




Widespread misperceptions of long-term attitude change

Adam M. Mastroianni^{a,1} , and Jason Dana^b

^aColumbia Business School, Columbia University, New York, NY 10027; and ^bSchool of Management, Yale University, New Haven, CT 06520

Edited by Elke Weber, Departments of Psychology & Public Affairs, Andlinger Center for Energy and Environment, Princeton University, Princeton, NJ; received May 4, 2021; accepted January 17, 2022

America is embroiled in cultural wars over abortion, immigration, gun control, climate change, religion, race, gender, and everything in between. Do people know how much attitudes have shifted on these contentious issues, or even which side is winning? Two pre-registered studies suggest they do not. In Study 1, we asked a nationally representative sample of participants to estimate how 51 different attitudes had changed over time and compared their estimates to actual polling data. Participants overestimated the amount of change on 29 attitudes (57%), underestimated change on 10 attitudes (20%), estimated change in the wrong direction on 10 attitudes (20%), and estimated change correctly on only two attitudes (4%). In most cases, participants did not know whether an attitude had grown to a majority or shrunk to a minority. These misperceptions had little to do with participants' demographics or ideologies and seemed instead to arise from a stereotype that the present is far more liberal than the past. Indeed, in Study 2, participants overestimated the liberal shift on most attitudes, believing that the liberal side had gained ground that it had in fact lost (e.g., gun control), or already held (e.g., climate change), or never held (e.g., religion). In three additional preregistered studies, we found that these misperceptions could justify policies that would otherwise seem objectionable. Overall, our findings suggest that widely shared stereotypes of the past lead people to misperceive attitude change, and these misperceptions can lend legitimacy to policies that people may not actually prefer.

attitudes | misperception | social change

American attitudes have shifted dramatically in the past 50 years, but perhaps not the way one might think. Consider the following claims: Most voters would have vehemently rejected a Black presidential candidate in the 1970s; a few decades later they elected Barack Obama twice. Climate change, denied and dismissed even 20 years ago, is finally accepted as a real and serious problem by most Americans. Belief in God has gone from ubiquitous to uncommon. A deluge of mass shootings has sapped what used to be near-universal support for gun rights. Meanwhile, Donald Trump's presidency reversed long-growing support for immigration.

Perhaps surprisingly, all of these claims are at odds with surveys of Americans' opinions. Most Americans have expressed willingness to vote for a Black president and concern about climate change since pollsters started asking in 1978 and 1990, respectively (1, 2). Belief in God has decreased only 10% over the past 60 years (3). Support for bans on assault rifles and handguns has decreased, not increased, since the 1980s and 1990s (4), and the years following Trump's election saw some of the highest support for immigration in decades (5).

Do Americans understand these changes? If they do not, it is concerning. People love to hop on a bandwagon: they are more likely to pick the vegetarian option, use less water, stop smoking, support women's rights, limit their screen time, avoid sugary drinks, and give more to charity when they believe more and more people are doing it too (6–9). Attitudes are invisible, however, which makes it easy to miss bandwagons that do exist and invent ones that do not. Imaginary

trends could thus spuriously shape people's real beliefs and behaviors.

Misperceptions may also distort people's policy preferences. For example, believing that the United States has made great strides toward racial equality undermines people's support for affirmative action (10, 11). White Americans tend to overestimate that progress in economic terms (12), so perhaps they also overstate declines in prejudice and bigotry. They might feel differently about affirmative action if they understood that some racial attitudes have stayed surprisingly stable.

Misperceiving attitude change could also mislead people about how the world works and how to change it. For example, Americans may assume that recent mass shootings have horrified people into supporting gun control. This may lead Americans who favor gun control to think there is no need for them to persuade their pro-gun neighbors or donate to anti-gun candidates, while Americans who oppose gun control may think persuasion and donation are futile. In fact, decades of gun tragedies have left the American public less in favor of gun control than ever, which means the losing side may think it is winning and the winning side may think it is losing. When people misperceive trends like these, they may never do what it takes to reverse them or capitalize on them.

Perhaps, however, Americans track changes in attitudes accurately on average. People's aggregate estimates tend to be accurate as long as they know more than nothing and have diverse and independent opinions (13, 14). Crowds can predict stock returns (15), integrate information (16), and accurately

Significance

People change when they think others are changing, but people misperceive others' changes. These misperceptions may bedevil people's efforts to understand and change their social worlds, distort the democratic process, and turn imaginary trends into real ones. For example, participants believed that Americans increasingly want to limit immigration, which they said justifies tighter borders. However, participants also said that limiting immigration would not be right if attitudes had shifted against it—which is what actually occurred. Our findings suggest that the national discourse around contentious social issues, policies resulting from that discourse, and perhaps the opinions that drive discourse in the first place would be very different if people better understood how attitudes have and have not changed.

Author contributions: A.M.M. and J.D. designed research; A.M.M. performed research; A.M.M. analyzed data; and A.M.M. and J.D. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: adam.m.mastroianni@gmail.com.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2107260119/-DCSupplemental>.

Published March 7, 2022.

order historical events (17). They might be similarly spot-on when it comes to shifts in attitudes.

On the other hand, people sometimes misperceive the attitudes around them. For example, college students overestimate their peers' comfort with drinking (18), taxpayers mistakenly believe their fellow citizens are more willing to commit fraud than themselves (19), Saudi Arabian men underestimate how much other men approve of women working outside the home (20), women believe that men desire a thinner body type than men really do (21), children assume that everyone accepts bullying except for themselves (22), and legislators overestimate how conservative their constituents are (23). These studies, however, did not select attitudes systematically or investigate multiple attitudes at once, so it is possible that people are generally well-informed and that their occasional misperceptions are simply more likely to garner scientific attention. Nevertheless, if people misperceive where attitudes stand today, perhaps they also misperceive how attitudes have changed.

Some people may be less prone to these misperceptions. Groups that have less power in society, like women (24) and racial minorities (25), may pay a higher price for going against prevailing attitudes and thus might be more attuned to changes in those attitudes. Educated people may simply be more informed. Older people may better track changes in attitudes because they witnessed more of the past firsthand, and they may especially outperform younger people when estimating changes across greater spans of time. On the other hand, political partisans may be more likely to misperceive attitude change because they seek out biased information (26) and apply more bias to it (27).

Bias and Inaccuracy. When estimating attitude change, people can make two different kinds of mistakes: 1) they may randomly miss the mark (*inaccuracy*) and 2) they may consistently over- or underestimate (*bias*). In wise crowds, people's random mistakes cancel out, leaving the average of their estimates close to the right answer. In unwise crowds, either overestimators or underestimators predominate, and the average of their estimates is too high or too low—that is, biased. Inaccurate estimates and biased estimates may fall equally far from the mark, on average, but biased estimates fall consistently in one direction.

Inaccuracy and bias are both problems, but bias is a bigger problem. Nobody is perfect, so we should always expect people to be at least a little inaccurate. As people become more inaccurate, however, we might become more concerned. People do not worry when an archer's salvo of arrows land a few inches left or right of the bullseye, but people do worry when those arrows land in the crowd. Bias, on the other hand, is not simply imperfection, noise, or error. When an archer's arrows all land north of the bullseye, it means something is amiss with her equipment or technique.

The Present Research. We seek to answer three questions. Do people know how important attitudes have changed in the long term? If not, what factors contribute to their bias and inaccuracy? And what might be the consequences of misunderstanding these changes? We began by asking a large, nationally representative sample to estimate how a diverse set of attitudes had changed over time and compared their estimates to actual polling data. This allowed us to measure their bias and inaccuracy, and investigate the factors underlying any misperceptions.

Study 1

We obtained 51 public opinion questions and associated data from national survey organizations and repositories, including the General Social Survey (1), Gallup (3–5, 28–31), Pew Research Center (32), the American National Election Studies (33), and the Roper Center for Public Opinion iPoll Database

(34, 35). These items were selected based on a pilot study, described in *SI Appendix*, in which participants nominated social issues about which opinions have changed in the last 50 years.

We presented these items to a sample of 943 participants assembled by the Prolific platform to be representative of the United States in terms of age, race, and gender. These participants viewed a random sample of 20 of the 51 items and estimated how people had responded at the earliest and latest time points for which data were available.

Results.

Estimated and actual change. We calculated estimated change for each participant by subtracting their estimate for the first time point from their estimate for the second time point. For instance, if a participant estimated that 30% of Americans favored legalizing marijuana in 1973 and 70% of Americans favored it in 2018, their estimated change was 40. We calculated actual change for each item in the same way using the item's polling data.

Analysis plan. Our analyses aimed to answer four questions about how well participants estimated attitude change. These analyses are illustrated in *SI Appendix, Fig. S2*.

Bias on individual items. First, were estimates biased for individual items? To find out, we compared estimated and actual change on each item using preregistered two-tailed, one sample *t* tests. Each result fell into one of four categories. Estimates were “correct” if the mean of estimated change did not differ significantly from actual change, “overestimated” if the mean was greater than a positive actual change or less than a negative actual change, “underestimated” if the mean was less than a positive actual change or greater than a negative actual change, and “wrong direction” if the mean was significantly different and opposite in sign from actual change. For instance, if the actual amount of change was –50, –50 would be a correct estimate, –75 would be an overestimate, –25 would be an underestimate, and +25 would be in the wrong direction.

Bias across all items. Second, did participants tend to be biased across items? This question requires a different analysis, for two reasons. First, the same estimated change can indicate different bias across different items—for instance, –25 is an underestimate when actual change is –50, but it is an estimate in the wrong direction when actual change is +50—so we cannot simply average all estimated changes, average all actual changes, and compare the two. Second, estimating public support of a position means different things for different questions. For example, one item asked about support for banning the death penalty and another asked about support for legalizing marijuana. We can, however, gauge whether people overestimated or underestimated the overall magnitude of change by comparing the absolute value of each estimated change to the absolute value of the actual change. Regardless of items and direction of change, this analysis shows us whether participants thought more or less change had occurred than had actually occurred. We tested this by fitting the following model using the *lme4* package in R (36) and calculating *P* values using the *lmerTest* package (37):

$$\begin{aligned} &|\text{estimated change}| - |\text{actual change}| \\ &\sim \text{intercept} + \text{participant random effects} \\ &+ \text{item random effects.} \end{aligned} \quad [1]$$

In this model, a positive and significant intercept would indicate that people tend to overestimate how much change has occurred; a negative and significant intercept would indicate underestimation.

Inaccuracy across all items. Third, how accurate were participants' estimates, on average? We calculated accuracy by taking the absolute value of the difference between estimated and actual change across all participants and items, and fitting the following model:

$$| \text{estimated change} - \text{actual change} | \sim \text{intercept} \\ + \text{participant random effects} + \text{item random effects.} \quad [2]$$

To make the intercept of this model more interpretable, we fit it separately for the items rated on a 100-point scale (such as percentage agreement, $n = 46$) and the items rated on a 7-point Likert scale ($n = 5$).

Were participants more accurate at Time 1 or Time 2? Fourth and finally, we tested whether participants were more inaccurate when estimating attitudes further in the past. We refit Model 2 using the absolute difference between each estimated and actual attitude at each time point, rather than the difference between estimated change and actual change, and included a fixed effect for time (Time 1 vs. Time 2).

Bias on individual items. Bias emerged on 96% of items. Out of 51 attitudes, participants overestimated 29 (57%, Figs. 1–3), underestimated 10 (20%, Fig. 4), estimated 10 in the wrong direction (20%, Fig. 5), and correctly estimated two (4%, Fig. 6). In *SI Appendix*, we show that these biases do not strongly depend on participants' own stance on each issue.

Bias across all items. When testing for bias across all items using Model 1, the intercept was 9.73 (95% CI = [6.07, 13.38], $t(52) = 5.26$, $P < 0.001$), indicating that participants significantly overestimated the magnitude of attitude change.

We explored this effect further by adding covariates to the model that could be related to participants' misperceptions: gender, race, education, political ideology, and whether participants agreed or disagreed with each item. We also included the interaction between participants' age and the interval between Time 1 and 2020 (the year in which Study 1 was run), since older participants may especially outperform younger participants when estimating changes across greater spans of time. We z-scored all continuous covariates and removed item effects because they account for the same variance as the time interval of each item, leaving us with this exploratory model:

$$| \text{estimated change} | - | \text{actual change} | \\ \sim \text{intercept} + \text{gender} + \text{race} + \text{education} \\ + \text{agreement} + \text{political ideology} + \text{age} \\ * \text{interval} + \text{participant random effects.}$$

(The reference level for gender was female, the reference level for race was White, and the reference level for agreement was "agree.")

Overestimation of change was even greater when taking account of these factors, $b = 12.55$, 95% CI = [11.69, 13.41], $t(1118) = 28.46$, and $P < 0.001$. People who agreed with an attitude overestimated change in that attitude more than people who disagreed with it, $M_{\text{agree}} = 10.82$, $M_{\text{disagree}} = 5.36$, $b = -5.48$, 95% CI = [-6.05, -4.90], $t(18551) = -18.81$, and $P < 0.001$. However, even those who disagreed still overestimated change. Participants also overestimated change more over longer intervals, $b = 4.08$, 95% CI = [3.80, 4.35], $t(18143) = 28.87$, and $P < 0.001$. No other effects were significant.

We also tested whether race and gender might matter for items directly related to race and gender, respectively. Exploratory models indicated no significant effect of gender on bias for the seven items related to gender and no significant effect of race on bias for the 11 items related to race (all P s > 0.05).

Inaccuracy across all items. For items rated on a 100-point scale, the intercept from Model 2 revealed that participants misestimated change by an average of 22.22 points, 95% CI = [20.37, 24.08], $t(49) = 23.72$, and $P < 0.001$. For items rated on a 7-point Likert scale, participants misestimated change by an average of 1.27 points, 95% CI = [1.06, 1.49], $t(4) = 12.79$, and $P < 0.001$.

We explored the factors that affect accuracy by adding the same covariates used in the exploratory model above. Lower numbers indicate greater accuracy. Participants were more accurate when they disagreed with an item, $b = -1.66$, 95% CI = [-2.18, -1.15], $t(18690) = -6.30$, and $P < 0.001$. More educated participants were slightly but significantly more accurate than less educated participants, $b = -0.50$, 95% CI = [-0.90, -0.09], $t(931) = -2.40$, and $P = 0.02$, and more liberal participants made smaller errors than more conservative participants, $b = 0.57$, 95% CI = [0.15, 0.98], $t(939) = 1.93$, and $P = 0.007$. Additionally, after conducting all pairwise comparisons between racial groups and applying a Holm-Bonferroni correction for multiple comparisons, White participants were slightly but significantly more accurate than Black participants, $M_{\text{White}} = 18.30$, $M_{\text{Black}} = 20.30$, $P = 0.02$. The intercept remained positive, significant, and essentially unchanged, $b = 20.80$, 95% CI = [18.35, 23.27], $t(56) = 16.59$, and $P < 0.001$, indicating that participants were still considerably inaccurate, even accounting for these factors. No other effects were significant.

We found limited evidence that participants' race and gender mattered for items directly related to race and gender, respectively. In exploratory models, there was no significant effect of gender on accuracy for the seven items related to gender (all P s > 0.05). For the 11 items related to race, White participants' estimates were slightly but significantly more accurate than Black participants' after correcting for multiple comparisons, $M_{\text{White}} = 16.00$, $M_{\text{Black}} = 18.40$, $b = -2.45$, and $P = 0.02$. No other contrasts were significant.

Were participants more accurate at Time 1 or Time 2? Participants were slightly but significantly more inaccurate at Time 1 ($M = 21.80$) compared to Time 2 ($M = 19.10$), $b = -2.75$, 95% CI = [-3.05, -2.44], $t(36726) = -17.77$, and $P < 0.001$.

Discussion. Participants had little idea how attitudes had changed. They most often overestimated change and sometimes even estimated change in the wrong direction. Their estimates were biased on 96% of items, and the average effect size of these biases was $d = 0.65$. Although this was likely a difficult task, and participants made inaccurate estimates—on average, they were 22 points away from the correct answer on a 100-point scale—inaccuracy alone cannot explain the results.

Many of these biases were remarkable. On a majority of items ($n = 28$, 55%), participants were mistaken about whether an attitude had shifted from a minority view to a majority view (or vice versa) or whether it had crossed the midpoint of the scale. For example, participants thought that only a minority of Americans (32%) were willing to vote for a woman for president in 1972 and a majority (70%) were willing in 2010. In fact, most Americans (74%) were willing in 1972 and virtually all (96%) were in 2010.

Why were misperceptions so prevalent and pervasive? Surprisingly, they had little to do with participants' age, race, gender, education, or political ideology. Participants were slightly more biased and inaccurate when they agreed with the attitude they were estimating, but both bias and inaccuracy remained after accounting for agreement. Motivated reasoning may have exacerbated these misperceptions, but it did not create them.

Instead, the results of Study 1 suggest that a stereotype about the past may have distorted participants' judgments. As seen in Figs. 1–6, participants often overestimated how much attitudes had shifted toward the putative liberal side of each

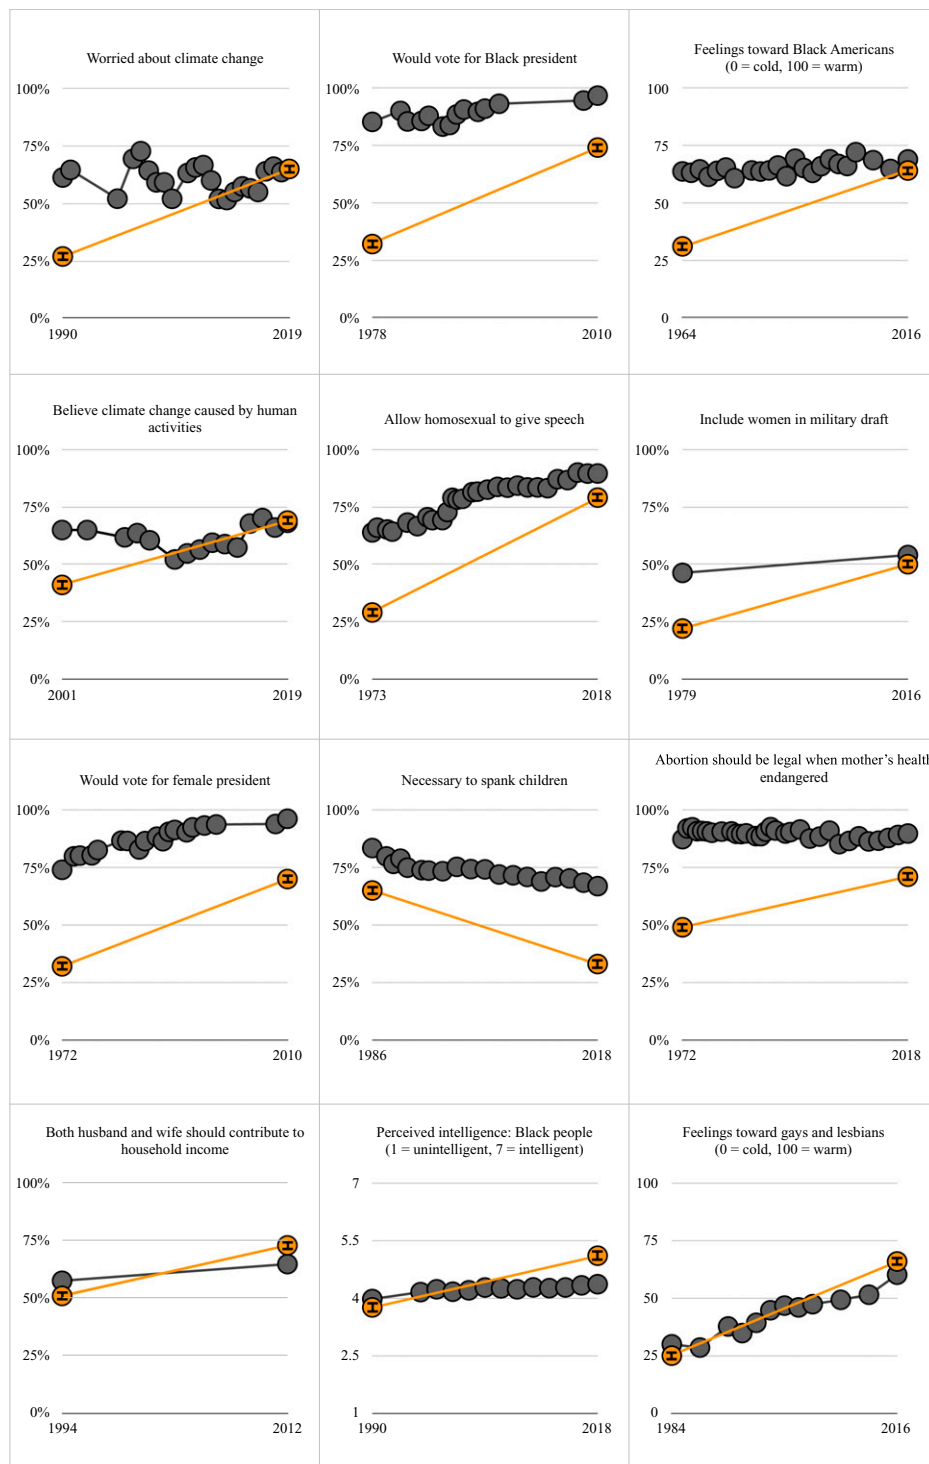


Fig. 1. All items overestimated in Study 1. Orange points represent participants' estimates, and error bars represent 95% confidence intervals. Gray points represent actual data. Graphs are ordered by the effect size of overestimation (Cohen's *d*) from largest to smallest.

issue. For instance, participants overestimated how positive attitudes had become toward racial minorities, abortion, and women's rights, all ostensibly liberal positions. Perhaps participants based their judgments on a stereotype that the present is far more liberal than the past, which could explain why participants often overestimated the amount of change and sometimes estimated change in the wrong direction. For example, participants thought that support had increased for gun control and decreased for the death penalty—both shifts toward traditionally

liberal positions—when in fact public opinion became more conservative on those issues.

To test this possibility, we need to know participants' beliefs about the political dimensions of each issue. Not everyone agrees which positions are liberal and which are conservative, so change in one direction may seem conservative to some and liberal to others. Accordingly, in Study 2, we measured how much participants thought attitudes had shifted in the direction they considered to be liberal.

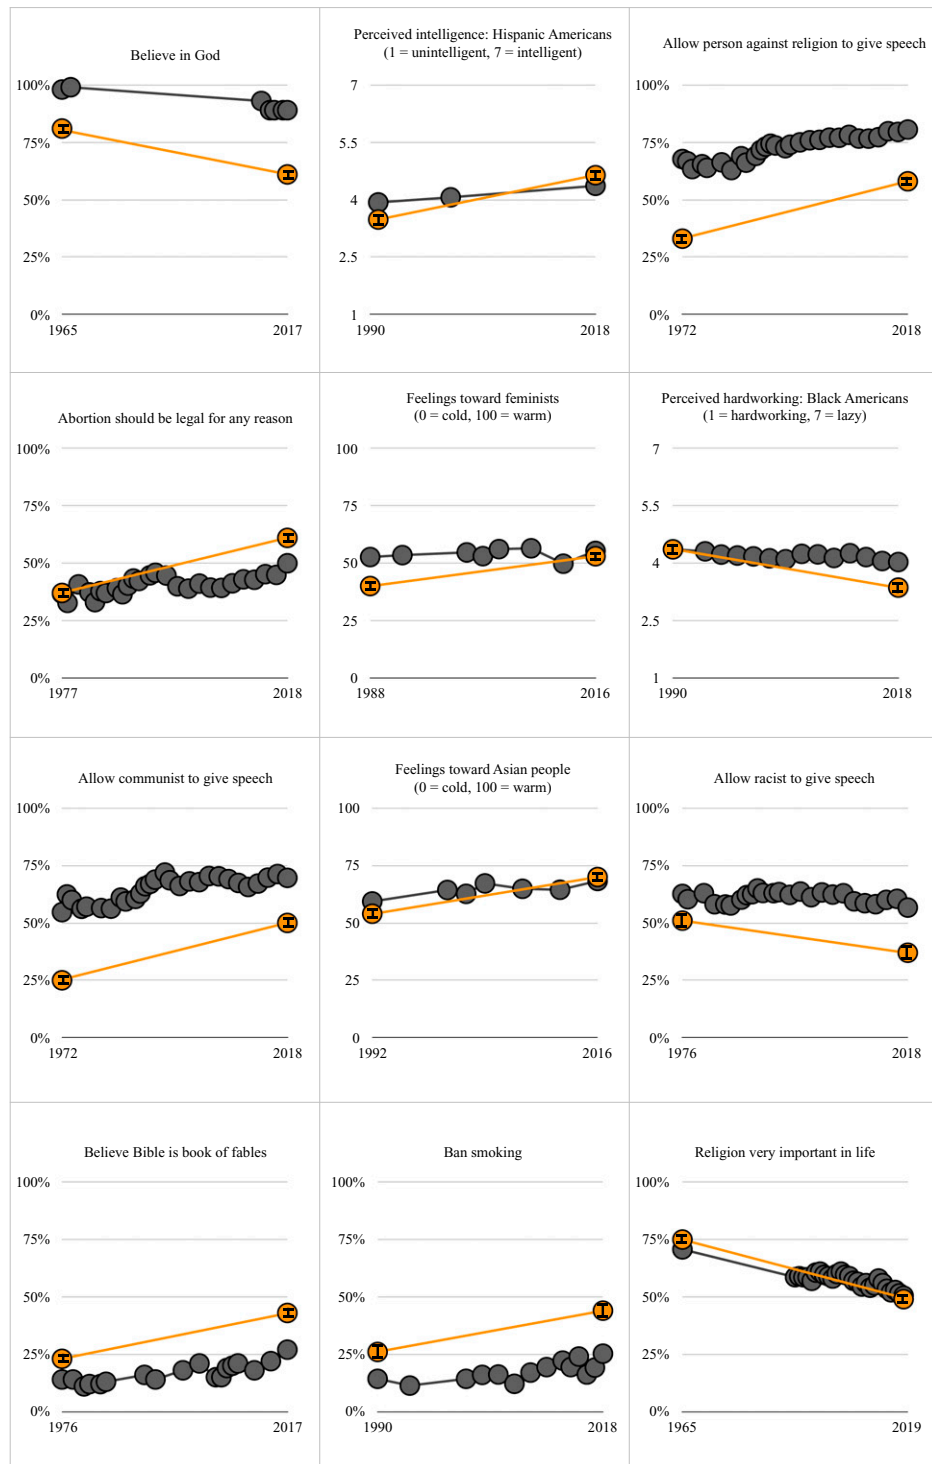


Fig. 2. All items overestimated in Study 1 (cont.). Orange points represent participants' estimates, and error bars represent 95% confidence intervals. Gray points represent actual data. Graphs are ordered by the effect size of overestimation (Cohen's *d*) from largest to smallest.

Study 2

In addition to completing the same procedure as in Study 1, participants in Study 2 made two additional ratings for each item: how liberal or conservative they would consider someone on each side of the issue. This allowed us to determine whether participants thought attitudes had become more liberal or less, and by how much.

Results.

Agreement on the direction of liberal change. On average, 76.36% of participants agreed on which side of each issue was more liberal and which side was more conservative (median = 79.17%).

Estimated and actual liberal change. After calculating estimated and actual change as in Study 1, we calculated estimated and actual

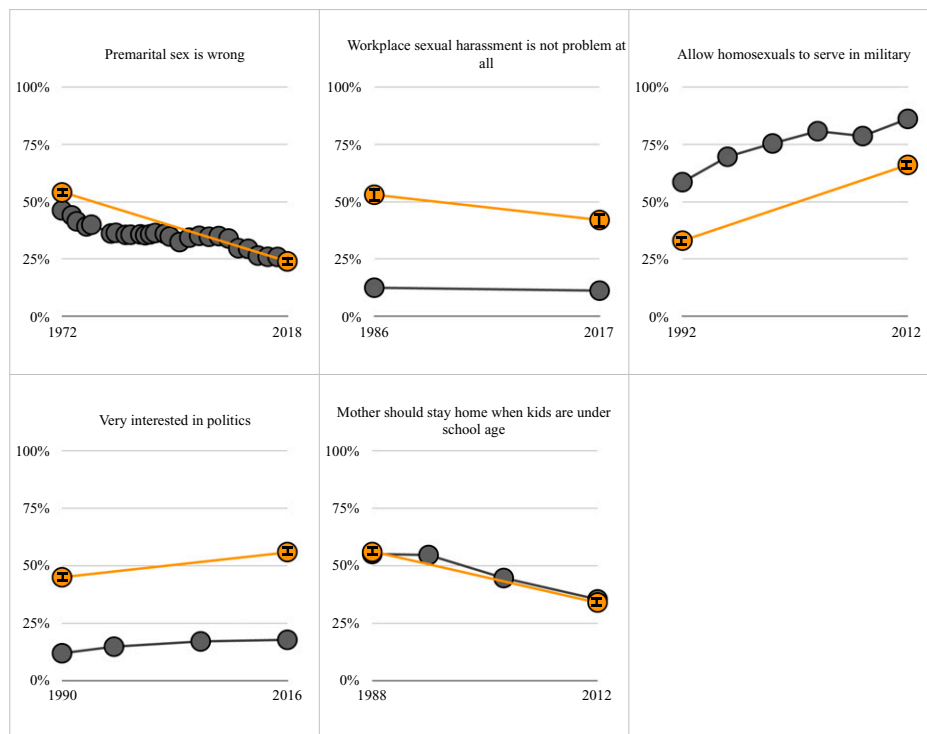


Fig. 3. All items overestimated in Study 1 (cont.). Orange points represent participants' estimates, and error bars represent 95% confidence intervals. Gray points represent actual data. Graphs are ordered by the effect size of overestimation (Cohen's *d*) from largest to smallest.

liberal change by transforming the sign of change so that positive scores indicated change in the direction the participant rated as more liberal. (For more detail on this process, see *SI Appendix*.)

Analysis plan. We followed an analysis plan similar to the Study 1 plan, this time using estimated and actual liberal change scores.

Bias on individual items. We first compared estimated vs. actual liberal change for each item using two-tailed one-sample *t* tests.

Bias across all items. Unlike in Study 1, analyzing bias across all items at once did not require taking the absolute value of estimated and actual change, because liberal change scores mean the same thing across each item. For example, an estimated liberal change of 50 is always a shift in the liberal direction, regardless of item. We tested for bias across all items at once using this model:

$$\begin{aligned} & \text{estimated liberal change} - \text{actual liberal change} \\ & \sim \text{intercept} + \text{participant random effects} \\ & + \text{item random effects.} \end{aligned} \quad [3]$$

Were misperceptions driven by perceptions of the past or the present? To compare estimates of attitudes at Time 1 and Time 2, we also calculated estimated and actual liberal attitudes at each time point by reverse-scoring items as necessary so that larger numbers indicated greater support for the liberal position. This process is described in detail in *SI Appendix*. We then used the following model to determine whether misperceptions of liberal change were driven more by misperceptions of the past or present, using dummy codes for actual vs. estimated change and Time 1 vs. Time 2:

$$\begin{aligned} & \text{liberal attitude} \sim \text{intercept} + (\text{actual vs. estimate}) \\ & * (\text{Time 1 vs. Time 2}) \\ & + \text{participant random effects} \\ & + \text{item random effects.} \end{aligned} \quad [4]$$

(The reference levels were actual and Time 1, respectively.)

What explained participants' misperceptions? Finally, we explored one possible source of participants' misperceptions. If participants overestimate liberal change because they are relying on a stereotype that the present is far more liberal than the past, they should estimate—and perhaps even overestimate—more liberal change on attitudes that they believe have more to do with being liberal or conservative. To test this possibility, we created a perceived partisan diagnosticity score for each participant on each item by calculating the absolute value between how liberal/conservative each participant rated someone on each side the issue. Attitudes earn high partisan diagnosticity scores when one side is perceived as very liberal and the other side is perceived as very conservative, like abortion ($M_{\text{diagnosticity}} = 4.20$, out of a possible six points). Attitudes earn low partisan diagnosticity scores when participants do not perceive a large partisan difference between the two sides of an issue, like being interested in politics ($M_{\text{diagnosticity}} = 0.89$). In an exploratory model, we regressed participants' liberal change estimates on perceived partisan diagnosticity scores with random intercepts for each participant and each item:

$$\begin{aligned} & \text{estimated liberal change} \sim \text{diagnosticity} \\ & + \text{participant random effects} \\ & + \text{item random effects.} \end{aligned} \quad [5]$$

We used another exploratory model to test whether diagnosticity predicted not just participants' estimates, but their overestimates:

$$\begin{aligned} & \text{estimated liberal change} - \text{actual liberal change} \\ & \sim \text{diagnosticity} + \text{participant random effects} \\ & + \text{item random effects.} \end{aligned} \quad [6]$$

Were estimates biased on individual items? Participants overestimated liberal change on 34 attitudes (67%), estimated liberal change correctly on 10 attitudes (20%), and underestimated

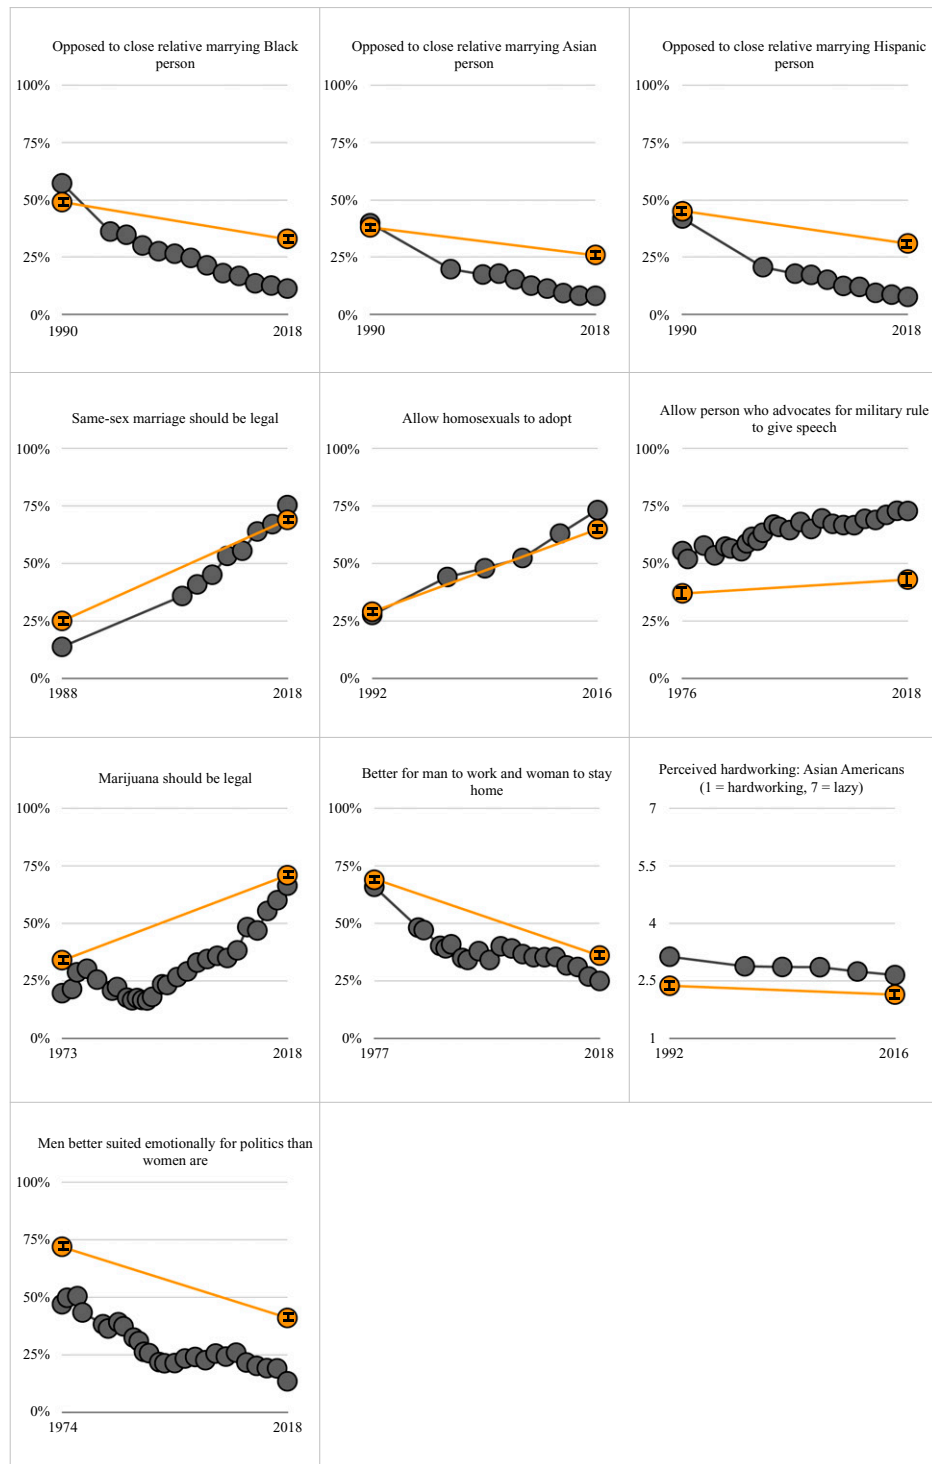


Fig. 4. All items underestimated in Study 1. Orange points represent participants' estimates, and error bars represent 95% confidence intervals. Gray points represent actual data. Graphs are ordered by the effect size of underestimation (Cohen's *d*) from largest to smallest.

liberal change on seven attitudes (14%). The full output of each test is reported in *SI Appendix*.

Did participants overestimate liberal change across all items? Participants significantly overestimated liberal change by almost 9 points on average ($b = 8.75$, 95% CI = [4.64, 12.85], $t(55) = 4.21$, and $P < 0.001$).

We then added the same covariates used in Study 1: gender, race, education, agreement with each item, political ideology,

and the interaction of age and the time interval of each item (the number of years between Time 1 and 2020). We again z-scored continuous variables. Older participants overestimated liberal change less, $b = -1.85$, 95% CI (-3.04, -0.67), $t(390) = -3.02$, and $P = 0.003$, as did participants who disagreed with the item they were estimating, $b = -7.07$, 95% CI (-8.75, -5.37), $t(3688) = -8.19$, and $P < 0.001$. When accounting for all of these factors, participants overestimated liberal change

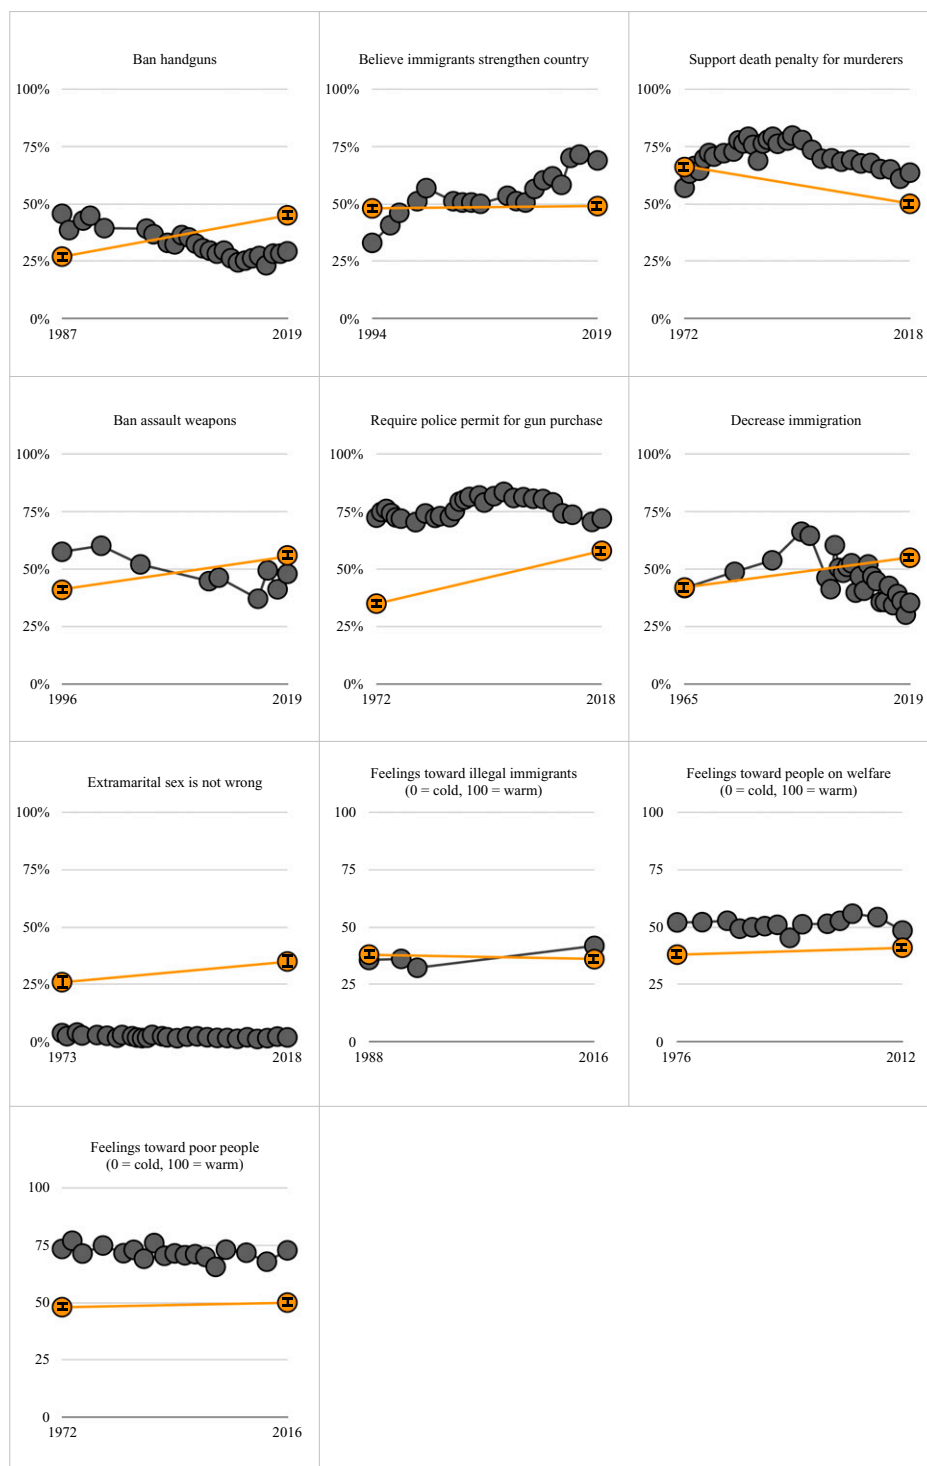


Fig. 5. All items estimated in the wrong direction in Study 1. Orange points represent participants' estimates, and error bars represent 95% confidence intervals. Gray points represent actual data. Graphs are ordered by the effect size of estimation in the wrong direction (Cohen's *d*) from largest to smallest.

even more, $b = 12.63$, 95% CI (10.69, 14.57), $t(546) = 12.61$, and $P < 0.001$. No other effects were significant.

This bias was not limited to liberals. We fit an exploratory model that included only conservative participants ($n = 124$) and found they also significantly overestimated liberal change, $b = 7.09$, 95% CI = [2.87, 11.31], $t(61) = 3.32$, and $P = 0.002$.

Were misperceptions of liberal change driven more by misperceptions of the past or the present? The interaction of rating type and time in Model 4 was significant, $b = 9.63$, 95% CI = [8.45, 10.80], $t(14,322) = 16.05$, and $P < 0.001$. Postestimation contrasts using a Holm-Bonferroni correction indicated that participants significantly underestimated liberal attitudes at both time points, but much more at Time 1 ($M_{\text{estimate}} = 32.80$,

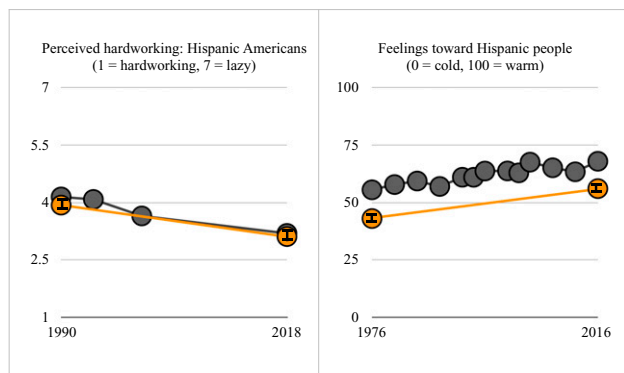


Fig. 6. Both items estimated correctly in Study 1. Orange points represent participants' estimates, and error bars represent 95% confidence intervals. Gray points represent actual data.

$M_{\text{actual}} = 44.70$, 95% $CI_{\text{diff}} = [10.96, 12.86]$, and $P < 0.001$) than at Time 2 ($M_{\text{estimate}} = 53.30$, $M_{\text{actual}} = 55.60$, 95% $CI_{\text{diff}} = [1.33, 3.23]$, and $P < 0.001$).

Although participants were less biased when estimating liberal attitudes at Time 2, they were only slightly more accurate. We refit Model 4 using the absolute difference between estimated and actual liberal attitudes as the outcome, and it indicated that participants misestimated by 23.30 points at Time 1, and by 18.50 points at Time 2, a small but significant difference, $b = -4.83$, 95% $CI = [-5.50, -4.16]$, $t(6948) = -14.14$, and $P < 0.001$.

Together, these results indicate that participants greatly underestimated liberal attitudes at Time 1, slightly underestimated liberal attitudes at Time 2, and estimated liberal attitudes quite inaccurately at both time points. These results are graphed in *SI Appendix, Fig. S3*.

What explained participants' perceptions of liberal change? Diagnosticity significantly predicted participants' estimates in Model 5, $b = 1.42$, 95% $CI = [0.84, 2.01]$, $t(3321) = 4.77$, and $P < 0.001$, meaning that participants estimated more liberal change on attitudes they thought were more diagnostic of being liberal or conservative.

Diagnosticity also significantly predicted participants' overestimation in Model 6, $b = 0.71$, 95% $CI = [0.16, 1.25]$, $t(3069) = 2.53$, and $P = 0.01$. When participants thought an attitude was more diagnostic of being liberal or conservative, they overestimated liberal change even more.

Discussion. Participants seemed to share a stereotype that the present is far more liberal than the past. They mistakenly believed the liberal side had gained ground that it had in fact lost (e.g., gun control), already held (e.g., climate change), or never held (e.g., religion). Consistent with this stereotype, participants overestimated liberal change more on items that they thought were better indicators of being liberal. Remarkably, even conservatives overestimated liberal change.

As expected, this stereotype explained several cases in which participants estimated change in the wrong direction in Study 1. Participants thought that more Americans now support gun control, feel warm toward the poor and those on welfare, approve of extramarital sex, and oppose the death penalty, changes that participants considered liberal. In fact, more Americans now support the putative conservative position on these issues.

Participants especially underestimated how liberal the past was, which may have happened for several reasons. One reason might be that participants knew most attitudes had become more liberal—not surprising, since some of these shifts have been seismic—but they did not know much more than that.

This stereotype may have led participants to exaggerate the difference between past and present, and it may have especially distorted their perceptions of the past simply because participants knew less about the past. Such stereotypes can arise when people rely too heavily on a “kernel of truth” in the absence of deeper knowledge (38, 39). Consistent with this explanation, participants overestimated liberal change the most on the issues they thought had the biggest gap between liberal and conservative positions—that is, the issues where their stereotype applied most.

To our knowledge, we are the first to document this stereotype, likely because previous research has rarely tapped people's perceptions of attitude change and has never investigated this many attitudes at once.

Studies 3a, 3b, and 3c

In Studies 3a, 3b, and 3c, we explored one possible consequence of these misperceptions: they may shape policy preferences. In Study 3a, participants imagined that support for lowering the number of immigrants admitted to the United States had either increased (the misperception from Studies 1 and 2) or decreased (the actual change). In Study 3b, participants imagined that support for an assault weapons ban had either increased (the misperception from Studies 1 and 2) or decreased (the actual change). In each case, they rated how justified Congress would be in changing policy on the issue.

In Study 3c, we also tested whether participants were sensitive to changes in attitudes or only sensitive to current attitudes. Participants imagined that support for banning the death penalty had either risen to 50% (the misperception from Studies 1 and 2) or fallen to 50% (the opposite of participants' estimates). In each case, they rated how justified Congress would be in banning the death penalty.

Results. All results were analyzed with preregistered two-tailed paired t tests.

Study 3a: Limits on immigration. In Study 3a, participants thought it would be more justifiable for Congress to limit immigration when support for doing so had increased from 42 to 56% ($M = 4.64$) than when it had decreased from 42 to 35% ($M = 3.03$), $t(229) = 14.89$, $P < 0.001$, and $d = 0.98$. This was true for those who supported the decrease ($M_{\text{increase}} = 5.72$, $M_{\text{decrease}} = 3.72$, $t(70) = 10.13$, $P < 0.001$, and $d = 1.20$), those who opposed it ($M_{\text{increase}} = 4.07$, $M_{\text{decrease}} = 2.66$, $t(122) = 9.35$, $P < 0.001$, and $d = 0.84$), and those who neither favored nor opposed it ($M_{\text{increase}} = 4.44$, $M_{\text{decrease}} = 2.89$, $t(35) = 6.65$, $P < 0.001$, and $d = 1.11$).

Study 3b: Assault weapons ban. In Study 3b, participants thought it would be more justifiable for Congress to ban assault weapons when support for the ban had increased from 48 to 58% ($M = 4.58$) than when it had decreased from 58 to 48%, ($M = 3.40$), $t(246) = 13.64$, $P < 0.001$, and $d = 0.87$. This was true for those who supported the ban ($M_{\text{increase}} = 5.63$, $M_{\text{decrease}} = 4.38$, $t(137) = 11.84$, $P < 0.001$, and $d = 1.01$), those who opposed the ban ($M_{\text{increase}} = 3.00$, $M_{\text{decrease}} = 2.06$, $t(86) = 5.90$, $P < 0.001$, and $d = 0.63$), and those who neither favored nor opposed the ban ($M_{\text{increase}} = 4.18$, $M_{\text{decrease}} = 2.64$, $t(21) = 5.58$, $P < 0.001$, and $d = 1.19$).

Study 3c: Death penalty ban. In Study 3c, participants thought it would be more justifiable for Congress to ban the death penalty when support for the ban had increased from 33 to 50% ($M = 4.52$) than when it had decreased from 67 to 50% ($M = 3.97$), $t(387) = 8.52$, $P < 0.001$, and $d = 0.43$. This was true for those who supported the death penalty ($M_{\text{increase}} = 3.91$, $M_{\text{decrease}} = 3.35$, $t(167) = 5.49$, $P < 0.001$, and $d = 0.42$) and those who opposed it ($M_{\text{increase}} = 5.21$, $M_{\text{decrease}} = 4.43$, $t(177) = 7.56$, $P < 0.001$, and $d = 0.57$), but not for those who neither supported

nor opposed it ($M_{\text{increase}} = 4.07$, $M_{\text{decrease}} = 4.12$, $t(41) = -0.26$, $P = 0.80$, and $d = 0.04$).

Discussion. Attitude change can legitimize policy change. According to participants, whether Congress could justifiably limit immigration, ban assault weapons, and eliminate the death penalty depended on whether public opinion had swung toward or away from those positions. Even supporters of these policies saw them as less justifiable when support was slipping away from their side, meaning that partisans were not merely capitalizing on favorable trends and ignoring unfavorable trends.

People paid attention not just to current attitudes but also to changes in attitudes. In Study 3c, participants judged banning the death penalty as acceptable when popular support had risen to 50% and objectionable when it had fallen to 50%, even though the current level of support was the same in both cases. (In *SI Appendix*, we report a study in which we replicate this effect using the immigration issue from Study 3a.)

These results make the misperceptions we identified in Studies 1 and 2 even more concerning. If shifts in attitudes can justify shifts in policy, then misperceiving shifts in attitudes could lead to policies that people do not actually prefer.

General Discussion

People do not seem to know how attitudes have changed. In fact, participants typically did not even know whether an attitude had gained or lost the support of the majority. Our work greatly expands on prior research, which has generally investigated people's perceptions of a single attitude at one moment in time (18–23).

Why This Happens. Though people may misperceive changes in attitudes for many reasons, our results provide direct evidence for two reasons and indirect evidence for a third. First, people may have misperceived attitude change because they believe what they want to believe. In Study 1, participants who agreed with an attitude were more likely to overestimate how much it had changed. However, as we show in *SI Appendix*, agreeing with an attitude made only small differences in people's estimates. In Study 2, even conservatives overestimated the liberal shift in attitudes. This finding is consistent with prior research on motivated reasoning, which shows that even the most motivated reasoners are still constrained by reality (40). No matter how much someone might oppose gay marriage, for example, it is hard for them to deny that support for gay marriage has increased.

Second, participants may have over-relied on a stereotype that the present is far more liberal than the past. Consistent with this possibility, participants overestimated liberal change more on items that they believed were more related to being liberal or conservative.

Third, our results hint that major events may mislead participants about attitude change. In one study, people surveyed immediately after the *Obergefell vs. Hodges* decision thought that more Americans supported gay marriage than did people surveyed immediately before the ruling (41). Our participants may have similarly used landmark events to infer changes in attitudes. For instance, the two most-overestimated changes were increases in concern about climate change and willingness to vote for a Black president. Perhaps participants assumed that the numerous recent wildfires, floods, and hurricanes made people worry more about the climate, and Barack Obama's 2008 election signaled new openness to African Americans in the Oval Office. Both assumptions are reasonable but wrong. Demographic factors did not moderate these mistakes, suggesting everyone was attending to the same misleading

information—perhaps momentous events like natural disasters and presidential elections.

Thus, misperceptions may arise when major events do not intuitively correspond to changes in actual attitudes or when attitudes change without corresponding major events. This seems broadly consistent with the results shown in Figs. 1–6. For instance, the percentage of Americans who would oppose a close friend or family member marrying someone who is Black, Hispanic, or Asian dropped from around 50% in 1990 to less than 10% in 2018. It is difficult to identify any landmark events related to both race and marriage that occurred in this period, which is perhaps why participants underestimated these shifts by about half, the largest underestimations we observed. The connection between salient events and the misperception of attitude change is speculative and does not fit every item, but it may be a fruitful area for further research.

Why This Matters. These misperceptions may have far-reaching consequences. Shifts in public opinion can lend legitimacy to laws, as Studies 3a through 3c showed, so misperceiving those shifts can justify policies that people do not actually prefer. For example, participants believed that more and more Americans favor decreasing immigration, which they said justifies tightening borders. However, participants also said more immigrants should be admitted if attitudes had shifted in favor of doing so—which is what actually occurred. This misperception may have boosted Donald Trump's anti-immigration agenda by emboldening his supporters and discouraging his opponents. Americans may have viewed Trump differently if they had known he was elected despite changes in attitudes toward immigration, not because of them.

In the long term, misperceiving attitude change could lead to self-fulfilling prophecies. People change when they think others are changing (6–9). If people think more and more Americans are against immigration, they might oppose it as well, creating a reality that was once imaginary. Beliefs about the future can also become self-fulfilling by spurring or stalling action today (7). In one study, participants donated less money to their favored candidate when they were told the candidate was likely to win (42). The misperceptions we identified could cause similar consequences: if people incorrectly assume their side is its way victory, they might not do what is necessary to secure victory in the first place.

Our work raises important questions to be answered in future research: Why do people overestimate the liberal shift in attitudes but underestimate liberal attitudes overall? How malleable are these misperceptions, and can they be leveraged to catalyze attitude change? (In *SI Appendix*, we report one study in which participants predicted a different future when shown accurate information about the past.) We investigated large changes over long periods of time; are these more or less powerful than small changes over short periods of time?

Conclusion. We live in a world beset by ideological conflicts over race, gender, immigration, free speech, abortion, climate change, and gun control, to name but a few issues. Both sides treat these struggles as life-and-death because they are—winning would earn them the right to make laws, control institutions, and set cultural norms. Our studies show, however, that people have little idea whether they are winning or losing, or even whether they have won or lost. People who believe in climate change think they have recently won the majority that they have, in fact, had for decades; people who favor gun control think they have recently gained the majority that they have, in fact, recently lost; and people who are both pro- and anti-immigration think they are in a stalemate when, in fact, the pro side is winning. These misperceptions may distort the

democratic process, hinder people's efforts to understand and change their social worlds, and turn imaginary trends into real ones.

Materials and Methods

All materials, measures, data, code, and preregistrations are available on the Open Science Framework (<https://osf.io/wud9a/>). All studies were preregistered; any deviations are discussed in *SI Appendix*. All studies were approved by the Institutional Review Board of Harvard University.

Study 1.

Materials. We selected 51 public opinion questions based on a pilot study (reported in *SI Appendix*) in which participants nominated topics on which they thought American public opinion had changed in the past 50 y. To be included, each item had to be administered to a nationally representative sample at least twice, with waves separated by at least 15 years and the most recent data having been collected in 2010 or later. The exploratory study, process of selecting items, and wording of each item are reported in *SI Appendix*.

Participants. We sought to recruit 1,000 participants using Prolific's nationally representative sample feature. A total of 1,036 people responded to an advertisement for our study. Of these, 33 did not complete the study, leaving 1,003 participants who completed all measures (514 female, 485 male, 4 "other," $M_{\text{age}} = 44.88$ y, 73% White, 13% Black, 7% Asian, 4% Hispanic, 2% "more than one of the above," <1% "other") in exchange for \$2.25. Our demographic measures differ from those the Prolific platform uses to assemble nationally representative samples and so do not perfectly reflect the demographics used to create the sample.

Procedure. After providing consent, participants reported their birth year. We later compared this to their reported age at the end of the study as a quality control check. Participants were then told they would estimate how other people answered various questions, which had all been asked of nationally representative groups of American adults. When estimating, they were instructed not to include anyone they thought said "do not know" or did not give an answer.

Participants then viewed a subset of 20 items drawn from the total set of 51 items. The selection was pseudorandom, as eight items—the intergroup feelings thermometers from the American National Election Studies—required additional instructions, so participants either saw all eight feelings thermometers or none of them. These items were randomized within their block of eight. If participants viewed the feelings thermometers, they also completed a random subset of 12 items selected from the rest of the items; the block of 12 and block of 8 were presented in random order.

For each item, participants read the public opinion question and then estimated how a nationally representative sample of Americans had answered at two time points. For most items, participants estimated the percentage of the sample that chose a certain answer by typing their estimate into a text box. For instance, participants read the following question: "Are you for or against a law which would make it illegal to manufacture, sell or possess semi-automatic guns known as assault rifles?" Participants were asked to estimate the percentage of respondents who responded "for" in 1996 and 2010 by typing a number between 0 and 100 in one text box labeled "1996" and another text box labeled "2010." The feelings thermometers were originally administered using a 0 to 100 scale, so participants used the same scale for those items, again typing their responses into text boxes. Five items were originally administered using Likert scales, so participants estimated the average response using the same Likert scales, one for each year. We include a screenshot of an example question in *SI Appendix*.

After estimating, participants reported how they would answer each question themselves. We coded participants as agreeing with the question if they chose the answer that they had just estimated and as disagreeing if they chose the opposite answer. For questions that used a 0 to 100 scale or a 1 to 7 Likert scale, we performed a median split and coded participants in the more positive half of responses (e.g., more hardworking, intelligent, warm) as agreeing with the item and the less positive half (e.g., less hardworking, intelligent, warm) as disagreeing with the item.

After completing 20 items, participants completed an attention check that resembled the items they had just completed but asked them to type the number one in each of the blanks provided. Participants then completed a set of demographic questions. To make models more interpretable, we recoded education as a continuous variable from 1 (did not complete high school) to 7

(graduate school) and recoded political ideology as a continuous variable from -2 (very liberal) to $+2$ (very conservative). Embedded in the demographics questions was an attention check that asked participants to select the option "other" and type the word "tree."

Exclusions. Sixty participants failed at least one of the three quality control checks and were excluded. The remaining 943 participants (488 female, 451 male, 4 "other," $M_{\text{age}} = 45.07$ y, 73% White, 13% Black, 7% Asian, 4% Hispanic, 2% "more than one of the above," 1% "other") were included in all analyses.

Study 2.

Replication of Study 1. Studies 2 and 3a-3c were run on Amazon Mechanical Turk (MTurk) instead of Prolific. To ensure that the platforms produce similar results, we also re-ran Study 1 on MTurk. The results, which are reported in the Supplement, were nearly identical. Thus, we felt confident running further studies on MTurk.

Methods. Participants. We aimed to recruit 500 people on MTurk. In order to participate, respondents had to pass a three-item test of American culture and the English language, which required them to know that students in their second year of high school are called sophomores, that granola would be an unusual dish to serve at a Fourth of July barbecue, and that dialing 911 connects the caller to emergency services. Five hundred and three people (261 female, 241 male, 1 "other," $M_{\text{age}} = 39.47$, 71% White, 10% Black, 9% Asian, 5% Hispanic, 1% Hawaiian or Pacific Islander, 4% "more than one of the above") passed this test, consented, and completed all measures in exchange for \$1.50.

Procedure. Participants completed the same procedure as in Study 1, except they completed 10 items instead of 20. After completing each item, they answered two additional questions: "How liberal or conservative would you consider someone if they answered [X]?" and "How liberal or conservative would you consider someone if they answered [Y]?" where [X] and [Y] were replaced with the two opposite answers for each item. These questions were altered slightly for the eight feelings thermometers and the items rated on a 7-point Likert scale: "How liberal or conservative would you consider someone if their answer was [below 10/above 90]?" and "How liberal or conservative would you consider someone if they answered [1/7]?" Participants answered using 7-point Likert scales with endpoints labeled 1 (very liberal) and 7 (very conservative) and midpoints labeled 4 (neither conservative nor liberal/cannot tell).

Exclusions. Fifty-six participants failed at least one attention check and were excluded, leaving 447 participants (233 female, 213 male, 1 "other," $M_{\text{age}} = 39.38$, 72% White, 9% Black, 9% Asian, 5% Hispanic, 1% Hawaiian or Pacific Islander, 3% "more than one of the above") in all analyses.

Studies 3a-3c.

Participants. We aimed to recruit 300 people from MTurk for Studies 3a, 300 people for Study 3b, and 500 people for Study 3c. Participants first had to complete the same test of American culture and the English language used in Study 2. All participants were compensated \$0.60.

3a. Three hundred people (145 female, 155 male, $M_{\text{age}} = 40.72$ y, 73% White, 11% Asian, 6% Black, 5% Hispanic, 4% "more than one of the above") participated.

3b. Three hundred and one people (160 female, 139 male, 2 "other," $M_{\text{age}} = 40.65$ y, 75% White, 8% Black, 6% Asian, 5% Hispanic, <1% Hawaiian or Pacific Islander, 4% "more than one of the above," 1% "other") participated.

3c. Five hundred people (259 female, 240 male, $M_{\text{age}} = 42.11$ y, 77% White, 9% Asian, 8% Black, 4% Hispanic, <1% American Indian or Alaska Native, 1% "more than one of the above") participated.

Procedure. 3a. In counterbalanced order, participants were asked to imagine either that support for lowering the number of legal immigrants annually admitted to the United States had increased from 42 to 56% from 1965 to today (corresponding to participants' estimates in Study 1) or that it had decreased from 42 to 35% (corresponding to actual survey data). For each hypothetical, participants rated how justified it would be for Congress to lower the number of legal immigrants admitted to the United States annually using 7-point Likert scales with endpoints labeled 1 (not justified at all) and 7 (very justified). They then indicated how much they would support or oppose this policy themselves using a 7-point Likert scale with endpoints labeled -3 (strongly oppose) and $+3$ (strongly support) and a midpoint labeled 0 (neither support nor oppose). They also completed all attention checks included in Study 2, as well as an additional attention check that appeared immediately after the main measures and asked them to recall the question they had just answered.

* Mean ages always exclude participants who reported being over 1,000 years old. All of these participants later failed attention checks and were excluded, perhaps because of their advanced age.

3b. Participants completed the same procedure, except they imagined that support for banning assault weapons had either increased from 48 to 58% from 1996 to today (roughly corresponding to participants' estimates in Study 1) or decreased from 58 to 48% (corresponding to actual survey data).

3c. Participants completed the same procedure, except they imagined that support for banning the death penalty had either increased from 33 to 50% from 1972 to today (corresponding to participants' estimates in Study 1) or decreased from 67 to 50% (the opposite of participants' estimates in Study 1).

Exclusions. 3a. Sixty-nine participants failed at least one attention check and were excluded, leaving 231 participants (111 female, 120 male, $M_{age} = 41.12$ y, 74% White, 11% Asian, 6% Black, 5% Hispanic, 3% "more than one of the above") in all analyses.

3b. Fifty-three participants failed at least one attention check and were excluded, leaving 248 participants (132 female, 114 male, 2 "other,"

$M_{age} = 40.93$ y, 76% White, 8% Black, 6% Hispanic, 5% Asian, 1% Hawaiian or Pacific Islander, 4% "more than one of the above") in all analyses.

3c. One hundred and eleven participants failed at least one attention check and were excluded, leaving 389 participants (202 female, 187 male, $M_{age} = 42.23$, 77% White, 8% Asian, 8% Black, 4% Hispanic, 2% "more than one of the above," 1% American Indian or Alaska Native) in all analyses.

Data Availability. Comma-separated values data have been deposited in Open Science Framework (<https://osf.io/wud9a/>). Previously published data were used for this work (1–5, 28–35).

ACKNOWLEDGMENTS. We thank Gus Cooney, Ethan Ludwin-Peery, Dan Gilbert, Priyanka Menon, and Jeremy Simon for their helpful comments. We also thank Ali Barker, Ali Javeed, Barbara Oedayrajsingh Varma, Dani Grant, Danny Blumenthal, Eric Neumann, Eric Wilson, Kara Xie, Jenna Lang, Ryan Barrett, Sandra Bajc-Dimitrov, and Tong Su for their assistance.

1. T. W. Smith, M. Davern, J. Freese, S. Morgan, *General Social Surveys, 1972–2018* (2018).
2. G. Organization, *Environment* (Gall. Hist. Trends, 2021). <https://news.gallup.com/poll/1615/environment.aspx>. Accessed 30 November 2021.
3. G. Organization, *Religion* (2021). <https://news.gallup.com/poll/1690/religion.aspx>. Accessed 30 November 2021.
4. G. Organization, *Guns* (2021). <https://news.gallup.com/poll/1645/guns.aspx>. Accessed 30 November 2021.
5. G. Organization, *Immigration* (2021). <https://news.gallup.com/poll/1660/immigration.aspx>. Accessed 12 August 2021.
6. G. Sparkman, G. M. Walton, Witnessing change: Dynamic norms help resolve diverse barriers to personal change. *J. Exp. Soc. Psychol.* **82**, 238–252 (2019).
7. G. Sparkman, G. M. Walton, Dynamic norms promote sustainable behavior, even if it is counternormative. *Psychol. Sci.* **28**, 1663–1674 (2017).
8. C. R. Mortensen et al., Trending norms: A lever for encouraging behaviors performed by the minority. *Soc. Psychol. Personal. Sci.* **10**, 201–210 (2019).
9. D. D. Loschelder, H. Siepelmeyer, D. Fischer, J. A. Rubel, Dynamic norms drive sustainable consumption: Norm-based nudging helps café customers to avoid disposable to-go-cups. *J. Econ. Psychol.* **75**, 102146 (2019).
10. A. B. Brodish, P. C. Brazy, P. G. Devine, More eyes on the prize: Variability in White Americans' perceptions of progress toward racial equality. *Pers. Soc. Psychol. Bull.* **34**, 513–527 (2008).
11. M. DeBell, Polarized opinions on racial progress and inequality: Measurement and application to affirmative action preferences. *Polit. Psychol.* **38**, 481–498 (2017).
12. M. W. Kraus, J. J. X. Tan, Americans overestimate social class mobility. *J. Exp. Soc. Psychol.* **58**, 101–111 (2015).
13. R. P. Larrick, A. E. Mannes, J. B. Soll, "The social psychology of the wisdom of crowds" in *Frontiers of Social Psychology. Social Judgment and Decision Making*, J. I. Krueger, Ed. (Psychology Press, 2012), pp. 227–242.
14. J. Surowiecki, *The Wisdom of Crowds* (Doubleday, 2004).
15. H. Chen, P. De, Y. Hu, B. H. Hwang, Wisdom of crowds: The value of stock opinions transmitted through social media. *Rev. Financ. Stud.* **27**, 1367–1403 (2014).
16. S. K. M. Yi, M. Steyvers, M. D. Lee, M. J. Dry, The wisdom of the crowd in combinatorial problems. *Cogn. Sci.* **36**, 452–470 (2012).
17. M. Steyvers, M. Lee, B. Miller, P. Hemmer, "The wisdom of crowds in the recollection of order information" in *Advances in Neural Information Processing Systems 22-Proceedings of the 2009 Conference* (2009), pp. 1785–1793.
18. D. A. Prentice, D. T. Miller, Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm. *J. Pers. Soc. Psychol.* **64**, 243–256 (1993).
19. M. Wenzel, Misperceptions of social norms about tax compliance: From theory to intervention. *J. Econ. Psychol.* **26**, 862–883 (2005).
20. L. Bursztyn, A. González, D. Yanagizawa-Drott, Misperceived social norms: Female labor force participation in Saudi Arabia (No. w24736) (National Bureau of Economic Research, 2018).
21. R. L. Bergstrom, C. Neighbors, M. A. Lewis, Do men find "bony" women attractive?: Consequences of misperceiving opposite sex perceptions of attractive body image. *Body Image* **1**, 183–191 (2004).
22. M. Sandstrom, H. Makover, M. Bartini, Social context of bullying: Do misperceptions of group norms influence children's responses to witnessed episodes? *Soc. Influ.* **8**, 196–215 (2013).
23. D. E. Broockman, C. Skovron, Bias in perceptions of public opinion among political elites. *Am. Polit. Sci. Rev.* **112**, 542–563 (2018).
24. P. Glick, S. T. Fiske, An ambivalent alliance. Hostile and benevolent sexism as complementary justifications for gender inequality. *Am. Psychol.* **56**, 109–118 (2001).
25. S. J. Spencer, C. Logel, P. G. Davies, Stereotype threat. *Annu. Rev. Psychol.* **67**, 415–437 (2016).
26. E. Greico, *Americans' Main Sources for Political News Vary by Party and Age* (Pew Res. Cent., 2020).
27. M. Lodge, R. Hamill, A partisan schema for political information processing. *Am. Polit. Sci. Rev.* **80**, 505–520 (2009).
28. G. Organization, *Gallup Organization Poll: March 2018 (Version 2) (Dataset)*. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research. 10.25940/ROPER-31114959.
29. G. Organization, *Gallup Organization Poll: March 2001, Question 49, USGAL-LUP.01MRH05.R36* (2001). Cornell University, Ithaca, NY: Roper Center for Public Opinion Research. 10.25940/ROPER-31108555.
30. L. Saad, *Support for Banning Public Smoking Holding Steady in U.S. Gallup*, (2017). <https://news.gallup.com/poll/214373/support-banning-public-smoking-holding-steady.aspx>. Accessed 30 November 2021.
31. G. Organization, *Gallup Poll # 1122G (1979) (Version 3) (Dataset)*. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research. 10.25940/ROPER-31087977.
32. B. Jones, *Majority of Americans continue to say immigrants strengthen the U.S.* (Pew Res. Cent., 2019). <https://www.pewresearch.org/fact-tank/2019/01/31/majority-of-americans-continue-to-say-immigrants-strengthen-the-u-s/>. Accessed 30 November 2021.
33. The American National Election Studies. These materials are based on work supported by the National Science Foundation under grant numbers SES 1444721, 2014–2017, the University of Michigan, and Stanford University. <http://www.electionstudies.org>. Accessed 23 April 2020.
34. C. N. Network, *CNN/ORC Poll: 2016 Presidential Election/Late Supreme Court Justice Antonin Scalia/Female President (2016). (Version 2) (Dataset)*. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research. 10.25940/ROPER-31095602.
35. Y. C. Shulman, *Time Soundings # 5772 (1986) (Version 2) (Dataset)*. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research. 10.25940/ROPER-31099181.
36. D. Bates, M. Mächler, B. Bolker, S. Walker, Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2014).
37. A. Kuznetsova, P. B. Brockhoff, R. H. B. Christensen, lmerTest Package: Tests in linear mixed effects models. *J. Stat. Softw.* **82**, 1–26 (2017).
38. J. L. Hilton, W. von Hippel, Stereotypes. *Annu. Rev. Psychol.* **47**, 237–271 (1996).
39. P. Bordalo, K. Coffman, N. Gennaioli, A. Shleifer, Stereotypes. *Q. J. Econ.* **2016**, 1753–1794 (2016).
40. Z. Kunda, The case for motivated reasoning. *Psychol. Bull.* **108**, 480–498 (1990).
41. M. E. Tankard, E. L. Paluck, The effect of a supreme court decision regarding gay marriage on social norms and personal attitudes. *Psychol. Sci.* **28**, 1334–1344 (2017).
42. T. Rogers, D. A. Moore, M. I. Norton, The belief in a favorable future. *Psychol. Sci.* **28**, 1290–1301 (2017).