ARTICLES
https://doi.org/10.1038/s41562-022-01327-5

nature human behaviour

Check for updates

Negativity bias, personality and political ideology

Christopher D. Johnston[1] ✉ and Gabriel J. Madson[2]

Research suggests that right-wing ideology is associated with negativity bias: a tendency to pay more attention and give more weight to negative versus positive stimuli. This work typically relies on either self-reported traits related to negativity bias in large, often-representative, samples or physiological and behavioural indicators of negativity bias in small convenience samples. We extend this literature and examine the relationship of negativity bias to political ideology using five distinct behavioural measures of negativity bias in four national samples of US residents with a total analytical sample size of about 4,000 respondents. We also examine the association of these behavioural measures to four of the most common self-report measures of personality in the literature on ideology. Across a wide range of tests, we find no consistent evidence for a relationship of negativity bias to either ideology or self-reported personality.

What factors influence attraction to the political left and right? Theories that highlight stable individual differences in personality now occupy an important place beside classic sociological and economic approaches focused on individual and group interests[1–6]. While diverse, much of this literature emphasizes dispositions and traits related to negativity bias—a tendency to pay more attention and give more weight to negative and threatening stimuli relative to positive and rewarding stimuli. In theory, a focus on negative potential outcomes creates a disposition to favour existing arrangements and certain outcomes over mixed gambles[7,8]. In the political realm, social order and the status quo provide predictability, while diversity, reform and change are risky prospects with the potential for either gain or loss. In turn, negativity bias should be associated with a preference for right-wing policies that promote stability in social and economic arrangements[4]. In this view, there is an enduring political conflict between two 'primal mindsets': one focused on order and stability and one on reform and progress[3].

Evidence for this theory tends to fall into one of two categories. First, many studies examine the association of political ideology to self-reported dispositions and traits related to negativity bias and its implications, such as the need for cognitive closure, threat and disgust sensitivity and openness to experience[5,9]. This work often, and increasingly, relies on large samples that are reasonably representative of the broader population of interest[1,2,6,10,11]. Yet such studies assume the exogeneity of traits to political ideology, as well as the unbiased reporting of such traits. If traits are instead endogenous to political preferences or if citizens misrepresent, or systematically misperceive, their own traits, the conclusions of such studies would be suspect. There is reason to be concerned. Recent work suggests that citizens exaggerate the degree to which they possess traits prototypical for their political in-group[12] and variation in traits over short periods of time may be driven by political variables[13–15]. It is thus critical to supplement existing work with alternative research designs that do not rely solely on self-reported personality.

Indeed, a growing literature examines the association of ideology to revealed indicators of negativity bias[3], such as patterns of decision-making under ambiguity[16,17], physiological responses to threatening stimuli[18,19] and neural activation in response to threat and uncertainty[20–22]. This work has the virtue of bypassing self-report. In many cases, there is also a strong argument for exogeneity, as indicators of negativity bias are derived from apolitical stimuli and often involve uncontrolled processes. On the downside, the procedures involved, especially in studies using psychophysiological or neuroscientific methods, often require substantial resources in time and money, and the physical presence of participants. This limits the size and representativeness of samples and creates concerns with replicability and generalizability. Consistent with these concerns, recent attempts at replicating a link between ideology and psychophysiological indicators of negativity bias (for example, changes in skin conductance) have largely failed[23–26].

Thus, the literature on the negativity bias hypothesis lacks studies with both of two critical characteristics: (1) large, diverse samples and (2) revealed, rather than self-reported, traits. The primary aim of the present study is to fill this gap. We examine the negativity bias hypothesis using five distinct behavioural measures in four national samples with an analytical sample size of about 4,000 respondents. As we rely on several different operationalizations of negativity bias, we ensure that any conclusions drawn are not specific to a single measurement approach. The large size of our sample also allows for more precise estimates of the relationship of negativity bias to ideology. Whereas small-sample studies tend to focus on statistical significance at the expense of effect size, we can more clearly rule out important subsets of the parameter space.

Our first three behavioural measures of negativity bias attempt conceptual replications of previous research using smaller and less representative samples. First, we measure the cognitive accessibility of threatening concepts (for example, disease and murder) using a lexical decision task in which respondents attempt to recognize positive, negative and neutral words as quickly as possible[27]. Second, we measure attentional biases to threatening images using a flanker task in which respondents attempt to categorize target images (for example, a seal pup and a gun) as positive or negative in the presence of positive or negative distractor images[28]. Third, we measure biases in learning negative over positive stimuli using a computer game called BeanFest[17]. Our fourth and fifth measures extend this research to examine the relationship between political ideology and individual differences in loss aversion—perhaps the most intuitive conceptualization of negativity bias and one that ties the literature to behavioural decision theory. We apply two different methods for eliciting individual differences in loss aversion under risk[29,30]. We also collected data for a measure of loss aversion under ambiguity[31]. However, for reasons discussed in the Supplementary Methods, we have less confidence in these data and thus do not report the

[1]Department of Political Science, Duke University, Durham, NC, USA. [2]RTI International, Durham, NC, USA. ✉e-mail: cdj19@duke.edu

results in the main text. Interested readers can find these results in Supplementary Tables 14, 21 and 28. They are consistent with the results for the other five measures.

All five of our measurement tasks are incentivized. That is, respondents had the opportunity to win monetary prizes (paid as gift cards) if they performed each task properly and well. This was done to encourage attentiveness, effort and truthful revelation of preferences.

We also contribute to the literature in two other ways. First, we provide a large-sample examination of the association between behavioural indicators of negativity bias and the most commonly studied self-report personality variables in the literature on ideology. These include the need for non-specific cognitive closure[32], openness to experience[33], conservation versus openness to change values[34] and authoritarian childrearing values[35]. Strong relationships between behavioural measures of negativity bias and self-report measures of personality would provide evidence for a possible pathway through which basic attentional and evaluative dispositions operate, namely, core values and traits[36].

Second, we explore the relationship of negativity bias to multiple dimensions of left–right political preferences and across levels of political engagement—that is, interest in and knowledge of politics. Recent work has argued that relationships between traits and ideology vary across ideological dimensions[9,11,37,38]. Despite growing constraint over time[39,40], mass belief systems are multidimensional in the sense that they are well-described by models positing distinct, correlated dimensions of values and policy preferences, with social (for example, abortion) and economic (for example, taxes and spending) dimensions most prominent[37,41]. Recent work finds that traits conceptually related to negativity bias (for example, need for closure) are more strongly associated with political identity and social policy preferences relative to economic policy preferences[9,37] and occasionally promote left-wing economic views[42–44]. Further, in models that condition on respondents' political engagement, these traits promote right-wing views among the engaged but often left-wing views among the unengaged[11,38,45,46]. We extend this line of work and examine the relationship of behavioural indicators of negativity bias to several measures of political preferences, including identification with the political left and right (Democrats and liberals versus Republicans and conservatives), general conservatism (average conservatism across all issues), economic conservatism and social conservatism. In all cases, we also test for interactions with political engagement.

On the basis of our review of the existing literature, we test the following hypotheses:

(1) Negativity bias will be associated with right-wing political preferences in general but will be more strongly associated with right-wing identity, general conservatism and social conservatism compared to economic conservatism.
(2) Negativity bias will be associated with self-reported personality traits commonly used to predict right-wing political preferences, including: (low) openness to experience, the prioritization of conservation values over openness to change values, authoritarian childrearing values and the need for non-specific cognitive closure.
(3) Political engagement will positively moderate the relationship between negativity bias and right-wing political preferences.

Results

To test for an unconditional relationship between negativity bias and right-wing political preferences (hypothesis 1), we report six sets of estimates: one set for each of our five measures of negativity bias and one for a dataset that combines the samples for all five measures. The combined dataset ignores differences across negativity bias measures, leveraging the gains in efficiency from the larger sample at the expense of conceptual clarity. Each set of models includes four ordinary least squares regressions—one for each measure of ideology. Each model regresses the respective dependent variable on the respective measure of negativity bias and a set of standard control variables, which include age, gender, race and ethnicity, education, household income and employment status. All statements below concerning statistical significance are based on whether or not the 95% confidence interval (CI) for the estimated quantity includes zero and are thus two-tailed tests. We report one-tailed significance tests in Supplementary Tables 38–41, along with two types of adjustments of P values to control the false discovery rate. As we show in Supplementary Figs. 6–11, our study is sufficiently powered to detect even small effect sizes.

In addition, we report Bayes factors in the direction of the null for all quantities of interest in Tables 1–3. Following a previous study[47], we calculate the Savage–Dickey density ratio for each estimated coefficient. This ratio estimates the factor change in the probability of the null from prior to posterior by taking the ratio of the two densities at the null value for the parameter of interest. For this analysis, we first standardize both independent and dependent variables and begin with a normal prior density for all coefficients with mean zero and standard deviation 0.50. For all coefficients of primary interest (those reported in Figs. 1–3), we then compute Bayes factors for

Table 1 | Bayes factors in the direction of the null (BF_{01}) for associations between operationalizations of negativity bias and political ideology

Dataset	DV	BF_{01} and qualitative interpretation
All	Identity	Strong evidence ($BF_{01}=62.81$) in favour of null
All	General	Strong evidence ($BF_{01}=81.76$) in favour of null
All	Social	Strong evidence ($BF_{01}=85.85$) in favour of null
All	Economic	Positive evidence ($BF_{01}=9.16$) in favour of null
Lexical	Identity	Strong evidence ($BF_{01}=40.20$) in favour of null
Lexical	General	Strong evidence ($BF_{01}=30.85$) in favour of null
Lexical	Social	Strong evidence ($BF_{01}=38.35$) in favour of null
Lexical	Economic	Positive evidence ($BF_{01}=11.00$) in favour of null
Flanker	Identity	Positive evidence ($BF_{01}=8.38$) in favour of null
Flanker	General	Positive evidence ($BF_{01}=12.83$) in favour of null
Flanker	Social	Strong evidence ($BF_{01}=23.74$) in favour of null
Flanker	Economic	Weak evidence ($BF_{01}=1/1.03$) against null
Toubia	Identity	Weak evidence ($BF_{01}=1.38$) in favour of null
Toubia	General	Positive evidence ($BF_{01}=5.65$) in favour of null
Toubia	Social	Positive evidence ($BF_{01}=3.29$) in favour of null
Toubia	Economic	Positive evidence ($BF_{01}=7.14$) in favour of null
Tanaka	Identity	Strong evidence ($BF_{01}=39.90$) in favour of null
Tanaka	General	Strong evidence ($BF_{01}=43.69$) in favour of null
Tanaka	Social	Strong evidence ($BF_{01}=42.36$) in favour of null
Tanaka	Economic	Strong evidence ($BF_{01}=27.38$) in favour of null
BeanFest	Identity	Positive evidence ($BF_{01}=17.99$) in favour of null
BeanFest	General	Strong evidence ($BF_{01}=28.45$) in favour of null
BeanFest	Social	Positive evidence ($BF_{01}=15.61$) in favour of null
BeanFest	Economic	Positive evidence ($BF_{01}=6.87$) in favour of null

All, includes samples for all negativity bias measures combined; Lexical, lexical decision task; Flanker, flanker task; Toubia, loss aversion under risk estimated using Toubia et al.[30] procedure; Tanaka, loss aversion under risk estimated using Tanaka et al.[29] procedure; BeanFest, differential memory for negative versus positive beans. Identity, right-wing identity; General, general conservatism; Social, social conservatism; Economic, economic conservatism.

Table 2 | Bayes factors in the direction of the null (BF_{01}) for associations between operationalizations of negativity bias and self-reported personality traits

Dataset	DV	BF_{01} and qualitative interpretation
All	NFC	Strong evidence ($BF_{01}=51.30$) in favour of null
All	Low openness	Strong evidence ($BF_{01}=36.34$) in favour of null
All	Conservation	Strong evidence ($BF_{01}=27.90$) in favour of null
All	Authoritarian	Strong evidence ($BF_{01}=87.03$) in favour of null
Lexical	NFC	Strong evidence ($BF_{01}=26.78$) in favour of null
Lexical	Low openness	Strong evidence ($BF_{01}=31.20$) in favour of null
Lexical	Conservation	Strong evidence ($BF_{01}=39.40$) in favour of null
Lexical	Authoritarian	Strong evidence ($BF_{01}=41.13$) in favour of null
Flanker	NFC	Strong evidence ($BF_{01}=41.37$) in favour of null
Flanker	Low openness	Strong evidence ($BF_{01}=31.06$) in favour of null
Flanker	Conservation	Positive evidence ($BF_{01}=15.61$) in favour of null
Flanker	Authoritarian	Strong evidence ($BF_{01}=32.05$) in favour of null
Toubia	NFC	Very strong evidence ($BF_{01}=1/640.04$) against null
Toubia	Low openness	Positive evidence ($BF_{01}=1/6.12$) against null
Toubia	Conservation	Positive evidence ($BF_{01}=1/15.35$) against null
Toubia	Authoritarian	Positive evidence ($BF_{01}=6.33$) in favour of null
Tanaka	NFC	Positive evidence ($BF_{01}=7.13$) in favour of null
Tanaka	Low openness	Weak evidence ($BF_{01}=1.54$) in favour of null
Tanaka	Conservation	Positive evidence ($BF_{01}=9.75$) in favour of null
Tanaka	Authoritarian	Strong evidence ($BF_{01}=49.33$) in favour of null
BeanFest	NFC	Strong evidence ($BF_{01}=50.52$) in favour of null
BeanFest	Low openness	Strong evidence ($BF_{01}=39.49$) in favour of null
BeanFest	Conservation	Strong evidence ($BF_{01}=23.11$) in favour of null
BeanFest	Authoritarian	Strong evidence ($BF_{01}=43.48$) in favour of null

All, includes samples for all negativity bias measures combined; Lexical, lexical decision task; Flanker, flanker task; Toubia, loss aversion under risk estimated using Toubia et al.[30] procedure; Tanaka, loss aversion under risk estimated using Tanaka et al.[29] procedure; BeanFest, differential memory for negative versus positive beans. NFC, need for closure; Low openness, reverse-coded openness to experience; Conservation, Schwartz' conservation versus openness to change value dimension; Authoritarian, authoritarian childrearing values.

Table 3 | Bayes factors in the direction of the null (BF_{01}) for interactions of operationalizations of negativity bias with political engagement

Dataset	DV	BF_{01} and qualitative interpretation
All	Identity	Positive evidence ($BF_{01}=15.20$) in favour of null
All	General	strong evidence ($BF_{01}=24.54$) in favour of null
All	Social	Positive evidence ($BF_{01}=14.84$) in favour of null
All	Economic	Strong evidence ($BF_{01}=24.03$) in favour of null
Lexical	Identity	Positive evidence ($BF_{01}=7.81$) in favour of null
Lexical	General	Positive evidence ($BF_{01}=14.00$) in favour of null
Lexical	Social	Positive evidence ($BF_{01}=11.66$) in favour of null
Lexical	Economic	Positive evidence ($BF_{01}=12.14$) in favour of null
Flanker	Identity	Positive evidence ($BF_{01}=6.54$) in favour of null
Flanker	General	Positive evidence ($BF_{01}=7.51$) in favour of null
Flanker	Social	Positive evidence ($BF_{01}=10.64$) in favour of null
Flanker	Economic	Positive evidence ($BF_{01}=3.55$) in favour of null
Toubia	Identity	Positive evidence ($BF_{01}=4.01$) in favour of null
Toubia	General	Positive evidence ($BF_{01}=5.42$) in favour of null
Toubia	Social	Positive evidence ($BF_{01}=6.41$) in favour of null
Toubia	Economic	Positive evidence ($BF_{01}=4.33$) in favour of null
Tanaka	Identity	Positive evidence ($BF_{01}=6.05$) in favour of null
Tanaka	General	Positive evidence ($BF_{01}=4.40$) in favour of null
Tanaka	Social	Weak evidence ($BF_{01}=2.27$) in favour of null
Tanaka	Economic	Positive evidence ($BF_{01}=7.47$) in favour of null
BeanFest	Identity	Positive evidence ($BF_{01}=5.98$) in favour of null
BeanFest	General	Positive evidence ($BF_{01}=7.25$) in favour of null
BeanFest	Social	Positive evidence ($BF_{01}=3.02$) in favour of null
BeanFest	Economic	Positive evidence ($BF_{01}=11.39$) in favour of null

See Table 1 footnote for explanations.

directional hypotheses (one-tailed) by setting order restrictions on their prior distributions using the R package bayestestR[48,49]. This approach works against the null hypothesis because a null of zero is a priori more likely with a directional relative to a non-directional prior. We also report qualitative interpretations of the estimated Bayes factors based on previous work[50].

Full regression tables can be found in the Supplementary Results. Figure 1 displays the key estimates and 95% CIs for each of the 24 models. Points in the figure represent the percentage point difference in right-wing preferences for a 1 s.d. difference in a particular measure of negativity bias (listed on the x axis). Hypothesis 1 thus predicts positive coefficients for all four dependent variables and all measures of negativity bias but predicts that these will be larger for right-wing identity, general conservatism and social conservatism, relative to economic conservatism. This hypothesis finds no consistent support in our data. Only one of the estimates attains statistical significance in the predicted direction—the flanker measure for economic conservatism ($B=0.02$, $P=0.04$, 95% CI $=0.00$, 0.03)—and this result does not survive P-value adjustments intended to control the false discovery rate (Supplementary Table 38).

Perhaps more importantly, the estimates generally rule out large, positive relationships between negativity bias and the four dependent variables. If we ignore conceptual differences in the five operationalizations—and thus pool all the data—the estimates for negativity bias are precisely estimated and of little substantive significance in each case (political identity: $B=-0.01$, $P=0.25$, 95% CI $=-0.01$, 0.00; general conservatism: $B=0.00$, $P=0.12$, 95% CI $=-0.01$, 0.00; economic conservatism: $B=0.00$, $P=0.22$, 95% CI $=-0.00$, 0.01; social conservatism: $B=-0.01$, $P=0.13$, 95% CI $=-0.01$, 0.00). Thus, our five measures—alone or combined—produce no consistent evidence for a meaningful relationship between negativity bias and right-wing political ideology. Further, with only one exception (the flanker measure for economic conservatism), the estimated Bayes factors for these models, reported in Table 1, indicate that the data provide evidence for the null hypothesis that the coefficient for negativity bias equals zero and this evidence is often strong.

We test hypothesis 2 in a similar fashion. We again estimate six sets of models: one set each for the five negativity bias operationalizations and one for the combined data. Each set consists of four ordinary least squares regressions, one each for the four self-reported personality traits in our study. Figure 2 displays the coefficient estimates and 95% CI for these models. All personality measures are coded so that higher values correspond with theoretically right-wing traits. Thus, hypothesis 2 predicts a positive relationship in all cases. This hypothesis finds no consistent support in our data. Only nine of 24 coefficients are statistically significant, six of these are in the opposite direction predicted by theory and only a few survive adjustments for multiple hypothesis tests (Supplementary Tables 38 and 40). Interestingly, the Toubia measure

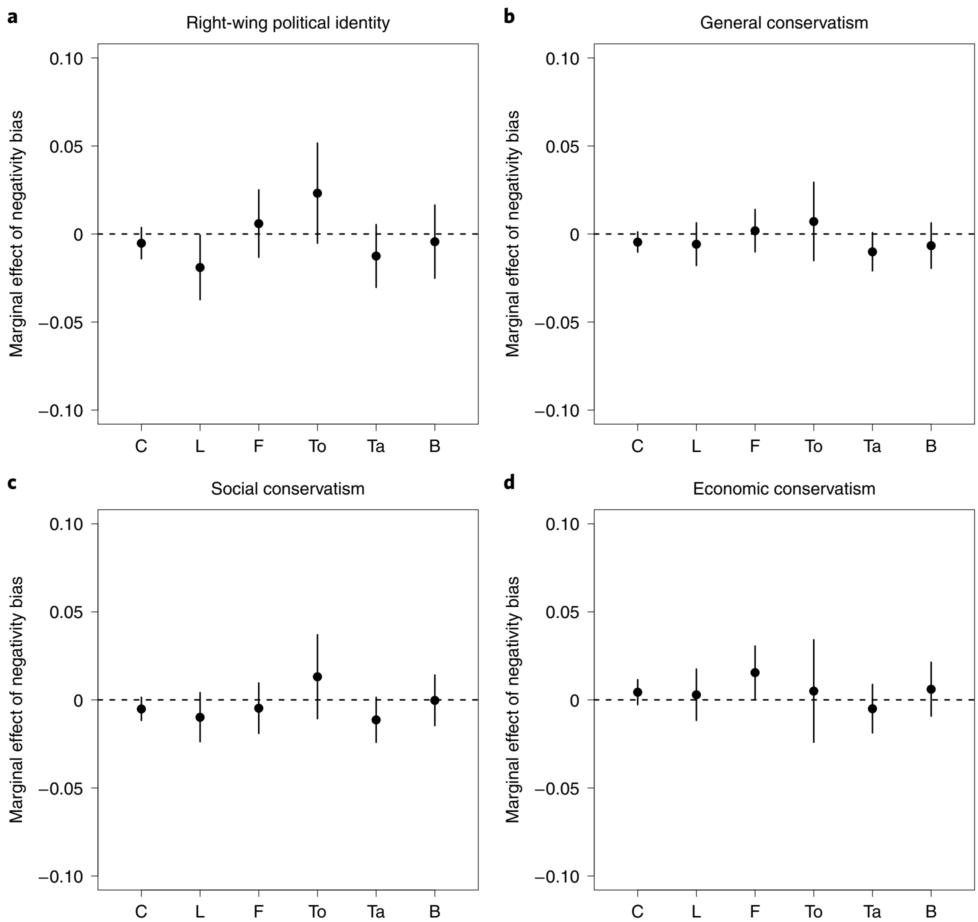

Fig. 1 | Associations between operationalizations of negativity bias and political ideology. a–d, Plots of ordinary least squares coefficients (the solid points) and their 95% CIs (vertical lines extended from points) for negativity bias for models in which right-wing political identity (**a**), general conservatism (**b**), social conservatism (**c**) and economic conservatism (**d**) were regressed separately on each measure of negativity bias (specified on the *x* axis) and a set of common controls. Each coefficient represents the difference in the respective dependent variable (coded 0 to 1) for a 1 s.d. difference in the respective measure of negativity bias. All variables are coded so that positive estimates are consistent with the negativity bias hypothesis. C, samples for all negativity bias measures combined (sample sizes for identity, general, social and economic conservatism models are, respectively, $n=4,055$, 3,934, 3,981 and 3,996); L, lexical decision task ($n=948$, 916, 926 and 932); F, flanker task ($n=881$, 861, 869 and 872); To, loss aversion under risk estimated using Toubia et al.[30] procedure ($n=379$, 367, 375 and 373); Ta, loss aversion under risk estimated using Tanaka et al.[29] procedure ($n=1,041$, 1,008, 1,017 and 1,028); B, differential memory for negative versus positive beans in BeanFest ($n=806$, 782, 794 and 791). Regression outputs, with all frequentist inferential statistics and the sample size for each estimated model, are provided in the Supplementary Results (Supplementary Table 8 for 'C'; Supplementary Table 9 for 'Ta'; Supplementary Table 10 for 'L'; Supplementary Table 11 for 'F'; Supplementary Table 12 for 'To'; and Supplementary Table 13 for 'B').

of loss aversion under risk[30] produces all three of the significant coefficients that are in the theoretically expected direction (the three dependent variables and their associated negativity bias estimates are low openness: $B=0.03$, $P=0.01$, 95% CI = 0.01, 0.05; need for closure: $B=0.04$, $P=0.00$, 95% CI = 0.02, 0.06; and conservation versus openness to change values: $B=0.03$, $P=0.00$, 95% CI = 0.01, 0.04). But given that these tests are not independent (the dependent variables are correlated and all tests use the same data), the result may be due to chance. Indeed, only the effect of need for closure remains significant after adjusting P values for multiple hypothesis tests with arbitrary dependence structures (Supplementary Table 40). In Table 2, we again show that, with few exceptions, the Bayes factors for these models indicate that the data provide support for the null hypothesis of no relationship. It is also important to note that the results for the Toubia measure[30] of negativity bias are less clear when we add controls for the other two prospect theory parameters (value curvature and probability weighting), which are highly correlated with loss aversion (Supplementary Table 35). The more important point is that, in all cases, we cast doubt on large, positive relationships. As in Fig. 1, points represent percentage point differences in the respective dependent variable for a 1 s.d. difference in negativity bias. Thus, even with the Toubia measure[30] the relationship is unlikely to exceed ten percentage points for a 2 s.d. difference in loss aversion.

Finally, we re-estimate all models from Fig. 1 with an added interaction between negativity bias and political engagement. Figure 3 displays these interaction coefficient estimates and 95% CI. As previous work has found[11] a positive interaction between self-reported traits related to negativity bias and engagement, hypothesis 3 predicts positive coefficients in all cases. We find only one significant interaction; it is in the opposite direction expected by theory and it does not survive adjustments to P values to control the false discovery rate (Supplementary Tables 39 and 41). We thus find no support for this hypothesis. The Bayes factors for these models also indicate that the data provide evidence for the null hypothesis of no relationship (Table 3).

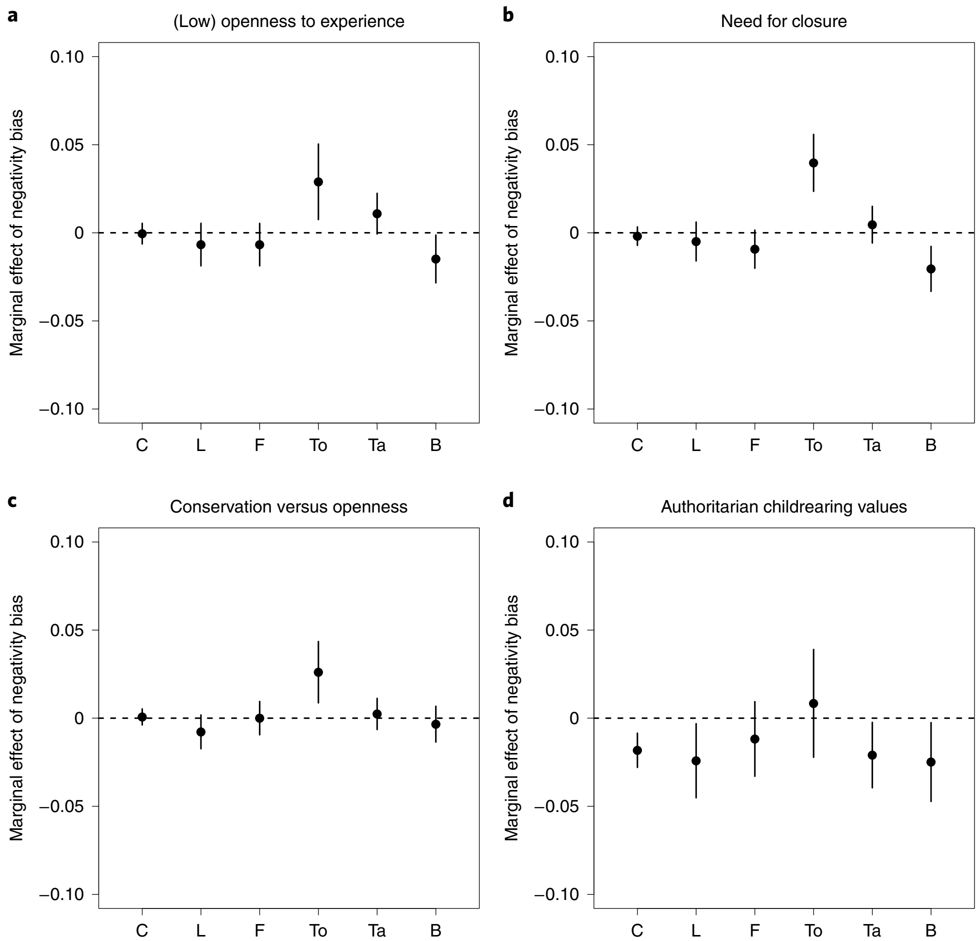

Fig. 2 | Associations between operationalizations of negativity bias and self-reported personality traits. a–d, Plots of ordinary least squares coefficients (the solid points) and their 95% CIs (vertical lines extended from points) for negativity bias for models in which reverse-coded openness to experience (**a**), need for closure (**b**), conservation versus openness to change (**c**) and authoritarian childrearing values (**d**) were regressed separately on each measure of negativity bias (specified on the *x* axis) and a set of common controls. Each coefficient represents the difference in the respective dependent variable (coded 0 to 1) for a 1 s.d. difference in the respective measure of negativity bias. All variables are coded so that positive estimates are consistent with the negativity bias hypothesis. C, samples for all negativity bias measures combined (sample sizes for openness, need for closure, conservation and authoritarianism models are, respectively, $n = 4{,}067$, 4,059, 4,065 and 4,058); L, lexical decision task ($n = 951$, 949, 951 and 950); F, flanker task ($n = 883$, 881, 881 and 881); To, loss aversion under risk estimated using Toubia et al.[30] procedure ($n = 377$, 379, 379 and 375); Ta, loss aversion under risk estimated using Tanaka et al.[29] procedure ($n = 1{,}046$, 1,042, 1,046 and 1,043); B, differential memory for negative versus positive beans in BeanFest ($n = 810$, 808, 808 and 809). Regression outputs, with all frequentist inferential statistics and the sample size for each estimated model, are provided in the Supplementary Results (Supplementary Table 15 for 'C'; Supplementary Table 16 for 'Ta'; Supplementary Table 17 for 'L'; Supplementary Table 18 for 'F'; Supplementary Table 19 for 'To'; and Supplementary Table 20 for 'B').

In summary, we find no consistent support for the hypotheses extracted from previous work. There is no consistent relationship of any of our five measures of negativity bias with any of our political dependent variables and there is no evidence for relationships conditional on political engagement. Further, we find strong evidence for the predicted relationship with personality for only one operationalization of negativity bias (Toubia measure of loss aversion[30]) and one measure of personality (need for closure). Bayes factors in the direction of the null for most tests indicate that the data provide support for the null hypothesis and this support is often strong. Even when significant, estimated relationships are substantively small and our confidence intervals rule out large, positive relationships between negativity bias and ideology or personality. Given a set of null results, ultimate conclusions rest on the quality of the data and research design. We thus undertake a series of data quality checks and robustness checks. These are reported in the Supplementary Methods and Supplementary Results and we summarize them here.

First, we are able to replicate well-established correlations in the American public opinion and political psychology literatures (Supplementary Table 3). Most relevant to the present paper, we find the expected relationships between our four self-reported personality indicators and right-wing political identity, general conservatism and social conservatism. These results suggest that our sampling frame can produce results similar to past work on personality and political ideology. Hypothesis 1 is supported when we use self-report measures.

Second, we examine our reaction time measures of negativity bias in greater depth. The lexical decision task and flanker task were modelled on two studies that were conducted in a controlled laboratory setting. Our studies, in contrast, were conducted over the internet and we have less control over the environment in which participants complete the task. While we provided a monetary incentive to take the task seriously, one may be concerned about the reliability of data such as these when collected over the internet.

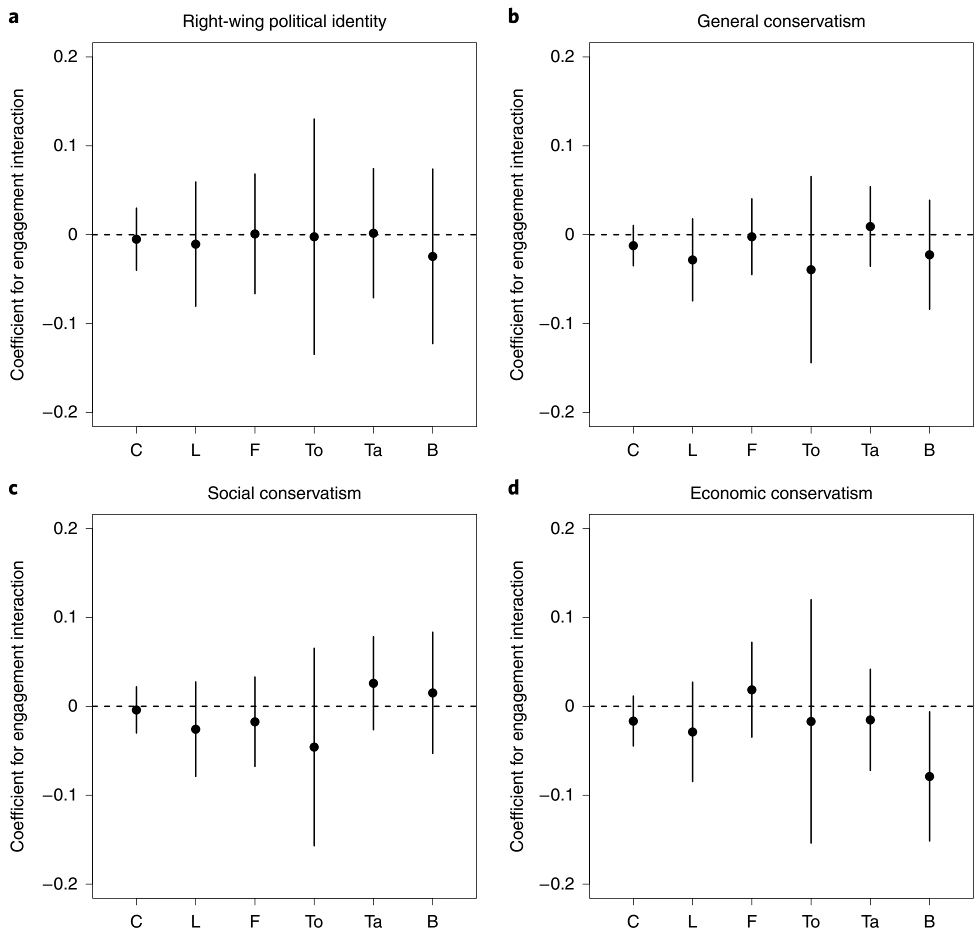

Fig. 3 | Interactions of operationalizations of negativity bias with political engagement. a–d, Plots of ordinary least squares coefficients (the solid points) and their 95% CIs (vertical lines extended from points) for models in which right-wing political identity (**a**), general conservatism (**b**), social conservatism (**c**) and economic conservatism (**d**) were regressed separately on each measure of negativity bias (specified on the *x* axis), political engagement and the interaction of negativity bias with political engagement, as well as a set of common controls. Each coefficient represents the difference in the coefficient for (standardized) negativity bias for a difference in political engagement comparing its minimum to its maximum value. All variables are coded so that positive estimates are consistent with the negativity bias hypothesis. C, samples for all negativity bias measures combined (sample sizes for identity, general, social and economic conservatism models are, respectively, $n=4{,}055$, 3,934, 3,981 and 3,996); L, lexical decision task ($n=948$, 916, 926 and 932); F, flanker task ($n=881$, 861, 869 and 872); To, loss aversion under risk estimated using Toubia et al.[30] procedure ($n=379$, 367, 375 and 373); Ta, loss aversion under risk estimated using Tanaka et al.[29] procedure ($n=1{,}041$, 1,008, 1,017 and 1,028); B, differential memory for negative versus positive beans in BeanFest ($n=806$, 782, 794 and 791). Regression outputs, with all frequentist inferential statistics and the sample size for each estimated model, are provided in the Supplementary Results (Supplementary Table 22 for 'C'; Supplementary Table 23 for 'Ta'; Supplementary Table 24 for 'L'; Supplementary Table 25 for 'F'; Supplementary Table 26 for 'To'; and Supplementary Table 27 for 'B').

Reassuringly, recent research strongly suggests that reaction time tasks conducted using web- and browser-based modes of delivery are comparable to laboratory-based studies in terms of their ability to reliably reproduce standard findings within the literature[51–54]. Looking at our own data, mean and median accuracy rates by respondent are >90% for both studies and we remove the small percentage of respondents with <80% accuracy (those who may be guessing on most trials). Reaction times also fall within a reasonable range and we remove the small number of trials that were extremely slow or fast (<200 or >5,000 ms).

Most importantly, however, we find theoretically expected relationships between reaction times and non-political variables. For the lexical decision task, we find a very strong relationship between mean reaction times for words and word frequency, which is perhaps the most important variable in the broader literature on word recognition (measured here as Lg10WF from the SUBTLEXUS database)[55,56]. A simple model with a linear and a quadratic term for word frequency accounts for 77% of the variance in mean reaction times across words in our study (Fig. 4) and the average difference in reaction times comparing low to high frequency words (~200 ms) is very similar to what has been found in recent published work using both laboratory and web-based delivery modes[52]. Thus, there does not seem to be anything intrinsically problematic about conducting this task over the internet and our reaction time data are not just noise.

Turning to the flanker task, we examine the flanker effect—the difference in mean reaction times comparing target-incongruent to target-congruent flanker trials—for both positive and negative targets separately (as in Table 1 of ref. [28]). As expected by past research on the flanker effect, when positive targets (for example, a seal pup) are flanked by negative images (for example, a spider), respondents are slower to respond than when they are flanked by positive images (mean difference = −18.55 ms, bootstrapped 95% CI = −38.02, 1.25). In contrast to this expectation, however, when negative targets are flanked by positive images, respondents are faster to respond than when they are flanked by valence-congruent negative images (mean difference = 35.57 ms, bootstrapped 95% CI = 5.51, 55.33). For both sets of targets—positive and negative—negative flankers capture attention and slow response times relative to positive flankers, which suggests a general negativity bias in our sample. While these magnitudes are smaller than a recent web-based study of the flanker effect using letters as targets and distractors (52 ms)[54], they are larger than the 10-ms flanker effect found by a previous study, which notes, 'Though a 10-millisecond difference may seem inconsequential, it is consistent with typical findings in the literature on attentional cueing'[28]. Thus, our data are reliable enough to detect flanker effects comparable in size to past research and we find a negativity bias in the sample as a whole—we simply do not find a difference in this bias across political ideology.

Third, we estimate alternative models for each operationalization of negativity bias to ensure that certain choices are not critical to the null results. For both the lexical and the flanker task, we re-estimate all models using only correct responses but the results do not support hypotheses 1–3 (Supplementary Tables 29 and 30). For the flanker task, we also estimate models identical to ref. [28], which examine the relationship between ideology and flanker effects for positive and negative targets separately. Again, we find no significant relationships (Supplementary Table 31). For both loss aversion measures, we re-estimate all models using alternative restrictions. Specifically, we first exclude respondents with extreme loss aversion coefficients (less than one-third or greater than three). This change produces similar results and no additional support for hypotheses 1–3 (Supplementary Tables 34 and 36). We also re-estimate all loss aversion models including the other two prospect theory parameters as controls (value function curvature and probability weighting parameters). Inclusion of these parameters, however, does not produce additional support for hypotheses 1–3 and reduces support for hypothesis 2 in the 2014 study (Supplementary Tables 35 and 37). For the study using BeanFest, we calculate two additional measures examined in past work using this research design: avoidance behaviour and negative valence weighting asymmetry (Supplementary Methods)[17]. Neither of these measures produces additional support for the hypotheses (Supplementary Tables 32 and 33).

Discussion

In four US datasets with a combined analytical sample size of about 4,000 respondents and across five distinct behavioural measures of negativity bias, we find no consistent evidence that negativity bias (1) is associated with right-wing ideology in terms of political identity or issue preferences; (2) is associated with 'closed' values or personality traits, such as need for closure or (low) openness to experience; or (3) interacts with political engagement to predict ideology. While we find a few statistically significant coefficients in the expected direction, these are not replicable across operationalizations and most do not survive P-value adjustments to control the false discovery rate. The more important implication of our results is that we rule out substantively large positive coefficients. That is, even if the predicted relationships exist in terms of direction, they are unlikely to be large in terms of magnitude. To support these null results, we conduct a series of data quality and robustness checks that suggest our data are comparable in quality to that of previous work and our results do not hinge on particular operationalizations or modelling choices. Moreover, our results are consistent with several recent studies that fail to replicate associations between physiological indicators of threat sensitivity and political ideology[23–26]. Our work extends this research on psychophysiology by broadening the set of measures of negativity bias, using large national survey samples in the United States, examining multiple measures of ideology, examining the link between negativity bias and a set of self-reported personality variables and examining the interaction of negativity bias with political engagement.

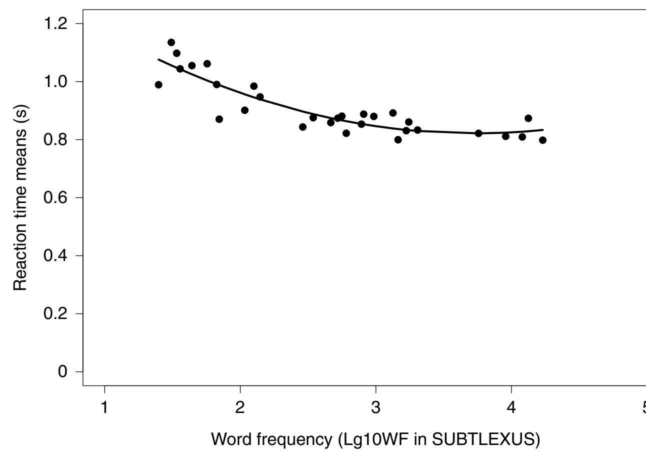

Fig. 4 | Mean reaction time to legal English words in the lexical decision task as a function of word frequency. Points are mean reaction times for target words in the lexical decision task as a function of the $\log_{10}$ of their frequency in the SUBTLEXUS corpus. The solid line contains the estimated values of mean reaction time as a function of word frequency based on coefficient estimates from an ordinary least squares regression ($n = 30$) of mean reaction time on two predictors: word frequency and its square. This model is statistically significant relative to an intercept-only model ($F(2, 27)$, $P < 0.001$, adjusted $R^2 = 0.77$).

These null findings present a puzzle. The relationship between self-reported personality traits theoretically related to negativity bias (for example, low openness to experience and disgust sensitivity) and right-wing political preferences is replicable and substantively meaningful. Our data also strongly support this pattern. It is thus surprising that seemingly straightforward behavioural measures of negativity bias are unrelated to ideology and self-reported personality. In the remainder of the paper we consider possible reasons for the conflict between results based on self-report and results using behavioural and physiological measures of negativity bias.

First, it is possible that common behavioural and physiological operationalizations of negativity bias are too unreliable to pick up these relationships. In this view, the virtue of such measures—that they bypass self-report—is inextricably linked to their major drawback—their low level of measurement reliability. A previous study, for example, find very low levels of reliability for a common measure of negativity bias based on electrodermal responses to threatening images[26]. This could explain the better performance of self-report measures, which are typically based on multi-items scales and which increase in predictive power as the number of items increases[57].

Reliability may be a particular problem for the lexical decision and flanker tasks, both of which rely on reaction time data to measure negativity bias. Moreover, since these were conducted over the internet, there may be an additional concern that respondents were inattentive or distracted while completing the tasks. As discussed above, however, data quality checks suggest that both measures reproduce well-established, non-political findings in the relevant literatures. Turning to our first measure of loss aversion under risk, previous work suggests that response error on the choice task is moderate: the average proportion of respondent choices inconsistent with their estimated preferences is 20–30% (ref. [30]). This work also suggests that the measure has out-of-sample predictive validity in both absolute terms and relative to a standard alternative measure of prospect theory preferences. Similarly, our second measure of loss aversion under risk shows predictive validity in a field study concerning real economic decisions[29]. One possible concern with this latter measure is that a substantial proportion of our sample displayed high levels of gain-seeking by choosing option B in the first paired lottery of the third sequential choice task (35%). This is reflected in a median loss aversion of slightly less than one. While we agree that further work with loss aversion would be useful, we note that a large (though smaller) percentage of respondents also made this choice in the original work using this measure which studied a much poorer population of respondents (20%)[29]. Further, excluding respondents with very high levels of gain-seeking does not change our results.

Another possibility is that the relationship between self-reported personality traits and political preferences is due to the influence of the latter on the former. That is, people adjust their self-reported traits to match the prototypes of their political identities. For example, a liberal may sort geographically and socially to be around liberal people and subsequently adopt the broader norms of that group, many of which are related to personality. There is some evidence for this. Another study, for example, find that liberals and conservatives report traits closer to their in-group's prototype relative to criterion measures of these same traits (for example, objective tests or third-party reports of traits)[12]. Other work suggests that traits may move as a function of political events[14], exert little causal effect on political ideology over time[58] and may even be endogenous to ideology and partisanship[13,15,59]. It is unlikely, however, that reverse causality can account for the entire association between self-reported traits and ideology. Indeed, this would raise questions about where such norms come from in the first place: if personality has no causal effect on political attitudes, why do personality-related partisan stereotypes exist at all? It is more likely that conformity to the political in-group strengthens the relationship than explains it away: people who are higher (lower) than average in (say) openness gravitate toward liberalism (conservatism) and the social connections formed through politics reinforce this disposition and extend it to new domains. For this reason, we do not expect reverse causality to explain away this puzzle.

Alternatively, it may be differences in beliefs rather than preferences that divide the left and right. That is, political groups may be similar in their attentiveness and reaction to negative potential outcomes but may differ in their beliefs about the likelihood of these outcomes. In a model of decision-making, for example, there would be no difference in loss aversion. Rather, right-wing 'negativity bias' would operate through the subjective probabilities assigned to negative and positive outcomes. Consistent with this idea, a substantial literature demonstrates that right-wing citizens are more likely to see the world as a threatening and dangerous place[5]. Since beliefs about the probability of negative outcomes are theoretically related to the same self-reported traits that robustly predict ideology, this is a worthwhile path for future research. Further, if differences in belief are important, left-wing citizens may be characterized by a 'negativity bias' on at least some issues. For example, one study finds that liberal identifiers self-report greater fear of climate change and overpopulation, while conservatives report greater fear of illegal immigration and terrorist attacks[60]. Similarly, the relationship between threat and ideology may depend on the type of threat, for example, economic threats (such as family in poverty) tend to be associated with left-wing political beliefs (such as government ownership of business and industry should be increased)[61].

Another possible resolution of this puzzle is that the search for a single individual difference variable that explains self-reported personality traits at the level of basic attentional and physiological processes (for example, negativity bias) is a misguided endeavour. It is a well-supported finding in behavioural genetics that complex traits are shaped by 'many genes of small effect'[62]. Analogously, complex personality traits, like openness to experience, may indeed be causally related to ideology but may be shaped by a diverse array of only weakly correlated antecedent factors, each of which has a small marginal effect on the overall trait level. Given the large number of diverse characteristics studied under the label 'negativity bias'[3], this seems like a viable, if speculative, hypothesis. The implication is that there may be a level of analysis, such as core values, beyond which it is no longer useful to theorize for the purpose of studying political behaviour because the antecedent causal factors become too weak and too numerous. While potentially dispiriting, this could also be fruitful by forcing the literature to focus on theoretical development at a single level of analysis.

Our study has methodological limitations that suggest the value of additional research. Most importantly, most of our data were collected via Lucid in the United States over a small period of time (2018–2020). Additional research with different populations is needed to determine if our null results generalize to other contexts and with alternative sampling frames. Further, while we rely on several distinct behavioural measures of negativity bias, we have not exhausted the realm of possible operationalizations. Finally, while we have tried to capture the most common ways in which political ideology is conceptualized and operationalized in the public opinion literature, we cannot speak to the relationship of negativity bias to other important dimensions, such as racial or foreign policy attitudes and we cannot make claims about opinion on specific issues.

Methods

All experiments were performed in accordance with relevant ethical guidelines and regulations. All participants in our studies provided informed and voluntary consent before beginning their respective study. All studies were approved by the institutional review board at Duke University (approved protocol nos. C0562, 2019-0038 and 2020-0523). As detailed in the Supplementary Methods, participants were compensated for their time through the relevant panel provider and had the opportunity to win additional bonus payments on the basis of their performance in the tasks to which they were assigned.

Our data consist of four US samples collected during four time periods (all n below represent the number of respondents with non-missing negativity bias measurements in each dataset but final analytical sample sizes vary by model on the basis of missingness across dependent variable and control measures—Supplementary Results): November/December 2014 ($n=381$, mean age = 50, female = 59%), August 2018 ($n=2{,}431$, mean age = 51, female = 57%), October/November 2019 ($n=1{,}056$, mean age = 50, female = 51%) and July 2020 ($n=821$, mean age = 51, female = 51%). The 2014 sample was collected by ClearVoice research through a contract with Qualtrics Panels (see Supplementary Methods for further information about all samples). Steps were taken to ensure data quality: respondents were removed from the survey if they failed either of two attention checks or if they failed to complete at least 90% of all survey items in the first half of the survey. The first attention check was a survey item that appeared to be a standard question about self-perceived knowledge in politics but embedded an instruction to select the response option 'No knowledge at all'. The second was a multiple-choice item concerning the position currently held by Barack Obama. The 2018, 2019 and 2020 samples were collected through Lucid's Marketplace (previously Fulcrum Exchange) with a restriction of the sample to only non-mobile-device users. We also removed respondents who failed either of two attention checks in the first part of the survey. In the first check, respondents had to accurately select two photos of roads (from a set of six) that contained a stop sign. The second, a multiple-choice question, asked respondents to identify the current President of the United States. We additionally removed any respondents from the data with no or duplicate Lucid identification numbers to avoid repeat survey takers.

Respondents completed one (and only one) of five different tasks to measure negativity bias, which are described in the following sections. All respondents to the 2014 sample completed measure 3; respondents to the 2018 sample were randomly assigned to either measure 1 or 2; all respondents to the 2019 sample completed measure 4; and all respondents to the 2020 sample completed measure 5. Information on background characteristics and political orientations of respondents in all samples is provided in Supplementary Table 1. Final analytical sample sizes for each regression model vary on the basis of data availability across respondents and are provided in the Supplementary Results. No statistical methods were used to predetermine sample sizes but our sample sizes are similar to, or larger than, those reported in previous publications addressing similar questions using behavioural or physiological measures of negativity bias[17,26–28,63].

We organize our discussion of these measures substantively rather than chronologically. After describing the measures in detail, we turn to our dependent variables, which are identical across the Lucid samples and very similar to the ClearVoice sample.

Measure 1: cognitive accessibility of threatening concepts. We adapt previous work on authoritarianism and use a lexical decision task (LDT) to measure negativity bias[27]. In an LDT, respondents are exposed to a series of trials. In each trial they are asked to decide, as quickly as possible while being accurate, whether a string of letters is a legal English word or a non-word. Speed and accuracy in the legal word trials are, theoretically, measures of the relative accessibility of the concepts represented by the words[64–66]. We use the difference in the average response time to positive relative to negative words as a measure of the relative accessibility of negative concepts and thus of negativity bias.

Respondents completed a total of 60 trials: 10 each for positive (for example, sunshine), negative (for example, die) and neutral words (for example, teletype) and 30 for non-words (for example, fusk). Words for each trial were randomly sampled without replacement. Our final measure was calculated similarly to the D-score in research on the implicit association test[67]. We exclude all responses <200 ms or >5,000 ms and all respondents with <80% correct responses. We then calculate, for each respondent, the difference between average response time to positive words and average response time to negative words and divide this difference by the respondent-specific standard deviation of latencies to both types of targets combined. The larger the value, the greater the variance-adjusted relative response time to positive words and thus the higher the relative accessibility of negative words. To incentivize effort and reduce measurement error, we informed respondents that the top ten performers in the LDT study would be eligible to receive a US$10.00 Amazon.com gift card and winners who provided an email address were paid after all data were collected. Additional information is available in the Supplementary Methods.

Measure 2: attentional biases to negative stimuli. Our second measure adapts recent work on political ideology using the flanker task[28]. In each of 30 trials, respondents are presented with a series of three images arrayed horizontally. The middle image is the target and the flanking images to the left and right are distractors. The goal of each trial is to identify, as quickly as possible while being accurate, whether the target image is positive or negative. The previous research uses angry and happy faces; however, we substitute three negative and three positive images drawn from the International Affective Picture System database (IAPS)[68] and the Geneva Affective Picture Database (GAPED)[69]. The images in these databases are normed for both valence and arousal and IAPS pictures have been used in other work on negativity bias and ideology[20]. The specific images used in our study were chosen for being extreme in terms of valence. We used three very positive images (flowers, a human baby and a baby seal) and three very negative images (an aimed gun, a dirty toilet and a spider).

A flanker effect is defined as the difference in average response times comparing trials where all three images are of the same valence (for example, all three are positive) to trials where the flanking images conflict in valence with the target (for example, a positive image flanked by two negative images), with conflict between target and flanker valence expected to slow response time[28]. However, if negative images strongly capture attention, this flanking effect should be small (or even reversed) for negative targets and large for positive targets. That is, negativity bias implies a comparative measure such that target-incongruent flankers are more distracting for positive relative to negative targets. A straightforward measure is thus the difference in the flanking effects comparing negative to positive targets. We again exclude all responses <200 ms or >5,000 ms and all respondents with <80% correct responses. We calculate this measure using a D-score-like procedure, as with measure 1, by dividing the difference in flanking effects for each respondent by the respondent-specific standard deviation of latencies to all types of targets combined. We again incentivized performance by informing respondents that the top ten performers would be eligible to receive a US$10.00 Amazon.com gift card. Winners who provided an email address were paid after the data were collected. Additional information is available in the Supplementary Methods.

Measure 3: loss aversion in decision-making under risk. Our third measure operationalizes negativity bias as the loss aversion parameter in prospect theory for decision-making under risk (with known outcome probabilities). Conceptually, loss aversion means that losses loom larger than nominally equivalent gains. Within cumulative prospect theory (CPT)[70], loss aversion is formalized as a weighting parameter (λ) on the value function for negative potential outcomes:

$$U\left(x_j\right) = \begin{cases} x_j^{\sigma} \text{ if } x_j > 0 \\ -\lambda(-x_j)^{\sigma} \text{ if } x_j < 0 \end{cases}$$

Values of λ above one mean that losses are worth more in terms of absolute utility than nominally equivalent gains, while values between zero and one imply the opposite (gain-seeking). We use the three-parameter version of CPT and estimate value curvature and probability weighting parameters along with loss aversion.

The questionnaire for eliciting CPT parameters consists of 16 questions (see a screenshot of a question in the Supplementary Methods). For each, the respondent is asked to choose between two prospects of the form $X=(x,P,y)$ where the probability of the outcome x is P and the probability of the outcome y is $(1-P)$. To incentivize truthful and accurate revelation of preferences, respondents were informed before completing the task that, following data collection, four respondents would be randomly selected, each of these four respondents would be given a US$20 endowment and one of their 16 choices would be selected and played for real stakes. Losses were thus paid out of the endowment and the selected respondents received US$20 plus the outcome of their selected gamble. The universe of prospects is as follows:

$$x \in (\text{US\$1, US\$30, US\$40, US\$100, US\$1,000})$$

$$y \in (-\text{US\$20}, -\text{US\$15}, -\text{US\$10}, -\text{US\$5}, \text{US\$5}, \text{US\$10}, \text{US\$30})$$

$$P \in (0.1, 0.3, 0.5, 0.7, 0.9)$$

The choice questionnaire is adaptive in the sense that the options in choices 2 through 16 are selected to minimize the expected uncertainty in the CPT parameters following that choice. This approach reduces the total number of choices necessary to achieve a given level of precision for each respondent's parameter estimates. After data collection, estimates are obtained via hierarchical Bayes, which partially pools information across respondents, again generating more efficient estimates. We provide additional details about the questionnaire and estimation process in the Supplementary Methods. Randomly selected winners were paid following data collection.

Measure 4: loss aversion in decision-making under risk. Our fourth measure also operationalizes negativity bias as loss aversion under risk but uses a different elicitation procedure adapted from past work in economics[29]. The procedure consists of three sequential tasks. In each, respondents are presented with a series of choices, each of which is presented on a separate screen (see Supplementary Table 6 for all stimuli). Each choice presents two gambles, A and B, with B increasing in relative attractiveness as the task proceeds. The goal is to determine the point at which the respondent switches from preferring A to preferring B. We operationalize this as the first choice at which the respondent chooses B (if at all)—that is, once the respondent chooses B, they move to the next task in the sequence. Our construction of the task thus imposes consistency on respondents' preferences (they cannot make incoherent choices within a task).

Switching points in the first two tasks uniquely determine the prospect theory value curvature and probability weighting parameters. Given values for these parameters, the third task determines a range of possible loss aversion values for that respondent. When this range is not bounded by zero or infinity, we calculate the midpoint of the range as the loss aversion estimate for the respondent. When one of the bounds is zero for a respondent, we use the upper bound as their estimate of loss aversion. When one of the bounds is infinity for a respondent, we use the lower bound as their estimate. We discuss the procedure for calculating parameter estimates in detail in the Supplementary Methods and provide the calculation tables in Supplementary Table 7. To incentivize truthful and accurate revelation of preferences, we informed respondents that five participants would be randomly selected to have one of their choices played out for real stakes with a US$21 initial endowment. Winners who provided an email were paid after data collection.

Measure 5: BeanFest. Our fifth measure builds on recent work in psychology using the BeanFest experimental framework[17]. On each trial of BeanFest, respondents are presented with a line drawing of a 'bean' which can vary along two dimensions: shape (from circular to oblong) and number of spots (few to many) (Supplementary Fig. 4). Their task is to choose whether to 'accept' the bean or 'reject' it. If they accept the bean, they receive a positive or negative payoff, with half the beans worth +10 points and half worth −10 points. If they reject the bean, they neither gain nor lose points but receive no information about the bean's value.

Respondents begin each round of the game with 50 points. They win the round if they get to 100 points and lose if they reach zero points (the round continues until they reach one or the other endpoint). Unknown to the respondents, the beans are drawn from a two-dimensional matrix within which clusters of spatially proximate beans share both attribute and point values (Supplementary Fig. 5). It is thus possible to learn about the characteristics of helpful and harmful beans and win the game with high probability. Each respondent played 6 practice trials followed by 108 trials in the game phase. Following previous work, there were 36 unique beans that were each viewed three times[17]. The first 12 beans viewed were in a predetermined order, with the remaining beans being randomly displayed without replacement (every respondent saw every bean exactly three times). We provided a monetary incentive to respondents. For each win, US$2 was added to the respondent's 'bonus payment total'. Each loss subtracted US$2 from this total. Respondents were told that, at the end of the study, we would randomly select ten respondents and those respondents would receive their bonus payment total as an Amazon.com gift card.

At the end of the game phase, respondents completed a test phase in which they were presented with 56 beans, 36 of which they had previously seen and 20 of which were new but were drawn from the same two-dimensional matrix as the game phase beans. Respondents' task in the test phase is to determine whether each bean is helpful (+10 points) or harmful (−10 points). They were told that for each correct answer we would add US$0.10 to their bonus payment total.

Following previous work, our primary measure of negativity bias is the respondent's 'learning asymmetry'—the difference in their accuracy during the test phase comparing negative to positive beans[17]. A high value of this measure suggests greater relative accuracy for negative beans and thus a stronger negativity bias. We consider two additional measures in the section on robustness. We provide additional details for this study in the Supplementary Methods. Randomly selected winners who provided an email were paid their bonus payment total after data collection was complete.

Dependent variables and controls. We examine four measures of political ideology. First, we operationalize right-wing political identity as the average of a seven-point, branching partisanship scale and a seven-point ideological identification scale. Second, we operationalize general conservatism as the factor scores estimated from one-dimensional principal factors analyses of several political value and policy items. These include: (1) four items tapping moral traditionalism, (2) three items tapping support for limited government and (3) ten policy items, including preferences on gay marriage, abortion, immigration, affirmative action for African-Americans, government health insurance, social security privatization, minimum wage, tax rates on wealthy Americans, military spending and import restrictions. Our third measure operationalizes social conservatism as the factor scores calculated from one-dimensional principal factors analyses of the four moral traditionalism items and policy preferences on gay marriage, abortion, immigration and affirmative action for African-Americans. Our fourth measure, economic conservatism, is operationalized similarly using the limited government items and policy preferences on insurance, Social Security, minimum wage and tax rates. In all cases, the factor analyses were run separately within each dataset, the factor scores were calculated for respondents within that dataset and the datasets were then merged. This allows for the covariance structure of the items to vary across datasets. Items for all dependent variables are provided in the Supplementary Methods.

We also consider four personality constructs that have been the focus of previous work on political preferences: openness to experience[2,6], the need for non-specific cognitive closure[10,71], conservation versus openness to change values[38] and authoritarian childrearing values[35,72]. Our measures are identical in all samples for the last two constructs but differ for the first two comparing the 2014 sample to the 2018–2020 samples. The changes were made to increase the validity and reliability of these two measures. All measures, however, use common items within the literature. For each personality construct, we average all available items for each respondent. All personality items are listed in the Supplementary Methods.

In all models we control for age, gender, race and ethnicity, education, household income and employment status. Following recent research, we measure political engagement as the average of political knowledge (on five objective knowledge questions) and both self-reported political interest and news consumption[11].

We recode all measures of negativity bias to have a mean of zero and standard deviation of one before analysis. All other variables, including all dependent variables, are coded from zero (minimum) to one (maximum). All regression models were estimated by ordinary least squares.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The data that support the findings of this study are publicly available in Harvard Dataverse with the identifier https://doi.org/10.7910/DVN/GRXTZY[73].

Code availability

The code necessary for reproducing the findings of this study are publicly available in Harvard Dataverse with the identifier https://doi.org/10.7910/DVN/GRXTZY[73].

Received: 1 July 2021; Accepted: 1 March 2022;
Published online: 9 May 2022

References

1. Aarøe, L., Petersen, M. B. & Arceneaux, K. The behavioral immune system shapes political intuitions: why and how individual differences in disgust sensitivity underlie opposition to immigration. *Am. Political Sci. Rev.* **111**, 277–294 (2017).
2. Gerber, A. S., Huber, G. A., Doherty, D., Dowling, C. M. & Ha, S. E. Personality and political attitudes: relationships across issue domains and political contexts. *Am. Political Sci. Rev.* **104**, 111–133 (2010).
3. Hibbing, J. R., Smith, K. B. & Alford, J. R. Differences in negativity bias underlie variations in political ideology. *Behav. Brain Sci.* http://digitalcommons.unl.edu/poliscifacpub/67 (2014).
4. Jost, J. T., Glaser, J., Kruglanski, A. W. & Sulloway, F. J. Political conservatism as motivated social cognition. *Psychol. Bull.* **129**, 339–375 (2003).
5. Jost, J. T. Ideological asymmetries and the essence of political psychology. *Political Psychol.* **38**, 167–208 (2017).
6. Mondak, J. J. *Personality and the Foundations of Political Behavior* (Cambridge Univ. Press, 2010).
7. Kahneman, D., Knetsch, J. L. & Thaler, R. H. Anomalies: the endowment effect, loss aversion, and status quo bias. *J. Econ. Perspect.* **5**, 193–206 (1991).
8. Tversky, A. & Kahneman, D. Prospect theory: an analysis of decision under risk. *Econometrica* **47**, 263–291 (1979).
9. Federico, C. M. & Malka, A. The contingent, contextual nature of the relationship between needs for security and certainty and political preferences: evidence and implications. *Political Psychol.* **39**, 3–48 (2018).
10. Federico, C. M. & Goren, P. in *Social and Psychological Bases of Ideology and System Justification* (eds Jost, J. T. et al.) 267–291 (Oxford Scholarship Online, 2009).
11. Johnston, C. D., Lavine, H. G. & Federico, C. M. *Open Versus Closed: Personality, Identity, and the Politics of Redistribution* (Cambridge Univ. Press, 2017).
12. Ludeke, S., Tagar, M. R. & DeYoung, C. G. Not as different as we want to be: attitudinally consistent trait desirability leads to exaggerated associations between personality and sociopolitical attitudes. *Political Psychol.* **37**, 125–135 (2016).
13. Bakker, B. N., Lelkes, Y. & Malka, A. Reconsidering the link between self-reported personality traits and political preferences. *Am. Political Sci. Rev.* **115**, 1482–1498 (2021).
14. Boston, J., Homola, J., Sinclair, B., Torres, M. & Tucker, P. D. The dynamic relationship between personality stability and political attitudes. *Public Opin. Q.* **82**, 843–865 (2018).
15. Hatemi, P. K., Crabtree, C. & Smith, K. B. Ideology justifies morality: political beliefs predict moral foundations. *Am. J. Political Sci.* **63**, 788–806 (2019).
16. Fiagbenu, M. E., Proch, J. & Kessler, T. Of deadly beans and risky stocks: political ideology and attitude formation via exploration depend on the nature of the attitude stimuli. *Br. J. Psychol.* **112**, 342–357 (2021).
17. Shook, N. J. & Fazio, R. H. Political ideology, exploration of novel stimuli, and attitude formation. *J. Exp. Soc. Psychol.* **45**, 995–998 (2009).
18. Oxley, D. R. et al. Political attitudes vary with physiological traits. *Science* **321**, 1667–1670 (2008).
19. Smith, K. B., Oxley, D., Hibbing, M. V., Alford, J. R. & Hibbing, J. R. Disgust sensitivity and the neurophysiology of left–right political orientations. *PLoS ONE* **6**, e25552 (2011).
20. Ahn, W.-Y. et al. Nonpolitical images evoke neural predictors of political ideology. *Curr. Biol.* **24**, 2693–2699 (2014).
21. Amodio, D. M., Jost, J. T., Master, S. L. & Yee, C. M. Neurocognitive correlates of liberalism and conservatism. *Nat. Neurosci.* **10**, 1246–1247 (2007).
22. Kanai, R., Feilden, T., Firth, C. & Rees, G. Political orientations are correlated with brain structure in young adults. *Curr. Biol.* **21**, 677–680 (2011).
23. Bakker, B. N., Schumacher, G., Gothreau, C. & Arceneaux, K. Conservatives and liberals have similar physiological responses to threats. *Nat. Hum. Behav.* **4**, 613–621 (2020).
24. Fournier, P., Soroka, S. & Nir, L. Negativity biases and political ideology: a comparative test across 17 countries. *Am. Political Sci. Rev.* **114**, 775–791 (2020).
25. Knoll, B. R., O'Daniel, T. J. & Cusato, B. Physiological responses and political behavior: three reproductions using a novel dataset. *Res. Politics* https://doi.org/10.1177/2053168015621328 (2015).
26. Osmundsen, M., Hendry, D., Laustsen, L., Smith, K. & Petersen, M. B. The psychophysiology of political ideology: replications, reanalyses and recommendations. *J. Politics* **84**, 50–66 (2022).
27. Lavine, H., Lodge, M., Polichak, J. & Taber, C. Explicating the black box through experimentation: studies of authoritarianism and threat. *Political Anal.* **10**, 343–361 (2002).
28. McLean, S. P. et al. Applying the flanker task to political psychology: a research note. *Political Psychol.* **35**, 831–840 (2014).
29. Tanaka, T., Camerer, C. F. & Nguyen, Q. Risk and time preferences: linking experimental and household survey data from Vietnam. *Am. Econ. Rev.* **100**, 557–571 (2010).
30. Toubia, O., Johnson, E., Evgeniou, T. & Delquié, P. Dynamic experiments for estimating preferences: an adaptive method of eliciting time and risk parameters. *Manag. Sci.* **59**, 613–640 (2013).
31. Abdellaoui, M., Bleichrodt, H., l'Haridon, O. & Van Dolder, D. Measuring loss aversion under ambiguity: a method to make prospect theory completely observable. *J. Risk Uncertain.* **52**, 1–20 (2016).
32. Kruglanski, A. W. & Webster, D. M. Motivated closing of the mind: seizing and freezing. *Psychol. Rev.* **103**, 263 (1996).
33. McCrae, R. R. & Costa, P. T. *Personality in Adulthood: A Five-Factor Theory Perspective* (Guilford Press, 2003).
34. Schwartz, S. H. in *Advances in Experimental Social Psychology* Vol. 25 (ed. Zanna, M. P.) 1–65 (Elsevier, 1992).
35. Stenner, K. *The Authoritarian Dynamic* (Cambridge Univ. Press, 2005).
36. Smith, K. B., Oxley, D. R., Hibbing, M. V., Alford, J. R. & Hibbing, J. R. Linking genetics and political attitudes: reconceptualizing political ideology. *Political Psychol.* **32**, 369–397 (2011).
37. Feldman, S. & Johnston, C. Understanding the determinants of political ideology: implications of structural complexity. *Political Psychol.* **35**, 337–358 (2014).

38. Malka, A., Soto, C. J., Inzlicht, M. & Lelkes, Y. Do needs for security and certainty predict cultural and economic conservatism? A cross-national analysis. *J. Pers. Soc. Psychol.* **106**, 1031 (2014).
39. Abramowitz, A. *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy* (Yale Univ. Press, 2010).
40. Freeze, M. & Montgomery, J. M. Static stability and evolving constraint:preference stability and ideological structure in the mass public. *Am. Politics Res.* **44**, 415–447 (2016).
41. Treier, S. & Hillygus, D. S. The nature of political ideology in the contemporary electorate. *Public Opin. Q.* **73**, 679–703 (2009).
42. Arikan, G. & Sekercioglu, E. Authoritarian predispositions and attitudes towards redistribution. *Political Psychol.* **40**, 1099–1118 (2019).
43. Crowson, H. M. Are all conservatives alike? A study of the psychological correlates of cultural and economic conservatism. *J. Psychol.* **143**, 449–463 (2009).
44. Johnston, C. D. Dispositional sources of economic protectionism. *Public Opin. Q.* **77**, 574–585 (2013).
45. Jedinger, A. & Burger, A. M. The role of right-wing authoritarianism and political sophistication in shaping attitudes toward redistribution. *Eur. J. Soc. Psychol.* **49**, 560–573 (2019).
46. Johnston, C. D. Authoritarianism, affective polarization, and economic ideology. *Political Psychol.* **39**, 219–238 (2018).
47. Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H. & Grasman, R. J. C. P. Bayesian hypothesis testing for psychologists: a tutorial on the Savage–Dickey method. *Cogn. Psychol.* **60**, 158–189 (2010).
48. Makowski, D., Ben-Shachar, M. S. & Lüdecke, D. bayestestR: describing effects and their uncertainty, existence and significance within the Bayesian framework. *J. Open Source Softw.* **4**, 1541 (2019).
49. Makowski, D., Ben-Shachar, M. S., Chen, S. & Lüdecke, D. Indices of effect existence and significance in the Bayesian framework. *Front. Psychol.* **10**, 2767 (2019).
50. Raftery, A. E. Bayesian model selection in social research. *Sociol. Methodol.* **25**, 111–163 (1995).
51. Barnhoorn, J. S., Haasnoot, E., Bocanegra, B. R. & van Steenbergen, H. QRTEngine: an easy solution for running online reaction time experiments using Qualtrics. *Behav. Res. Methods* **47**, 918–929 (2015).
52. Hilbig, B. E. Reaction time effects in lab- versus web-based research: experimental evidence. *Behav. Res. Methods* **48**, 1718–1724 (2016).
53. Kim, D., Lowder, M. W. & Choi, W. Effects of print exposure on an online lexical decision task: a direct replication using a web-based experimental procedure. *Front. Psychol.* https://doi.org/10.3389/fpsyg.2021.710663 (2021).
54. Semmelmann, K. & Weigelt, S. Online psychophysics: reaction time effects in cognitive experiments. *Behav. Res. Methods* **49**, 1241–1260 (2017).
55. Brysbaert, M. & New, B. Moving beyond Kučera and Francis: a critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behav. Res. Methods* **41**, 977–990 (2009).
56. Brysbaert, M. et al. The word frequency effect. *Exp. Psychol.* **58**, 412–424 (2011).
57. Bakker, B. N. & Lelkes, Y. Selling ourselves short? How abbreviated measures of personality change the way we think about personality and politics. *J. Politics* **80**, 1311–1325 (2018).
58. Smith, K. B., Alford, J. R., Hibbing, J. R., Martin, N. G. & Hatemi, P. K. Intuitive ethics and political orientations: testing moral foundations as a theory of political ideology. *Am. J. Political Sci.* **61**, 424–437 (2017).
59. Ciuk, D. J. Assessing the contextual stability of moral foundations: evidence from a survey experiment. *Res. Politics* https://doi.org/10.1177/2053168018781748 (2018).
60. Jost, J. T., Stern, C., Rule, N. O. & Sterling, J. The politics of fear: is there an ideological asymmetry in existential motivation? *Soc. Cogn.* **35**, 324–353 (2017).
61. Brandt, M. J. et al. The association between threat and politics depends on the type of threat, the political domain, and the country. *Pers. Soc. Psychol. Bull.* **47**, 324–343 (2021).
62. Plomin, R., DeFries, J. C., Knopik, V. S. & Neiderhiser, J. M. Top 10 replicated findings from behavioral genetics. *Perspect. Psychol. Sci.* **11**, 3–23 (2016).
63. Bakker, B. N., Schumacher, G., Gothreau, C. & Arceneaux, K. J. N. H. B. Conservatives and liberals have similar physiological responses to threats. *Nat. Hum. Behav.* **4**, 613–621 (2020).
64. Higgins, E. T. in *Social Psychology: Handbook of Basic Principles* (eds Higgins, E. T. & Kruglanski, A. W.) (Guilford Press, 133–168 1996).
65. Lodge, M. & Taber, C. S. The automaticity of affect for political leaders, groups, and issues: an experimental test of the hot cognition hypothesis. *Political Psychol.* **26**, 455–482 (2005).
66. Meyer, D. E. & Schvaneveldt, R. W. Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *J. Exp. Psychol.* **90**, 227–234 (1971).
67. Greenwald, A. G., Nosek, B. A. & Banaji, M. R. Understanding and using the implicit association test: I. An improved scoring algorithm. *J. Pers. Soc. Psychol.* **85**, 197–216 (2003).
68. Lang, P. J., Bradley, M. M. & Cuthbert, B. N. *International Affective Picture System (IAPS): Affective Ratings of Pictures and Instruction Manual* Technical Report A-8 (Univ. Florida, 2008).
69. Dan-Glauser, E. S. & Scherer, K. R. The Geneva affective picture database (GAPED): a new 730-picture database focusing on valence and normative significance. *Behav. Res. Methods* **43**, 468 (2011).
70. Tversky, A. & Kahneman, D. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* **5**, 297–323 (1992).
71. Van Hiel, A., Pandelaere, M. & Duriez, B. The impact of need for closure on conservative beliefs and racism: differential mediation by authoritarian submission and authoritarian dominance. *Pers. Soc. Psychol. Bull.* **30**, 824–837 (2004).
72. Hetherington, M. J. & Weiler, J. D. *Authoritarianism and Polarization in American Politics* (Cambridge Univ. Press, 2009).
73. Johnston, C. & Madson, G. Reproduction materials for negativity bias, personality, and political ideology. *Harvard Dataverse* https://doi.org/10.7910/DVN/GRXTZY (2022).

Acknowledgements

Funding for this project was provided by Duke University and the Worldview Lab at the Kenan Institute for Ethics. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

C.D.J. and G.J.M. contributed to research design, data collection, analyses and write-up.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41562-022-01327-5.

Correspondence and requests for materials should be addressed to Christopher D. Johnston.

Peer review information *Nature Human Behaviour* thanks Bert Bakker and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

nature portfolio

Corresponding author(s): Christopher Johnston

Last updated by author(s): Mar 1, 2022

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. F, t, r) with confidence intervals, effect sizes, degrees of freedom and P value noted *Give P values as exact values whenever suitable.*
☐	☒	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's d, Pearson's r), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Survey data was collected using Qualtrics survey software. Custom JavaScript was included in the Qualtrics survey (.qsf) files for the 2018, 2019, and 2020 Lucid studies. The incentivized decision making task in the 2014 study was run in a website built by another research group. Final estimates of prospect theory parameters for the 2014 study, which are based on the decisions of the respondents in the incentivized task, were calculated by this other research group and shared with us. The current version of their software is available at the following website: https://sites.google.com/a/decisionsciences.columbia.edu/cds-wiki/deep-software.
Data analysis	Data was analyzed using code written in R (3.6.2) and Stata (17).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

The data that support the findings of this study are publicly available in Harvard Dataverse with the identifier https://doi.org/10.7910/DVN/GRXTZY

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	The study is observational and quantitative.
Research sample	Our data consist of four U.S. samples collected during four time periods: November/December 2014, August 2018, October/November 2019, and July 2020. The 2014 sample was collected by ClearVoice research through a contract with Qualtrics Panels. The 2018, 2019, and 2020 samples were collected through Lucid's Marketplace (previously Fulcrum Exchange) with a restriction of the sample to only non-mobile-device users. Quotas for age, sex, race and ethnicity, and region, based on the 2016 American Community Survey, were used for the Lucid samples. The mean age for each of these four samples is, respectively, 50, 50, 50, and 52. The percent female in each is 59%, 55%, 52%, and 52%. Additional demographic information is available in the supplemental materials. These organizations were used to obtain our samples because we believed that, among the options available to us given the budget constraints of our project and our sample size goals, they would provide maximally representative samples of U.S. adults. Nonetheless, respondents opt-in to participation and thus likely differ from a truly representative sample in unobserved ways.
Sampling strategy	Samples were obtained through intermediaries (Qualtrics Panels and Lucid Marketplace) that contract with survey panel providers. The panel providers use a variety of strategies to recruit participants (e.g., online advertisements) and typically provide incentives for participation in surveys (e.g., gift cards). The sample size for the 2014 study was the maximum size possible given the funding received for the study. The 2018 and 2019 studies had a target sample size of 1,000 per measure of negativity bias (so 3,000 total for 2018 and 1,000 for 2019). The 2020 study had a target sample size of 1,500. These sample sizes were chosen to balance budget constraints at the time of data collection with the our desire to have sufficient power to detect small effect sizes. Final analytical sample sizes are smaller than our targets due to respondent drop off, exclusions (see below), and missing data on model variables.
Data collection	All data was collected over the internet using Qualtrics survey software and, in one case (the incentivized decision making task in the 2014 study), a website built by another researcher group.
Timing	Our data consist of four U.S. samples collected during four time periods: November/December 2014, August 2018, October/November 2019, and July 2020.
Data exclusions	For the lexical decision task and the flanker task, we exclude all responses below 200 milliseconds or above 5000 milliseconds and all respondents with less than 80% correct responses. 45 respondents were excluded from the lexical decision task on this basis, while 85 were excluded from the flanker task. These exclusions are intended to remove respondents who are not taking the task seriously, for example, by speeding or straight-lining. For model estimation, we used listwise deletion for respondents with unavailable data on model variables. These exclusion conditions were pre-established.
Non-participation	It is not possible to determine how many potential participants declined an offer to participate as these offers are made by the panel providers and we only receive data for participants who choose to enter our survey. 746 respondents to the 2014 sample were forcibly dropped from the survey for failing either of two attention checks. In the 2018, 2019, and 2020 samples, some respondents dropped out without completing the survey while others were forcibly dropped for failing either of two attention checks. The total number failing to complete the survey for these three samples are 835, 120, and 853, respectively.
Randomization	Allocation to one of three measures of negativity bias in the 2018 sample was random.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study	
☒	☐	Antibodies
☒	☐	Eukaryotic cell lines
☒	☐	Palaeontology and archaeology
☒	☐	Animals and other organisms
☐	☒	Human research participants
☒	☐	Clinical data
☒	☐	Dual use research of concern

Methods

n/a	Involved in the study	
☒	☐	ChIP-seq
☒	☐	Flow cytometry
☒	☐	MRI-based neuroimaging

Human research participants

Policy information about studies involving human research participants

Population characteristics	See above.
Recruitment	Participants were recruited through intermediaries (Qualtrics Panels and Lucid Marketplace) which contract with survey panel providers for participants. The panel providers use a variety of strategies to recruit participants (e.g., on line advertisements) and typically provide incentives for participation in surveys (e.g., gift cards). Since these are opt-in samples, the sampling frame cannot be assumed representative of the target population under study (U.S. adults). Individuals who self-select into participation in our studies are likely to be more experienced at taking surveys than the average U.S. adult and are likely to be more interested in politics. They are also likely to have higher levels of unemployment. Higher levels of political interest may bias results in the direction of stronger estimated relationships between measures of negativity bias and political ideology.
Ethics oversight	The studies reported in our manuscript were approved by the Duke University Institutional Review Board.

Note that full information on the approval of the study protocol must also be provided in the manuscript.