Errata

On Baxter-King, Coppock, Straus and Vavreck (2025), *PNAS Nexus* 4(6)

The maintained replication code behind these errata is at https://github.com/active-maintenance-aec/baxter-king_etal_2025 and is archived at Zenodo at https://doi.org/10.5281/zenodo.24168270, a DOI that always resolves to the latest version.

This note lists 44 corrections, and four items it flags without correcting, in "Endorsements vs. information: experimental evidence of backlash and parallel persuasion during the COVID-19 public health crisis," published in *PNAS Nexus* 4(6):pgaf185, doi 10.1093/pnasnexus/pgaf185, and its online appendix. **None of them changes a conclusion of the article.** Two of the corrections concern claims the article makes. On p. 4 the article says that the sign and significance of the interaction terms match those of the differences-in-CATEs in all cases but one, and that equivalence can be affirmed at 10 points in two of the four mandate vignettes. By the printed estimates the first count is four, not one (entry 5), and under the correct standard error for the difference the second is one, not two (entry 6), because the deposit's code mistypes that standard error. Two more concern what a printed figure and an appendix table show: the series labels of Fig. 2 are swapped (entry 25), and the G-1 and G-3 rows of the appendix's guidance equivalence table carry each other's numbers (entry 26). No estimate in the article's main tables or figures changes value, so the Discussion's findings about partisan endorsers, expert guidance, and nonpartisan information stand. The appendix's equivalence tables print standard errors and equivalence p-values from the mistyped formula, and they are not reprinted here.

Four corrections are counts of respondents or arms: the headline number of respondents and the size of survey wave five (entries 1 and 2), the number of estimates in the appendix's comparison of weighted and unweighted models (entry 3), and one arm of a CONSORT diagram (entry 4). The rest are descriptions of stimuli, samples, and outcomes that do not match the appendix, labels and captions, cross-references to appendix figures and tables, wording and spelling in the article and the appendix, and details of the reference list. The four items without a correction are cases where the page or the deposit cannot say what the right value is, or where a repair would rewrite the authors' sentence (entries 7, 12, 18, and 48).

All were found by reading the article and its appendix word by word, together with the maintained replication, which recomputes every number the article prints from the deposited replication archive (osf.io/bwv3c) and checks each against the printed value. Each quotation was checked against an image of the published page, and the references were checked against

the Crossref record named in each entry's note. The counts in entries 1 and 6 are computed from the maintained output at the time this document is rendered. The code supporting the maintained analysis is at github.com/active-maintenance-aec/baxter-king_etal_2025.

1. The headline count of survey respondents

Published, p. 1, abstract; repeated in the Significance Statement on p. 1 and the Introduction on p. 2:

> Using data from 10 experiments with **85,191** survey respondents conducted over a 2-year period during the COVID-19 pandemic, we assess the effectiveness of these three types of persuasive messages.

Corrected:

> Using data from 10 experiments with **96,567** survey respondents conducted over a 2-year period during the COVID-19 pandemic, we assess the effectiveness of these three types of persuasive messages.

The same figure appears three times: here, in the Significance Statement ("Ten experiments conducted between 2020 and 2022 with 85,191 respondents"), and in the Introduction ("conducted among 85,191 survey respondents interviewed between 2020 and 2022"). The ten deposited analysis files, one per experiment in Table 1, hold 127,439 rows and 96,567 distinct respondent identifiers. Every identifier carries its wave as a prefix and no identifier appears in two waves, so the overlap is entirely within waves, whose experiments draw on the same respondents: the mandate vignettes and the first mask-guidance experiment share wave six, the two conversation experiments and the second mask-guidance experiment share wave seven, and the three booster experiments share wave eight. The analysis files run 0.1 to 0.7 percent short of the enrollment Table 1 reports, so the count of enrolled respondents is larger still: Table 1's own N column, with the nested experiments counted once as the appendix's CONSORT diagrams show them, gives at least 96,840. No sum of the N figures the article prints produces 85,191, and the deposit contains no file from which it could have been counted. The figure given here is the deposit's.

2. Wave five holds half the respondents the Materials and methods give it

Published, p. 2, Materials and methods:

> Waves five through eight, which consist of 30,000 respondents each, were conducted 2021 March 25 to April 13, 2021 June 17 to July 6, 2021 September 3 to October 4, and 2022 October 24 to December 20.

Corrected:

> **Wave five, which consists of 15,000 respondents, and waves six** through eight, which consist of 30,000 respondents each, were conducted 2021 March 25 to April 13, 2021 June 17 to July 6, 2021 September 3 to October 4, and 2022 October 24 to December 20.

The deposit ships experiment subsets rather than waves, but three of its wave-five files (the masking, sick-leave, and incentive experiments, none of them in this article) each hold 14,557 respondents, which is the N the appendix's CONSORT diagram gives for the wave in which E-2 was fielded. The largest deposited files from waves six, seven, and eight hold 30,857, 32,943, and 29,448 respondents against CONSORT Ns of 30,857, 33,088, and 29,448, so those three waves are the 30,000 the sentence describes and wave five is not. The corrected sentence rounds as the article does, which gives wave three's 14,946 enrolled respondents as 15,000 and wave eight's 29,448 as 30,000. The field dates are unchanged.

3. The count of matching weighted and unweighted estimates is 51 of 60, and Figure S46 holds 52 estimates

Published, app. p. 42, Section E.1:

> That said, the **signs and significances** of our estimates match in most (**51 of 60**) cases.

Corrected:

> That said, the **significances** of our estimates match in most (**45 of 52**) cases.

Figure S46 plots 52 estimates, the average treatment effect and the Treatment x Party ID interaction for each of 26 experiment arms. The next paragraph lists seven that differ in significance, five where only the unweighted estimate is significant and two where only the weighted estimate is, so 52 less 7 gives 45. A further six differ in sign without differing in significance, so signs and significances together match in 39 of 52. Both counts come from refitting the weighted and unweighted models on the deposited data. The corrected sentence keeps the paragraph's own criterion, significance, and the sign count is given here.

4. A control-arm count in the I-1 diagram repeats the Friend arm's total

Published, app. p. 13, I-1 diagram, the Doctor control box:

> N: 1,426/**1,488**

Corrected:

> N: 1,426/**1,436**

The Doctor group has N: 2,958, and its treatment box prints 1,518/1,522, so the control box must be 1,436 at the larger count. The other groups balance the same way (Friend 1,488 + 1,439 = 2,927 and CDC 1,416 + 1,409 = 2,825), and the three groups sum to the 8,710 unvaccinated respondents. The Friend control box, directly beside it, prints 1,479/1,488.

5. The interaction terms and the differences-in-CATEs agree in significance in all but four cases, not one

Published, p. 4, Materials and methods:

> the **sign and significance** of the interaction terms in the main text **match** the **sign and significance** of the difference-in-CATEs in all cases but **one**.

Corrected:

> the **significance** of the interaction terms in the main text **matches** the **significance** of the difference-in-CATEs in all cases but **four**.

Counted on the stars the appendix prints at the 5 percent level, the interaction term and the difference-in-CATEs disagree in significance in four of 34 comparisons: E-1 Trump + Fauci and E-1 Spiritual leader (Table S1 against appendix p. 45), where the difference-in-CATEs is significant and the interaction term is not, G-2 Trip (Friend), where the same holds (Table S3 against p. 47), and G-2 Restaurant (Self), where the interaction term is significant and the difference-in-CATEs is not. The 34 comparisons are the 7 E-1 endorsers, the 8 E-2 endorsers, G-1, G-3, the 5 information experiments, and the G-2 vignettes pooled, in the friend arm, and in the solo arm (4 each). The sign also differs, by the printed estimates, in eight further cells in which neither estimate is significant: E-1 Pharmacy, E-2 Fauci, E-2 Obama, E-2 James, G-1, G-2 Concert (Friend), G-2 Trip (Self), and G-2 Restaurant (pooled). The corrected sentence applies one rule, significance at the 5 percent level, and the sign differences are reported here. For the 26 arms the maintained replication refits, the corrected standard error described in entry 6 changes none of the printed stars.

6. Only one of the four G-2 vignettes affirms equivalence at 10 points

Published, p. 4, Materials and methods:

> G-2 is more nuanced; we can affirm equivalence at 10 points in **two cases** but cannot in **two cases**.

Corrected:

> G-2 is more nuanced; we can affirm equivalence at 10 points in **one case** but cannot in **three cases**.

The deposit's dic_estimator() in code/helpers.R computes the standard error of the difference-in-CATEs as sqrt(std.error_2^2 + std.error_2^2), squaring the Republican standard error twice and never using the Democratic one. The correct formula is sqrt(std.error_1^2 + std.error_2^2). Under the corrected standard error the pooled Restaurant vignette affirms equivalence at 10 points (p = 0.043) and the Concert vignette does not (p = 0.116; the appendix prints 0.038), nor do Team (0.791) and Trip (0.324). The same error affects the standard errors and the equivalence p-values in all of the appendix's equivalence tables (pp. 45, 47, and 62), so the values printed there differ from the corrected ones, for example the I-1 p-value at 10 points, printed as 0.006 and 0.046 when corrected. No other count in the article changes: both guidance experiments affirm equivalence at 5 points, and three of the five information experiments affirm it at 10 points, under either standard error.

7. Two exclusion counts in the I-3 diagram print the same value

Published, app. p. 15, I-3 diagram:

> Yes; Exclude N: 4,380 … Not sure (I got a booster shot since August 31, but don't know if it was the new one or the original); Exclude N: 4,380

The diagram starts from N: 29,448, excludes 12,951 respondents who are not fully vaccinated, and includes 10,700 who answered No, which leaves 5,797 for the Yes and Not sure boxes together. Two boxes of 4,380 sum to 8,760, so at least one is wrong, and the page cannot say which. If Yes is right, Not sure is 1,417, and if Not sure is right, Yes is 1,417. The deposit does not settle it. The I-3 analysis file holds the 10,700 respondents who have an assignment, and the other 18,748 rows are 12,951 plus 5,797, but the file does not record the screener answer, so it cannot split the 5,797.

8. Appendix F.4 swaps the control and treatment labels on the I-2 stimuli

Published, app. p. 67, Section F.4:

> [**CONTROL**] **Imagine** you're having a conversation with your doctor about the way businesses in your neighborhood are reacting to the spread of the highly contagious Delta variant of the virus. … [**TREATMENT**] **Imagine** you're having a conversation with your doctor about the way businesses in your neighborhood are reacting to the state of the virus.

Corrected:

> [**TREATMENT**] **Imagine** you're having a conversation with your doctor about the way businesses in your neighborhood are reacting to the spread of the highly contagious Delta variant of the virus. … [**CONTROL**] **Imagine** you're having a conversation with your doctor about the way businesses in your neighborhood are reacting to the state of the virus.

Only the two labels move. The text that mentions the Delta variant is the treatment, and the text that does not is the control. Main p. 8 says so in prose ("Respondents assigned to control were told that their doctor was urging them to get the vaccine. Treatment respondents were told to imagine they were having a conversation with their doctor in which the doctor was urging them to get the vaccine and that their doctor was emphasizing the increased contagiousness of the Delta variant of the virus"), and the I-2 diagram on appendix p. 14 labels the version without the Delta text "Control" and the next box "Treat; Delta". Section F.4 disagrees with both. The deposit records which arm each respondent was in but not the text shown, so it cannot settle the question independently.

9. The party shares among unvaccinated E-2 respondents do not say where leaners are counted

Published, p. 5, Results:

> Among the 7,249 unvaccinated people we interviewed in Spring of 2021, **79%** of the Democrats said they were likely to get the vaccine, whereas only 45% of Republicans said this. Independents were in the middle at 50%.

Corrected:

> Among the 7,249 unvaccinated people we interviewed in Spring of 2021, **counting leaners as independents, 79%** of the Democrats said they were likely to get the vaccine, whereas only 45% of Republicans said this. Independents were in the middle at 50%.

In the deposited E-2 analysis file the weighted shares likely to vaccinate are 79.1 percent of Democrats, 49.6 percent of independents, and 44.5 percent of Republicans when leaners are counted as independents, which reproduces the printed 79, 50, and 45. Page 4 defines partisans as including leaners. Counted that way, the shares are 77.1, 45.7, and 44.1 percent, and independents are no longer in the middle. The correction names the grouping the printed figures use.

10. The I-5 outcome categories differ from those described above them

Published, p. 9, Information experiments:

> Outcome categories **were** identical to those detailed above.

Corrected:

> Outcome categories **in I-4 were** identical to those detailed above.

I-3 and I-4 use a five-point scale running from "I will definitely get it" to "I will not get it" (appendix pp. 15, 16, and 69). I-5 asks "How likely are you to vaccinate the eligible, unvaccinated children in your household for COVID-19?" with the options Extremely likely, Very likely, Somewhat likely, Not at all likely, and Not sure, the first three coded 1 and the last two coded 0 (appendix pp. 17 and 69).

11. The unboosted sample is defined by a bivalent booster, and the text says booster

Published, p. 3, Table 1 note c, and p. 8:

> Unboosted means the respondent indicated they have completed a full course of a primary vaccine but have not yet received a **booster**. … The (I-3) Bivalent booster experiment was fielded on people who had received a full vaccine course but had not yet received a **booster shot**.

Corrected:

> Unboosted means the respondent indicated they have completed a full course of a primary vaccine but have not yet received a **bivalent booster**. … The (I-3) Bivalent booster experiment was fielded on people who had received a full vaccine course but had not yet received a **bivalent booster shot**.

The screener on appendix p. 15 asks "Have you received a dose of a"bivalent" COVID-19 booster since August 31st, 2022?" and includes the respondents who answer No (N: 10,700). A respondent who had received an earlier booster is therefore in the sample.

12. I-1 and I-2 are described as doctor's office vignettes that give information about vaccine effectiveness

Published, p. 8:

> In September 2021, we conducted two vignette experiments that placed people in an imaginary doctor's office and gave respondents information about the contagiousness of the recent variant of the virus and the vaccine's effectiveness in preventing serious illness (I-1 and I-2).

I-1 places the respondent in a conversation with a friend, a doctor, or the CDC (appendix p. 13), and I-2 in a conversation with the doctor (appendix p. 14). Neither gives information about the vaccine's effectiveness: I-1 gives information about contagiousness, and I-2 about the Delta variant, with the doctor adding that the risks are minimal. The sentence is not corrected here, because a repair would rewrite it rather than change a word.

13. The appendix says only half of the unvaccinated respondents were assigned to the mandate experiment

Published, app. p. 63, Section F.1:

> **Only half** of the unvaccinated survey respondents were assigned to these experimental conditions.

Corrected:

> **All** of the unvaccinated survey respondents were assigned to these experimental conditions.

The G-2 diagram on appendix p. 11 includes all 10,298 unvaccinated respondents, 5,153 in the friend vignettes and 5,145 in the solo vignettes, and Table 1 gives N = 10,298 for G-2.

14. The appendix's activity-mandate response options say buy something else where the experiment says stay home

Published, app. p. 64, Section F.1:

> I would probably not get vaccinated and **buy something else**. … I would definitely not get vaccinated and **buy something else**.

Corrected:

> I would probably not get vaccinated and **stay home**. … I would definitely not get vaccinated and **stay home**.

Main p. 7 and the response options on appendix p. 11 both read "stay home" for these two options.

15. The vignette quoted on p. 7 omits the sentence about the friend's vaccination status

Published, p. 7, Results, the friend vignette:

> … proof of a COVID-19 vaccination ['is' or 'is NOT'] required to enter the venue. **Which** of the following best describes what you would do in this situation?"

Corrected:

> … proof of a COVID-19 vaccination ['is' or 'is NOT'] required to enter the venue. **You believe your friend has been vaccinated. Which** of the following best describes what you would do in this situation?"

Appendix pp. 11, 63, and 64 give the friend vignettes with the sentence "You believe your friend has been vaccinated." before the question.

16. Table 1 note a does not describe the three-option outcomes or the coding of G-1

Published, p. 3, Table 1, note a:

> We convert **4- or 5-point Likert scales** into binary measures, with the **prosocial option** coded as the higher value.

Corrected:

> We convert **three-option scales and 4- or 5-point Likert scales** into binary measures, with the **prosocial option (in the guidance experiments, the option the CDC recommended)** coded as the higher value.

The guidance experiments G-1 and G-3 have three response options (p. 6), and appendix p. 46 says each is collapsed to a binary outcome "based on which recommendation is CDC-backed at the time". In June (G-1) that is the option under which vaccinated people need not wear masks, so the coded option is the CDC-backed one rather than the prosocial one.

17. The I-4 diagram dates the booster from August 31, where its question says October 12

Published, app. p. 16, I-4 diagram:

> Not sure (They got a booster shot since **August 31**, but don't know if it was the new one or the original);

Corrected:

> Not sure (They got a booster shot since **October 12**, but don't know if it was the new one or the original);

The question above the boxes asks about a bivalent booster "since October 12, 2022", and the survey instrument on appendix p. 68 uses October 12 in the children's question. The adult question (I-3, appendix p. 15) correctly uses August 31.

18. The field dates in the Materials and methods and the appendix differ by a day

Published, p. 2, Materials and methods, and app. pp. 9, 12, and 13:

> conducted 2021 March 25 to April 13, 2021 June 17 to July 6, 2021 September 3 to October 4, and 2022 October 24 to December 20. … Dates Fielded: 03/24/2021 - 04/14/2021 … Dates Fielded: 09/03/2021 - 10/05/2021

Appendix p. 9 gives the E-2 field dates as March 24 to April 14, 2021, a day earlier at the start and a day later at the end than the Materials and methods' March 25 to April 13. Appendix pp. 12 and 13 give the September wave (G-3 and I-2) as September 3 to October 5, a day later at the end than October 4. None of the deposited analysis files has a date, time, or start or end field, so the deposit cannot say which set of dates is right. Entry 2 leaves the dates in the Materials and methods as printed.

19. Table 1 describes the I-3 outcome as a children's outcome

Published, p. 3, Table 1, row I-3, Outcome column:

> Likelihood of **children getting** the **Bivalent** booster

Corrected:

> Likelihood of **getting** the **bivalent** booster

I-3 is the adult experiment (Sample: Unboosted, N 10,700). The children's experiment is I-4, whose row reads "Likelihood of children getting the bivalent booster", with a lowercase b.

20. The Fig. 4 caption calls the mandates endorsements

Published, p. 8, Figure 4 caption:

> … and the coefficient on the treatment × party interaction term describes how much more effective the **endorsement** is for subjects who are more Republican on the 1–7 party identification scale.

Corrected:

> … and the coefficient on the treatment × party interaction term describes how much more effective the **mandate** is for subjects who are more Republican on the 1–7 party identification scale.

Fig. 4 reports the activity-specific mandate vignettes (experiment G-2). The caption of Fig. 2, which reports the endorsement experiments, contains the same sentence.

21. The Fig. 5 caption names an information treatment that I-3 does not give

Published, p. 9, Figure 5 caption:

> Fig. 5. Effects of information about **contagiousness** on intentions to get booster shot (October 2022, experiment I-3)

Corrected:

Fig. 5. Effects of information about **a winter surge and the effectiveness of boosters** on intentions to get booster shot (October 2022, experiment I-3)

The I-3 treatment, reproduced on appendix p. 15, warns that another COVID-19 surge is likely this winter and reports that the CDC says vaccines and boosters are the best protection.

22. The pointer to the endorsement regression tables includes Table S3

Published, p. 5, Results:

> Corresponding regression tables are presented in Appendix Tables S1–**S3**.

Corrected:

> Corresponding regression tables are presented in Appendix Tables S1–**S2**.

Tables S1 and S2 report the endorsement experiments E-1 and E-2, which the paragraph describes. Table S3 reports the mandate experiment G-2.

23. The appendix figures for respondents predicted to be unvaccinated are described as being for those predicted to be vaccinated, and cited from S23

Published, p. 4, Materials and methods:

> … estimates among the "predicted to be **vaccinated**" (Appendix Figs. **S23**–S45), we also present estimates of conditional average treatment effects

Corrected:

> … estimates among the "predicted to be **unvaccinated**" (Appendix Figs. **S24**–S45), we also present estimates of conditional average treatment effects

Figs. S24 to S45 are titled "among never vaxxers", and the same page's later sentence correctly speaks of the "predicted to be un-vaccinated at the final wave" group. Fig. S23 is the variable-importance plot for the random forest that makes the prediction. The comparison of weighted and unweighted estimates is Fig. S46.

24. The differences-in-CATEs for all experiments are cited to figures that show other estimates

Published, p. 4, Materials and methods:

> … and the corresponding differences-in-CATEs for all experiments (Appendix **Figs. S48–S61**).

Corrected:

> … and the corresponding differences-in-CATEs for all experiments (Appendix **pp. 45, 47, and 62**).

Figs. S48 to S61 plot the G-2 vignette effects and the I-1 source comparisons. The CATEs and differences-in-CATEs for every experiment are in the unnumbered tables on appendix p. 45 (E-1 and E-2), p. 47 (G-1, G-2, and G-3), and p. 62 (I-1 to I-5).

25. Figure 2 plots each framing's estimates under the other framing's label

Published, p. 6, Figure 2, October 2020 panel, legend, and p. 5, Results:

> Legend: Pooled, **Social** (orange triangles), **Personal** (teal circles). … Estimates from the **personal arms** are plotted with **circles** while estimates from the **social arm** that included additional information that the vaccine would protect other people are plotted with **triangles**.

Corrected:

> Legend: Pooled, **Personal** (orange triangles), **Social** (teal circles). … Estimates from the **social arm** that included additional information that the vaccine would protect other people are plotted with **circles** while estimates from the **personal arms** are plotted with **triangles**.

The deposit's code/analysis/run_models_w3_endorse_by_arm.R fits the models named *_prosocial on experiment_arm == "A (Personal)" and the models named *_not_prosocial on experiment_arm == "B (Social)". code/figures/figure_2.R then relabels "Prosocial" as "Social" and "Not Prosocial" as "Personal", so each framing is drawn under the other's label. The arm variable agrees with the appendix: Arm A (Personal) on appendix p. 8 is the version without the sentence about protecting others, and its Trump and Fauci cells (N: 937/940 and 879/882) are the counts of the deposit's A (Personal) rows, 937 and 879 treated respondents. The maintained replication labels each series by the arm it was fitted on. The article's prose about Fig. 2 describes the series by the printed labels and so carries the same swap, which the corrected sentence reverses. No printed sentence changes direction: the text says only that the Personal versus Social variation does not matter (joint test $p = 0.37$), and the estimates themselves are unchanged.

26. The G-1 and G-3 rows of the appendix's guidance equivalence table carry each other's numbers

Published, app. p. 47, Guidance Experiments CATEs, Difference-in-CATEs, and Equivalence Tests:

> **G-1**: CDC Mask Guidance 1: **0.015 (0.009)**, **0.028 (0.013) ***, **0.013 (0.018)**, **0.022 ***, **0.000 ***. **G-3**: CDC Mask Guidance 2: **0.054 (0.011) ***, **0.064 (0.012) ***, **0.010 (0.016)**, **0.007 ***, **0.000 ***.

Corrected:

> G-1: CDC Mask Guidance 1: **0.054 (0.011) ***, **0.064 (0.012) ***, **0.010 (0.016)**, **0.005 ***, **0.000 ***. G-3: CDC Mask Guidance 2: **0.015 (0.009)**, **0.028 (0.013) ***, **0.013 (0.016)**, **0.011 ***, **0.000 ***.

The deposit's code/tables/equiv_test_guidance.R maps the label "G-1: CDC Mask Guidance 1" to the model object w7_fit_cdcmask and "G-3: CDC Mask Guidance 2" to w6_fit_cdcmask. G-1 is the June 2021 experiment, whose analysis file is w6, and G-3 is the September 2021 experiment, whose file is w7, so the two rows carry each other's numbers. Two errors are separate here. The row swap is one. The other is the mistyped standard error of entry 6, which affects the printed difference-in-CATEs standard errors and the equivalence p-values in both rows. The corrected row values given above fix both, and they are the maintained replication's. Correcting only the swap would leave the G-1 row at 0.010 (0.016), 0.007 *, 0.000* and the G-3 row at 0.013 (0.018), 0.022 *, 0.000* . Both experiments affirm equivalence at 5 and 10 points under any of these values, so the claim on p. 4 is unchanged, and so are the counts in entries 5 and 6.

27. Tables 2 and 3 print thousands separators as decimal points in the observation counts

Published, p. 7, Table 2, and p. 9, Table 3:

> Num.Obs. **30.751** | **32.933** … Num.Obs. **8.659** | **8.690** | **10.677** | **1.626** | **1.711**

Corrected:

> Num.Obs. **30,751** | **32,933** … Num.Obs. **8,659** | **8,690** | **10,677** | **1,626** | **1,711**

The counts are 30,751 and 32,933 in Table 2, and 8,659, 8,690, 10,677, 1,626, and 1,711 in Table 3. All seven match the regression tables the maintained replication fits on the deposited data (maintained/output/table_2_guidance_experiments.csv and table_3_information_experiments.csv).

28. A transposed pair of words and a stray word in the introduction

Published, p. 2, Introduction:

> To be clear, **we while** find that some information treatments have effects that are positive for all subgroups … endorsements by others **either often** had null or even negative effects on vaccine intentions

Corrected:

> To be clear, **while we** find that some information treatments have effects that are positive for all subgroups … endorsements by others **often** had null or even negative effects on vaccine intentions

In the first sentence the clause that follows ("we also find that other information treatments have effects that are close to zero") completes the "while" construction. In the second, "either" has no clause to pair with.

29. The first experiments' endorsers are described as political figures and political figures

Published, p. 2, Introduction:

> As we will show, our first experiments on endorsements by **political figures** and political figures underlined the dangers associated with that approach

Corrected:

> As we will show, our first experiments on endorsements by **public health officials** and political figures underlined the dangers associated with that approach

The same noun phrase is printed twice. The endorsers in the first experiments, listed on p. 4, include President Trump, Dr. Fauci, a personal physician, a pharmacy, and a health insurance company.

30. A missing article and a misspelling on p. 4

Published, p. 4, Materials and methods:

> **In E-1 set** of endorsers included (i) their health insurance company, … we have not masked important **nonlinearites** (14).

Corrected:

> **In E-1, the set** of endorsers included (i) their health insurance company, … we have not masked important **nonlinearities** (14).

31. A verb that does not agree with its subject on p. 6

Published, p. 6, Results:

> The panels **looks** different, of course, as everyone moves in the same direction and by roughly the same amounts

Corrected:

> The panels **look** different, of course, as everyone moves in the same direction and by roughly the same amounts

32. A verb that does not agree with its subject and an unmatched parenthesis on p. 7

Published, p. 7, Results:

> baseline levels of willingness to get vaccinated in order to participate in these activities **was** low, around 20% … did not consistently vary depending on whether the activity was for the respondent or their friend**).**

Corrected:

> baseline levels of willingness to get vaccinated in order to participate in these activities **were** low, around 20% … did not consistently vary depending on whether the activity was for the respondent or their friend**.**

The subject of the first sentence is "levels". In the second, no parenthesis is open when the closing one is printed.

33. A missing preposition in the Discussion

Published, p. 10, Discussion:

> Even so, we find just one instance **cues** working as intended across all audiences (the E-1 Fauci endorsement).

Corrected:

> Even so, we find just one instance **of cues** working as intended across all audiences (the E-1 Fauci endorsement).

34. Dr. Fauci's title is misspelled in the E-2 diagram, and his institute is misnamed

Published, app. pp. 8 and 9:

> Dr. Anthony Fauci (**Dirctor** of the U.S. National **Isnstitute** of Allergy and Infectious **Disease**), and other scientific sources

Corrected:

> Dr. Anthony Fauci (**Director** of the U.S. National **Institute** of Allergy and Infectious **Diseases**), and other scientific sources

Page 9 prints the string with "Dirctor" and "Isnstitute" in three boxes. Page 8 spells those two words correctly and prints the same "Infectious Disease". The agency is the National Institute of Allergy and Infectious Diseases.

35. A misspelled label in the G-2 diagram

Published, app. p. 11:

> Solo **Concrert** Control:

Corrected:

> Solo **Concert** Control:

36. A comparison in the I-1 CDC treatment box reads as contagious than

Published, app. p. 13, I-1 diagram, the CDC treatment box:

> that it is as contagious **than** the Chicken Pox.

Corrected:

> that it is as contagious **as** the Chicken Pox.

The Friend and Doctor treatment boxes on the same page read "as contagious as Chicken Pox".

37. Vaccine is misspelled in two I-5 figure titles

Published, app. p. 29, Figure S22, and p. 41, Figure S45:

> Figure S22: Information experiment I-5: **Vaccince** - Child … Figure S45: Information experiment I-5: **Vaccince** - Child among "never vaxxers"

Corrected:

> Figure S22: Information experiment I-5: **Vaccine** - Child … Figure S45: Information experiment I-5: **Vaccine** - Child among "never vaxxers"

38. Children's is misspelled in a heading

Published, app. p. 1, the contents, and p. 69, Section F.5.3:

> F.5.3 Step 3: Measure Intentions Towards **Childrens'** Vaccinations

Corrected:

> F.5.3 Step 3: Measure Intentions Towards **Children's** Vaccinations

39. The Concert prompt in the appendix drops a word

Published, app. p. 64, Section F.1:

> … ['Because' or 'Even though'] **will be** lots of people together, proof of a COVID-19 vaccination ['is' or 'is NOT'] required to enter the venue.

Corrected:

> … ['Because' or 'Even though'] **there will be** lots of people together, proof of a COVID-19 vaccination ['is' or 'is NOT'] required to enter the venue.

The Team and Restaurant prompts below it, and main p. 7, read "there will be lots of people together".

40. A sentence in the appendix's weighting discussion is garbled

Published, app. p. 42, Section E.1:

> Indeed** that it is** what we find: the weighted standard errors are on average 1.6 times their unweighted counterparts.

Corrected:

> Indeed**, that is** what we find: the weighted standard errors are on average 1.6 times their unweighted counterparts.

41. SARS-CoV-2 and the 1980s are misprinted on the first page

Published, p. 1, Introduction:

> In 2020, during the **SARS-CoV2** COVID-19 pandemic … antidrug messaging from the **1980's** and concluded there was "no good evidence" any of it changed behavior

Corrected:

> In 2020, during the **SARS-CoV-2** COVID-19 pandemic … antidrug messaging from the **1980s** and concluded there was "no good evidence" any of it changed behavior

42. The appendix spells LeBron James as Lebron

Published, app. pp. 5, 9, 25, 37, and 72:

> **Lebron** James

Corrected:

> **LeBron** James

Main p. 4 spells the name correctly ("NBA star LeBron James"). The appendix misspells it in the Table S2 header (p. 5), the E-2 diagram (p. 9), and the E-2 figures (pp. 25, 37, and 72).

43. Fig. S47 prints a p-value of exactly zero

Published, app. p. 46, Figure S47, panel "Less restrictive (G-1, June 2021)":

> $2 = 111.78$, $\mathbf{p = 0}$

Corrected:

> $2 = 111.78$, $\mathbf{p < .001}$

A p-value is not exactly zero. A chi-square of 111.78 on two degrees of freedom has a p-value far below .001. The panel below it prints "p = .016" in the same format.

44. The y-axis label of 29 appendix figures ends in a stray bracket

Published, app. pp. 19 to 25 and 31 to 38, Figures S2 to S15 and S24 to S38, left panel, y-axis label:

> to get the vaccine = 1, 0 **otherwise**]

Corrected:

> to get the vaccine = 1, 0 **otherwise**

The label opens no bracket. Figure S1 (appendix p. 18) prints the same label without it.

The reference list

45. Two references run words together

Published, p. 11, refs, references 7 and 8:

> 7 Coppock A. **2022.Persuasion** in parallel: how information changes minds about politics. … 8 Tappin BM, Berinsky AJ, Rand DG. 2023. **Partisans'receptivity** to persuasive messaging is undiminished by countervailing party leader cues.

Corrected:

> 7 Coppock A. **2022. Persuasion** in parallel: how information changes minds about politics. … 8 Tappin BM, Berinsky AJ, Rand DG. 2023. **Partisans' receptivity** to persuasive messaging is undiminished by countervailing party leader cues.

Crossref, doi 10.1038/s41562-023-01551-7, gives the Tappin, Berinsky, and Rand year as 2023, volume 7, issue 4, pages 568 to 582, as printed. Crossref's record of Persuasion in Parallel (doi 10.7208/chicago/9780226821832.001.0001) gives 2022 as the print year, which agrees with the printed year, so the year in reference 7 is not corrected.

46. The Nationscape reference lowercases the project name and misplaces a space

Published, p. 11, refs, reference 11:

> 11 Holliday D, et al. 2021. Democracy **fund+ UCLA nationscape** methodology and representativeness assessment.

Corrected:

11 Holliday D, et al. 2021. Democracy **Fund + UCLA Nationscape** methodology and representativeness assessment.

The project is Democracy Fund + UCLA Nationscape. The printed title lowercases “Fund” and “Nationscape” and omits the space before the plus sign.

47. The Lin reference lowercases Freedman’s name

Published, p. 11, refs, reference 15:

15 Lin W. 2013. Agnostic notes on regression adjustments to experimental data: reexamining **freedman’s** critique. Ann Appl Stat. 7(1):295–318.

Corrected:

15 Lin W. 2013. Agnostic notes on regression adjustments to experimental data: reexamining **Freedman’s** critique. Ann Appl Stat. 7(1):295–318.

Crossref, doi 10.1214/12-AOAS583, gives volume 7, issue 1, pages 295 to 318, issued March 2013, as printed. The lowercase “reexamining” after the colon follows the journal’s sentence-case style, so only the surname is corrected.

48. Reference 1 names no outlet

Published, p. 11, refs, reference 1:

1 Balonon-Rosen P. 2019. From cringeworthy to scary: a history of anti-drug PSAs.

The other references in the list name a journal, a publisher, or a URL. This one names none, so a reader cannot locate it. A Crossref search for the title finds no matching record, and we have not been able to verify the outlet independently, so the entry is not corrected here.