

Supporting Information for
Split Feelings: Understanding Implicit and Explicit Political Persuasion

Timothy J. Ryan and Yanna Krupnikov

Table of Contents

1. Sample characteristics	1
2. Ad script.....	2
3. Ad Pre-Test and Ad Post-Hoc Check.....	4
4. Study 2: A Second Test of the APE Model.....	10
5. Results Including International Students	14
6. Subgroup Analyses for Study 3b.	16
7. Pre-registration for Study 1	18
8. Pre-registration for Study 2.....	19
9. Pre-registration for Study 3a	22
10. Pre-registration for Study 3b.....	25
11. Pre-registration for Study 4.....	28
12. Procedure to Ensure Data Integrity in Study 3b	30
13. Adherence to Principles and Guidance for Human Subjects Research.....	31

1. Sample characteristics

Table SI-1: Sample Characteristics

	Study 1	Study 2	Study 3a	Study 3b	Study 4
Age					
Under 30	99.3%	36.2%	99.6%	20.8%	98.8%
30-39	0.7	37.2	0.4	35.7	1.0
40-49	0.0	12.8	0.0	21.5	0.2
50-59	0.0	8.5	0.0	12.9	0.1
Over 60	0.0	5.3	0.0	9.1	0.0
Education ¹					
No Diploma	0.0%	--	0.0%	0.6%	0.0%
Diploma only	0.0	--	0.0	13.0	0.0
Some college	100.0	--	100.0	22.1	100.0
Associate's Degree	0.0	--	0.0	12.5	0.0
BA	0.0	--	0.0	37.5	0.0
Advanced Degree	0.0	--	0.0	14.3	0.0
Gender					
Female	58.1%	56.2%	41.8%	51.3%	56.1%
Male	41.3	43.8	57.1	48.1	42.7
Non-binary	0.7	0.0	1.1	0.6	1.3
Partisanship					
Democrat	44.0%	36.6%	44.9%	42.1%	43.5%
Republican	20.2	19.0	16.4	27.9	14.2
Independent	33.2	26.6	34.6	21.3	31.5
Other / Don't know	2.6	17.7	4.2	8.6	10.8

¹ For studies 1, 3a, and 4, we infer respondent education from the fact that they are undergraduate samples. Study 2 is an MTurk sample, but unfortunately we did not measure respondents' education.

2. Ad script

The following script was held constant across all advertisement conditions.

Context	Narration
Mike Harper’s school: the outside of a classroom, with student projects visible in the background.	Hi, I’m Mike Harper. I’m a dad, a teacher, and a soccer coach. And I’d like to be your next representative to the United States Congress. My dad had an expression. He said, “Start small, go far.” That’s the kind of stick-to-it attitude I’d take to Washington. I want to be your voice against the incredible egomania that drives both parties. I want to make the changes we need to keep our communities safe, protect America’s status as a world leader, and make sure the next generation has all the opportunities we had. But that’s the end of the story. Let’s start at the beginning.
Driving around town in Mike Harper’s car.	I was born in Ellerbe and raised right here in Statesville. Actually grew up in that house [right here / right there]. Every morning in the summer at 5am I’d ride my bike to Umstead Park. That’s where I worked at the concession stand, selling coffee and newspapers in the morning, hotdogs and ice cream to the little league spectators later in the afternoon. The local kids still work there, just like I did. I like to think they’re still learning the value of a dollar, just like I did.
Mike Harper’s kitchen, with a cutaway shot to Peter Mueller in his high school yearbook.	When I was sixteen, my parents died in a car accident and I was totally devastated. My high school history teacher, Peter Mueller, helped me and my brother get through it. He always used to tell us that when things get tough, people pull together. Sometimes I think the politicians in Washington could use a reminder of that. There’s a lot more Republicans and Democrats could agree on if we forgot about silly partisan hang ups and focused on solving problems.
Mike Harper’s classroom, with six students visible	I went on to become a history teacher myself, and was honored to be awarded the Proust National Award for Education Excellence in 2014. It’s a privilege to come into the classroom each day. I never cease to be amazed at what our students can do. It makes me confident in our future.
Outside school building	(Spoken by Student A) Mr. Harper inspires us every day. Learning about US history made me want to get more involved in the community. He also taught me to be a good person and hard worker.”
School hallway	(Spoken by Student B) He’s a really great teacher. He says his motto is, “Start small, go far,” which is an easy idea to remember.
Faculty lounge	Now I know what you’re thinking. What does a simple school teacher from Statesville know about fixing what’s broken in Washington? Well I’ll tell you, the problem in Washington these days is that people are too focused on getting to the next level and taking credit whenever they can. Both parties do it. The last time Republicans were in charge, they kicked all the hard problems down the road. Then Democrats came in and did the same thing. Me? I don’t care about credit. I don’t need to take meetings with lobbyists. I don’t want to run for president. I want to work to solve problems, and then come back home to Statesville, the community I love. That’s the mindset that’s going to keep America the most powerful and respected nation in the world.

<p>Mike Harper's living room</p>	<p>This is my grandfather's compass, which he used as a paratrooper in World War II. He passed it on to me, and I still carry it around in my pocket. For me, it's a reminder to stay on course. To keep true to my values, and not get sidetracked by silly distractions. I've made it my campaign logo because that's my promise to you: to go to DC, and to stay on course for as long as I'm there.</p> <p>Now if you think another lawyer or lobbyist is what we need to fix DC, I'm not your guy. But if you want someone who is determined to put old fights aside and focus on solving problems, I'd sure appreciate your support. if we start small, we can go far.</p>
<p>Long, still image of Harper's campaign logo and slogan</p>	<p>(None)</p>

3. Ad Pre-Test and Ad Post-Hoc Check

We conducted two tests designed as checks on our advertisements. The first was a pre-test and was conducted prior to fielding Study 1. The second was a post-hoc check, and was conducted after the completion of the studies at the suggestion of a reviewer of this manuscript. We discuss each of the checks in turn.

3.1 Ad Pre-Test

Prior to fielding Study 1, we conducted a pre-test to ensure that our advertisement did not produce any unexpected confounds. In particular, we wanted to ensure that people perceived the ad script as not clearly partisan. The test was conducted on MTurk in September 2017 (N=204). In this test, participants were assigned to either the high or low quality version of our ad (just as they were in the studies). After exposure to the ad, we asked participants a variety of questions about the candidate in the advertisement and the advertisement itself. We present these results below.

Table SI-2: Political Characteristics of the Advertisement

	Perceived Ideology of Candidate (1-7, higher values = more conservative)	Perceived Partisanship of Candidate (1-5, higher values, more Democratic)
High Quality	4.24	2.82
Low Quality	4.08	2.98
Difference	0.15 (p=0.43)	0.15 (p=0.25)

We note that in our measures also had options that allowed participants to note that they are uncertain. In the ideology measure, 7.61% of participants in the high quality condition, and 10% in the low quality condition reported that they could not tell what ideology the candidate had. In the partisanship condition, 40.91% in the low quality condition and 46.7% in the high quality condition reported that they were unsure about the partisanship.

We also asked participations about various possible characteristics that the candidate might have. In these measures the participants rated the candidate on a scale of 1 to 5, where 1 meant that the term described the candidate very well, and 5 meant that it did not describe the candidate at all. We note that on two characteristics, we see significant differences: the candidate is perceived as taller and more attractive in the high quality conditions. We believe this is a signal of the differences in ads – the high quality condition had better lighting deliberately, which may have led to the candidate appearing more attractive.

We also asked the participants to estimate the candidate’s age. Again, we see that likely due to the more flattering lighting, the candidate appears slightly younger in the high quality ad.

The differences that do emerge on these measures suggest that our goal to manipulate purely visual components, we were successful in doing so.

Table SI-3: Perceptions of Candidate’s Characteristics

	High Quality	Low Quality	Diff
Educated	2.09	2.2	p=0.41
Community leader	2.37	2.59	p=0.15
Experienced	3.47	3.47	p=0.97
Healthy	2.24	2.41	p=0.19
Tall	2.73	3.05	p=0.02
Attractive	2.90	3.35	p=0.002
Age	41.24	43.17	p=0.008

We also measured general perceptions of the ad as a whole. We find that people believed the ads were equally memorable (identical distribution over 5 possible categories, based on a Kolmogorov-Smirnov test, $p=0.88$; similar means 2.46 versus 2.66, $p=0.17$, based on a t-test). We also find that a nearly identical number of participants perceived that this could possibly be a real ad (group difference $p=0.95$).

Finally, we also included a number of open-ended measures to track whether there were differences in how memorable these advertisements were. If participants had an easier time remembering the candidate in, say, a high quality versus a low quality ad, this would cause difficulties for the measurement of implicit attitudes.

We find very few differences in these open-ended responses. In the low quality ad condition, 84.4% of participants were able to recall the candidate’s name exactly, and 5.5% were close. In the high quality ad condition, 83.7% stated the name exactly, and 5.43% were close. The difference in exact recall between conditions was 0.7 percentage points, $p=0.89$. Next, 53.7% in the low quality condition and 45.7% in the high quality condition successfully recalled the candidate’s logo; the difference was not significant at $p=0.33$. Finally, in the low quality condition 63.6% recalled the candidate’s slogan, while 67.4 did so in the high quality condition, again this was not a significant difference at $p=0.57$. In sum, across all the open-ended measures, we see no indication that one of our ads was more memorable than the other.

3.2 Post-Hoc Check

In addition to the pre-test that we fielded prior to completing, we also fielded a second check after the completion of our studies. This was fielded at the suggestion of a reviewer to ensure that our treatments had functioned as we theorized: the high-quality advertisement led people to have a positive experience, and the low-quality advertisement led people to have a negative experience.

The post-hoc check was conducted in December 2020. We used CloudResearch to recruit MTurk participants (329 participants entered the study). CloudResearch increases data integrity by weeding out workers who were flagged for fraudulent responding in the past (see Kennedy et al. 2020 on the issues). Additionally, we included several attention check questions to help ensure the responses we collected were high-quality.

The check randomly assigned participants to one of two conditions: (1) high quality ad first, low quality ad second or (2) low quality ad first, high quality ad second. This approach means that we can analyze the data as a within-subject experiment or a between-subject experiment (based on assignment to the first ad).

Prior to assignment to treatment, participants were asked a series of demographic measures, along with a question measuring their feelings about music. We will use these measures to consider whether a personal experience with music affects response to treatment:

Music Measure 1: How many hours do you spend listening to music on an average day?

<1> I do not listen to any music on an average day

<2> Less than 1 hour

<3> 1 to 2 hours

<4> 2 to 3 hours

<5> 3 to 4 hours

<6> 4 to 5 hours

<7> 5 to 6 hours

<8> More than 6 hours per day

Music Measure 2: How well does the phrase "appreciator of music" describe you?

<1> Very well

<2> Somewhat well

<3> Neither well or not well

<4> Somewhat not well

<5> Not well at all

Post-treatment, we asked a variety of questions designed to track the extent to which the participants found the ad pleasant and the extent to which they reported a positive mood post-exposure. The pleasantness questions are based on a measure by Madigan and Bollenbach (1986) and the mood questions are based on the Brief Mood Introspection Scale by Mayer and Cavallaro (2019).

The measures were as follows:

Enjoyment: How much did you enjoy watching the ad? A score of 1 would mean you did not enjoy the ad at all and a score of 7 would mean you enjoyed the ad very much.

Pleasantness: How pleasant or unpleasant do you find this ad? Scores closer to 0 mean that you found the ad mostly unpleasant and scores closer to 10 mean that you found the ad mostly pleasant. A score of 5 means that you found the ad largely neutral.

Visual Pleasantness: How pleasant or unpleasant do you find **the visual components** (e.g. lighting, backgrounds, and so forth) in this ad? Scores closer to 0 mean that you found these visual components mostly unpleasant and scores closer to 10 mean that you found the visual components mostly pleasant. A score of 5 means that you found the ad largely neutral.

Audio Pleasantness: How pleasant or unpleasant do you find **audio components** (e.g. music, narration, and so forth) in this ad? Scores closer to 0 mean that you found the audio components mostly unpleasant and scores closer to 10 mean that you found the audio components mostly pleasant. A score of 5 means that you found the ad largely neutral.

Mood: How well does each of the words below describe your current mood? Scores closer to 1 mean that the adjective does not describe your current mood and scores closer to 7 mean that they do describe your current mood [presented in random order]

- Lively
- Happy
- Sad
- Gloomy
- Grouchy
- Drowsy
- Cheery

- Active

Our study also included checks for data validity — i.e. to avoid fraudulent responding. In the end, we found that 25 responses could be flagged as potentially suspicious (based on methods suggested by CloudResearch and Kenney et al. 2020). We present the results without these 26 responses, though we note that the results do not substantively change if these responses are retained. Three additional participants were excluded from the analysis because they reported they could not watch the videos all the way through.

Within-Subject Results, Post-Hoc Check

The within subject results reflect the difference between watching a high-quality ad and a low-quality ad. Positive differences mean that a person perceived a high-quality ad more positively (e.g. enjoyed it more, was more pleasant) than a low quality ad. On the mood measures, positive differences mean that a person was more likely to state this mood in the high quality ad, negative differences mean that they were more likely to have the mood in the low quality ad. All p-values are two-tailed.

Table SI-4: Effects of Treatment Condition on Affect and Mood (Within-subjects)

	High Quality	Low Quality	Difference, p-value
Enjoyed Ad (1-7 scale)	4.49	3.06	1.43 (p<0.0001)
Pleasant (1-10 scale)	6.63	4.44	2.19 (p<0.0001)
Visually Pleasant (1-10 scale)	6.96	4.13	2.82 (p<0.0001)
Auditory Pleasant (1-10 scale)	6.78	4.59	2.20 (p<0.0001)
<i>Mood Measures, all 1-7 scale</i>			
Lively	4.25	3.49	0.76 (p<0.0001)
Happy	4.59	4.00	0.59 (p<0.0001)
Sad	1.89	2.05	-0.16 (p=0.0052)
Gloomy	2.01	2.37	-0.36 (p<0.0001)
Grouchy	1.93	2.28	-0.35 (p<0.0001)
Drowsy	2.28	2.79	-0.52 (p<0.0001)
Cheery	4.21	3.58	0.63 (p<0.0001)
Active	4.40	3.83	0.57 (p<0.0001)

These results confirm that the high-quality worked as intended, increasing positive affect (and moods), relative to the low-quality ad.

Between-Subject Check

Since the order in which participants saw the ads was assigned randomly, we can also analyze the data as a between-subject study, using only the first ad to which the participant was assigned. We present these results below.

Table SI-5: Effects of Treatment Condition on Affect and Mood (Between-subjects)

	High Quality	Low Quality	Difference, p-value
Enjoyed Ad (1-7 scale)	4.21	3.46	0.75 (p=0.0008)
Pleasant (1-10 scale)	6.43	5.29	1.14 (p=0.0002)
Visually Pleasant (1-10 scale)	6.69	4.82	1.87 (p<0.0001)
Auditory Pleasant (1-10 scale)	6.59	5.15	1.44 (p<0.0001)
<i>Mood Measures, all 1-7 scale</i>			
Lively	4.12	3.80	0.32 (p=0.132)
Happy	4.67	4.19	0.48 (p=0.023)
Sad	1.89	2.15	-0.25 (p=0.151)
Gloomy	2.06	2.40	-0.35 (p=0.080)
Grouchy	1.98	2.32	-0.33 (p=0.073)
Drowsy	2.36	2.66	-0.30 (p=0.147)
Cheery	4.29	3.75	0.54 (p=0.011)
Active	4.19	4.31	-0.12 (p=0.55)

These results likewise confirm that the high-quality worked as intended, increasing positive affect (and moods), relative to the low-quality ad.

The Role of Personal Experience

An additional possible concern about our instrumentation that we considered is that some pre-existing personal characteristic might moderate how participants responded to the treatments, leading some to experience negative affect even while others experience positive affect. While it is not possible to investigate all possible moderating factors, we identified one that seemed more likely than most: a respondent’s appreciation for music. This factor might moderate reactions, since our instrumentation is substantially music-based. Since music appreciation is among the most likely moderating factors, if we find null heterogeneous effects in this dimension, we suspect we are unlikely to find them in others.

Music appreciation was measured with the two “Music Measures” described above. In the table below, we use each of these questions (one in each column) to estimate differences in participants’ responses to each of the outcome variables (a within-subject analysis). (The underlying models also control for the order in which participants saw the two ads.)

The table shows that experience variables do not have any effect on the difference between the ratings. Furthermore, in the between-subject comparisons, when we interact treatment with the music variables and find no significant interactions. These results serve to corroborate that the ads’ effects on affect are reasonably uniform.

Table SI-6: Music Appreciation Does Not Moderate Reactions to Ads

	Time Spent on Music Coefficient, SE, p-value	Appreciator of Music Coefficient, SE, p-value
Outcome in OLS model:		
Enjoyed Ad (1-7 scale)	0.032 (0.062, p=0.606)	-0.003 (0.104, p=0.973)
Pleasant (1-10 scale)	0.034 (0.094, p=0.723)	0.157 (0.156, p=0.314)
Visually Pleasant (1-10 scale)	0.061 (0.093, p=0.516)	-0.005 (0.155, p=0.975)
Auditory Pleasant (1-10 scale)	0.025 (0.100, p=0.795)	0.256 (0.156, p=0.106)
<i>Mood Measures (1-7 scale)</i>		
Lively	0.028 (0.055, p=0.612)	0.137 (0.091, p=0.131)
Happy	0.067 (0.049, p=0.172)	0.083 (0.081, p=0.304)
Sad	-0.014 (0.036, p=0.704)	-0.083 (0.060, p=0.165)
Gloomy	-0.058 (0.048, p=0.224)	-0.028 (0.079, p=0.721)
Grouchy	-0.010 (0.040, p=0.795)	0.005 (0.067, p=0.936)
Drowsy	0.022 (0.055, p=0.689)	-0.089 (0.091, p=0.333)
Cheery	0.061 (0.053, p=0.255)	0.052 (0.089, p=0.562)
Active	0.55 (0.049, p=0.263)	-0.006 (0.082, p=0.939)

Cell entries are regression coefficients predicting differences in responses to the high-quality and low-quality ads. Standard errors in parentheses.

References for this section:

- Kennedy, Ryan et al. 2020. "The Shape of and Solutions to the MTurk Quality Crisis." *Political Science Research and Methods* 8(4): 614–29.
- Madigan, R. J., and A. K. Bollenbach. 1986. "The Effects of Induced Mood on Irrational Thoughts and Views of the World." *Cognitive Therapy and Research* 10(5): 547–62.
- Mayer, John D. and Rachel Cavallaro. 2019. "Brief Mood Introspection Scale (BMIS): Technical and Scoring Manual (3rd Edition)." Available online at https://scholars.unh.edu/personality_lab/13/

4. Study 2: A Second Test of the APE Model

As the text notes, the purpose of Study 2 was to separately target propositional and associational processing related to Mike Harper. The structure of Study 2 was much the same as Study 4 in the main text, and we refer readers to that study for a discussion of the theoretical grounding for the approach and stimuli. Here, we focus on highlighting a few important differences between Study 2 and Study 4.

The first difference is that Study 2 took place about one year before Study 4, at which point we still adopted the theoretical perspective of the APE Model. Thus, and as the pre-registration for Study 2 reflects, we predicted the audio/visual manipulation in Study 2 to influence implicit attitudes, and the text-based manipulation to influence explicit attitudes.

The second difference is that Study 2 relied on a different data source: Amazon.com's Mechanical Turk (MTurk) crowdsourcing service. As we elaborate below, Mechanical presented unexpected difficulties. Mechanical Turk workers proved to be noticeably less diligent in compliance with the IAT protocol. (They had a considerably higher failure rate than students.) Additionally, shortly after our study was fielded, it came to light that surveys on MTurk were at least temporarily afflicted by fraudulent responding (Kennedy et al., in press). We took extra steps to ensure the integrity of the results we present below. Still, these issues, plus the benefits of a uniform data source, are what caused us to focus exclusively on student samples for the remaining studies.

Third, Study 2's Scandal manipulation has a different control condition than Study 4. After conducting Study 2, we wondered if the control article displayed below was too upbeat and enthusiastic in its tone. As such, we substituted a different (more neutral) condition in Study 4.

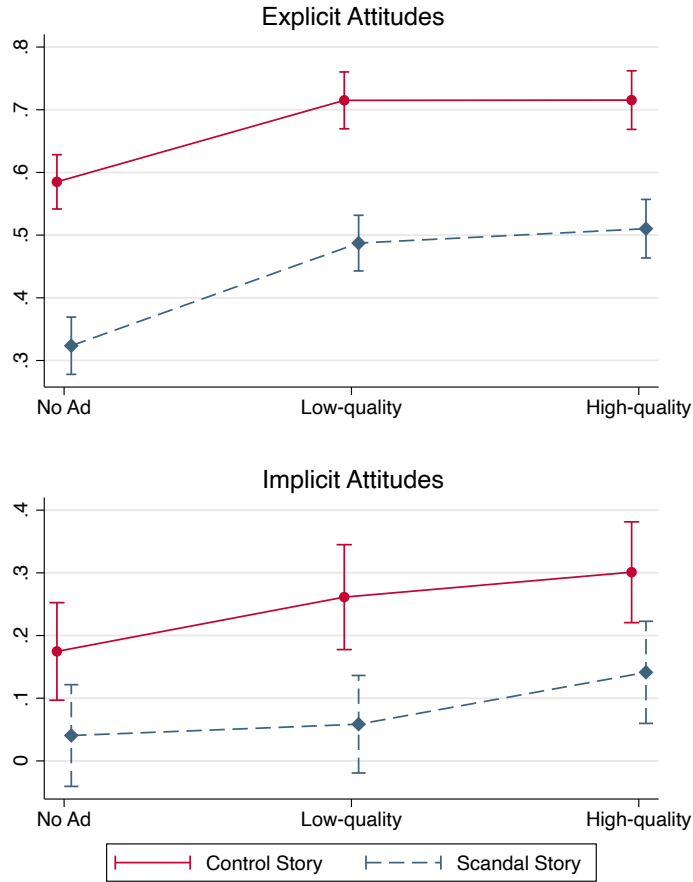
Procedure. Study 2 was conducted in July of 2018. We recruited 1,088 MTurk workers to participate. From this pool, we excluded participants for two reasons. First, we excluded responses (N=120) who exhibited patterns almost certainly indicative of the fraudulent responding we allude to above. Specifically, we excluded respondents whose ISP geographic information (automatically captured by our survey software) is one known to generate fraudulent responses. We also exclude respondents whose answers to open-ended questions have clear indications of fraud.² This procedure was not part of our pre-registration (indeed, we could not have anticipated the potential contamination of MTurk data) but we consider it fully justified, in light of clear evidence of fraudulent responding.

This said, we do not expect that our results will be sensitive to slightly different exclusion criteria. For one thing, the fraudulent responses should be evenly spread across randomly assigned conditions, and thus should mostly add statistical noise. Further, 75.8% of the responses we identify as fraudulent had unusable IAT scores, on the basis of "button-mashing" through the procedure (see a footnote in Study 1). Given that usable IAT scores are necessary for our main analysis, this means that the potentially fraudulent responses had a low likelihood of being part of our sample of analysis in the first place.

After excluding these 120 responses, we find that an additional 120 respondents (14.48% of the remaining data) had unusable IAT scores on the basis of the button-mashing. This represents a more substantial data loss than we experienced in the student samples, where only 3% of responses exhibited button mashing. This is a cautionary tale for researchers conducted implicit attitude studies with online convenience samples. We cannot include button-mashing respondents in analyses focused on implicit attitudes, though we include them in analyses of other measures.

² As has been discussed in online forums, the fraudulent responders commonly provide one-word answers that do not address the question being asked ("good," "nice.")

Figures SI-2: Study 2 Results



Whiskers represent 95% confidence intervals.

Results. Figure SI-2 displays participants’ explicit and implicit affect toward Mike Harper, depending on the treatment condition. A between-subjects ANOVA reveals significant effects for both manipulations. (For explicit attitudes, $F_{ad}(2,828) = 22.62, p < .001$; $F_{scandal}(1,828) = 136.94, p < .001$. For implicit attitudes, $F_{ad}(2,709) = 3.70, p < .03$; $F_{scandal}(1,738) = 24.06, p < .001$.) As the figure reveals, these effects are such that the scandal story had a significant negative effect on both implicit and explicit attitudes, while the ad quality manipulation had a generally positive (but less consistent) positive effect on both implicit and explicit attitudes. Thus, there is little evidence for the domain-specific effects that the APE predicts.

An especially helpful prediction for assessing the APE Model is, as we stated in Study 2’s pre-registration, that the scandal “[will] have a significantly smaller effect (or no effect) on implicit attitudes when the scandal story is preceded by an advertisement.” This is the prediction focused on the possibility that different kinds of content could have reasonably compartmentalized effects on implicit and explicit attitudes. Visually, it would manifest as a narrowing of the distance between the blue and red lines, as one reads from left to right in the bottom panel. As one can see, there is no evidence of this.

Discussion. Study 2 confirms that campaign messages have the power to influence both explicit and implicit attitudes about an unfamiliar candidate. However, it does not uncover support for the reasonably segmented, domain-specific effects predicted by the APE Model.

References for Study 2

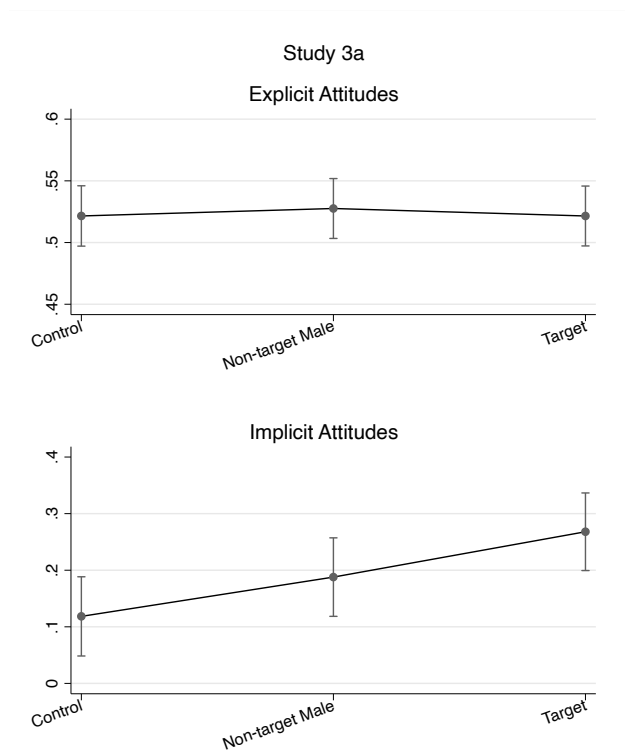
Kennedy, Ryan, et al. In press. "The Shape of and Solutions to the MTurk Quality Crisis." *Political Science Research and Methods*. Forthcoming.

5. Results Including International Students

The analyses reported in the main text follow our pre-registered plan to exclude international students from the analyses for Studies 3a and 4. However, it arguably makes sense to include these respondents. They might be less immersed in the dynamics of partisan competition in the United States and thus might speak more clearly to the psychological relationships that are our main focus.


Figure SI-3 below is parallel to Figure 5 (left column) in the text, except international students are now included. Figure SI-4 below is parallel to Figure 6 in the text—again with international students included. As can be seen, the results are nearly identical.

Figure SI-3: Results from Study 3a, with international students included



Whiskers represent 95% confidence intervals.

Figure SI-4: Results from Study 4, with international students included



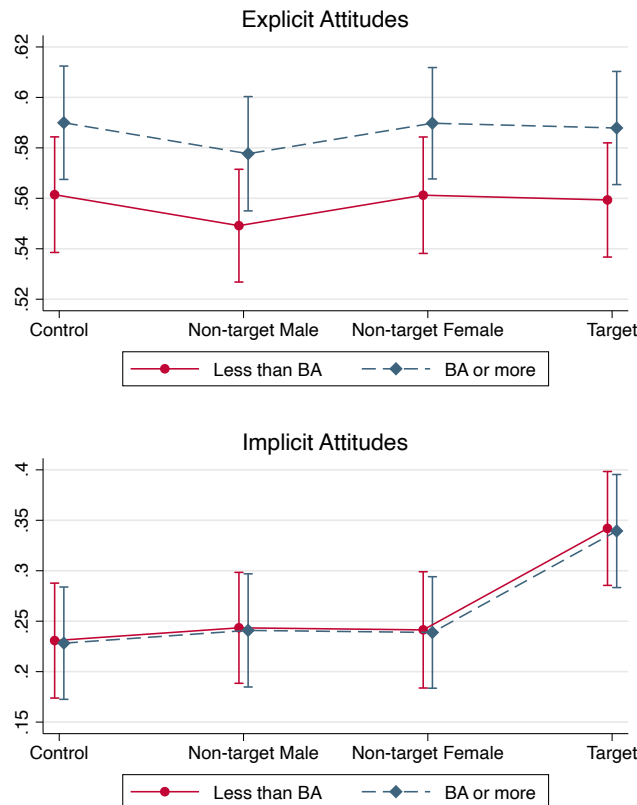
Whiskers represent 95% confidence intervals.

6. Subgroup Analyses for Study 3b.

As we note in the main text, Study 3b affords us opportunities to assess treatment effect heterogeneity, since this sample is diverse with respect to age and education (see Table SI-1). We report the relevant results here.

First, we stratify the analysis by education, separating participants who have obtained a college degree (48.2% of our sample) from those who have not (51.8%).³ Figure SI-5 presents these results. Education has a main effect on explicit attitudes. But critically, it does not appear to moderate effects on either explicit or implicit attitudes.⁴ As concerns implicit attitudes, education has neither a main nor interactive effect.⁵

Figure SI-5: Effects by Education in Study 3b.



Whiskers represent 95% confidence intervals.

Next, we stratify by age, separating participants under 40 years old (56.5% of the sample) from those 40 and over (43.5%). Figure SI-6 presents these results. This analysis reveals main effects of age, but not evidence that age moderates treatment effects.⁶

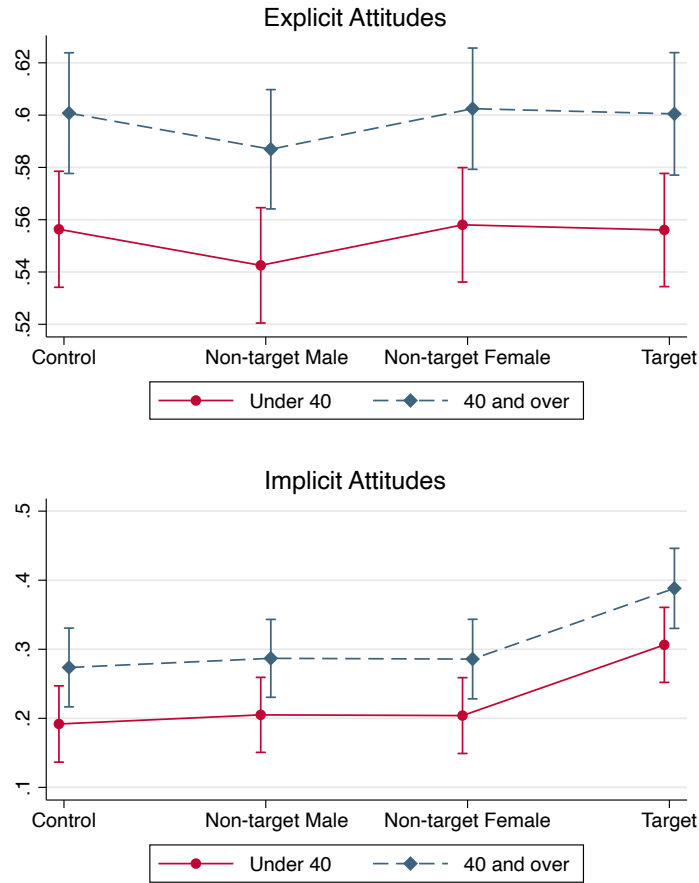
³ Recall that we applied a recruitment quota in Study 3b that limited individuals with a BA to no more than 50% of our sample. The slight imbalance reported here arises because our planned exclusions (see Section 11) applied more heavily to low-education participants than high-education participants.

⁴ In a two-way ANOVA, the F -statistic for the education factor is 7.70 ($p < .01$). However, for the treatment \times education interaction, $F = 0.43$; $p = 0.73$.

⁵ In a two-way ANOVA, $F_{\text{education}} = 0.01$, $p = .94$. $F_{\text{education} \times \text{condition}} = 0.26$, $p = .24$.

⁶ For explicit attitudes, $F_{\text{age}} = 18.37$, $p < .01$; $F_{\text{age} \times \text{condition}} = 0.39$, $p = .76$. For implicit attitudes, $F_{\text{age}} = 10.22$, $p < .01$; $F_{\text{age} \times \text{condition}} = 0.17$, $p = .92$.

Figure SI-6: Effects by Education in Study 3b.



Whiskers represent 95% confidence intervals.

In summary, these analyses show that age and education are predictive of implicit and explicit attitudes, but not in the interactive way that would necessarily be necessary for our reliance on student samples to undermine the conclusions in the main text. Such results are in step with the finding that pathological treatment effect heterogeneity is rare (Coppock et al. 2018), though of course any particular study can deviate from the general trend.

References for this section:

Coppock, Alexander, Thomas J Leeper, and Kevin J Mullinix. 2018. "Generalizability of Heterogeneous Treatment Effect Estimates across Samples." *Proceedings of the National Academies of Science* 115(49): 12441–46.

7. Pre-registration for Study 1

1) Have any data been collected for this study already?

No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?

All hypotheses are based on a "bland" ad (devoid of music and rich imagery) and a "rich" ad (with music and imagery)

(1) The bland ad will improve implicit attitudes relative to control, but not as much as the rich ad.

(2) None of the ads will affect explicit attitudes

RQ: Is the effect on implicit attitudes moderated by political engagement and sophistication, such that people low in engagement/sophistication are most affected?

3) Describe the key dependent variable(s) specifying how they will be measured.

Implicit measures: IAT

Explicit measures: how much S likes candidate in ad, feeling thermometer of candidate in ad, open-ended response (likes/dislikes) to candidate in ad, political leanings of candidate in ad (ideology/partisanship)

Checks on ad response: characteristics of candidate in ad (educated, community leader, experienced, healthy, tall, attractive); guess at candidate age.

4) How many and which conditions will participants be assigned to?

Three conditions: control (no ad), bland ad (ad without music/rich imagery), rich ad (ad with music/imagery)

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

Main analyses will compare implicit attitudes and explicit attitudes in each group to the control. Analyses will focus on overall attitudes, but will also consider how implicit and explicit attitudes affect each other. Analyses will also focus on particular subgroups -- as a research question considers how engagement and knowledge affect the role the ad plays in implicit attitudes.

6) Any secondary analyses?

We will consider how partisanship of subject affects responses.

7) How many observations will be collected or what will determine sample size?

No need to justify decision, but be precise about exactly how the number will be determined.

Study will be sent out to approximately 400 undergraduate students in the Fall 2017 semester and another 250 in the Spring 2018 semester.

8) Anything else you would like to pre-register?

(e.g., data exclusions, variables collected for exploratory purposes, unusual analyses planned?)

Participants may be excluded under the following conditions: (1) there is evidence in the study that they didn't watch the ad. (2) they did not complete full IAT.

8. Pre-registration for Study 2

Have any data been collected for this study already?

No, no data have been collected for this study yet

What's the main question being asked or hypothesis being tested in this study?

The current pre-registration relates to As Predicted pre-registration #6198. That pre-registration was deposited before an experiment we conducted in the Fall of 2017, with an exact replication in the Spring of 2018. The present pre-registration records how we interpreted results from those initial tests and how they informed the design of a new test to be conducted in the Summer of 2018. Thus, our answer in Question 1 relates specifically to the Summer 2018 round of data collection.

As pre-registration #6198 documents, we conducted an experiment in which participants (undergraduates partaking in the study for course credit) were randomly assigned to one of three experimental conditions. They watched an advertisement for a fictional political candidate that was high in production quality (the High-quality Condition), an advertisement with an identical script that was low in production quality (the Low-quality Condition), or no ad at all (a pure control condition). After ad exposure, participants completed an Implicit Association Test (IAT) focused on the fictional candidate. They also recorded (via a self-report) their liking or disliking of the candidate, as well as a number of secondary measures: the perceived ideology of the candidate in the ad, the candidate's perceived partisanship, their perception of the candidate's personal traits, and the candidate's age.

As Pre-registration #6198 records, we expected 1) the low-quality ad to “improve implicit attitudes relative to control, but not as much as the rich [high-quality] ad.” We also expected 2) neither of the ads to influence explicit attitudes. And we posed a research question: Are people low in political engagement and sophistication more responsive to political advertising (by either implicit or explicit measures) than people who are high in political sophistication.

Expectation (1) was based on our reading of the literature on implicit attitude change. In particular, we drew on Gawronski & Bodenhausen's (2006) suggestion that implicit attitudes become more favorable following the pairing of an attitude object (the candidate in our case) with positive stimuli (our advertisements). We saw the low-quality ad as providing a minimally positive stimulus (given the positive influence of “mere exposure” (cf. Zajonc 1968)). But we saw the high-quality ad as providing a more powerfully positive association, and thus likely to induce greater change in implicit attitudes.

Expectation (2) similarly drew from Gawronski & Bodenhausen's (2006) model of attitude change. The model these authors advance suggests that explicit attitude change is induced by propositional reasoning—in particular persuasive arguments that are relevant to the evaluative judgment (see especially p. 701). The script in our ad provided no concrete policy-focused information, so we expected neither ad to improve explicit attitudes relative to control. More importantly, the two versions of the advertisement were identical in terms of the persuasive arguments they offered—they had an identical script—and as such we expected them to be identical in their influence of explicit (but not implicit) attitudes.

In the initial two experiments, these expectations were supported only in part. Consistent with expectations, the low-quality ad improved implicit attitudes relative to the control condition. Pooling the two identical experiments together, the mean D-score in the control condition was 0.154 (SE=0.037), while in the Low-quality condition, it was 0.314 (SE=0.037), a significant difference ($p < .01$). However, contra expectations, the High-quality condition did not result in improved implicit attitudes relative to the Low-quality condition: the mean D-Score in the High-quality condition was 0.319 (SE=0.037), an insignificant difference ($p = .92$). As

concerns explicit attitudes, the Low-quality condition did not improve explicit attitudes relative to the control: the mean in the control condition was 0.483 (SE=0.017), while in the Low-quality condition it was 0.490 (SE=0.017, $p=.76$). However, the High-quality ad did improve explicit attitudes relative to control (mean=0.608, SE=0.017, $p<.01$).

In sum, we expected the effect on implicit attitudes to increase as the ad quality increased, with no clear pattern for the effects on explicit attitudes. Instead, we found that the ads improved implicit attitudes to approximately the same extent irrespective of their quality. We also found that the ad improved explicit attitudes—but only if it was high-quality.

While the results did not corroborate our *ex ante* expectations, we found no reason to think the tests we conducted were not valid. As such, we returned to the literature to reconsider what model the pattern of results might support. We see our results as broadly consistent with the patterns described by Gregg et al. (2006). These authors find that implicit attitudes about a previously unknown attitude object (which our fictional candidate is) can readily be induced by minimal stimulus. They also find that, once formed, implicit attitudes are difficult to change. This pattern is consistent with what we found: any positive candidate exposure can improve implicit attitudes (irrespective of whether the exposure is high-quality or low-quality). But the high-quality ad offers little extra improvement relative to the low-quality ad. Similarly, explicit attitudes about unknown attitude objects can readily be formed, which is consistent with the pronounced positive effect of the high-quality ad on explicit attitudes. (This framework does not readily explain why the effect of the Low-quality ad on explicit attitudes was null, though at this point we conjecture that it could be due to social desirability pressures: participants might feel they do not have a rational basis via which to explain liking of a candidate with a poorly-produced advertisement.)

We designed an extension to our initial set of tests to examine an additional implication of the Gregg et al. (2006) framework. In these authors' understanding, implicit attitudes, once formed, are difficult to change. On the other hand, explicit attitudes, even if they are already formed, can change rapidly when the propositions that originally supported the attitude are sharply undermined (see especially Studies 3 and 4).

Our new test seeks to induce this pattern in the context of political advertising about an unfamiliar candidate. Specifically, we plan to manipulate whether participants are told that the candidate featured in the ad has been implicated in a political scandal. Participants will either read a brief news account stating this to be the case (the Scandal condition), or a nonpolitical story about making cheese (a control condition). This manipulation will be fully crossed with the original manipulation, resulting in six conditions: No ad, Low-quality ad, or High-quality ad X nonpolitical story or scandal story. Our expectations follow from Gregg et al. (2006). Specifically, we hypothesize that the scandal will negatively influence explicit attitudes—irrespective of the advertisement condition. We also expect the scandal to have a negative effect on implicit attitudes if participants saw no candidate advertisement. However, we expect the scandal to have a significantly smaller effect (or no effect) on implicit attitudes when the scandal story is preceded by an advertisement. This constellation of results would be consistent with Gregg et al.'s (2006) notion that implicit attitudes, once formed, are difficult to change—especially by administering new propositional information.

One question that might arise in this setup is whether the design should randomize the order in which participants are exposed to the candidate advertisement and news stories. We view this elaboration as a reasonable design extension, since it would help elucidate what kinds of stimuli can influence implicit and explicit attitudes when they are being formed anew. However, we believe the additional conditions would strain our resources too much, and would represent too large a deviation from the more narrow theoretical objectives laid out above.

Describe the key dependent variable(s) specifying how they will be measured.

As described above, the key dependent variables in our study are implicit and explicit attitudes toward the fictional candidate described in the study. These will be measured as in past studies. We also intend to carry forward our secondary measures (e.g. the candidate's perceived partisanship and ideology).

How many and which conditions will participants be assigned to?

As described above, there are six conditions. In one randomization, participants are assigned to view a low-quality ad, a high-quality ad, or not ad at all. In the second (fully crossed) random assignment, participants are assigned to read a short news article suggesting that the candidate featured in the ad was implicated in a scandal, or a short recipe.

Specify exactly which analyses you will conduct to examine the main question/hypothesis.

The main analysis will focus on difference in means of implicit and explicit attitudes, by condition. (Implicit attitude scoring will follow standard procedures for calculating a D-score.)

Any secondary analyses?

Our previous instruments included instrumentation to assess whether treatment effects differed by subjects' level of the personality trait Need for Cognition—a possibility explicitly discussed in previously literature (such as Gawronski & Bodenhausen 2006). The initial tests provided little evidence that treatment effects differ by Need for Cognition, but we intend to carry forward the relevant instrumentation to assess the evidence further.

How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.

We intend to collect 600 responses via Amazon.com MTurk. Then, we intend to collect 1,200 responses via Survey Sampling International, though we may modify the second sample size on the basis of power calculations from the MTurk sample.

Anything else you would like to pre-register? (e.g., data exclusions, variables collected for exploratory purposes, unusual analyses planned?)

Similar to our previous pre-registration, participants may be excluded if (1) there is evidence in the study that they did not watch the advertisement or attend to the news story, or (2) if they did not complete the full instrument (including the Implicit Association Test).

9. Pre-registration for Study 3a

Have any data been collected for this study already?

No, no data have been collected for this study yet

What's the main question being asked or hypothesis being tested in this study?

The current pre-registration relates to a separate pre-registration filed via OSF on July 7, 2017 at 16:22pm, as well as As Predicted pre-registration #6198 (submitted in October of 2017). We conducted studies as described in those previous documents, and now are executing a follow-up study based on what we found. Thus, our answer in Question 1 relates specifically to the current (Fall 2018) round of data collection.

As we discuss in a conference proceeding (blinded), the existing studies uncover two results that appear inconsistent with the Associational-Propositional Evaluation (APE) Model as described by Gawronski & Bodenhausen (2006). First, shifting from exposure to a low-quality advertisement (no music, and bland imagery) to a high-quality advertisement (professional music and imagery) is associated with an improvement in explicit, but not implicit attitudes. (As previous pre-registrations document, our interpretation of the APE model implies the opposite pattern.) Second, we find that a proposition-focused message (succinct, text-based information that a candidate was involved in embezzling funds) is associated with a negative effect on both implicit and explicit attitudes. (The APE model predicts effects localized in explicit attitudes, especially if implicit attitudes are already formed.)

The purpose of the present experiment is to assess whether an idiosyncrasy in our experimental materials could be responsible for these results. In particular, we are concerned that manipulating ad quality influences participants' propositional (and not just association-based) processing. This would be the case, for instance, if participants made a reasoned, deduction-based inference about our hypothetical candidate's competence for public office on the basis of the quality of his advertisement.

Thus, we propose a test of the APE model that does not rely on the previous stimuli—and one where it is difficult to attribute treatment effects to deductive reasoning about the target candidate. Specifically, the current experiment manipulates mere exposure to the target candidate—without providing any information about him whatsoever.

There are three conditions. In one condition, participants are exposed to the target candidate (Mike Harper) in a way that provides no substantive information about Mike Harper. They evaluate two still photographs of him, guessing his age and reporting how well various traits (e.g. intelligent, wealthy) describe him. They also rate the attractiveness of his campaign logo and are induced to repeated his campaign slogan five times. Then, the procedure measures implicit and explicit attitudes toward Mike Harper.

In a randomly assigned contrast condition, the procedure induces mere exposure to a different candidate: Paul Coleman. Participants rate pictures of a different person, evaluate a different logo, and repeat a different campaign slogan. (But implicit and explicit attitudes are then measured about Mike Harper, not Paul Coleman.)

Finally, in a pure control condition, participants skip the mere exposure induction and proceed directly to the implicit and explicit attitude measures about Mike Harper.

The APE Model predicts that a mere familiarity induction as we employ here should have a positive effect on implicit attitudes, but no effect on explicit attitudes (expect perhaps as mediated by implicit attitudes). In other words, implicit attitudes about Mike Harper should be most positive in the first condition described

above, and less positive in the other conditions, whereas differences in explicit attitudes should be muted across conditions. These are the expectations we test, though as the background above implies, we see this test as addressing a potential limitation of our existing studies, and not one that we undertake because our existing studies substantially corroborate the APE Model.

We are submitting this pre-registration document just after the proposed study was fielded at one of the two planned universities. We have not yet analyzed or even downloaded the data. (In fact, the study will be open for another week before we can do so.)

Describe the key dependent variable(s) specifying how they will be measured.

As described above, the key dependent variables in our study are implicit and explicit attitudes toward the fictional candidate described in the study. These will be measured as in past studies. We also intend to carry forward our secondary measures (e.g. the candidate's perceived partisanship and ideology).

How many and which conditions will participants be assigned to?

The three conditions for this study are fully describe above.

Specify exactly which analyses you will conduct to examine the main question/hypothesis.

The main analysis will focus on difference in means of implicit and explicit attitudes, by condition. (Implicit attitude scoring will follow standard procedures for calculating a D-score.)

Any secondary analyses?

Our previous instruments included instrumentation to assess whether treatment effects differed by subjects' level of the personality trait Need for Cognition—a possibility explicitly discussed in previously literature (such as Gawronski & Bodenhausen 2006). The initial tests provided little evidence that treatment effects differ by Need for Cognition, but we intend to carry forward the relevant instrumentation to assess the evidence further.

Additionally, we added measurement of a third personality trait to the current study: Need to Evaluate (NTE). We conjecture that participants high in this trait will be most influenced to the treatments described above. We will treat NTE as a moderating variable.

How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.

We expect to collect approximately 250 responses by conducting this study on the participant pool at [blinded], and up to an additional 200 respondents via a participant pool at [blinded].

Anything else you would like to pre-register? (e.g., data exclusions, variables collected for exploratory purposes, unusual analyses planned?)

Similar to our previous pre-registration, participants may be excluded if (1) there is evidence in the study that they did not watch the advertisement or attend to the news story, or (2) if they did not complete the full instrument (including the Implicit Association Test).

[Thirteen days after the above pre-registration, we filed a separate registration that was blank except for the following response:]

Anything else you would like to pre-register? (e.g., data exclusions, variables collected for exploratory purposes, unusual analyses planned?)

The registration relates to [blinded] , which is a registration for the same study. As that registration indicates, we intend to conduct the experiment we describe at two universities. We have already conducted it at one university (University A) and are preparing to field it at the second (University B). In doing so, we realized that University B has a far larger share of international students, for whom a study that is in English and which focuses on American politics might have peculiar properties. For this reason, we are pre-registering a small change in the protocol we intend to implement in University B. The instrument will measure whether the respondent is an international student or not, and the randomization will be blocked on this response. We intend to conduct analyses excluding the international students, and report their results separately if they seem to respond to the experiment differently than non-international students.

10. Pre-registration for Study 3b

Have any data been collected for this study already?

[The following text is one of three closed-ended response options, the other two being “Yes” and “No”:]
“It's complicated. We have already collected some data but explain in Question 8 [the Sample Size question] why readers may consider this a valid pre-registration nevertheless.”

Hypothesis.

This is the fifth in a series of related studies on the antecedents of implicit and explicit attitudes. At the time of writing, we have accumulated evidence that increased familiarity with an attitude object improves implicit (but not explicit) attitudes toward that object. We are seeking to replicate and extend one of our existing studies (Study 3 in our most recent manuscript draft, on mere familiarity effects), with three objectives in mind.

First, we seek to examine whether the same instrumentation leads to the same conclusions, thereby testing their reliability.

Second, we seek to examine the extent to which our existing conclusions might depend on the demographic composition of our initial study. Study 3 was conducted on student samples at two universities. It was diverse with respect to race, gender, and political orientation, but not age or education. Therefore, we seek to examine the extent to which treatment effects differ along these lines.

Third, we introduce a new treatment condition intended to extend the initial study. While Study 3 had only one non-target induction, the present study will have two: the one included in the initial study (focused on a male political candidate named Paul Coleman), plus a new one focused on a female candidate of approximately the same age (named Paula Coleman).

The new condition has two purposes. First, we seek to examine an attitude object that is more dissimilar from the target object (a politician named Mike Harper) than Paul Coleman was. Both Harper and Paul Coleman were white male politicians of approximately the same age, and this similarity might explain why exposure to Paul Coleman would improve implicit attitudes about Harper (Figure 5 in our most recent draft): the two objects might be related closely enough that positive affect toward one spills over to the other. In contrast, because, as a matter of social cognition, another person's sex is encoded automatically (Stangor et al. 1992), we expect such possible spillover effects to be smaller in the case of Paula Coleman. Thus, we expect exposure to a non-target female to improve implicit attitudes toward Mike Harper less than exposure to a non-target male. (And we expect induced familiarity with Harper to improve implicit attitudes about Harper most of all.) We do not hypothesize any effects on explicit attitudes about Mike Harper.

The second purpose of the additional treatment condition is to broaden the scope of this research. So far, the studies in this series have focused only on a generic white, male candidate. This was a design choice intended to create control across studies and help us isolate posited antecedents to implicit attitudes. But, it is eminently reasonable to examine how the relationships we have identified so far relate to other aspects of social cognition. The addition of this treatment arm is a (very initial) foray into such topics.

Dependent variable.

As in Study 3, there are two primary dependent variables: implicit attitudes toward Mike Harper (measured with an IAT), and explicit attitudes toward Mike Harper (the same question as in Study 3).

Conditions.

The four conditions for this study are fully described above.

Analyses

The main analysis is a difference-in-means by experimental condition, for each of the dependent variables mentioned above.

We also intend to examine possible heterogeneity in treatment effects by education, comparing people without a college degree to people with a college degree (and smaller groupings, if there are sufficient cases). We will also examine differences by age, with the expected groupings being people under 30, people between 30 and 50, and people over 50.

Secondary analyses

Because this study includes a manipulation of the sex of an evaluative target, we will also conduct analyses stratified by respondent gender, as a check for heterogeneous effects by gender. (For instance, a female evaluative target might generate more positive affect among male respondents than a male evaluative target.)

Anything else you would like to pre-register?

As is standard, we plan to code IAT scores as missing values for respondents for whom more than 10% of IAT trials are faster than 300 milliseconds. (These respondents will still be included in the analysis of explicit attitudes).

We also intend to take extra precautions in light of the fact that this study will be conducted via Amazon Mechanical Turk, wherein it is more difficult to monitor participant attention compliance than among university students. In some respects, MTurk respondents are impressive in terms of attention and care (Peer et al. 2014). But, social scientists have shown that a small-but-consequential proportion of MTurk responses are fraudulent (Kennedy et al. 2020). We will take steps to exclude such problematic responses. In particular:

- Participants will be asked in what state they reside, twice. First, they will be asked to choose the state from a drop-down menu. Later and on a separate screen, they will be asked to write the state in an open-ended text box. We will confirm that the two answers are the same.
- Participants will be asked the following question, “For you, what is the most important meal of the day, and why? Please write one sentence.” We will examine responses to this question to confirm that they 1) represent a coherent answer to the question, 2) do not appear to be plagiarized from some other website, and 3) show facility in written English.
- We will enable Qualtrics features to capture fraudulent responding via Recaptcha and RelevantID. We will exclude responses with Recaptcha scores below 0.5, as well as respondents with RelevantID scores above 30 (guidelines recommended by the software providers).

Respondents who fail any of these attention checks will be excluded from analysis. However, for transparency we will retain data collected from them, such that it can be reanalyzed should people wish to revisit our decisions. Additionally, we will report (in supporting materials) the open-ended responses that we excluded

(subject to redactions for privacy), so readers can evaluate the basis for our decisions. Insofar as time and funds allow, we will replace excluded responses with valid ones, up to our target N of 1,200.

Sample size

We intend to recruit 300 responses per condition, for a total of 1,200.

Note that above, we chose the option for “It’s complicated. We have already collected some data but explain in Question 8 why readers may consider this a valid pre-registration nevertheless.” Here is the explanation: We collected 15 responses (1.2% of the intended total) to confirm correct programing of our survey instrument. Otherwise, no data have been collected.

References

- Kennedy, Ryan et al. 2020. “The Shape of and Solutions to the MTurk Quality Crisis.” *Political Science Research and Methods* 8(4): 614–629.
- Peer, Eyal, Joachim Vosgerau, and Alessandro Acquisti. 2014. “Reputation as a Sufficient Condition for Data Quality on Amazon Mechanical Turk.” *Behavior Research Methods* 46: 1023–31.
- Stangor, Charles, Laure Lynch, Changming Duan, and Beth Glas. 1992. “Categorization of Individuals on the Basis of Multiple Social Features.” *Journal of Personality and Social Psychology* 62(2): 207.

11. Pre-registration for Study 4

Data collection. Have any data been collected for this study already? Note: 'Yes' is a discouraged answer for this preregistration form.

No, no data have been collected for this study yet.

Hypothesis

This study follows from several others we have conducted in a series of investigations on how political campaign messages influence implicit and explicit attitudes about candidates. As separate pre-registrations—beginning with registration #6198 on AsPredicted.org—document, our initial expectation in this project derived from the Associational-Propositional Evaluation (APE) Model. However, in our two initial studies, we found mixed support for this model. In particular, under our reading, the APE Model expects negative propositional information to have a reduced negative effect on implicit attitudes when preceded by positive associations (induced, in our studies, by a simulated television advertisement). As we documented in a 2019 American Political Science Association conference paper, these patterns did not consistently play out in two studies we ran.

However, we noticed that the first two studies we conducted are consistent with an alternative understanding of the antecedents of implicit attitudes. Where these studies manipulated candidate messaging without also providing negative propositional information—that is, in Study 1 and the No-scandal condition in Study 2—any advertisement for a political candidate (whether high- or low-quality) appeared to improve implicit attitudes toward the candidate. Elsewhere, we have hypothesized and tested that such a pattern can be explained by “mere exposure” effects, wherein almost any familiarity with an attitude object improves attitudes toward that object. (See [blinded] and Study 3 in a related APSA paper.)

The present pre-registration is for a revised iteration of Study 2. We are re-conducting Study 2 for three main reasons. First, trends in the initial run of Study 2 were part of what motivated the mere exposure study referenced above, but the trends were subtle and we want to be more confident in them. Second, as our conference draft discusses, the initial run had to contend with fraudulent responses on Amazon Mechanical Turk, and also found that our IAT measure is difficult to administer via MTurk. We wished to improve on this. Third, we are striving to accumulate an evidence base that relies on a consistent data source (college students).

Given the history above, we are approaching this iteration of Study 2 with different expectations than in the initial run. As in previous studies, we expect explicit attitudes to respond markedly to substantive information about a political candidate. Thus, explicit attitudes toward the candidate (Mike Harper) should be more favorable when respondents view an advertisement for Harper than when they view no ad. But more importantly, we expect that when there is no scandalous information about Harper, exposure to Harper advertisement—whether low- or high-quality—should improve implicit attitudes toward Harper. This expectation does not derive from the APE Model, but rather our own work on this topic.

It bears emphasis that we do not expect ad exposure to improve implicit attitudes toward Harper when the advertisement is followed by a scandalous news story. We had this expectation when we were adopting the perspective of the APE Model, but no longer. As our APSA paper discusses, we suspect that the effects scandalous information are not limited to propositional mental processes, and can influence (indeed overwhelm) the effects of mere familiarity. Thus, we expect the candidate advertisement manipulation’s effects on implicit attitudes to manifest primarily in the no-scandal condition.

More generally, we should acknowledge that the current study is not an exacting test of mere exposure effects. That was the purpose of the mere-exposure study referenced above, but it is not our purpose here. Rather, we are out to develop a more uniform basis of evidence (student samples), address a limitation (fraudulent responding) in the series of findings that motivated the mere-exposure study, and see if we can use the knowledge we have accumulated so far to make an out-of-sample prediction.

Dependent variable

The primary dependent variables are implicit and explicit attitudes toward Mike Harper, measured exactly as in previous studies.

Conditions. How many and which conditions will participants be assigned to?

The conditions are the same as in the study we are replicating, with one minor change. On reflection, we had some concern that the control story we used might have too enthusiastic a tone. (See materials.) We updated it to be a bland recipe for making cheese pastry, rather than a news story about making cheese. See this previous pre-registration for a fuller description of conditions: [blinded]

Analyses

As in previous studies, the primary analysis will be means of the two dependent variables, broken down by random assignment.

Outliers and Exclusions

No response

Sample Size

The precise sample size will be determined by student availability and participation, but we expect approximately 900 students from University A, and a maximum of 570 students from University B (though likely somewhat fewer based on past participation patterns).

Other

We expect to analyze international students at University B separately if they appear to respond to the experiment differently than non-international students. See our previous pre-registration for more details: [blinded]

12. Procedure to Ensure Data Integrity in Study 3b

As we discuss in the our pre-registration, we were concerned that a study conducted on Amazon Mechanical Turk would be afflicted by fraudulent responses (Kennedy et al. 2020; see also a discussion of this issue in Study 2 above). Our pre-registration for Study 3b lays out our plan for addressing this problem. Here, we report how our pre-registered exclusion plan unfolded.

We invited MTurk respondents to participate in our study via the CloudResearch platform. CloudResearch (formerly known as TurkPrime) is an interface that facilitates survey research conducted on MTurk (see Litman et al. 2017). Aside from some helpful programming tools, CloudResearch maintains a database of pre-approved MTurk workers who, in previous studies, provided high-quality (i.e. attentive) responses. We limited data collection to this pre-approved panel.

From there, we implemented our pre-registered exclusions as follows:

- We excluded respondents who provided nonsensical answers to our open-ended question about what is the meal they consider most important. There were only two such responses: “very important and enjoy this day” and “ghjghjhgj”.
- We excluded respondents who had Recaptcha scores below 0.5. There were six additional exclusions on this basis.
- We excluded respondents who had RelevantID Fraud scores above 30. There were 40 such exclusions.
- We excluded respondents who did not consistently report their state of residence during the study. There was only one additional exclusion on this basis—a respondent who reported living in the state “12065” in response to an open-ended question.

Additionally, we noticed that a small number of respondents began the study, stopped, and then began it again (in a new browser window). In all likelihood, these respondents had some unexpected interruption, such as a power outage or accidentally closing the browser window. In such cases, to preserve respondent naivete, we retain only the earliest response for which the respondent entered into one of the treatment assignments. This step is a small but justified departure from our pre-registration. It results in six additional exclusions.

It bears note that, in examining our data, we grew to think that the Recaptcha and Relevant ID exclusions above are overly conservative. (I.e., they result in the unnecessary exclusion of valid responses.) This assessment is based on reviewing open-ended responses for these respondents. The response appear to be perfectly cogent and valid. However, we did not wish to deviate from our pre-registered exclusion protocol on this basis.

References for this section:

- Kennedy, Ryan et al. 2020. “The Shape of and Solutions to the MTurk Quality Crisis.” *Political Science Research and Methods* 8(4): 614–29.
- Litman, Leib, Jonathan Robinson, and Tzvi Abberbock. 2017. “TurkPrime.Com: A Versatile Crowdsourcing Data Acquisition Platform for the Behavioral Sciences.” *Behavior research methods* 49(2): 433–42.

13. Adherence to Principles and Guidance for Human Subjects Research

Here, we discuss ethical considerations related to the human subjects research described above. We cover several matters emphasized in the American Political Science Association's [Principles and Guidance for Human Subjects Research](#).

Researcher responsibility. The APSA Guidelines state that “have an individual responsibility to consider the ethics of their research related activities and cannot outsource ethical reflection to review boards.” Although the research we conducted was approved by the relevant institutional review boards, the researchers take the responsibility for conducting ethical studies as our own. We thought carefully about ethical matters in all phases of the described research.

Potential for harm. We assess the potential for direct harm coming from participating in this research to be low. (It conforms to the “no more than minimal risk” standard.) The only techniques or measures used are unobtrusive survey-based measures. We considered whether any parts of our survey instruments could induce anxiety or feelings of trauma for participants, and did not identify any for which this was likely. (We discuss potential harm related to breaches of confidentiality below.)

Potential benefits. We also considered any potential benefits that might redound to study participants (aside from any compensation related to participation). The benefits are small, though it is possible that being induced to think about their own political views and psychological processes could be a fulfilling reflection task.

Privacy and Confidentiality. We took steps to address risks related to breaches of confidentiality. First, we considered if any information we collected would be harmful if disclosed. This seemed unlikely, as the studies we report focus on opinions about a fictional political candidate. Nevertheless, we collected data anonymously: individually identifiable participant information was not part of the data collection. We also considered the possibility of deductive identification of respondents. This too seemed unlikely, since we asked about only a few pieces of core demographic information.

Informed consent. All the studies described above included an informed consent disclosure that stated the true objectives of the proposed research, as well as compensation. The consent forms also covered all considerations standardly required by governmental oversight. The consent clarified, per IRB guidelines, the precise process of the studies.

Deception. The studies we report involved did not involve deception.

Power. Because some study participants are undergraduate students at the investigators' institutions, we considered the possible power relationships carefully and took steps to address them. These steps depended on which of the two universities students came from.

At University A, students were recruited via various courses' “research activity requirements.” Although students were recruited in this way, no student was required to participate: there were alternative (and comparably onerous) ways to fulfill the course requirement. Specifically, students could avoid participating in studies for the entire semester by writing a short research paper under the guidance of a faculty member (who was not an investigator). Second, the investigator did not receive information on which students participated in the study.

At University B, participants received extra credit for participation. In the case that they did not want to take a study, they were offered an alternative task through which they could earn the same amount of extra credit.

Just as in University A, the investigator was not the instructor of record for any of the students who took part in the study. The extra credit amount is standardized and approved by the IRB based on the length of participation.

Power differentials are less evident for the MTurk studies we report herein. MTurk is an opt-in participant pool, and investigators would not normally have any personal relationship with the study participants.

Compensation. Student participations was facilitated by a course requirement, as described above. For MTurk respondents, compensation was set to be at a level consistent with the minimum wage in the U.S. state where one of the investigators resides. A response field at the very end of MTurk study allowed participants to submit feedback on the study. We inspected this feedback for complaints about compensation. Less than 1% of participants registered any complaints about the level of compensation.

Impact. We considered the possible beneficial impacts of this research. While they are abstract, we believe they have the potential to deepen understanding of persuasion and political campaigning—and thereby help elucidate ways to bring about a better political system and world.

Laws and regulations. The research herein conformed with all local, state, and national laws.