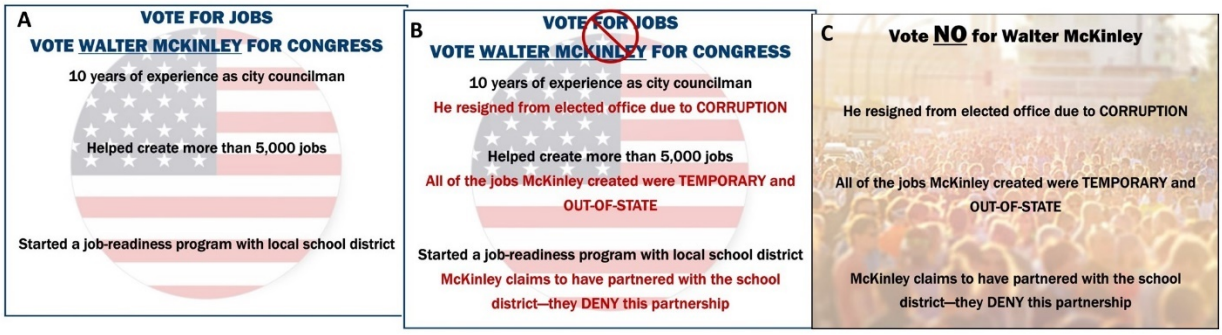

Fig. 1. Walter McKinley ads used in Studies 1 and 2. (A) Rival ad; (B) PPC ad; (C) Traditional Response ad

At the end of the survey, participants were told that they would be shown one of the candidate ads from the first section, chosen arbitrarily. All participants were again shown the rival McKinley ad and were subsequently asked:

1. If Walter McKinley was running in your state, how likely is it that you would vote for him?
2. How honest do you think Walter McKinley is?

Both outcomes were ranked on a five-point scale, from extremely unlikely to extremely likely, and extremely dishonest to extremely honest, respectively.

Analysis and Results

We hypothesized that by directly challenging the claims in the rival ad and creating mnemonic links between the response ad and the rival ad, the PPC procedure would result in recurring, cue-based recall of the counter-messages whenever the rival's ad was re-encountered, making the counterargument resistant to normal memory degrading processes. This would, in turn, lead participants to view McKinley as more dishonest and reduce their willingness to vote for him compared to a traditionally presented response ad.

Using a standard linear model, we regressed each of our primary outcomes on an indicator for treatment assignment, as well as participants' gender, age, party affiliation (Democrat, Republican, Independent, or Other), and an indicator for whether they were college-educated. To facilitate comparability across studies and outcomes, we also standardize each of our outcome measures.

As predicted, exposure to the PPC ad subsequently reduced participants' reported likelihood of voting for McKinley by 0.61 SD ($SE = 0.11$, $p < .001$; Fig. 2) relative to participants who had seen the Traditional Response ad. This is equivalent to a 0.62 point reduction in likelihood of voting on a 5-point scale (see SOM). In a similar pattern, participants who saw the PPC ad rated McKinley as 0.70 SD ($SE = 0.11$, $p < .001$) or 0.59 points ($SE = 0.09$, $p < .001$) less honest than participants in the Traditional Response condition.

We also conducted a mediation analysis to examine the extent to which perceptions of a candidate's dishonesty mediates the relationship between the PPC procedure and voting

preference. As shown in the SOM, perception of McKinley's honesty mediates approximately 46% percent of the total effect of assignment to the PPC condition on voting preference (z = 33.25, $p < .001$). This suggests that, as hypothesized, emphasizing a rival candidate's dishonesty and duplicity via the PPC procedure may produce an especially potent "poison" to undermine the rival's communication.

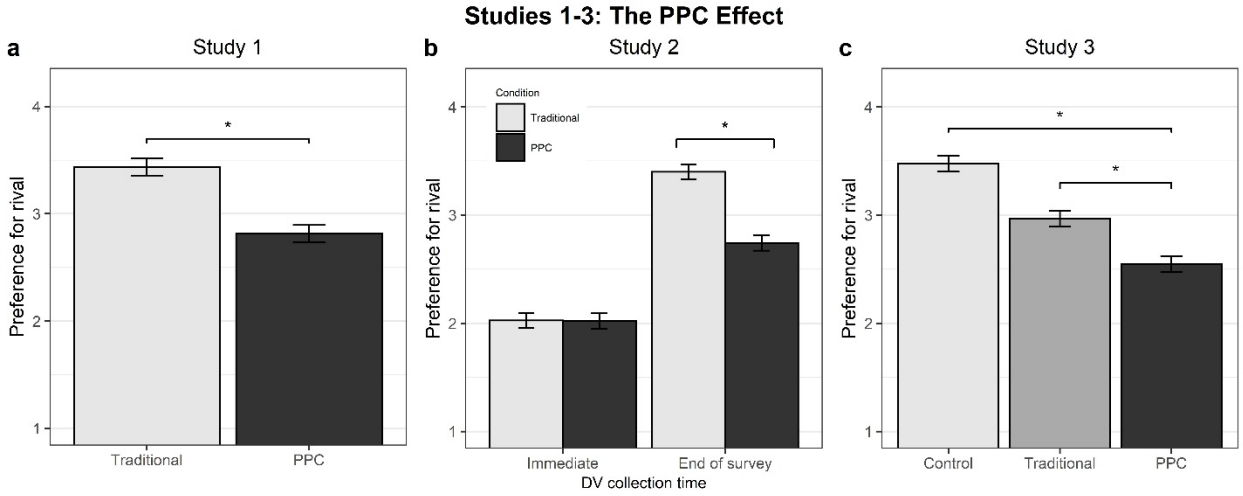

Fig. 2. Results from **(a)** Study 1; **(b)** Study 2; and **(c)** Study 3. Outcomes reflect regression-adjusted unstandardized means, measured on 1-5 scale where 5 represents "extremely likely to vote." Error bars reflect ±1 standard error * reflects differences significant at $p < .01$.

Study 2: PPC Mechanism

The PPC and Traditional Response ads in Study 1 had identical content and were from the same source, which ensures that any difference between conditions was not the result of source sleeper effects (Kumkale & Albarracín, 2004). Yet, it is possible that the effects were only driven by differential ad quality. If the PPC ad was simply a superior ad, we would expect to see that superiority over the Traditional Response ad immediately after raters had seen one or the other ad. Instead, if as hypothesized, the superiority of the PPC is due to reduced decay of its impact on the rival communication, its negative effect on participants' perceptions should appear only upon participants' re-exposure to the rival communication—after they had seen interference-producing filler and decoy materials. Study 2 tests this directly using the same materials as Study 1 in a pre-registered factorial design.

Method

Participants. We recruited a total of 713 MTurk workers (mean age = 36 years, *SD* = 10.8; 52% female) to complete a 15-minute online survey for which they were each paid $1.20. Only MTurk workers who were located in the United States, who had an MTurk approval rating of at least 95%, who had not participated in previous similar studies, and who passed an initial attention check were eligible to complete the survey. The necessary sample size was decided ex ante based on expected effect sizes from pilot testing.

Procedure. The content and outcomes for this study were identical to those used in Study 1 (Fig. 1), with Walter McKinley as the focal candidate. In a 2 x 2 factorial design, all participants who consented to participate and passed the attention check were first randomly assigned to either the PPC or Traditional Response condition, and were then subsequently randomly assigned to either an Immediate Outcome or End-of-survey Outcome condition. Participants were not aware of their condition assignment.

The study procedures were identical to Study 1, with the exception of the timing of the outcome questions. In the Immediate Outcome condition, participants answered the two dependent variable questions immediately after seeing the PPC or Traditional Response ad. In the End-of-survey Outcome condition, participants answered the two dependent variable questions at the end of the 15-minute long survey, after being re-exposed to the rival McKinley ad. In both conditions, the dependent variable questions were the same as in Study 1.

In order to keep the survey the same length for participants in all experimental conditions, those assigned to the End-of-survey Outcome condition were also asked two questions—one recall, and one filler—immediately after viewing the PPC or Traditional Response ad:

1. [Recall] Which of the following was NOT mentioned on the previous ad? *[School district corruption; raised taxes; out-of-state jobs]*
2. [Filler] Please rank the following policy areas in order of importance to you, with 1 being the most important and 5 being the least important. *[Immigration; national security; military spending; national debt; international diplomacy]*

Analysis and Results

This study was pre-registered on OSF (osf.io/aps8h), and analyses procedures followed those described in Study 1. If the PPC and Traditional Response ads are of similar quality, participants' reactions to McKinley should be the same across both conditions when assessed immediately after viewing the counter ad. As expected, when measured immediately after viewing the PPC ad or Traditional Response ad, participants' likelihood of voting for McKinley was equivalent across conditions ($F(1,348) = 0.00$, $p = 0.95$), as was perceived honesty ($F(1,348) = 1.28$, $p = 0.26$). Additionally, immediate recall was also equivalent across conditions: when participants in the End-of-survey Outcome condition were to identify the counterclaims they had just seen immediately after viewing the PPC or Traditional Response ad, 72% of participants in both the PPC condition and the Traditional Response condition were able to correctly do so ($\chi^2(1) = 0.02$, $p = 0.89$). Thus, at the time of exposure, the two ads spurred equal recall and produced equally strong counter-responses to the rival ad, making it implausible that the PPC ad was merely a superior ad or that the qualities of the PPC procedure inherently facilitate a more fluid comparison with the rival material.

When participants were asked to rate McKinley's honesty and their likelihood of voting after being re-exposed to the rival McKinley ad at the end of the survey—after interference-producing filler and decoy materials—we find results that parallel those of Study 1 (Fig. 2). As would be expected given evidence of quick decay of persuasive effects and the memory interference effects of interpolated materials, overall favorability of McKinley was higher among participants in both conditions at the end of the survey compared to immediately after exposure

to the treatment ad. Yet, as we hypothesized, the PPC procedure mitigated this decay: upon re-exposure to the rival's communication, participants in the PPC condition who responded to the outcome measures at the end of the survey were significantly less likely to vote for McKinley (0.63 SD, $SE = 0.10$, $p < .001$) and rated him as significantly less honest (0.41 SD, $SE = 0.10$, $p < .001$) than those who saw the Traditional Response ad.

Study 3: PPC Versus No Counter-Message Control

Studies 1 and 2 utilized PPC and Traditional Response ads that offered identically parallel content (Fig. 1); only the visual presentation differed between the two ads. Study 3 tested the effect of the PPC ad against a more externally valid Traditional Response ad, which offered the same counterarguments as the PPC ad, but presented in a more realistic narrative form. Study 3 also added a control group, in which participants saw an ad for an entirely different fictional candidate, in order to compare the effects of the PPC and Traditional Response ads relative to no counter-message at all.

Method

Participants. Participants were recruited through MTurk to complete a 10-minute online survey for which they were paid $.60 each. Only MTurk workers who were located in the United States, who had an MTurk approval rating of at least 95%, who had not participated in previous similar studies, and who passed an initial attention check were eligible to complete the survey. A total of 602 participants (mean age = 40 years, $SD = 12.4$; 58% female) completed the study. The necessary sample size was decided ex ante based on expected effect sizes from pilot testing.

Procedure. All workers who consented to participate and passed the attention check were randomly assigned to one of three conditions: Control, Traditional Response, or PPC. Participants were not aware of their condition assignment.

In Studies 1 and 2, the counter-messages presented in the Traditional Response ad were identical to the messages presented in PPC ad (see Fig. 1). In reality, response ads often encompass cohesive narratives rather than individual and fragmented counterclaims. As such, Study 3 tested the effect of the PPC procedure against a more externally valid Traditional Response ad, as shown in Fig. 3. While the Traditional Response ad offered the same counterarguments as the PPC ad, it was presented in a more realistic form. In addition, we added a pure control group, in which participants saw an ad for an entirely different fictional candidate, in order to compare the effect sizes of the PPC and Traditional Response ads.

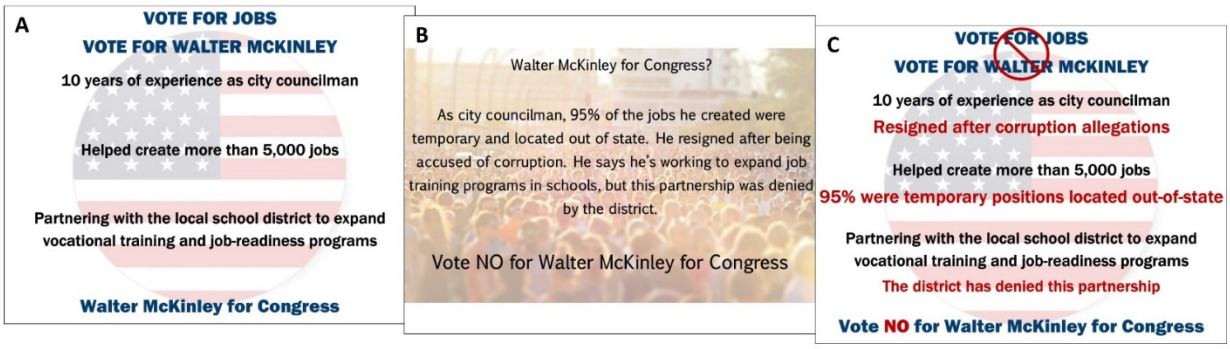

Fig. 3. Walter McKinley ads used in Studies 3, 4, and 7, as well as Supplement Studies I, II, and III. (A) Rival ad; (B) Traditional Response ad; (C) PPC ad.

In the first section of the survey, all participants were shown a series of 10 political ads for five fictional candidates, with Walter McKinley again acting as the focal candidate. All participants saw the rival McKinley ad on the second screen, and the PPC, Traditional Response, or Control ad fourth screen. The other eight ads were presented in random order. After viewing the 10 political ads, participants were asked a set of five filler questions designed to distract them from the focus of the study and interfere with memory processes. Then, in the third and final section of the survey, participants were told that they would be shown one of the candidate ads from the first section, chosen arbitrarily. All participants were shown the rival McKinley ad and were subsequently asked the same dependent variable questions as in Studies 1 and 2.

Analysis and Results

This study was pre-registered on OSF (osf.io/aps8h), and analysis procedures followed those described in Study 1. As shown in Fig. 2, exposure to the PPC ad significantly reduced participants' subsequent likelihood of voting for McKinley relative to the Control ad (0.85 SD, $SE = 0.10$, $p < .001$), and relative to the Traditional Response ad (0.38 SD, $SE = 0.09$, $p < .001$). In a similar pattern, participants who saw the PPC ad also rated McKinley as significantly less honest than participants in both the Control condition (0.68 SD, $SE = 0.10$, $p < .001$), as well as in the Traditional Response condition (0.30 SD, $SE = 0.10$, $p = .002$). As in Study 1, honesty mediated approximately 46% percent of the total effect of condition assignment on voting preference ($z = 28.77$, $p < .001$; see SOM).

Relative to the Control condition, the Traditional Response ad also reduced participants' likelihood of voting for McKinley (0.46 SD, $SE = 0.09$, $p < .001$), and perceptions of his honesty (0.38 SD, $SE = 0.09$, $p < .001$), although the effect sizes are significantly smaller than those produced by the PPC ad.

Discussion

Although the Traditional Response ad also reduced participants' likelihood of voting for McKinley and their perceptions of his honesty compared to the control, it did so to a significantly lesser extent than the PPC ad. Together, the results of Studies 1, 2, and 3

demonstrate that the PPC procedure slows the decay of persuasive effects over a short period of time. In Supplement Studies I and II, we replicated these results with variations in the frequency of exposure to both the rival McKinley ad and the Traditional Response ad.

Study 4: PPC Durability Over Time

Study 4 tests whether the effect of the PPC procedure is durable over the passage time and in the presence of more significant memory interference.

Method

Participants. We recruited an initial sample of 557 participants (mean age = 37 years, *SD* = 11.7; 67% female) via Amazon's MTurk. The sample was again restricted to participants who were located in the United States, had an MTurk approval rating of at least 95%, and passed an initial attention check. All participants who had taken previous surveys as part of this study were excluded. Participants were compensated $1.10 for the each of the first and second surveys; $0.70 each for the third and fourth surveys; and $1.00 for the fifth survey. Additionally, all participants who completed the first four surveys received a $0.75 bonus as an incentive to return for the fifth and final wave.

Of the 557 participants who completed the first day of the study, 330 completed all five days (mean age = 39 years, *SD* = 12.6; 69% female). Attrition was balanced evenly across conditions ($\chi^2(2) = 0.80$, $p = 0.67$; see SOM). Younger participants, participants whose party affiliation was Independent, and non-college educated participants were all more likely to attrit. A joint significance test shows that we cannot reject the null hypothesis that attrition was balanced across conditions and all covariates ($\chi^2(8) = 23.05$, $p = 0.003$). All analyses control for party affiliation, college education, and age, as well as participant gender.

Procedure. All workers who consented to participate and passed the attention check were randomly assigned to one of three conditions: Control, Traditional Response, or PPC. Participants were not aware of their condition assignment. Participants were told in wave 1 (day 1) that the study entailed five separate surveys that would be conducted over a span of two weeks, and were asked to signal their intention to complete all five parts before proceeding to the survey. Each of the four follow-up waves was open for 24 hours to all participants who had completed the preceding survey. Reminder emails were sent to all eligible workers when each follow-up survey opened, as well as at the 12-hour mark during each survey window.

This study used the same materials as Study 3 (Fig. 3), with Walter McKinley again serving as the rival candidate. The study was run over a period of 17 days, with waves conducted on days 1, 3, 6, 9, and 16. Each wave followed the same procedures as previous studies: participants were shown a series of fictional political ads that were interspersed with news article excerpts and filler questions to produce memory interference. To ensure that we did not unintentionally provide associative links between the McKinley ad and the PPC or Traditional Response ads, participants viewed the rival McKinley ad during the first wave on day one, and the PPC, Traditional Response, or control ads during the second wave which was administered on day three. In all other waves, participants were only shown the rival McKinley ad. By the end

of the five wave study, all participants had seen the rival McKinley ad nine times, and the PPC, Traditional Response, or Control ad once.

At the end of every wave, participants were asked the two dependent variable questions about two "arbitrarily" selected candidates—one of whom was always Walter McKinley, while the second subject differed in each wave. As in all earlier studies, page timers were used to ensure that participants spent a minimum amount of time on each article and ad. Overall, the first and second waves were each about 15 minutes long, while the third, fourth, and fifth waves took about 10 minutes each.

Analysis and Results

The results from the second wave, the only wave in which participants saw either the PPC, Traditional Response, or Control ads, parallel those of Study 3. As expected, the PPC procedure had a large and significant effect on the likelihood of voting for and perceived honesty of McKinley relative to both the Traditional Response and the Control ads. In the waves on all later days, the effect of the PPC procedure remained significantly superior to the Traditional Response ad (see Fig. 4, which reports unstandardized means). By the fifth wave (on day 16), participants who had seen the PPC ad during the second wave (on day 3) were still significantly less likely to vote for McKinley (0.74 SD, $SE = 0.13$, $p < .001$), and rated him as significantly less honest (0.43 SD, $SE = 0.13$, $p = .001$) relative to those who had seen the Traditional Response ad (on day 3).

Despite large frequency-of-exposure disparities and time-induced memory interference, the PPC procedure continued to undercut the rival ad throughout the two-week period. In contrast, the initial effectiveness of the Traditional Response ad waned relative to the Control condition, in keeping with the characteristic decay of political ad effectiveness (Gerber et al., 2011; Mutz & Reeves, 2005; Kalla & Broockman, 2018) over time: in waves three (on day 6), four (on day 9), and five (on day 16), there was no significant difference in likelihood of voting or perceived honesty for McKinley between participants who saw the Control ad and those who saw the Traditional Response ad during wave two (on day 3; see SOM). These findings were replicated in Supplement Study III.

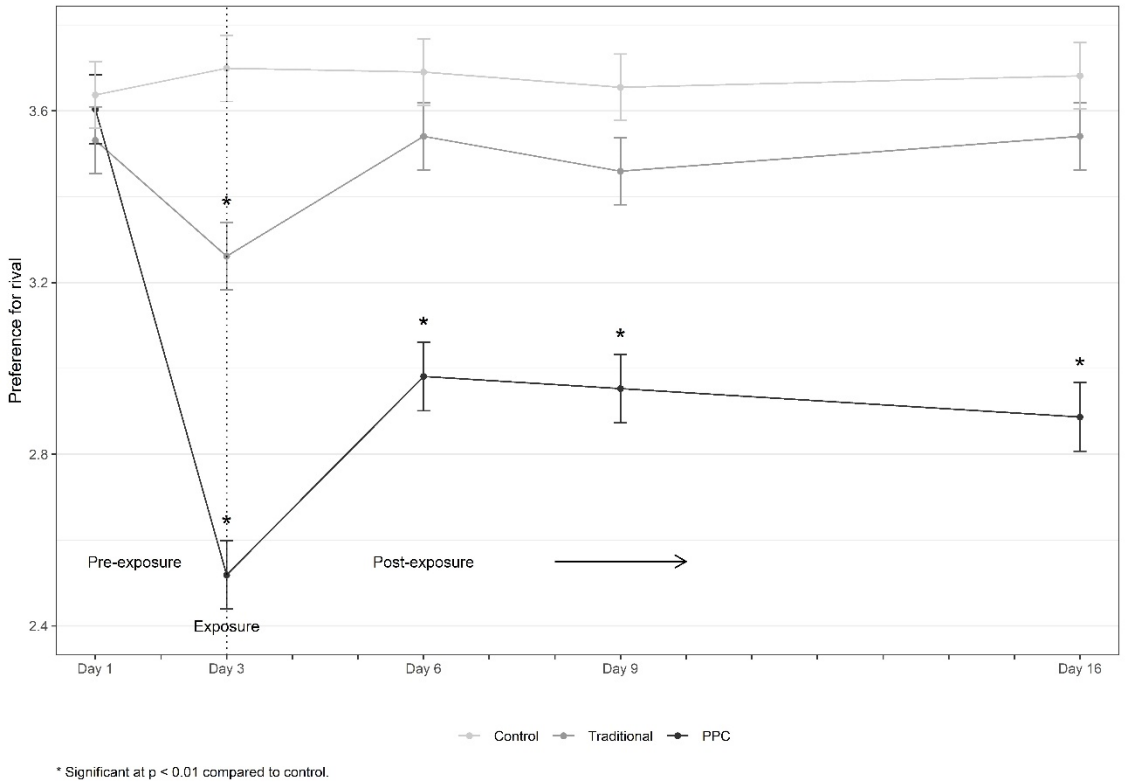

Fig. 4. Preference for rival each day in study 4. All participants saw Control, Traditional Response, or PPC ad a single time on day 3. In all other waves (on days 6, 9, and 16), participants saw the rival ad interspersed with decoy ads and filler material. Outcomes reflect regression-adjusted unstandardized means, measured on 1-5 scale where 5 represents "extremely likely to vote." Error bars reflect ±1 standard error. *p < .01 relative to Control condition.

Study 5: PPC Effect on Real-World Political Donations

Our first four studies provide strong evidence of the efficacy and durability of the PPC procedure in a hypothetical setting—a contest between fictional political candidates. It remains possible that the PPC is ineffective when the focus is a real political candidate and election for which opinions may be less malleable. Study 5 aimed to demonstrate that the PPC procedure can affect a consequential behavior in a real campaign with real political ads. Relying on the same paradigm as Study 1, we replaced the fictional ads with real political communications from an election that was on-going at the time of the study—the 2018 Democratic gubernatorial primary election in Michigan. Critically, we also added a consequential outcome measure that matters to political campaigns—political donations.

Method

Participants. Participants were 299 Amazon MTurk workers (mean age = 36 years, *SD* = 10.7; 55% female) who received $1.30 compensation for completing the survey. The sample was again restricted to participants who were located in the United States, had an MTurk approval rating of at least 95%, and passed an initial attention check. All participants who had taken previous surveys as part of this study were excluded. Because this study used actual campaign ads from the Michigan Democratic primary election, we also excluded all workers in the state of Michigan per IRB rules so as not to influence any potential voters' opinions prior to the election. Finally,

since our primary outcome measure asked about voting in a Democratic primary election, we also excluded all self-identified Republicans. The necessary sample size was decided ex ante based on expected effect sizes from pilot testing.

Procedure. As the rival ad, we used a real print ad produced by Gretchen Whitmer's campaign, the leading 2018 Democratic gubernatorial candidate in Michigan. The Traditional Response ad was an actual ad produced and circulated during the election by the campaign of Shri Thanedar, one of her main Democratic opponents. Thanedar's ad provided strong and explicit counter-messages, but no associative links to the rival Whitmer ad. We modified this response ad to create two new versions that used the PPC procedure: a "full" and a "tailored" PPC ad (Fig. 5). In the "full" application, we placed the rival Whitmer ad and the Traditional Response ad side-by-side with a line down the middle and the respective headers: "Typical Gretchen Whitmer ad" and "Here's what we say in our ad." In the "tailored" application, we took the exact counter-messages from Thanedar's response ad and embedded them in the rival Whitmer ad. Both ads were created purely for research purposes; neither was actually used in the campaign or circulated to prospective voters. We had no a priori hypothesis as to which manifestation of the PPC would be more effective, and thus tested each relative to the actual (traditional) response ad.

The study was run prior to the primary election on August 6, 2018. The procedures for this study paralleled those of Studies 1 and 2. All workers who consented to participate and passed the attention check were randomly assigned to one of four conditions: Control, Traditional Response, Full PPC, or Tailored PPC. Participants were not aware of their condition assignment.

All participants were told that they would see a series of ads for real political candidates currently running for office, as well as excerpts from actual news articles. All participants first saw the rival Whitmer ad and then, depending on their experimental condition, either one of the PPC ads, the Traditional Response ad, or a Control ad for a different candidate in a different election. The ads were interspersed with eight decoy ads, all of which were real campaign ads from other current political races across the United States, as well as four news article excerpts and related filler questions. The rival Whitmer ad appeared third, and the Full PPC, Tailored PPC, Traditional Response, or Control ad appeared fifth. At the end of the roughly 15-minute survey, all participants were asked:

1. If you lived in Michigan, how likely would you be to vote for Gretchen Whitmer in the upcoming Democratic primary election?
2. How honest do you think Gretchen Whitmer is?
3. You have a chance to allocate real resources. We are donating $0.10 on behalf of every worker who takes our survey. We can either donate this $0.10 to Gretchen Whitmer's campaign or to the campaign of Shri Thanedar, her opponent. Who would you like us to donate this $0.10 to?

In response to question 3, a total of $21.20 was donated to Whitmer's campaign, and $8.70 to Thanedar's campaign. All donations were made by the authors prior to the primary election.

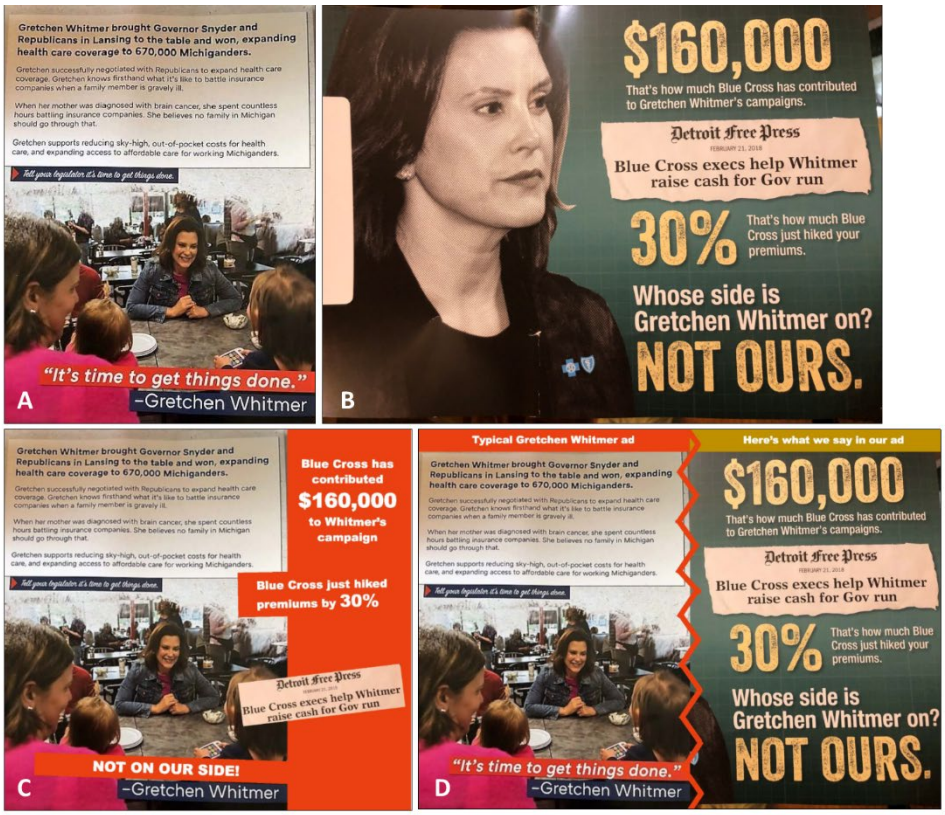

Fig. 5. Gretchen Whitmer ads. (A) Rival ad; (B) Traditional Response ad; (C) Tailored PPC ad; (D) Full PPC ad. Used in Study 5.

Analysis and Results

Consistent with results from the previous studies, participants in both PPC conditions were significantly less likely to express willingness to vote for Whitmer and rated her as significantly less honest than participants in the Traditional Response and Control conditions (Fig. 6). Exposure to the Tailored PPC ad reduced participants' subsequent likelihood of voting for Whitmer by 0.40 SD ($SE = 0.16$, $p = .01$) and 0.50 SD ($SE = 0.16$, $p = .002$) relative to the Traditional Response and Control ads, respectively. Similarly, exposure to the Full PPC ad reduced participants' likelihood of voting by 0.60 SD ($SE = 0.16$, $p < .001$) and 0.71 SD ($SE = 0.16$, $p < .001$) compared to the Traditional Response and Control ads, respectively.

Participants in the Tailored PPC condition rated Whitmer as 0.38 SD ($SE = 0.16$, $p = .02$) less honest than participants in the Traditional Response condition, and 0.43 points ($SE = 0.16$, $p = .008$) less honest than those in the Control conditions. Meanwhile, participants in the Full PPC condition rated Whitmer as 0.62 SD ($SE = 0.16$, $p < .001$) less honest than participants in the Traditional Response condition, and 0.67 points ($SE = 0.16$, $p < .001$) less honest than those in the Control conditions. As in the previous two studies, honesty mediated approximately 44% of the total effect of assignment to the PPC condition on voting preference ($z = 23.88$, $p < .001$; see SOM).

Meanwhile, there was no significant difference in voting likelihood or perceived honesty between participants who saw the Traditional Response ad and those who saw the Control ad (see Fig. 6 and SOM), demonstrating that the mere presence of counterarguments is not necessarily sufficient to produce meaningful resistance to a rival ad. Additionally, there was no significant difference between the two PPC conditions on either of the outcome measures, although the Full PPC ad produced stronger directional results.

After the questions on voting and perceived honesty, participants were told that a $0.10 donation would be made on their behalf to either Gretchen Whitmer or Shri Thanedar, her opponent, and were asked to direct the donation to their preferred candidate. Fifty-seven percent of participants in the Full PPC condition and 61% in the Tailored PPC condition directed the donation to Whitmer's campaign (instead of Thanedar's), compared to 75% in the Traditional Response condition and 90% in the control condition. The percentage of participants who directed the donation to Whitmer was significantly lower in both PPC conditions than in the Traditional Response condition (both $\chi^2(1) > 3.6$, $p < .10$) and the control condition (both $\chi^2(1) > 15.1$, $p < .001$).

Discussion

The relative equivalency of the PPC versions suggests there are multiple ways to implement the PPC procedure effectively. Via either method, not only does the PPC procedure reduce the viability of a rival's message, it also influences related behavior in the form of financial support for that candidate.

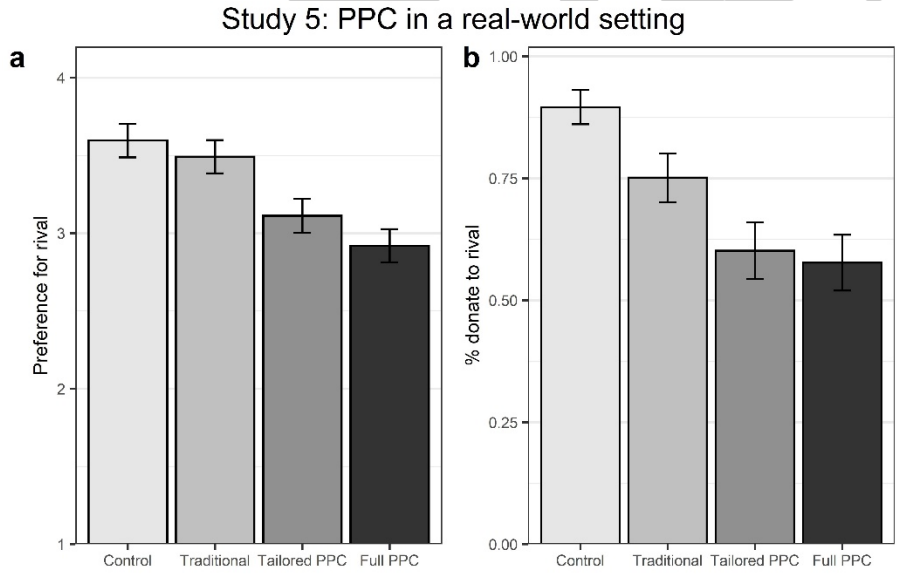

Fig. 6. (a) Preference for rival, and (b) donations to rival in study 5. Outcome for (a) reflects regression-adjusted unstandardized means, measured on 1-5 scale where 5 represents "extremely likely to vote." Error bars reflect ±1 standard error.

Study 6: PPC within Super Bowl Video Ads

Given that video ads comprise a growing share of advertising, including commercial advertising, showing the PPC procedure can be effectively implemented via video and for a nonpolitical, commercial entity would extend its applicability. Study 6 tested this applicability.

Method

Participants. On the Friday before the 2020 Super Bowl, we recruited an initial sample of 2,429 participants (mean age = 41 years, *SD* = 12.8; 50% female) via Amazon's MTurk for a three-wave study that was run over nine days. The sample was again restricted to participants who were located in the United States, had an MTurk approval rating of at least 95%, and passed an initial attention check. All participants who had taken previous surveys as part of this study were excluded.

Of the 2,429 participants who completed the first wave of the study (day 1), 2,152 (89%) completed the second wave (day 3), balanced evenly across conditions ($\chi^2(1) = 1.71$, $p = .43$). Then, anyone who completed the second wave and reported having watched the Super Bowl was invited to complete a third wave exactly one week later (day 9). Of the 2,152 participants who completed the second wave of the study, 1,463 reported watching the Super Bowl between wave 1 and wave 2, and 1,172 (80%) of these participants completed the third wave. As such our final analytic sample is comprised of 1,172 participants who watched the 2020 Super Bowl and completed all three surveys over the nine-day period (mean age = 41 years; *SD* = 12.6; 44% female).

Overall attrition across all three surveys was balanced evenly across conditions ($\chi^2(2) = 0.41$, $p = 0.82$), as was self-reported Super Bowl viewing ($\chi^2(2) = 1.31$, $p = 0.52$). However, across demographics, women were less likely to report having watched the game, while those with higher incomes were more likely to have watched the game. We thus control for both, as well as a range of other covariates, in all analyses. Participants were compensated $1.60 for the first survey; $0.70 each for the second survey; and $1.00 for the third survey.

Procedure. During the 2020 Super Bowl, TurboTax—one of the largest online tax preparation companies—ran a 45-second ad highlighting the simplicity and benefits of their software. They released the ad five days before the game, which afforded a unique opportunity to test the PPC procedure knowing that participants who watched the game would be subsequently re-exposed to the rival's communication. Prior to the game, we developed three response ads. The PPC ad overlaid a counter message on the exact TurboTax ad that was to run during the Super Bowl. This message stated: "TurboTax says they work to make filing taxes easy for us. Yet, they've spent $10 million lobbying lawmakers to prevent free automatic filing. This makes filing harder and more expensive for us, so they can make money." In the PPC ad, this text scrolled twice across the screen over the course of the 45-second video, and then ended with a static screen that displayed this message for an additional three seconds. In the Poison Only ad, the exact same scrolling text was overlaid on a different TurboTax commercial, which was of an equivalent length. And in the Pure Counterargument condition, the same scrolling text was presented with a solid black screen as the background.[2]

[2] The ads created for this study are at: https://bit.ly/poisonparasite.

Participants were told in wave 1 (day 1) that the study would two parts, and were asked to signal their intention to complete both parts before proceeding to the survey. The first wave took place during the 24 hours preceding the Super Bowl. In a design that mimicked that of our previous studies (see SOM), we randomly assigned participants to see the PPC ad, the Poison Only ad, or the Pure Counterargument ad. In a 15-minute survey, we interspersed the assigned treatment ad with the rival TurboTax ad—the one that was to run during the Super Bowl—and eight other television ads for a variety of services and products. The rival TurboTax ad was shown first, followed by the assigned treatment ad approximately four minutes later. At the end of the survey, participants were re-exposed to the rival TurboTax ad and asked a series of four questions to gauge how favorably they viewed TurboTax:

1. How positively or negatively do you view TurboTax?
2. If you were to use an online tax filing service, how likely would you be to use TurboTax?
3. There are many competing tax preparation companies. Imagine they all offer tax filing for the same price. Would you choose to file your taxes through TurboTax or one of its comparable competitors?
4. If a friend asked you for a recommendation on online tax filing services, which company would you be most likely to recommend?

Then, we re-recruited the same sample of participants on the Monday after the Super Bowl—beginning about 12 hours after the game. The rival TurboTax ad was aired during the second quarter of the Super Bowl. This constituted the second exposure to the rival ad for all participants who watched the game. We hypothesized that, as in prior controlled experiments with static ads, the PPC procedure would mitigate decay of the counterclaims all participants saw in wave 1 (day 1), thereby neutralizing the persuasive effects of being re-exposed to the rival TurboTax ad during the Super Bowl. As such, in wave 2, we solicited opinions on TurboTax using the same outcome measures as in the first wave, but without explicit re-exposure to the rival ad since the game served as the second re-exposure for those who watched.

Finally, we followed-up with the same sample again for a third wave seven days later, offered a third re-exposure to the rival ad, and again re-assessed their attitudes toward TurboTax. Participation in the third wave was limited to only those participants who had completed the second wave and reported watching the Super Bowl. In this wave, we again showed participants a series of video ads, including the rival TurboTax Super Bowl ad, and asked the same four outcome questions as in prior waves.

Analysis and Results

We limit our analysis to those participants who reported watching the Super Bowl (N = 1,172) and thus had the opportunity to be re-exposed organically to the rival TurboTax ad during the game. If the PPC procedure is effective, it should undercut the persuasive effects of the rival TurboTax ad upon each subsequent re-exposure—first during the Super Bowl, and then during the third wave. Our analysis focuses on the effect of the PPC procedure relative to both the Poison Only ad and the Pure Counterargument ad. In the SOM, we also explore how removing the distraction of a background video (as in the Pure Counterargument ad) may affect recall and influence the efficacy of the ad, but that is not the main emphasis here.

As our primary outcome of interest, we create a standardized aggregate "favorability index" comprised of our first two outcome measures. In a linear model, we analyzed the effect of treatment assignment on Turbo Tax favorability, controlling for participant race, age, party affiliation, gender, and income, as well as indicators for college education, prior Turbo Tax use, and whether self-reported weekly hours of television watched was above the median for this sample.

We find that exposure to the PPC ad in the first wave (before the Super Bowl) effectively reduced TurboTax favorability during each subsequent wave, with the largest effects seen during the first wave. This aligns with the results of Study 3, which also showed the strongest effects of the PPC procedure at the time at which participants were initially exposed to the counter messages. In the first wave (day 1), TurboTax favorability among participants who had seen the PPC ad was significantly lower than favorability among participants who had seen the Poison Only ad (0.32 SD; $SE = 0.07$, $p < .001$), or the Pure Counterargument ad (0.30 SD, $SE = 0.07$, $p < .001$).

At the second wave (day 3)—after being re-exposed to the rival TurboTax ad during the Super Bowl—participants in the PPC condition rated TurboTax 0.31 SD ($SE = 0.07$, $p < .001$) less favorably than participants in the Poison Only condition, and 0.21 SD ($SE = 0.07$, $p = .002$) less favorably than participants in the Pure Counterargument condition. Effect sizes decreased slightly by the third wave (day 9), but remained highly significant: participants who had seen the PPC ad in wave 1 still rated TurboTax 0.30 SD ($SE = 0.07$, $p < .001$) and 0.24 SD ($SE = 0.07$, $p < .001$) less favorably than participants who had seen the Poison Only or Pure Counterargument ad, respectively. Similar patterns can be seen across all four outcome measures independently (see SOM).

Discussion

In a setting in which participants were organically and then experimentally exposed repeatedly to a rival communication—the TurboTax Super Bowl ad—the PPC lastingly undercut the rival communication's persuasive effects. This meaningfully extends the applicability and utility of the PPC procedure, especially for less prominent communicators whose messages may be infrequently seen and whose communication channels may rely heavily on either TV or online video appeals.

Study 7: Test of Cue-Based Recall

Our first set of studies established the effectiveness of the PPC procedure across domains and modalities. We hypothesized the PPC would be effective against a more frequently encountered rival ad because its parasitic component spurs recall of the embedded counter-messages when the rival ad is subsequently encountered, thereby hindering memory decay. Study 7 tested this mechanism by adding two questions to measure participants' recall of the counter-messages.

Method

Participants. Participants were 266 Amazon MTurk workers (mean age = 35 years, $SD = 11.5$; 51% female) who received $1.80 for completing the survey. As in the previous studies, the sample was again restricted to participants who were located in the United States, had an MTurk

approval rating of at least 95%, and passed an initial attention check. All participants who had taken previous surveys as part of this study were excluded. The necessary sample size was decided ex ante based on expected effect sizes from pilot testing.

Procedure. Study 7 utilized the same materials as Studies 3 and 4 (Fig. 3), with McKinley serving as the rival candidate. All workers who consented to participate and passed the attention check were randomly assigned to one of three conditions: Control, Traditional Response, PPC. Participants were not aware of their condition assignment.

The study procedures followed those described in Studies 1 and 2, except we added two additional outcome questions asked at the end of the survey, upon re-exposure to the rival ad:

1. When you see the claims in McKinley's ad, how clearly do you recall the specifics of any arguments you may have viewed against those claims?
2. Which of the following anti-McKinley claims do you recall seeing, if any? Participants were asked to choose from a list of eight options. Three of the choices were the exact counter-messages included in both the Traditional Response ad and the PPC ad.

Analysis and Results

Replicating the effects seen in previous studies, exposure to the PPC advertisement reduced participants' likelihood of voting for McKinley and reduced perceived honesty. Following the same analysis procedures as in the previous studies, we find exposure to the PPC ad reduced participants' likelihood of voting for McKinley by 0.70 SD (SE = 0.14, $p < .001$) relative to exposure to the Control condition, and by 0.44 SD (SE = 0.14, $p = .002$) relative to exposure to the Traditional Response condition. Similarly, participants in the PPC condition rated McKinley as 0.62 SD (SE = 0.14, $p < .001$) less honest than Control condition participants, and 0.44 SD (SE = 0.14, $p = .003$) less honest than those in the Traditional Response condition.

On the first recall measure, 27% of participants in the PPC condition reported they recalled seeing specific arguments against McKinley's claims "extremely clearly," versus 15% of participants in the Traditional Response condition ($\chi^2(1) = 3.71$, $p = .05$) and 4.6% of participants in the Control condition ($\chi^2(1) = 13.25$, $p < .001$). Corroborating these reports, when asked to identify which anti-McKinley messages they saw—from a list of eight possible choices—PPC condition participants correctly identified more of the three counter-messages than Traditional Response condition or Control condition participants. Relatedly, 24% of participants in the PPC condition correctly identified all three counter-messages, compared to 15% of participants in the Traditional Response condition ($\chi^2(1) = 2.57$, $p = .11$) and 0% of participants in the control condition.

Tellingly, the effect of assignment to the PPC condition on perceived honesty was mediated by participants' recall of the specific counter-messages against McKinley ($z = -4.55$, $p < .001$). In turn, honesty mediated the effect of assignment to the PPC condition on participants' likelihood of voting for McKinley ($z = -4.14$, $p < .001$).

Discussion

Study 7 supported our hypothesis that the PPC spurs recall of its counter-messages upon repeated exposures to a rival's ad, and that this approach is most potent when the counter-messages offer evidence of a rival's duplicity.

General Discussion

The PPC procedure leverages associative memory and retrieval-cue-based recall to structure one-time communications that lastingly undermine the claims of a more frequently encountered disputable rival. Across the seven studies reported here (and three additional studies in the SOM), the PPC procedure effectively and enduringly undercut the persuasive effects of a rival's communication by using retrieval cues to activate memory and slow decay of its countervailing messages. Our studies examined one prevalent form of communication, paid advertisements. Future research should test the generality of the PPC's effectiveness to earned media and direct forms of communication as well. Our studies employed the PPC only in the context of asymmetries in frequency of exposure to an unfamiliar rival communication and communicator. Future research should explore the extent of the PPC effect on advantaged rival communicators who are also more familiar to audience members.

In all, the PPC procedure's effectiveness represents an important step toward redressing at least some imbalanced information environments in which legitimate critiques by disadvantaged voices can be drowned out by more advantaged communicators who exploit the advantage to disseminate misleading information.

References

Argo, J. J., and Main, K. J. (2004) Meta-analyses of the effectiveness of warning labels. *Journal of Public Policy and Marketing*, *23*, 193-208.

Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture,* J. H. Barkow, L. Cosmides, J. Tooby, Eds. Oxford Univ. Press, pp. 163-228.

Darke, P. R., & Ritchie, R. B. (2007). The defensive consumer: Advertising deception, defensive processing, and distrust. *J. Market. Res. XLIV*, 114-127.

Eagly, A. H., Kulesa, P., Brannon, L. A., Shaw, K., & Hutson-Comeaux, S. (2000). Why counterattitudinal messages are as memorable as proattitudinal messages: The importance of active defense against attack. *Personality and Social Psychology Bulletin, 26,* 1392-1408.

Fehr, E., & Gachter, S. (2002). Altruistic punishment in humans. *Nature, 415,* 137-140.

Farrelly, M. C., Davis, K. C., Haviland, M. L., Messeri, P., & Healton, C. G. (2005). Evidence of a dose-response relationship between "truth" antismoking ads and youth smoking prevalence. *American Journal of Public Health, 95*, 425-431.

Fazio, L. K., & Sherry, C. L. (2020). The effect of repetition on truth judgments across development. *Psychological Science, 31*, 1150-1160.

Gerber, A. (1998). Estimating the effect of campaign spending on senate election outcomes using instrumental variables. *American Political Science Review, 92*, 401-411.

Gerber, A. (2004). Does campaign spending work?: Field experiments provide evidence and suggest new theory. *American Behavioral Scientist, 47*.

Gerber, A. S., Gimpel, J. G., Green, D. P., & Shaw, D. R. (2011). How large and long-lasting are the persuasive effects of televised campaign ads? Results from a randomized field experiment. *American Political Science Review, 105*, 135-150.

Hill, S. J., Lo, J., Vavreck, L., & Zaller, J. (2013). How quickly we forget: The duration of persuasion effects from mass communication. *Political Communication, 30*, 521-547.

Jang, J., Lee, E., & Shin, S. Y. (2019). What debunking of misinformation does and doesn't. *Cyberpsychology, Behavior, and Social Networking, 22*, 423-427. http://doi.org/10.1089/cyber.2018.0608

Johar, G.V. and Simmons, C.J. (2000) The use of concurrent disclosures to correct invalid inferences. *Journal of Consumer Research*, *26*, 307-22.

Kalla, J., & Broockman, D. (2018). The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments. *American Political Science Review, 112*, 148-166.

Koerth-Baker, M. (2018, Sep. 10). How money affects elections. *FiveThirtyEight*. Retrieved from: https://fivethirtyeight.com/features/money-and-elections-a-complicated-love-story/

Koriat, A., Goldsmith, M., & Pansky, A. (2000). Toward a Psychology of Memory Accuracy. *Annual Review of Psychology, 51*, 481-537.

Kumkale, G. T., & Albarracín, D. (2004). The sleeper effect in persuasion: A meta-analytic review. *Psych. Bulletin, **130***, 143–172.

Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K.M., Menczer, F., et al. (2018). The science of fake news. *Science, 359*, 1094-1096. DOI: 10.1126/science.aao2998.

Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest, 13*, 106–131.

Lowery, W. (2014, April 4). 91% of the time the better-financed candidate wins. Don't act surprised. *The Washington Post*. Retrieved from: https://www.washingtonpost.com/news/the-fix/wp/2014/04/04/think-money-doesnt-matter-in-elections-this-chart-says-youre-wrong.

Maaravi, Y., Ganzach, Y., & Pazy, A. (2011). Negotiation as a form of persuasion: Arguments in first offers. *Journal of Personality and Social Psychology, 101,* 245-255.

Mutz, D. C., & Reeves, B. (2005). The new videomalaise: Effects of televised incivility on political trust. *American Political Science Review, 99*, 1-15.
New steps to protect US elections (2020, September, 3).
https://about.fb.com/news/2020/09/additional-steps-to-protect-the-us-elections/.

OpenSecrets.org. (2018). Did money win? Retrieved from: https://www.opensecrets.org/elections-overview/did-money-win.

Petty, R. E., & Brinol, P. (2010). Attitude change. In R. F. Baumesiter & E. J. Finkel (Eds.), *Advanced social psychology: The state of the science* (pp. 217-259). New York: Oxford University Press.

Petrova, P. K., & Cialdini, R. B. (2011). New approaches toward resistance to persuasion. In G. Hastings, C. Bryant., & K. Angus (Eds.). *The Sage handbook of social marketing* (pp. 107-122). London: Sage Publications.

Rogers, T., & Milkman, K. (2016). Reminders through association. *Psychological Science, 27*, 973-986.

Roediger, H. L., Tekin, W., & Uner, O. in *Learning and Memory: A Comprehensive Reference*, H. L. Roediger, J. H. Byrne, Eds. (Elsevier, 2017), vol. 2, chap. 2, pp. 5-26.

Sagarin, B. J., Cialdini, R. B., Rice, W. E., & Serna, S. B. (2002). Dispelling the illusion of invulnerability: The motivations and mechanisms of resistance to persuasion. *Journal of Personality and Social Psychology, 83*, 526-541.

Schwarz, N., Sanna, L., Skurnik, I., and Yoon, C. (2007). Metacognitive experiences and the intricacies of setting people straight: Implications for debiasing and public information campaigns. *Advances in Experimental Social Psychology, 39,* 127-161.

Schwarz, N., Newman, E., & Leach, W. (2016). Making the truth stick & the myths fade: Lessons from cognitive psychology. *Behavioral Science & Policy 2*, 85-95. doi:10.1353/bsp.2016.0009

Siegel, J. T., Alvaro, E. M., Crano, W. D., Lac, A., Ting, S., & Jones, S. P. (2008). A quasi-experimental investigation of message appeal variations on organ donor registration rates. *Health Psychology, 27*, 170–178. https://doi.org/10.1037/0278-6133.27.2.170

Skurnik, I., Yoon, C., Park, D.C., and Schwarz, N. (2005). How warnings about false claims become recommendations. *Journal of Consumer Research, 31*, 713-724.

Steinmetz, K. (2018, August 20). The real news, fake news crisis. *Time*, 26-31.

Truth vs. Lies (2020, Fall). Special edition, *Scientific American, 29*.

Tulving, E. (1983). *Elements of Episodic Memory*. Oxford Univ. Press.

Unkelbach, C., Koch, A., Silva, R. R., & Garcia-Marques, T. (2019). Truth by repetition: Explanations and implications. *Current Directions in Psychological Science*, *28*, 247-253. doi:10.1177/0963721419827854

Vosoughi, S., Roy, D, & Aral, S. (2018). The spread of true and false news online. *Science, 359*, 1146-1156. DOI: 10.1126/science.aao9559.

Wahlheim CN, Alexander TR, Peske CD. (2020). Reminders of everyday misinformation statements can enhance memory for and beliefs in corrections of those statements in the short term. *Psychological Science*, 1325-1339. doi:10.1177/0956797620952797

Wixted, J.T. (2004). The psychology and neuroscience of forgetting. *Ann. Rev. Psychol. 55*, 235-269.

Wood, W., & Quinn, J. M. (2003). Forewarned and Forearmed? Two meta-analysis syntheses of forewarnings of influence appeals. *Psychological Bulletin, 129,* 119-138.

Zuckerberg, M. (2020). https://www.facebook.com/zuck/posts/10112270823363411.