

The heart trumps the head: Desirability bias in political belief revision

Running Head: BIAS IN POLITICAL BELIEF REVISION

Ben M. Tappin^{a, 1, *}, Leslie van der Leer^{b, 2}, & Ryan T. McKay^{a, 3}

^a ARC Centre of Excellence in Cognition and its Disorders, Department of Psychology, Royal Holloway, University of London, Egham, Surrey TW20 0EX, UK

Institutional telephone number: +44 (0) 1784 434455

^b Regent's School of Psychotherapy and Psychology, Regent's University London, London NW1 4NS, UK

Institutional telephone number: +44 (0) 207 487 7700

Email: ¹ Ben.Tappin.2015@live.rhul.ac.uk; ² leslievdl@gmail.com; ³ Ryan.McKay@rhul.ac.uk

* Corresponding author

ARTICLE TYPE: Article

Word count: 3823

In press at Journal of Experimental Psychology: General

© 2017, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors permission. The final article will be available, upon publication, via its DOI: 10.1037/xge0000298

Acknowledgements: This research was supported by a grant from the Economic and Social Research Council awarded to the first author, reference number ES/J500148/1.

Data availability: For access to the raw data supporting this article please contact Research Data Management at Royal Holloway, University of London: rdm@rhul.ac.uk.

Abstract

Understanding how individuals revise their political beliefs has important implications for society. In a pre-registered study (N=900) we experimentally separated the predictions of two leading theories of human belief revision—desirability bias and confirmation bias—in the context of the 2016 US presidential election. Participants indicated who they desired to win, and who they believed would win, the election. Following confrontation with evidence that was either consistent or inconsistent with their desires or beliefs, they again indicated who they believed would win. We observed a robust desirability bias—individuals updated their beliefs more if the evidence was consistent (versus inconsistent) with their desired outcome. This bias was independent of whether the evidence was consistent or inconsistent with their prior beliefs. In contrast, we find limited evidence of an independent confirmation bias in belief updating. These results have implications for the relevant psychological theories and for political belief revision in practice.

Keywords: Confirmation bias; desirability bias; motivated cognition; belief updating; politics

Introduction

People are routinely exposed to a bewildering array of information relevant to their political beliefs. Whether and how they incorporate this information has profound consequences for society. The belief that vaccines have harmful side effects (Moritz, 2011), or that climate change is a hoax (Lewandowsky et al., 2013), can reduce people's intentions to vaccinate (Gangarosa et al., 1998; Jolley & Douglas, 2014a; Horne et al., 2015), or to minimize their carbon footprint (Douglas & Sutton, 2015; Jolley & Douglas, 2014b). Even simple infographics displayed during live televised election debates can meaningfully shape beliefs about debate outcome, potentially influencing the voting intentions of millions of viewers (Davis et al., 2011). A clear understanding of how people incorporate information into their political beliefs is thus of considerable practical importance.

Two prominent theories offer similar yet distinct predictions regarding when and how people incorporate new information into their beliefs. One theory contends that individuals assign greater weight to information that is desirable versus undesirable—i.e., a *desirability bias*. This bias is reported to underlie an asymmetry whereby people update their prior beliefs to incorporate new and desirable information more than new but undesirable information (Sharot & Garrett, 2016). The other theory, *confirmation bias*, contends that people preferentially search for, evaluate and incorporate new information that confirms their prior beliefs (Nickerson, 1998). This bias is reported to underlie an asymmetry whereby people update their prior beliefs to incorporate new and confirming information more than new but disconfirming information—even if they receive a balanced set of both types of information (Lord et al., 1979; Taber & Lodge, 2006; Taber et al., 2009).

Unfortunately, the predictions of desirability bias and confirmation bias are often conflated. In the domain of self-belief, the tendency for people to believe desirable things about themselves and their futures (Sedikides & Strube, 1997; Weinstein, 1980) means that new information is typically either confirming *and* desirable, or disconfirming *and* undesirable (Eil & Rao, 2011). In the domain of political belief, rigorous separation of desirable and confirming information is similarly difficult. Of the few experiments that *are* appropriately designed to disentangle them, group identity is taken as a proxy for the desirability of information—that is, whether the information is consistent (i.e., desirable) or inconsistent (undesirable) with the position of an individual’s cultural group—and belief updating is not the target outcome measure (e.g., see Kahan, 2016a; 2016b).

Here we experimentally separated desirability bias and confirmation bias in political belief updating. To do so, we capitalized on the political context prior to the 2016 US presidential election. To illustrate the advantage of this context, consider that many supporters of candidate Donald Trump may have believed Hillary Clinton would win the election—owing to her establishment support (Green & Kapur, 2016) or, more conspiratorially, a rigged ballot (Graham, 2016). In such circumstances, new information may have been simultaneously confirming but undesirable—for instance, polls indicating a Clinton win—or disconfirming but desirable—polls indicating a Trump win: causing desirability bias and confirmation bias to yield divergent predictions for belief updating.

We exploited the profusion of close polling results¹ to credibly suggest to individuals that either Donald Trump or Hillary Clinton would become the next president, and

¹ At the time of study (data collection commenced 26th September 2016); see http://www.realclearpolitics.com/epolls/2016/president/us/general_election_trump_vs_clinton-5491.html

measured how individuals with congruent (i.e., same candidate) desire-belief profiles, and incongruent (different candidate) desire-belief profiles updated their beliefs following receipt of this information. We thus independently manipulated whether information was consistent or inconsistent with (a) who individuals *desired* to win the election, or (b) who they *believed* would win the election.

Methods

Participants

We collected data from 900 participants online via Amazon's Mechanical Turk (59% female; $M_{\text{age}} = 37.89$ $SD = 12.91$). Participants were US residents as determined by IP address (IP addresses located outside of the US were blocked prior to the start of the experiment). We required 779 participants to attain greater than 80% power ($\alpha = .05$) to detect a small effect of partial eta squared (η_p^2) = .01 in our primary analyses of covariance. We added approximately 15% to this number to guard against power loss due to planned data exclusions. Following these data exclusions, we retained 811 participants for analyses. The study hypotheses, design, data collection, and analysis plan were pre-registered (see <https://aspredicted.org/idxgj.pdf>).

Procedure & design

At the beginning of the survey, participants completed a brief screening questionnaire designed to determine who they (a) *desired* to win, and (b) *believed* would win the 2016 US presidential election. Responses to (a) were provided in a nominal choice format: Participants selected "Donald Trump", "Hillary Clinton", or "neither". Responses to (b) were provided on a bipolar sliding scale from 0-100 with "Hillary Clinton" (0) at one end,

and “Donald Trump” (100) at the other (the numerical values were hidden from participants). Participants were instructed that the more confident they were that a candidate would win, the closer they should slide the pointer to that candidate’s name. Those who responded with scores greater than 50 were categorized as *believing* Trump would win, and scores less than 50 as *believing* Clinton would win. Participants selecting “neither” for (a), or exactly 50 for (b), were directed to an end-of-survey message and were unable to continue. This yielded two quasi-experimental groups; those whose *desire-believe* candidates were congruent, and those whose *desire-believe* candidates were incongruent. We balanced these condition assignments to obtain approximately 450 in each quasi-experimental condition (final condition samples after data exclusions: congruent desire-belief: $n=406$ [$\text{desire}_{\text{Trump}}/\text{believe}_{\text{Trump}}$: $n=127$, $\text{desire}_{\text{Clinton}}/\text{believe}_{\text{Clinton}}$: $n=279$]; incongruent desire-belief: $n=405$ [$\text{desire}_{\text{Clinton}}/\text{believe}_{\text{Trump}}$: $n=91$, $\text{desire}_{\text{Trump}}/\text{believe}_{\text{Clinton}}$: $n=314$])².

Participants in both conditions then completed a filler task (the 16-item Balanced Inventory of Desirable Responding; Hart et al., 2015) before being randomly presented with evidence either consistent, or inconsistent, with who they believed would win the election. Specifically, participants read a short passage about nationwide polling results, which emphasized either that Hillary Clinton or Donald Trump was likely to win the upcoming election. Participants were also presented with a bar graph figure illustrating such an outcome (study materials are available in the Materials Supplement). Evidence presentation was balanced within each specific candidate that participants initially believed would win the election. For example, of those participants who initially *believed* Trump would win, half

² The substantial variance in condition sizes per candidate reflects the fact that approximately three-quarters of our sample initially believed Clinton would win (see Figure 1 in the results section)—explaining the smaller number of individuals in the “ $\text{believe}_{\text{Trump}}$ ” condition(s).

received the polling manipulation suggesting Clinton would win, and half received the polling manipulation suggesting Trump would win (likewise for those who initially *believed* Clinton would win). Thus, collapsing over specific candidates, this yielded four between-subjects conditions in a 2 x 2 design: Evidence consistent or inconsistent with who the participant initially *believed* would win (Confirmation: Confirmatory or Disconfirmatory) and consistent or inconsistent with who they *desired* to win (Desirability: Desirable or Undesirable). Following the evidence presentation participants responded to several filler questions about polling data—e.g., “*To what extent have you been following the polling data for the upcoming US presidential election?*”—before again indicating who they believed would win the election, on the same bipolar scale used initially.

Belief updating

We calculated how much participants updated their confidence in who they believed would win the election in the following steps. First, we converted both the participants’ initial confidence (T1), and their subsequent confidence (T2), onto a comparable scale indicating the *absolute* confidence they had in the candidate they initially believed was most likely to win. Thus, for those who initially believed Trump would win we subtracted 50 from T1 and T2 scores, whereas for those who initially believed Clinton would win we subtracted T1 and T2 scores from 50. Next, we computed the absolute difference between these newly converted T1 and T2 scores for each participant. Finally, we multiplied this difference by either 1 (if the participant updated towards the presented evidence) or -1 (if the participant updated away from the presented evidence); meaning that higher numbers represented greater belief updating towards the presented evidence.

Results

Data exclusions

Participants were excluded from all analyses for fulfilling one or more of the pre-registered criteria: Failing an attention check embedded in the filler task (n=22, 2.44% of sample), answering “yes” to a question asking them if they responded dishonestly or mistakenly during the survey (n=48, 5.33%), or recording a belief update score of greater than the mean \pm 3SD in their respective condition (n=26, 2.89%). We excluded 1 (0.11%) further participant for taking the survey more than once (identified via their unique Amazon Mechanical Turk ID). Following these exclusions, 811 participants were retained for analyses.

Descriptives

Figure 1 displays the proportion of participants reporting who they (a) desired to win and (b) initially believed would win the election (for these results split by gender, age group and ethnicity, see figures S1-S3 in the Analysis Supplement).

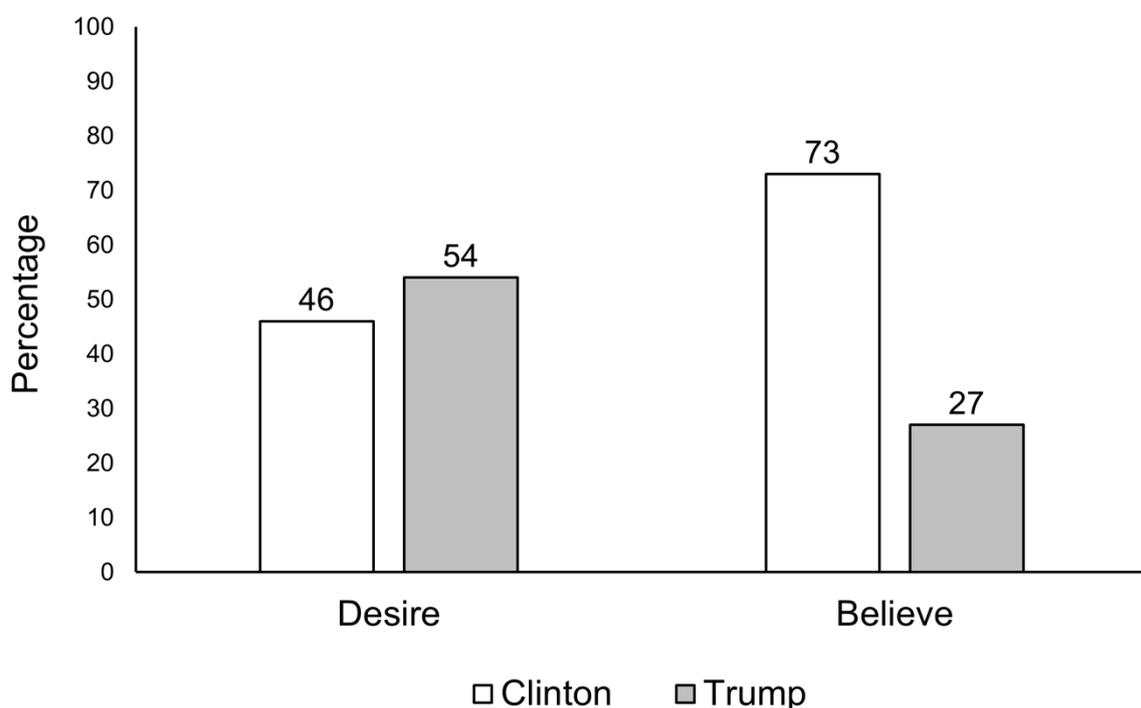


Figure 1. Percentage of participants reporting which candidate they (a) desired to win and (b) initially believed would win the 2016 US presidential election. $N=811$.

Pre-registered analyses

We conducted an analysis of covariance (ANCOVA) to investigate the effect of Desirability and Confirmation factors on belief updating, adjusting for absolute T1 confidence scores³ (Figure 2 displays the adjusted mean update in each condition⁴). There was a main effect of Desirability: $F(1, 806) = 32.81, p < .001, \eta_p^2 = .04$ 90% CI [0.02, 0.06], such that participants updated more towards the evidence when it was consistent (versus inconsistent) with the candidate they desired to win. There was also a main effect of Confirmation: $F(1, 806) = 76.63, p < .001, \eta_p^2 = .09$ [0.06, 0.12], but in this case participants updated more towards the evidence when it was *inconsistent* (versus consistent) with the candidate they initially believed would win. In other words, we observed a *disconfirmation* bias. Finally, we observed a small interaction between Desirability and Confirmation: $F(1, 806) = 7.15, p = .008, \eta_p^2 = .01$ [0.00, 0.02]. To decompose this interaction, we conducted planned ANCOVAs comparing updating in each condition—while adjusting for absolute T1 confidence scores.

³ This prevents regression to the mean spuriously affecting belief updating.

⁴ The raw means and distributions of update scores are reported in the Analysis Supplement.

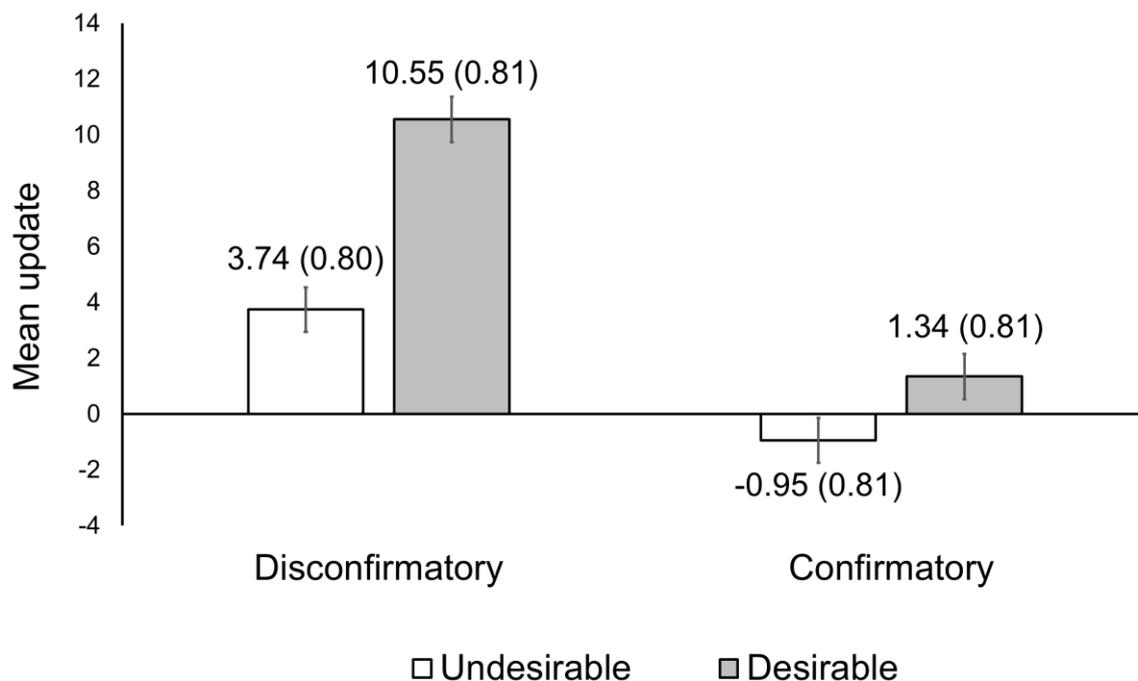


Figure 2. Mean update by condition. Error bars and parentheses denote standard error of the mean. Note: Means are adjusted for absolute T1 confidence and based on the 2x2 ANCOVA model. One unit of update corresponds to a 1% adjustment on the bipolar scale used to measure belief. $N=811$.

For those participants receiving disconfirming information, updating was greater if that information was desirable (versus undesirable): $F(1, 407) = 36.58, p < .001, \eta_p^2 = .08$ 90% CI [0.04, 0.13]. This pattern was the same for those receiving confirming information, albeit less pronounced: $F(1, 398) = 20.62, p < .001, \eta_p^2 = .05$ [0.02, 0.09]. Next we examined those participants who received undesirable information—updating was greater for *disconfirming* (versus confirming) information: $F(1, 406) = 23.76, p < .001, \eta_p^2 = .06$ [0.02, 0.09]. This disconfirmation pattern was the same, yet more pronounced, for those receiving

desirable information: $F(1, 399) = 47.72, p < .001, \eta_p^2 = .11 [0.06, 0.16]$. Finally, directly comparing the unique effect of desirable information (disconfirming-desirable condition) against the unique effect of confirming information (confirming-undesirable condition) revealed that updating was greater for the former: $F(1, 402) = 75.26, p < .001, \eta_p^2 = .16 [0.11, 0.21]$.

In the following sections, we report a series of exploratory analyses to examine (a) the robustness of our results, and (b) extant debates in the field of politically motivated cognition.

Robustness

Prior exposure. It is likely that participants had different amounts of prior exposure to the election polls. Examination of the distribution of one of our filler questions— “*To what extent have you been following the polling data for the upcoming US presidential election?*”—suggested this was the case (see Figure S4 in the Analysis Supplement). It is possible this affected our manipulation and subsequent results. We thus repeated our pre-registered ANCOVA with the addition of this variable as a covariate. However, the pattern of results remained the same.

Initial confidence. Participants’ initial (T1) confidence scores were negatively skewed—in particular, a substantial number reported complete (or strong) confidence in their initial belief regarding which candidate would win (see Figure 3). This constrains belief updating for those receiving confirming information because they are unable to update towards the new information (i.e., increase their confidence). In contrast, those receiving

disconfirming information *can* update towards the new information (i.e., decrease their confidence). This may account for the disconfirmation bias we observed⁵.

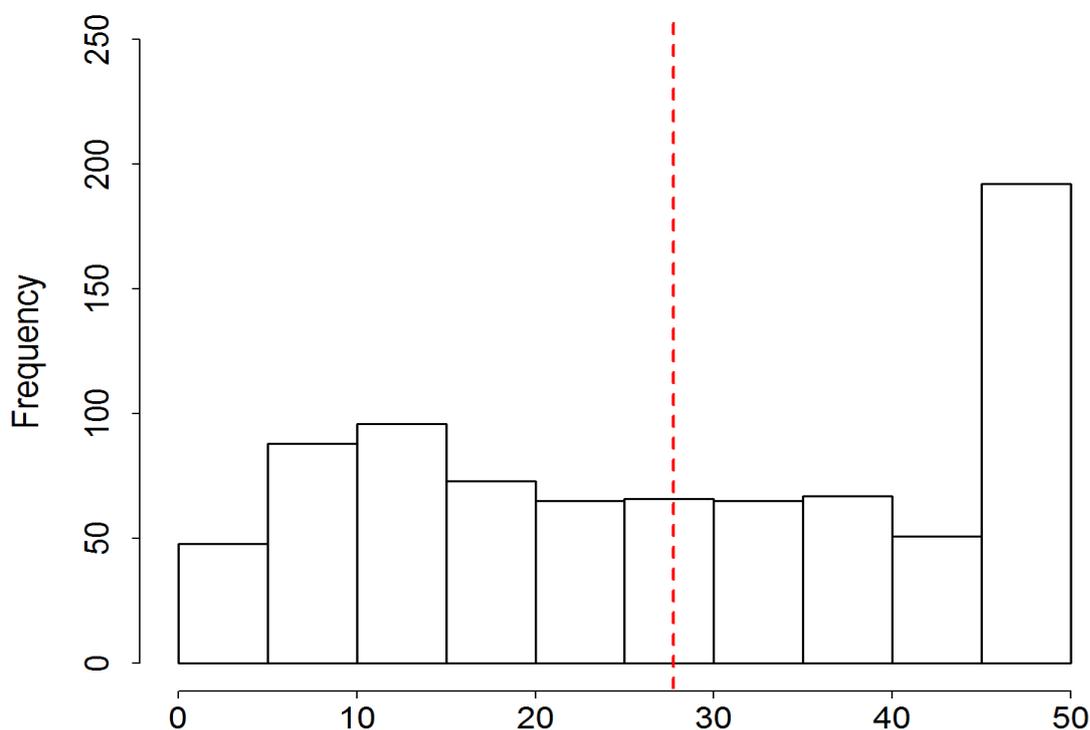


Figure 3. Distribution of absolute T1 confidence in belief about which candidate would win the election. Note: The dashed line denotes the median. $N=811$.

To explore this possibility, we selected a subset of participants ($N = 370$)—excluding those with high initial confidence (absolute T1 confidence scores > 25 , $N_{\text{excluded}} = 441$)—and recomputed the mean update in each condition (Figure 4 displays the results). The pattern of means in this truncated sample indicated a diminished disconfirmation bias, but an enduring desirability bias. To confirm this statistically we conducted separate Kruskal-Wallis

⁵ We are grateful to two anonymous reviewers for emphasizing this point.

tests on the distribution of belief updating in the Confirmation and Desirability conditions, respectively⁶. As suspected, there was now only a trivial difference in updating for participants who received disconfirmatory (Median = 2.01, IQR = 11.69) versus confirmatory (Median = 1.78, IQR = 6.58) information: $\chi^2(1, N = 370) = 2.01, p = .156$. In contrast, participants receiving desirable information updated more (Median = 3.08, IQR = 11.61) than those receiving undesirable information (Median = 0.71, IQR = 6.16): $\chi^2(1, N = 370) = 25.84, p < .001$.

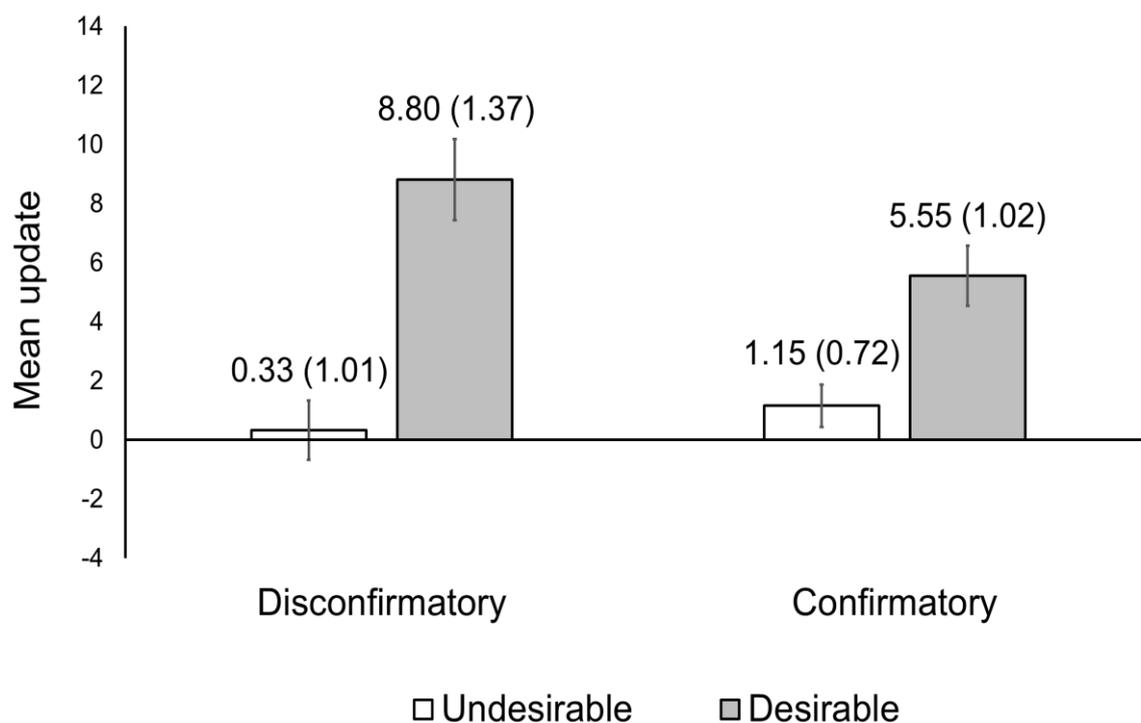


Figure 4. Mean update by condition following sample truncation. Error bars and parentheses denote standard error of the mean. Note: Means are unadjusted. One unit of update corresponds to a 1% adjustment on the bipolar scale used to measure belief. $N=370$.

⁶ Parametric analyses were inappropriate as the cell N 's across conditions were unequal following sample truncation.

To supplement this analysis, we also specifically examined updating among those with *weak* confidence in their initial belief. This is worthwhile because participants with particularly low confidence may have been (a) less constrained by the upper limit of the confidence scale, or (b) simply more receptive to confirming information, compared to their higher confidence counterparts. Thus, we selected those participants with low confidence (absolute T1 confidence scores ≤ 12.5 , $N_{\text{excluded}} = 622$) and again recomputed the mean update in each condition. Because the resultant N was small ($N_{\text{low-confidence}} = 189$) and unevenly distributed across conditions, we simulated belief updating scores using the parameters from the low confidence sample. Specifically, for each of the four conditions, we drew 500 scores from a random normal distribution centred on the respective condition mean, and the pooled SD (i.e., computed across the four conditions) (the simulation script and data are available in the Simulation Supplement).

This simulated sample conferred greater than 99% power to detect small effects ($\eta_p^2 = 0.01$, $\alpha = 0.05$). Conducting an ANOVA on this data revealed a main effect of Desirability, $F(1, 1996) = 53.73$, $p < .001$, $\eta_p^2 = .03$ 90% CI [0.02, 0.04], similar in size and equivalent in direction to that observed in the preceding empirical analyses. The main effect of Confirmation was trivial in size, $F(1, 1996) = 2.06$, $p = .151$, $\eta_p^2 = .001$ [0.000, 0.005], as was the interaction between the two factors, $F(1, 1996) = 1.54$, $p = .215$, $\eta_p^2 = .001$ [0.000, 0.004].

Ideological Asymmetry Hypothesis

There is ongoing debate over whether motivated cognition is more pronounced among individuals on the political right than the political left (Jost et al., 2003; Kahan, 2016b). We thus explored whether supporters of Donald Trump demonstrated greater

desirability bias than supporters of Hillary Clinton. We conducted an ANCOVA (adjusting for absolute T1 confidence as before) with two factors: Desirability, and a dummy coded variable denoting which candidate the participant desired to win (“Supporter”). There was a small Desirability by Supporter interaction, $F(1, 806) = 8.58, p = .004, \eta_p^2 = .01$ 90% CI [0.00, 0.03]. Separate ANCOVA models revealed a stronger desirability bias among supporters of Donald Trump, $F(1, 438) = 34.07, p < .001, \eta_p^2 = .07$ [0.04, 0.11], than supporters of Hillary Clinton, $F(1, 367) = 2.54, p = .112, \eta_p^2 = .01$ [0.00, 0.03].

Further exploration, however, revealed this asymmetry was due to the previously identified ceiling effect in initial (T1) confidence. First, a large number of participants supported Clinton and also believed she would win ($n=279$)—whereas fewer than half this number supported Trump while also believing he would win ($n=127$). Second, these participants (i.e., those with *congruent* desire and prior belief) had strong negative skew in their initial confidence, with many believing that their desired candidate was certain to win (see Figure S5 in the Analysis Supplement). Taking these facts together implies that supporters of Clinton were more numerous among those who received desirable information but were constrained (by virtue of their extreme initial confidence) in updating their belief towards this information.

This was confirmed by examining participants who (i) had a congruent desire-belief profile, (ii) received desirable information, and (iii) reported extreme initial confidence (absolute T1 confidence > 45). Of these participants ($n=69$), 67% supported Clinton ($n=46$) and 33% supported Trump ($n=23$). This discrepancy may have disproportionately suppressed desirability bias among Clinton supporters. Indeed, truncating the sample to exclude those with extreme initial confidence (absolute T1 confidence > 45 , $N_{\text{excluded}} = 192$,

$N_{\text{included}} = 619$), and repeating the ANCOVA analysis, diminished the size of the Desirability by Supporter interaction, $F(1, 614) = 1.08$, $p = .298$, $\eta_p^2 = .002$ 90% CI [0.000, 0.012]. Supporters of Donald Trump and supporters of Hillary Clinton demonstrated similar desirability bias in this sample: $F(1, 348) = 27.19$, $p < .001$, $\eta_p^2 = .07$ [0.03, 0.12] and $F(1, 265) = 10.55$, $p = .001$, $\eta_p^2 = .04$ [0.01, 0.08], respectively.

Discussion

Understanding how people revise their political beliefs has important implications for society. In the context of the 2016 US presidential election, we observed a robust desirability bias: individuals incorporated information more if it was consistent (versus inconsistent) with their desired outcome. This bias was independent of whether the information was consistent or inconsistent with individuals' prior beliefs. In contrast, we found limited evidence of an independent confirmation bias in belief updating. These results have implications for the underlying psychological theories and for political belief revision in practice.

A substantial body of work spanning neuroscience, economics, and clinical psychology reports an asymmetry in the updating of self-beliefs whereby desirable information is incorporated more than undesirable information. This asymmetry has been observed when individuals receive information about their personality traits (Korn et al., 2016; Korn et al., 2012), abilities and attractiveness (Eil & Rao, 2011; Mobius et al., 2011), or risk of experiencing future negative life events (Moutsiana et al., 2013; Sharot et al., 2011; but see Shah et al., 2016; Garret & Sharot, in press). A similar yet distinct asymmetry has been reported in the updating of political beliefs whereby individuals become more confident in their prior beliefs despite receiving a balanced set of confirming and

disconfirming information. When two individuals with conflicting prior beliefs are thus exposed to the same stream of information, polarization of political beliefs is an oft-observed outcome (e.g., Lord et al., 1979; Taber & Lodge, 2006; Taber et al., 2009).

The present study advances this work twofold. First, we find a robust asymmetry in political belief updating that is consistent with desirability bias, *independent* of individuals' prior beliefs. In contrast, we find little independent effect of prior beliefs on belief updating. This suggests that the belief polarization reported in previous studies may be due to individuals' conflicting desires, not their prior beliefs per se. Second, whereas past investigations of political belief updating have mainly focused on political *attitudes* (e.g., support for or against a policy), here we examined belief updating about political *reality*—specifically, individuals' belief about which presidential candidate was going to be elected. Though one might expect biased belief updating in the former case—after all, attitudes are guided by preferences and desires—it is somewhat more surprising to find that individuals' desires biased their belief updating over a question of fact (Kahan, 2016a).

A recent study reported that individuals updated their beliefs about the facts of global warming asymmetrically, but that the specific pattern depended upon whether they were weak or strong believers in anthropogenic climate change (Sunstein et al., 2016). Particularly, when confronted with new information regarding global temperature increase, strong believers updated their beliefs more upon receipt of ostensibly *undesirable* information (i.e., a faster temperature increase than expected), whereas weak believers updated their beliefs more upon receipt of ostensibly *desirable* information (a slower increase than expected). Though this pattern appears consistent with an independent confirmation bias, such an outcome may emerge when individuals are personally invested in

“being right”—indeed, for many climate change activists a belief that the world is warming constitutes a core part of their identity (Stern et al., 1999). For such people, objectively undesirable (but confirming) information about the rate of global warming may be subjectively *desirable*: vindicating their commitment to combatting climate change (Sunstein et al., 2016) and affirming their cultural group identity (Kahan et al., 2012).

It is unlikely that our design inadvertently conflated confirmation with desirability in this way. Ahead of an election, it is difficult to imagine an individual being personally invested in the belief that their desired candidate would *not* get into office. Indeed, in the domain of self-belief updating, rigorous separation of confirming and desirable information yields identical results to those reported in the present study—namely, a robust desirability bias but limited evidence of confirmation bias (Eil & Rao, 2011). We note the important distinction, however, between (a lack of) confirmation bias observed in *belief updating* as measured here, and confirmation bias observed in measures of *information search* and *evaluation* (e.g., Ditto & Lopez, 1992). We did not directly examine the latter, which may yet manifest independent of information desirability. Additional exploration of our own data lent support to this distinction (see “informational value of polls” available in the Analysis Supplement).

Finally, our results offer a mechanistic explanation for why impassioned political disagreements in the US, such as those over gun control or immigration, appear increasingly polarized and intractable (Pew Research Center, 2016). Insofar as individuals have strong preferences concerning these issues (Koleva et al., 2012), our findings suggest they selectively incorporate new evidence into what they believe to be true regarding the relevant facts—provided it is consistent with what they *desire* to be true. Polarization over

factual beliefs is inimical to the effective functioning of democratic society (Kahan et al., 2012); it is thus a priority to continue exploring which interventions ameliorate the motivated integration of evidence (Lewandowsky & Oberauer, 2016).

References

- Davis, C. J., Bowers, J. S., & Memon, A. (2011). Social influence in televised election debates: A potential distortion of democracy. *PLoS One*, *6*, e18154. Doi: 10.1371/journal.pone.0018154
- Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, *63*, 568-584. Doi: 10.1037/0022-3514.63.4.568
- Douglas, K. M., & Sutton, R. M. (2015). Climate change: Why the conspiracy theories are dangerous. *Bulletin of the Atomic Scientists*, *71*, 98-106. Doi: 10.1177/0096340215571908
- Eil, D., & Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, *3*, 114-138. Doi: 10.1257/mic.3.2.114
- Gangarosa, E. J., Galazka, A. M., Wolfe, C. R., Phillips, L. M., Gangarosa, R. E., Miller, E., & Chen, R. T. (1998). Impact of anti-vaccine movements on pertussis control: The untold story. *The Lancet*, *351*, 356-361. Doi: 10.1016/S0140-6736(97)04334-1
- Garret, N. & Sharot, T. (in press). Optimistic update bias holds firm: Three tests of robustness following Shah et al. *Consciousness and Cognition*. Doi: 10.1016/j.concog.2016.10.013
- Graham, D. A. (2016, 17 October). *Republicans have been 'rigging' elections for years*. Retrieved from <http://www.theatlantic.com/politics/archive/2016/10/trump-election-rigged/504347/>
- Green, J. & Kapur, S. (2016, 22 June). *Nearly half of Sanders supporters won't support Clinton: Poll*. Retrieved from <http://www.bloomberg.com/politics/articles/2016-06-22/nearly-half-of-sanders-supporters-won-t-support-clinton>
- Hart, C. M., Ritchie, T. D., Hepper, E. G., & Gebauer, J. E. (2015). The Balanced Inventory of Desirable Responding Short Form (BIDR-16). *SAGE Open*, *5*, 1-9. Doi: 10.1177/2158244015621113
- Horne, Z., Powell, D., Hummel, J. E., & Holyoak, K. J. (2015). Countering antivaccination attitudes. *Proceedings of the National Academy of Sciences*, *112*, 10321-10324. Doi: 10.1073/pnas.1504019112
- Jolley, D. & Douglas, K. M. (2014a). The effects of anti-vaccine conspiracy theories on vaccination intentions. *PLoS One*, *9*, e89177. Doi: 10.1371/journal.pone.0089177
- Jolley, D. & Douglas, K. M. (2014b). The social consequences of conspiracism: Exposure to conspiracy theories decreases intentions to engage in politics and to reduce one's carbon footprint. *British Journal of Psychology*, *105*, 35-56. Doi: 10.1111/bjop.12018
- Jost, J. T., Glaser, J., Kruglanski, A. W., & Sulloway, F. J. (2003). Political conservatism as motivated social cognition. *Psychological Bulletin*, *129*, 339-375. Doi: 10.1037/0033-2909.129.3.339

- Kahan, D. M. (2016a). The Politically Motivated Reasoning Paradigm, Part 1: What politically motivated reasoning is and how to measure it. In R. Scott & S. Kosslyn (Eds.), *Emerging Trends in Social and Behavioral Sciences* (pp. 1-16). Doi: 10.1002/9781118900772.etrds0417
- Kahan, D. M. (2016b). The Politically Motivated Reasoning Paradigm, Part 2: Unanswered questions. In R. Scott & S. Kosslyn (Eds.), *Emerging Trends in Social and Behavioral Sciences* (pp. 1-15). Doi: 10.1002/9781118900772.etrds0418
- Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., & Mandel, G. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, 2, 732-735. Doi: 10.1038/nclimate1547
- Koleva, S. P., Graham, J., Iyer, R., Ditto, P. H., & Haidt, J. (2012). Tracing the threads: How five moral concerns (especially Purity) help explain culture war attitudes. *Journal of Research in Personality*, 46, 184-194. Doi: 10.1016/j.jrp.2012.01.006
- Korn, C. W., La Rosée, L., Heekeren, H. R., & Roepke, S. (2016). Social feedback processing in borderline personality disorder. *Psychological Medicine*, 46, 575-587. Doi: 10.1017/S003329171500207X
- Korn, C. W., Prehn, K., Park, S. Q., Walter, H., & Heekeren, H. R. (2012). Positively biased processing of self-relevant social feedback. *The Journal of Neuroscience*, 32, 16832-16844. Doi: 10.1523/JNEUROSCI.3016-12.2012
- Lewandowsky, S., & Oberauer, K. (2016). Motivated Rejection of Science. *Current Directions in Psychological Science*, 25, 217-222. Doi: 10.1177/0963721416654436
- Lewandowsky, S., Oberauer, K. & Gignac, G. E. (2013). NASA Faked the Moon Landing-- Therefore, (Climate) Science Is a Hoax: An Anatomy of the Motivated Rejection of Science. *Psychological Science*, 24, 622-633. Doi: 10.1177/0956797612457686
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37, 2098-2109. Doi: 10.1037/0022-3514.37.11.2098
- Mobius, M. M., Niederle, M., Niehaus, P., & Rosenblat, T. S. (2011, May). *Managing self-confidence: Theory and experimental evidence*. National Bureau of Economic Research Working Paper No. 17014. Doi: 10.3386/w17014
- Moritz, A. (2011). *Vaccine-nation: Poisoning the population, one shot at a time*. Enerchi.com.
- Moutsiana, C., Garrett, N., Clarke, R. C., Lotto, R. B., Blakemore, S. J., & Sharot, T. (2013). Human development of the ability to learn from bad news. *Proceedings of the National Academy of Sciences*, 110, 16396-16401. Doi: 10.1073/pnas.1305631110
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2, 175-220. Doi: 10.1037/1089-2680.2.2.175

- Pew Research Center (2016, 26 August). *Gun Rights vs. Gun Control*. Retrieved from <http://www.people-press.org/2016/08/26/gun-rights-vs-gun-control/#party>
- Sedikides, C., & Strube, M. J. (1997). Self-evaluation: To thine own self be good, to thine own self be sure, to thine own self be true, and to thine own self be better. *Advances in Experimental Social Psychology, 29*, 209-269. Doi: 10.1016/S0065-2601(08)60018-0
- Shah, P., Harris, A. J. L., Bird, G., Catmur, C., & Hahn, U. (2016). A pessimistic view of optimistic belief updating. *Cognitive Psychology, 90*, 71-127. Doi: 10.1016/j.cogpsych.2016.05.004
- Sharot, T. & Garrett, N. (2016). Forming beliefs: why valence matters. *Trends in Cognitive Sciences, 20*, 25-33. Doi: 10.1016/j.tics.2015.11.002
- Sharot, T., Korn, C. W., & Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality. *Nature Neuroscience, 14*, 1475-1479. Doi: 10.1038/nn.2949
- Stern, P. C., Dietz, T., Abel, T. D., Guagnano, G. A., & Kalof, L. (1999). A value-belief-norm theory of support for social movements: The case of environmentalism. *Human Ecology Review, 6*, 81-97. Available at: https://works.bepress.com/troy_abel/3/
- Sunstein, C. R., Bobadilla-Suarez, S., Lazzaro, S. C., & Sharot, T. (2016). *How People Update Beliefs about Climate Change: Good News and Bad News*. Available at SSRN: 10.2139/ssrn.2821919
- Taber, C. S., Cann, D., & Kucsova, S. (2009). The motivated processing of political arguments. *Political Behavior, 31*, 137-155. Doi: 10.1007/s11109-008-9075-8
- Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science, 50*, 755-769. Doi: 10.1111/j.1540-5907.2006.00214.x
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology, 39*, 806-820. Doi: 10.1037/0022-3514.39.5.806