

Patterns of Implicit and Explicit Attitudes: IV. Change and Stability From 2007 to 2020



Tessa E. S. Charlesworth^{ID} and Mahzarin R. Banaji

Department of Psychology, Harvard University

Psychological Science
2022, Vol. 33(9) 1347–1371
© The Author(s) 2022
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/09567976221084257
www.psychologicalscience.org/PS



Abstract

Using more than 7.1 million implicit and explicit attitude tests drawn from U.S. participants to the Project Implicit website, we examined long-term trends across 14 years (2007–2020). Despite tumultuous sociopolitical events, trends from 2017 to 2020 persisted largely as forecasted from past data (2007–2016). Since 2007, all explicit attitudes decreased in bias between 22% (age attitudes) and 98% (race attitudes). Implicit sexuality, race, and skin-tone attitudes also continued to decrease in bias, by 65%, 26%, and 25%, respectively. Implicit age, disability, and body-weight attitudes, however, continued to show little to no long-term change. Patterns of change and stability were generally consistent across demographic groups (e.g., men and women), indicating widespread, macrolevel change. Ultimately, the data magnify evidence that (some) implicit attitudes reveal persistent, long-term change toward neutrality. The data also newly reveal the potential for short-term influence from sociopolitical events that temporarily disrupt progress toward neutrality, although attitudes eventually return to long-term homeostasis in trends.

Keywords

implicit attitude change, explicit attitude change, Implicit Association Test (IAT), long-term change, time-series analysis, autoregressive-integrated-moving-average (ARIMA) model, open data, open materials, preregistered

Received 4/27/21; Revision accepted 2/3/22

The 4-year period from 2016 to 2020 saw unpredicted and tumultuous events in the United States, from the election of Donald Trump and the rise of White supremacy groups to the emergence and growth of progressive movements for social justice. In the face of such shocks to the sociopolitical landscape, we examined how societal-level collective attitudes have persisted or responded to change: Did social-group attitudes—both explicit and implicit—continue along their predicted paths from the past? Or did they alter course and, if so, for how long?

Previously, we reported that across 10 years (2007–2016), biases related to six explicit attitudes (race, skin tone, sexuality, age, disability, and body weight), as well as implicit sexuality, race, and skin-tone attitudes, had decreased and were forecasted to continue decreasing toward neutrality (Charlesworth & Banaji, 2019). On the other hand, implicit age, disability, and body-weight attitudes had been largely stable and were forecasted to

remain so. These data covered the period of 2007 to 2016, and during the following 4 years (2017–2020), there was much anticipation about the new state of social-group attitudes in the United States.

Here, capitalizing on the continuous online collection of explicit and implicit attitudes through Project Implicit, we examined an additional sample of more than 3.4 million tests from 2017 to 2020 to investigate whether attitude trends persisted or altered course. Although the data have limitations (e.g., sample representativeness; see discussion below), they nevertheless are the first large-scale attitude data collected at high temporal granularity across a 14-year period that included a particularly tumultuous period in U.S.

Corresponding Author:

Tessa E. S. Charlesworth, Harvard University, Department of Psychology

Email: tessa_charlesworth@fas.harvard.edu

history. The data thus facilitated both confirmatory tests of past forecasts as well as exploratory investigations of any unexpected deviations in trends during 2017 to 2020.

Attitude Trends May Alter Course

An observer of the years 2017 to 2020 might reasonably expect that attitude trends of the past would have substantively altered following the many sociopolitical changes of this time. For one, widespread reversals in trends (i.e., backsliding from past progress toward attitude neutrality) could be expected in response to the endorsement, election, and/or presidency of Donald Trump, whose presidential communications deviated from those of past presidents in containing uncivil and even insulting comments toward specific social groups (Dimock & Gramlich, 2021). Could such public displays of explicitly stated prejudice have emboldened other individuals to explicitly express similarly hostile attitudes (e.g., Newman et al., 2020)? Additionally, could the repetition and amplification of these negative comments through news and social media have generated more negative implicit attitudes (Gawronski & Bodenhausen, 2006), even among people who might not endorse explicit bias?

At the same time, and perhaps in response, this period (2017–2020) also witnessed the growth of progressive social movements such as #metoo and Black Lives Matter (BLM), prompting national conversations about implicit and explicit biases toward a broad range of group identities. Could increased awareness of these issues, however contentious, have accelerated change, leading bias to decrease toward neutrality, possibly at even faster rates than before? Indeed, recent polls have documented widespread, rapid approval of BLM in mid-2020 (Cohn & Quealy, 2020) and particularly notable decreases in explicit racial bias since 2016 (Hopkins & Washington, 2020). Perhaps such decreases may have extended to other topics beyond race.

Finally, it is possible that attitude trends in 2017 to 2020 altered course within some demographic groups but not others (i.e., conservative respondents, most attuned to the words and actions of Trump and his supporters, may have shown unique increases). Previously, we found that demographic differences were surprisingly absent for most attitude trends, observing parallel change across respondent gender, race/ethnicity, religion, and education (Charlesworth & Banaji, 2021). However, in moments of sociopolitical upheaval, group differences are often magnified, especially across the political spectrum (Pew Research Center, 2017). Thus, with a larger time span of data over a more tumultuous

Statement of Relevance

The years 2016 to 2020 brought a whirlwind of social transformations, from unpredictable elections to widespread social movements. In the face of such events, did long-term trends of societal-level implicit and explicit attitudes continue along their forecasted paths, or did they alter course? In the present study, we tracked implicit and explicit attitudes across 14 years (2007–2020) using more than 7.1 million tests collected continuously through the Project Implicit educational website. Despite many societal changes, recent attitude trends (2017–2020) generally followed past trends (2007–2016). All explicit attitudes, as well as implicit race, skin-tone, and sexuality attitudes, continued to decrease in bias, whereas implicit disability and body-weight attitudes continued to remain mostly stable. Nevertheless, some implicit attitudes also revealed temporary increases in bias around 2015 to 2017, possibly aligned with societal events (e.g., Trump's presidential campaign). Such nuanced patterns reveal how implicit attitudes are capable of both maintaining long-term trends and responding to temporary shocks in the world.

period, attitude trends may reveal new demographic divides in their rate or even direction of change.

Persistence of Attitude Trends Over the Long Term

Despite the intuition that social-group attitudes will alter course following sociopolitical events, social scientists have often shown that, perhaps surprisingly, attitudes rarely experience sharp rises, falls, or reversals in trends, even in the face of significant events (Marsden et al., 2020; Page & Shapiro, 1992; Stimson, 2015). For instance, reviewing more than 100 explicit opinions from the General Social Survey, Smith (1994) found that approximately 75% of trends either remained stable or shifted at a slow, predictable pace. In the current data, we expected similar patterns of such attitude homeostasis, whereby societal attitudes maintain predictable trends over the long term. Homeostasis may be especially apparent for implicit attitudes, long thought to be slow to change (Bargh, 1999). Indeed, implicit attitudes have previously revealed no impacts from events as significant as the election and presidency of Barack Obama (Schmidt & Axt, 2016; Schmidt & Nosek, 2010).

Alternatively, a more nuanced pattern of both long-term persistence and temporary disruption following events could emerge in the current data, which were collected at a much higher temporal granularity than in any previous investigation over this period. After all, a major contribution of social-psychological theory has been the demonstration of short-term malleability or “context effects” on attitudes (e.g., Blair, 2002; Schwarz, 2007). Recent widespread societal events could act as large-scale context effects, prompting temporary alterations in trends that disappear as the events fade from public memory. Again, this pattern of temporary impact may be more notable on implicit attitudes, which have been theorized to be more attuned to the temporarily accessible associations in the environment (Fazio & Olson, 2003; Payne et al., 2017). This research newly tested and explored such possibilities of the persistence of trends or responsiveness to sociopolitical change for implicit and explicit attitudes toward multiple social groups over a turbulent period in recent U.S. history.

Method

The design and data-analysis plans for this study were preregistered at <https://aspredicted.org/mh958.pdf>. Deviations from the preregistration plan are noted below and summarized in the Supplemental Material available online. All data and analysis scripts are provided at the project’s OSF page (<https://osf.io/qywh4/>). Below, we review the data source, sample characteristics, collected measures (e.g., the Implicit Association Test [IAT]), and the analytic strategy of the primary change analyses. Additional analytic details are reported in the Results section, as necessary.

Data source

Data were retrieved from the Project Implicit demonstration website (<https://implicit.harvard.edu/>) from open data archived at OSF (<https://osf.io/qywh4/>). Respondents came to the website as volunteers from diverse sources (e.g., assignments from work or school, links in news articles, word of mouth), provided informed consent, and selected either the sexuality, race, skin-tone, age, disability, or body-weight IAT from among the available tests. Data across all 14 years began January 1, 2007, and ended December 31, 2020; the additional data reported in this article span from January 1, 2017, to December 31, 2020, and were merged with past data that we reported for the period January 1, 2007, to December 31, 2016 (Charlesworth & Banaji, 2019). For all data, participants were excluded if their IAT D scores fell outside of the conditions in the revised scoring algorithm (Greenwald et al., 2003).

We also excluded participants from outside the United States as well as any participants without complete explicit measures or demographic information on age, gender, race, political ideology, and education. After applying these additional demographic restrictions within the 2017 to 2020 data, we retained an average of 77% of complete U.S. sessions across tests (the Supplemental Material provides test-specific retentions).

Sample demographics

Across all attitudes, the sample from 2017 to 2020 consisted of approximately 3.4 million completed tests from U.S. respondents, with the full combined sample of continuous data from 2007 to 2020 now totaling over 7.1 million tests. Table 1 reports the sample demographics for data combined across 2007 to 2020. Overall, the majority of the sample was female, White, college-educated, and liberal. Additionally, for comparison, Tables S3.1 and S3.2 in the Supplemental Material provide sample demographics separated across previous data from 2004 to 2016 (the full data previously reported by Charlesworth & Banaji, 2019) and 2017 to 2020. Sample demographics were generally similar across both sets of data (e.g., both samples were overall approximately 66% female, 72% White, 9–10% Black, and 5–6% Asian). However, as we report in the Supplemental Material, over time, the sample also became slightly less educated (86% of people in the previous sample had a college degree or more education, whereas 79% of the most recent sample had a college degree or more), less politically conservative (from 26% to 21%), less politically neutral (from 28% to 25%), and more liberal (from 46% to 53%). Crucially, we controlled for such changes through raking and weighting approaches (described below) that adjust attitude estimates so that demographic compositions are consistent over time.

Materials

Implicit Association Test. The IAT is a computerized task that compares respondents’ reaction times when categorizing social groups (e.g., thin vs. fat) and attributes (good vs. bad). For all tests, respondents categorized target stimuli (e.g., images of fat and thin body silhouettes, good words such as “joyful” and “friend,” and bad words such as “evil” and “poison”) into four categories (e.g., *fat*, *thin*, *good*, *bad*). Average response latencies were compared across two critical blocks in which groups and attributes were sorted to the same computer key. For example, in the body-weight IAT, response latencies were compared across blocks in which (a) thin **Z** good was sorted with the “E” key and

Table 1. Sample Demographics Across All Attitudes From 2007 to 2020

Attitude	N	Age (years)		Gender (%)		Race (%)			College educated	Politics (%)		
		M	SD	Female	Male	White	Black	Asian	(%)	Liberal	Neutral	Conservative
Age	851,955	28.41	13.10	69.23	30.28	74.04	8.42	5.89	87.48	37.70	36.65	25.65
Disability	366,889	29.91	12.93	73.59	25.34	77.73	7.32	4.72	89.57	46.06	30.41	23.53
Skin tone	803,532	29.55	12.64	69.80	29.39	62.41	16.17	5.76	88.29	50.93	29.76	19.30
Race	2,962,461	29.67	13.35	62.47	36.78	71.21	11.24	5.37	78.89	50.23	28.04	21.73
Sexuality	1,128,350	26.78	11.92	66.09	32.15	73.83	8.57	5.22	87.19	53.36	27.81	18.83
Body weight	999,710	27.92	12.15	71.86	27.59	75.78	6.99	5.64	78.87	46.18	27.27	26.56
Total	7,112,897	28.81	12.88	66.57	32.55	71.95	10.23	5.46	82.84	48.52	29.24	22.23

Note: For sample demographics across separated data sets (2007–2016 vs. 2017–2020), see Tables S3.1 and S3.2 in the Supplemental Material available online. Sample demographic representations were generally consistent over time.

fat \bar{Z} bad was sorted with the “I” key, whereas versus (b) thin \bar{Z} bad was sorted with the “E” key and fat \bar{Z} good was sorted with the “I” key. The assumption is that faster responses in these paired blocks reflect a stronger association between the group and evaluative attribute (e.g., thin–good/fat–bad). In all tests, positive IAT D scores indicate a relatively positive association toward the typically preferred group (i.e., young, White, light-skinned, able-bodied, straight, thin). All IATs used a standard seven-block format; the order of the two paired blocks was randomized across respondents.

Issues concerning the test-retest reliability and predictive validity of implicit measures continue to be debated (for discussions, see Gawronski, 2019; Jost, 2019). Here, we note that these concerns are generally raised in the context of the stability and predictive validity of individual attitudes and are less applicable to aggregated IAT scores taken as an indicator of collective attitudes (Payne et al., 2017). That is, aggregating measures of implicit attitudes to provide population estimates—as we do in the current article—advantageously reduces measurement error and addresses many of the common psychometric questions concerning the IAT, including both the magnitude of attitude–behavior correlations and test-retest stability.

Explicit preference. Explicit attitudes were assessed on a 7-point Likert scale ranging from $\bar{1}3$ to $\bar{Z}3$, where positive scores indicate the typical preferences in society (e.g., a score of $\bar{Z}3$ indicates “I *strongly* prefer thin people to fat people,” a score of $\bar{Z}2$ indicates “I *moderately* prefer thin people to fat people,” and a score of $\bar{Z}1$ indicates “I *slightly* prefer thin people to fat people”), and negative scores indicate atypical preferences in society (e.g., “I *strongly/moderately/slightly* prefer fat people to thin people”). The midpoint (0) represents equal liking of both groups (i.e., no self-reported preference between the groups).

Analytic strategy: primary analyses of change over time

Autoregressive-integrated-moving-average (ARIMA) models. In a previous study (Charlesworth & Banaji, 2019), we used ARIMA time-series models (Cryer & Chan, 2008) to model long-term attitude data because ARIMA models are designed to model the unique features of time-series data (i.e., autocorrelations, seasonality, and forecasts). Specifically, ARIMA models explicitly account for temporal autocorrelations (i.e., model the fact that a measure at time t is dependent on, and correlated with, a measure at time $t - 1$). ARIMA models also capture non-linearity and seasonality in attitude trends, thereby offering a more fine-grained understanding of the underlying profiles of attitude change. Finally, and most relevant for the current article, ARIMA models can offer forecasts of future patterns of change, inferred from underlying patterns in past data (e.g., the patterns of trend, autocorrelations, seasonality). Such forecasts provide an opportunity to quantitatively test whether long-term attitude trends either persist predictably over time (and remain within the boundaries set by forecasts) or reveal altered trends in response to new social shocks unforeseen by forecasts (and thus move off from the forecasted trends).

In this project, ARIMA models were fitted to the aggregated monthly time series for each of the six attitudes using an automated forecasting algorithm implemented through the `auto.arima` function in the *forecast* package (Version 8.13; Hyndman & Khandakar, 2008) in the R programming environment (R Core Team, 2020). Models are specified with (a) an *autoregressive* (p) parameter, which specifies the number (order) of autoregressive terms to explain the autocorrelation structure in the time series (e.g., the dependencies between successive time points); (b) the *differencing* (d) parameter, which specifies the number of differencing parameters to explain the differences in magnitude

between consecutive time points and render the time series flat or stationary; and (c) the *moving-average* parameter (q), which specifies the number of moving-average parameters used to explain the lagged errors or random shocks in the series. These three parameters are applied to both the general trend of the series as well as to the seasonal component of the series. Finally, in some instances in which the time-series trend is clearly linear, the best-fitting ARIMA model may also specify a *drift* parameter (essentially a slope parameter). Thus, in the end, any ARIMA model can have the structure $(p, d, q)(P, D, Q)Z^k$ drift, in which the first three values specify the order of nonseasonal parameters, the second three values specify the order of seasonal parameters, and the drift parameter specifies any slope.

When ARIMA parameters are interpreted, the order of the d parameter is informative in revealing whether the time series is already stable (when $d = 0$) or is changing over time (when d is nonzero). Here, we inferred that an attitude trend was stable when the ARIMA model included no differencing parameter (because this indicates the trend is already flat) and when the ARIMA model forecasts indicated that they would not touch neutrality in more than 200 years, should past trends continue (because this indicates high confidence in stability going forward). The order of autoregressive and moving-average parameters is generally not intuitively interpretable on its own. Instead, the parameters can simply be understood as capturing the underlying structure of the time series (e.g., whether the underlying structure includes seasonal components or not) used to generate forecasts.

Comparing 2017 to 2020 trends to ARIMA model forecasts. Using ARIMA models, we previously generated forecasts on the basis of attitude trends from 2007 to 2016 to predict the mean future trends should the same underlying time-series structure continue (Charlesworth & Banaji, 2019). Confidence intervals for the forecasts were also calculated (using the `auto.arima` function) by (a) computing the standard deviation of the forecast distribution (a transformation of the ARIMA model's fitted residuals) and then (b) multiplying that standard deviation by the z -score values for a given confidence interval (specifically by 1.96 for a 95% confidence interval and by 1.28 for an 80% confidence interval, as is standard for the z -score distribution). Following our preregistration, we tested the accuracy of these past forecasts using a set of four commonly used accuracy statistics (Hyndman & Athanasopoulos, 2013): the mean absolute error (MAE), root-mean-square error (RMSE), mean absolute scaled error (MASE), and mean absolute percentage error (MAPE). The computation and interpretation of each statistic is defined briefly below and elaborated in the Supplemental Material.

Briefly, the MAE gives the average absolute amount (in raw-score units) that the observed values differ from the mean forecasted values; lower MAEs indicates better forecasts. The RMSE, although commonly used for assessing model accuracy, is not as intuitively interpretable, but smaller values can be similarly understood as indicating greater accuracy. The MASE is the recommended approach for assessing forecast accuracy and can be conveniently interpreted such that values greater than 1 indicate that a naive forecast (i.e., estimating that future data will persist at the last value observed) would be better than the forecast generated by the ARIMA model. In contrast, MASE values less than 1 indicate that the forecasts from the ARIMA models are better than the naive forecast. Finally, the MAPE gives the average percentage by which the forecast deviates from observed values and is scale invariant. However, the MAPE has been criticized because it cannot be accurately estimated if any observed value is zero, and it overpenalizes negative errors; consequently, we computed the MAPE (following our preregistration) but report the results only in the Supplemental Material.

Robustness analyses.

Addressing sample change. To control for confounds of sample change over time, we adjusted the data using raking and weighting with the *anesrake* package (Version 0.8; Pasek, 2018) in the R programming environment. We describe these raking and weighting approaches in detail in the Supplemental Material. Briefly, we first computed the target weights of sample demographics across the full study period (from 2007–2020) on the proportions of gender (female, male), race (White, American Indian, Asian, Black, biracial, multiracial, other), education (less than high school, high school graduate, college graduate, graduate school, professional degree, advanced degree), age (0–24, 25–34, 35–44, 45–54, 55+), and politics (liberal, neutral, conservative; collapsed from a 7-point Likert scale ranging from *strongly conservative* to *strongly liberal*). Next, we reweighted the individual data in each year to create the same demographic compositions across time for all 14 years.

For illustration, suppose the representation of liberals in a given data set across the full study period (2007–2020) was approximately 60% (the target weight), but in 2007 it was 65% and in 2020 it was 55%. Because, relative to the whole sample across time, liberals were overrepresented in 2007 and underrepresented in 2020, we would reweight the data so that liberals in 2007 had a weight less than 1 (decreasing their contribution in that year), whereas liberals in 2020 had a weight greater than 1 (increasing their contribution in that year). By applying this logic across the intersection of gender, age, race, education, and politics, we would essentially

be equating the demographic composition of the sample across the full study period. Similar approaches are routinely used in public polling (such as the General Social Survey) to reweight data in accordance with a given target demographic composition.

Addressing alternative weighting, repeat test takers, and sample sources. We also ensured that these trends were generally consistent across various analytic decisions. First, we ensured robustness to our decisions around the weighting approach by recomputing analyses across unweighted and weighted data. Indeed, we found that the general conclusions were equivalent across both unweighted and weighted data (see the Supplemental Material).

Second, we ensured robustness across both novice and repeat IAT test takers. That is, a common concern is that some respondents come back to take the IAT for a second time (or more), which may reduce the IAT scores through practice effects and result in regression to the mean artificially showing decreases over time. In contrast to this potential critique, we found that similar trends were observed in novice test takers alone (see the Supplemental Material) who reported never having visited Project Implicit or taken an IAT before. This consistency across the overall sample and novice test takers lends confidence that the reported results were not merely due to repeat test takers.

Third, we ensured robustness across participants coming to the website through different sources, whether from (a) assignments (e.g., assignment for school or work; approximately 76% of the sample), (b) self-directed searches and word of mouth (e.g., clicking on a link in a news article; approximately 6% of the sample), or (c) other sources (as reported by the remaining 18% of the sample). Intuitively, the latter group of self-directed participants might be more motivated to change and have lower overall biases to begin with because they are intrinsically interested in the website and in addressing implicit bias. The group of assigned participants may, in contrast, have more test hesitation or skepticism, be less motivated to change, and have higher overall biases to begin with. If there are such differences across sample sources in either the overall magnitude of bias or the rate of attitude change, then it is crucial to examine any changes in the relative proportions of these two sources of participants (e.g., with more self-directed participants in later years relative to earlier years) to ensure that observed attitude trends are not due to this specific type of nondemographic sample change.

Eliminating both concerns around the impact of sample sources, we found that (a) patterns of change and, for the most part, overall magnitudes of bias, were similar across both sources of participants (assigned to

participate or self-directed) and (b) there has been little meaningful change in the relative proportions of participants coming from assignments or self-directed sources between 2017 and 2020; see the Supplemental Material). Thus, it is unlikely that the observed trends are artifacts of the self-selected participants alone or of differences in the proportion of self-selected participants across time.

Results

We will summarize the results in five sections. First, we will report the means and correlations collapsed across the full study period in the overall data of 14 years (2007–2020). Second, we will turn to the primary questions of change, reporting overall trends of change across all 14 years (2007–2020), using ARIMA models, as well as an exploratory analysis of the relationship between implicit and explicit attitude trends. We will report trends across all available data, rather than across only a small portion of years (e.g., only recent years), to provide the most complete understanding of contemporary long-term attitude change and stability. Additionally, including earlier data (from 2007–2016) is necessary to demonstrate the accuracy of the ARIMA method, as we can directly compare forecasts generated from past data (2007–2016) with the observed trends from 2017 to 2020.

In the third section, we will investigate demographic differences and similarities in the overall trends (2007–2020), including a discussion of both the elimination and emergence of differences since 2017. In the fourth section, after observing striking patterns of temporary increases in some implicit attitudes, we will provide post hoc exploratory analyses using segmented regressions to identify the timing, magnitude, and duration of such increases. Here, we will also provide a speculative exploration of explanations for the uptick in implicit race attitudes using difference-in-differences analyses across participant political orientation and geography. In the fifth section, we will close with a high-level summary of the key results.

Means and correlations for implicit and explicit attitudes

Averaging all available data from 2007 to 2020, we found that explicit and implicit attitudes were significant in the expected, positive direction for all topics (i.e., pro-straight/anti-gay, pro-White/anti-Black, pro-light skin/anti-dark skin, pro-young/anti-old, pro-able-bodied/anti-disabled, and pro-thin/anti-fat; Table 2). The strongest implicit and explicit biases were observed for disability, body-weight, and age attitudes, and robust but

Table 2. Means and Correlations for the Six Implicit and Explicit Social-Group Attitudes (2007–2020)

Attitude	Implicit attitudes				Explicit attitudes				Implicit–explicit correlation	
	<i>M</i>	<i>SD</i>	95% CI	<i>d</i>	<i>M</i>	<i>SD</i>	95% CI	<i>d</i>	<i>r</i>	95% CI
Sexuality	0.22	0.50	[0.21, 0.22]	0.43	0.36	1.24	[0.36, 0.36]	0.29	.44	[.44, .45]
Race	0.29	0.44	[0.29, 0.29]	0.65	0.17	1.03	[0.17, 0.17]	0.17	.31	[.31, .31]
Skin tone	0.29	0.43	[0.29, 0.30]	0.69	0.18	0.92	[0.18, 0.19]	0.20	.23	[.22, .23]
Age	0.43	0.39	[0.43, 0.43]	1.11	0.41	1.17	[0.41, 0.41]	0.35	.12	[.12, .12]
Disability	0.51	0.45	[0.51, 0.51]	1.13	0.40	0.90	[0.39, 0.40]	0.44	.14	[.14, .15]
Body weight	0.48	0.41	[0.48, 0.48]	1.18	0.79	1.06	[0.79, 0.79]	0.75	.21	[.20, .21]

Note: Means for implicit attitudes are *D* scores; means for explicit attitudes were obtained from Likert scales. All means and correlations are significantly different from zero ($p < 0.001$). CI = confidence interval.

relatively weaker implicit biases were observed for race, skin-tone, and sexuality attitudes. Significant positive correlations, ranging in magnitude from small to medium effects, were observed between implicit and explicit attitudes for all six topics. The strongest (medium-sized) correlations were observed for sexuality and race attitudes and the weakest (small) correlations for disability and age attitudes.

Overall trends (2007–2020) and accuracy of past forecasts

Sexuality attitudes. Among all implicit attitudes, the fastest and most consistent change from 2007 to 2020 was observed in implicit sexuality attitudes, which decreased in implicit anti-gay/pro-straight bias by 65% (Table 3). ARIMA model forecasts project that implicit sexuality attitudes could reach neutrality as early as mid-2022, should past trends continue (Fig. 1; see the Supplemental Material for a full report of ARIMA model forecasts). Similarly, explicit sexuality attitudes have decreased in bias by a comparable 75% over the past 14 years and are forecasted to reach explicit attitude neutrality as early as mid-2021, if the trends persist.

The consistency of sexuality attitude change was also revealed when we examined only recent years: For 2017 to 2020, both implicit and explicit sexuality attitudes decreased in bias at rapid rates (27% and 41%, respectively, over 4 years) that even outpaced the mean forecasted change based on previous ARIMA models (Table 4). This particularly rapid change meant that implicit sexuality attitudes had the lowest accuracy of any attitude (i.e., the largest MAE and RMSE results and a MASE result that indicated the ARIMA model forecasts were not better than a naive forecast; Table 4). Further inspection showed that the months that fell outside the confidence intervals (and thus the months driving the low accuracy) fell below the lower bounds,

thus indicating that recent change was relatively faster than would be expected from past trends.

Race and skin-tone attitudes. Since 2007, implicit race and skin-tone attitudes have decreased in bias by 26% and 25%, respectively (Table 3; Fig. 1), notably, less than half the percentage of change observed in sexuality attitudes. The lower bounds of the ARIMA model forecasts suggest that these attitudes could first touch neutrality in approximately 15 years (implicit race attitudes) and 20 years (implicit skin-tone attitudes; see the Supplemental Material), assuming past trends continued as before.

The transformation of explicit race and skin-tone attitudes was much greater than for implicit attitudes: Since 2007, explicit race and skin-tone attitudes have dropped in bias by 98% and 79%, respectively (Table 3). In fact, by the end of 2020, the mean of explicit race attitudes had passed neutrality in the full sample, meaning that, today, large numbers of Americans express no preference for White/light-skinned over Black/dark-skinned groups. Indeed, most demographic subgroups showed this reduction in prejudice, and even the furthest groups from expressing neutral attitudes (e.g., White Americans, conservatives) were forecasted to reach neutrality within the next few years, should past trends continue (see the Demographic Differences section). Only future analyses will tell us whether these explicit attitudes now reach asymptote at a neutral attitude (zero bias) or whether there is a possibility of the population expressing Black-favoring explicit attitudes in the future. Ultimately, we note that this arrival at explicit race neutrality is an important milestone in measuring race prejudice: It is the first moment that a large sample has shown zero explicit bias since surveys of race attitudes began in the early years of the 20th century (see Banaji & Greenwald, 2013, Appendix 2).

Turning to just the 2017 to 2020 trends, we found that, as with sexuality attitudes, change in implicit and

Table 3. Overall Patterns of Change Across the Full Study Period (2007–2020) and Autoregressive-Integrated-Moving-Average (ARIMA) Model Forecasts

Attitude	Start	End	% 8	Raw 8	ARIMA model parameters (<i>p</i> , <i>d</i> , <i>q</i>) (<i>p</i> , <i>d</i> , <i>q</i>)	ARIMA model forecast		
						95% CI lower limit	<i>M</i>	95% CI upper limit
Implicit attitudes								
Sexuality	0.33	0.12	!65	0.22	(0, 1, 2) (2, 0, 0) Z drift	1 year, 6 months (<i>n</i>)	6 years, 9 months (<i>n</i>)	26 years, 6 months (<i>n</i>)
Race	0.33	0.24	!26	0.09	(1, 1, 1) (2, 0, 0)	15 years (<i>n</i>)	2 200 years	70 years 8 months (<i>d</i>)
Skin tone	0.33	0.24	!25	0.08	(0, 1, 1) Z drift	20 years, 9 months (<i>n</i>)	41 years, 4 months (<i>n</i>)	81 years, 6 months (<i>n</i>)
Age	0.45	0.41	!8	0.04	(1, 1, 1) (1, 0, 1)	2 200 years	2 200 years	2 200 years
Disability	0.50	0.49	!2	0.01	(1, 0, 1)	2 200 years	2 200 years	2 200 years
Body weight	0.47 ^a	0.47	!1	0.01	(1, 0, 0) (1, 0, 2)	2 200 years	2 200 years	2 200 years
Explicit attitudes								
Sexuality	0.68	0.17	!75	0.51	(1, 1, 1) (2, 0, 0)	6 months (<i>n</i>)	2 200 years	68 years, 10 months (<i>d</i>)
Race	0.32	0.01	!98	0.31	(1, 1, 2) (1, 0, 1) Z drift	Already neutral	Already neutral	4 years, 2 months (<i>n</i>)
Skin tone	0.27	0.06	!79	0.21	(0, 1, 1) (1, 0, 1)	Already neutral	2 years, 11 months (<i>n</i>)	30 years, 5 months (<i>d</i>)
Age	0.52	0.40	!22	0.11	(0, 1, 1) (0, 1, 2)	3 years, 7 months (<i>n</i>)	27 years, 7 months (<i>n</i>)	12 years, 4 months (<i>d</i>)
Disability	0.54	0.34	!37	0.20	(0, 1, 1) Z drift	16 years, 5 months (<i>n</i>)	21 years, 6 months (<i>n</i>)	26 years, 9 months (<i>n</i>)
Body weight	0.96	0.67	!31	0.30	(0, 1, 2) (2, 0, 0) Z drift	11 years, 3 months (<i>n</i>)	26 years, 6 months (<i>n</i>)	64 years (<i>n</i>)

Note: Start and end values are D scores. Start and end values, as well as percentage of change (% **8**) and raw change (raw **8**), are calculated from the start and end points of the decomposed trend line (removing seasonality and noise). Using these decomposed trend values rather than raw monthly estimates ensures that results from an outlier month (unusually high or low in bias) does not disproportionately affect the reported change. The first three parameters of the ARIMA model are nonseasonal; the second three values are seasonal; and, in some cases, a drift parameter (slope) is included. In each set of parameters, *p* specifies the number of autoregressive parameters used to explain the autocorrelations in the data, *d* specifies the number of differencing parameters necessary to make the series stationary, and *q* specifies the number of moving-average parameters used to explain the lagged forecast errors. Forecasts indicate the number of years and months it could take for the attitude to reach neutrality (*n*) or to double in magnitude (*d*) from January 2021, should past trends continue. CI **1** confidence interval. ^aThe Implicit Association Test stimuli for implicit body-weight attitudes changed in April 2010 from face images to body silhouettes; to facilitate inferences, we report the current data only for the continuous test (the body-silhouette stimuli), starting in April 2010. The start value, percentage of change, and raw-change values for the implicit body-weight results reflect the later start date.

explicit race and skin-tone attitudes outpaced the mean forecasted trends from past ARIMA models but nevertheless largely remained within the forecasted confidence intervals (Table 4). Indeed, the RMSE and MAE indicated good forecast accuracy with relatively low error, although the MASE indicated that forecasts were not better than a naive forecast because of the underestimation in the observed rate of change. This more rapid change is particularly notable in mid-to-late 2020: Implicit race attitudes consistently fell below the lower limit of the forecasted 80% confidence interval for June to December 2020, and implicit skin-tone attitudes fell below the forecasted 80%-confidence-interval limit for 9 of the 12 months of 2020. Interestingly, this moment of faster change coincides with the growth of the BLM

movement in summer 2020 (Cohn & Quealy, 2020). Although the recency of these drops makes it difficult to model the possible impact of events, such findings nevertheless appear in line with work suggesting that earlier BLM protests coincide with decreases in implicit race bias (Sawyer & Gampa, 2018).

Age attitudes. In contrast to the three consistently changing attitudes (sexuality, race, and skin tone), implicit age attitudes have been relative stability over the past 14 years (2007–2020), decreasing in anti-old/pro-young bias by 8% (Table 3). Although an 8% change for implicit age attitudes is greater than what was reported over the 10 years of 2007 to 2016 (by Charlesworth & Banaji, 2019), the change nevertheless remains sufficiently slow that

ARIMA confidence intervals indicate we cannot expect attitude neutrality even within the next 200 years, should past trends continue. Indeed, looking at only the 2017 to 2020 trends, we see evidence for such continued stability, because implicit age attitudes decreased by less than 1% over the 4 years of 2017 to 2020 (Table 4).

Explicit age attitudes have, across the full study period (2007–2020), decreased in bias by 22%, the slowest change of all explicit attitudes (Table 3). Even though change was relatively slow, the ARIMA model forecasts from the 14-year trend nevertheless suggest that explicit attitude neutrality could be possible as early as mid-2024 (see the Supplemental Material), on the basis of the lower bound of the forecasted confidence interval. However, the most recent trends paint

a less optimistic picture: Between 2017 and 2020, explicit anti-old/pro-young attitudes have switched direction over the past 4 years to be stable and even hint at increasing bias (increasing by about 4%; Table 4). This is the only case of any attitude (whether implicit or explicit) changing in the opposite direction of the forecast and the only case of any explicit attitude changing toward increasing bias.

Why might explicit age attitudes be uniquely increasing in 2017 to 2020? One possibility is the increase in negative information that has emerged regarding the elderly and their vulnerability during the COVID-19 pandemic (Fraser et al., 2020), but this is implausible given that attitude change commenced in late 2016, prior to the onset of the pandemic. Instead, we

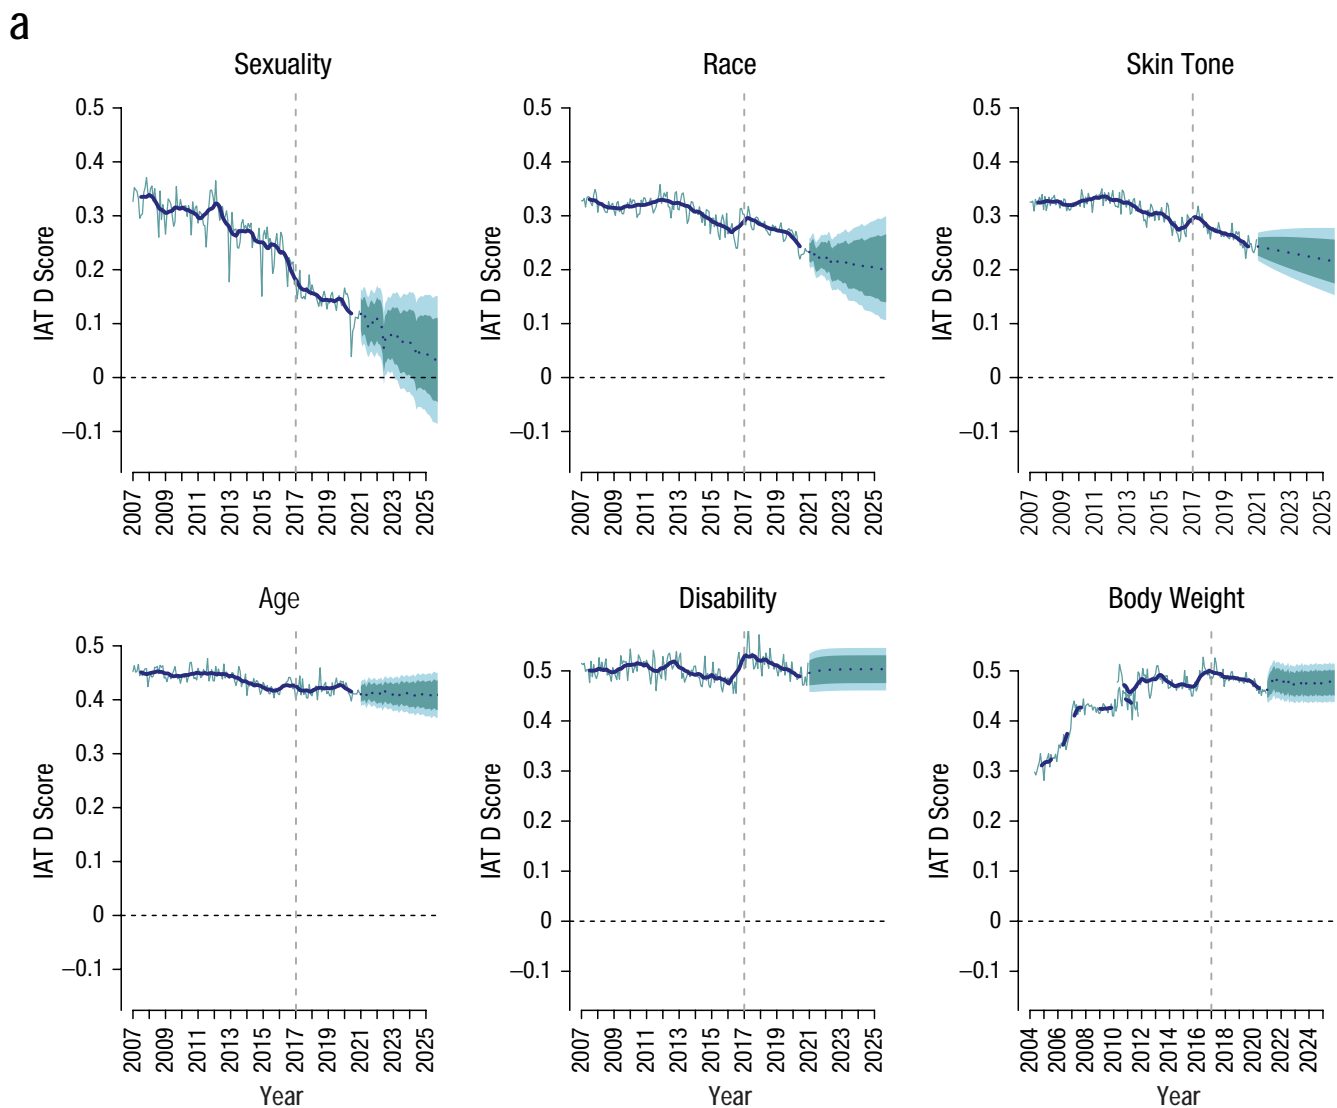


Fig. 1. (continued on next page)

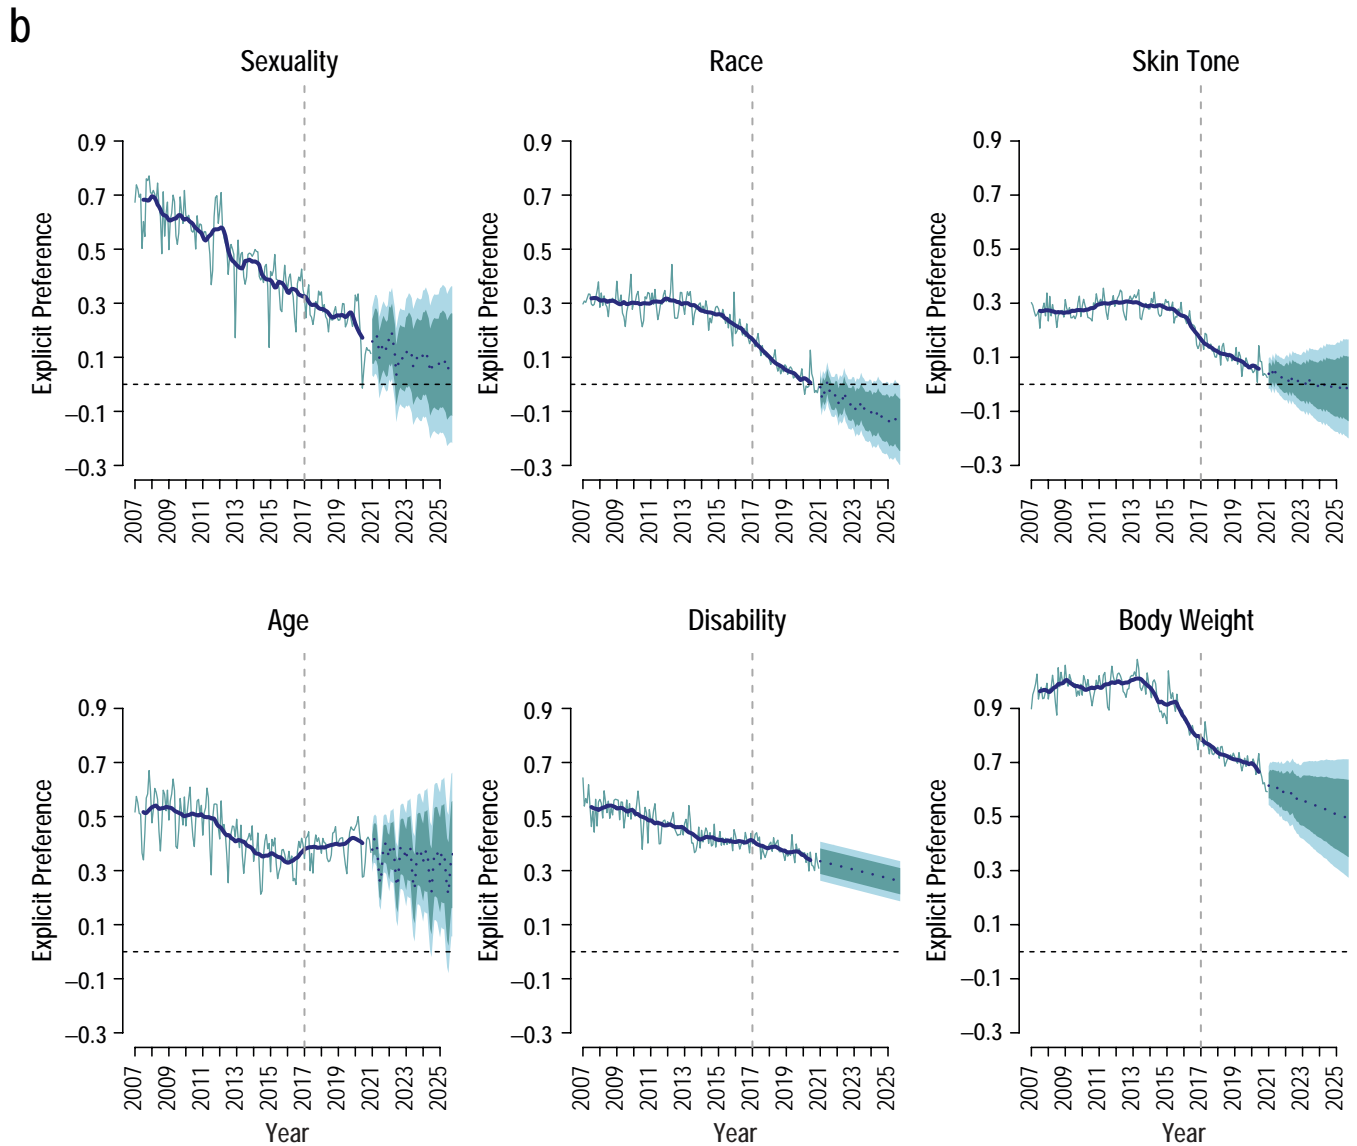


Fig. 1. Implicit (a) and explicit (b) attitude trends and forecasts from 2007 to 2020. Vertical dashed gray lines indicate the onset of the 2017 to 2020 data. The thick, dark-blue lines indicate the decomposed trend (removing seasonality and noise) for the raw observed data (thin, light-blue lines). Dark-shaded areas indicate 80% forecasted confidence intervals, and light-shaded areas represent 95% forecasted confidence intervals. Implicit body-weight attitudes included two tests (differentiated by stimuli); the early test (using face stimuli) is plotted with a dashed line and the later test (using silhouette stimuli) with a solid line.

speculate that the shift in anti-elderly explicit attitudes may be spurred, in part, by a growing generational war between younger (e.g., millennial, Gen Z) and older (e.g., baby boomer) age groups—a conflict that increased throughout 2016 to 2019 and culminated in memes such as “OK boomer” (a retort to perceived traditional beliefs of the baby-boomer generation; Mueller & McCollum, 2022). This generational conflict incorporates the belief among younger people (and particularly among younger liberals) that elderly people stand in the way of necessary progress on issues such as economic inequality, climate change, and racial

injustice. Frustration with elderly people’s perceived resistance to change could be the cause of negative explicit anti-elderly attitudes. Exploratory and admittedly post hoc support for this interpretation comes from initial tests of the differences in explicit age attitude trends between younger and older respondents and liberal and conservative respondents, reported below.

Disability attitudes. Across 14 years, implicit disability attitudes have shifted by only 2%, and the best-fitting ARIMA model indicates no evidence of past or forecasted

Table 4. Trends From 2017 to 2020 and Accuracy of Autoregressive-Integrated-Moving-Average (ARIMA) Model Forecasts Based on Past Data (2007–2016)

Attitude	Trends from 2017–2020				Accuracy of ARIMA model forecast						
	Start	End	% 8	Raw 8	ARIMA model parameters (<i>p</i> , <i>d</i> , <i>q</i>) (<i>p</i> , <i>d</i> , <i>q</i>)		Forecasted % 8	Forecasted raw 8	MAE	RMSE	MASE
Implicit attitudes											
Sexuality	0.16	0.12	!27	!0.04	(1, 1, 1)		!21	!0.04	0.05	0.05	1.47
Race	0.29	0.24	!17	!0.05	(2, 1, 0) (1, 1, 0)		O!1	O!0.01	0.02	0.02	1.06
Skin tone	0.29	0.24	!17	!0.05	(0, 1, 1)		O!1	O!0.01	0.02	0.03	1.47
Age	0.42	0.41	O!1	!0.00	(0, 0, 0)		O!1	O!0.01	0.01	0.02	1.20
Disability	0.53	0.49	!8	!0.04	(0, 1, 1) (0, 0, 1)		O!1	O!0.01	0.02	0.03	0.95
Body weight	0.49	0.47	!6	!0.03	(0, 1, 1) Z drift		O!1	O!0.01	0.01	0.01	0.50
Explicit attitudes											
Sexuality	0.29	0.17	!41	!0.12	(0, 1, 0)		!7	!0.02	0.09	0.12	1.19
Race	0.13	0.01	!96	!0.13	(4, 1, 1) Z drift		!57	!0.09	0.05	0.06	1.49
Skin tone	0.14	0.06	!59	!0.08	(0, 1, 1)		!33	!0.05	0.04	0.05	1.47
Age	0.39	0.40	Z 4	Z 0.01	(2, 0, 0) (0, 1, 0)		!2	!0.01	0.06	0.07	1.25
Disability	0.39	0.34	!14	!0.05	(0, 1, 1) (1, 0, 0)		!13	!0.05	0.02	0.03	0.61
Body weight	0.77	0.67	!13	!0.10	(0, 1, 1) Z drift		!6	!0.04	0.06	0.07	1.13

Note: Start and end values, as well as percentage of change (% **8**) and raw change (raw **8**), are calculated from the start and end points of the decomposed trend line (removing seasonality and noise). Using these decomposed trend values rather than raw monthly estimates ensures that results from an outlier month (unusually high or low in bias) does not disproportionately affect the reported change. The first three parameters of the ARIMA model are nonseasonal; the second three values are seasonal; and, in some cases, a drift parameter (slope) is included. In each set of parameters, *p* specifies the number of autoregressive parameters used to explain the autocorrelations in the data, *d* specifies the number of differencing parameters necessary to make the series stationary, and *q* specifies the number of moving-average parameters used to explain the lagged forecast errors. Forecasts indicate the number of years and months it could take for the attitude to reach neutrality (*n*) or to double in magnitude (*d*) from January 2021, should past trends continue. Mean absolute error (MAE) refers to the actual scale-dependent difference between the observed and forecasted mean, calculated as $\text{mean}(|e|)$. Root-mean-square error (RMSE) refers to the square root of the scale-dependent amount that the mean of the forecast was off from the observed mean, calculated as $\sqrt{\text{mean}(|e^2|)}$. Mean approximate scaled error (MASE) refers to the relative accuracy of the ARIMA model forecast compared with a standard naive forecast based on the training data (i.e., from a simple prediction of y_t from y_{t-1}); values greater than 1 indicate that the errors from the ARIMA model forecast are larger than the errors from the naive forecast (i.e., the naive forecast would have been just as efficient).

change (i.e., the ARIMA model has no differencing parameter; Table 3). In sharp contrast, explicit disability attitudes have consistently decreased in anti-disabled/pro-able-bodied bias by 37% between 2007 and 2020, and forecasts suggest that they could reach no explicit bias by approximately 2037, should past trends persist (see the Supplemental Material). Although the long-term trends of implicit disability attitudes show stability, looking only at more recent data reveals a temporary uptick starting in early 2017 and ending about a year later. Below, we provide a post hoc exploration of this increase to better identify the timing, magnitude, and duration, as well as speculate on possible coinciding events. Despite this noticeable increase in bias, both implicit and explicit disability attitudes revealed high accuracy in ARIMA model forecasts, as revealed by both low RMSE and MAE results as well as MASE results indicating that ARIMA model forecasts were significantly more accurate than forecasts from a naive model (Table 4).

Body-weight attitudes. Since 2010, when the continuous body-weight test (using body-silhouette stimuli) became available, implicit body-weight attitudes have decreased by only 1%, and the data suggest no future change (i.e., the ARIMA model includes no differencing parameter; Table 3). Thus, although we previously reported a unique increase in implicit body-weight attitudes that occurred from 2004 to 2010 (Charlesworth & Banaji, 2019), the trend now appears to have leveled off. Indeed, inspecting only the most recent data reveals that the trend now even hints toward decreases in bias (a drop of about 6% from 2017 to 2020; Table 4).

In contrast to the high and stable implicit body-weight bias, explicit body-weight attitudes have decreased in bias over all 14 years (2007–2020), dropping in anti-fat/pro-thin attitudes by 31% (Table 3). Within 2017 to 2020, the decrease even slightly outpaced the forecast (attitudes dropped by 13% in 4 years, despite being forecasted to drop only 6%; Table 4).

Finally, both implicit and explicit body-weight attitudes showed high forecast-accuracy statistics: MAE and RMSE were both low, and body weight had the lowest MASE for any implicit attitude. Thus, as with nearly all implicit and explicit attitudes, the trends of past body-weight attitudes generally predicted the observed patterns in 2017 to 2020, despite the tumultuous socio-political changes of this period.

Exploratory analysis: relationship between implicit and explicit attitude change. Among the many interesting questions that can be posed about this data is whether, and to what extent, implicit attitude change precedes or follows the changes observed in explicit attitudes. This question is of particular interest for thinking about possible mechanisms of change, since most existing theories (albeit proposed to explain individual-level rather than societal-level change) describe mechanisms in terms of the mediation and relationships between implicit and explicit attitudes (e.g., Gawronski & Bodenhausen, 2006). We offer an additional, nonpreregistered exploratory analysis of this question in the Supplemental Material using Granger causality analyses. In brief, Granger causality investigates whether a time-series x (e.g., Google searches) can help predict a second time-series y (e.g., implicit attitudes) at a lag of t time steps (e.g., at a lag of 1 month), thus providing initial insight into the time-locked relationships between two series.

As reported in the Supplemental Material, Granger causality models revealed no systematic relationship between implicit and explicit attitudes across topics or across time lags. That is, although there were some idiosyncratic relationships (especially for race attitudes), explicit attitude change did not systematically precede implicit attitude change, nor did implicit attitude change precede explicit attitude change. Thus, it appears that the two measures of attitudes are most likely sensitive to different sources of influence and societal transformations. In fact, this finding is in line with conclusions we will discuss below that show implicit attitudes as more attuned to widespread, macrolevel change in the societal environment, whereas explicit attitudes may respond more to demographic-specific forces and long-term consistency in attitudes. Future work is positioned to continue exploring the societal-level mechanisms that may help explain the relationships between implicit and explicit attitude trends.

Demographic differences (age, race, sexuality attitudes)

Analytic strategy. We tested demographic similarities against demographic differences in attitude trends across

14 years (2007–2020). Demographic subgroups were defined by respondents' religion (Christian, Jewish, non-religious, other), education (college educated, non-college educated), race (White American, Asian American, Black American), gender (male, female), politics (liberal, conservative), age (younger ≤ 20 years, older ≥ 40 years, determined on the basis of the youngest and oldest age cutoffs with sufficient sample sizes), and, for sexuality attitudes only, sexual orientation (straight, gay/lesbian, bisexual). The attitude time series for each demographic subgroup was weighted (see above) to help ensure that the comparison subgroups (e.g., male and female) approximated each other on all relevant demographics except the one of interest (e.g., attitudes for male and female respondents were weighted to match on the representation of race, age, politics, and education). Additionally, because of the smaller sample sizes when respondents were separated by both month and demographic subgroup, we used only the three largest tests—age, race, and sexuality—as in previous work (Charlesworth & Banaji, 2021).

The trends of two demographic subgroups could be interpreted as (a) *converging* (i.e., the trends were moving toward each other over time), such as if men, who started out as initially more biased, decreased in bias faster than women, essentially “catching up”; (b) *diverging* (i.e., the trends were moving away from each other over time), such as if women, who already started out as initially less biased, decreased in bias faster than men, essentially “falling away”; or (c) *parallel change*, such as if men and women decreased in bias at the same rate and direction. The interpretation of converging, diverging, or parallel change was informed by examining the differenced series between the two demographic groups (e.g., subtracting the trajectory of men's attitudes from that of women's attitudes), among other criteria (see Charlesworth & Banaji, 2021). If the differenced series is moving toward neutrality, that indicates the trends are converging, because the gap between the two groups' attitudes is reducing over time as the attitudes of one group move toward the attitudes of the other group. If the differenced series is moving away from neutrality, that indicates the trends are diverging because the gap between the two groups is growing over time as one group pulls away from the other. Finally, if the differenced series is stable (e.g., there is no d parameter in the ARIMA model), that indicates the trends are parallel, because both groups must be changing at similar rates in the same direction to keep the gap equivalent over time. Interpretations for all demographic differences are reported in Table 5; results from the differenced series are reported in the Supplemental Material.

Table 5. Demographic Differences and Similarities Across the Full Study Period (2007–2020) for Implicit and Explicit Age, Race, and Sexuality Attitudes

Attitude and group	Implicit attitudes					Explicit attitudes				
	Start	End	% 8	Raw 8	Interpretation (2007–2020)	Start	End	% 8	Raw 8	Interpretation (2007–2020)
Age										
Liberals (L)	0.43	0.42	!4	!0.02	Parallel	0.58	0.57	!2	!0.01	Diverging
Conservatives (C)	0.46	0.42	!9	!0.04	(L → 1 C →)	0.50	0.31	!37	!0.18	(C ↘ 2 L ↘)
Race										
Liberals (L)	0.31	0.20	!37	!0.11	Diverging	0.25	!0.10	2 –100	!0.35	Diverging
Conservatives (C)	0.37	0.32	!14	!0.05	(L ↘ 2 C ↘)	0.52	0.22	!58	!0.30	(L ↘ 2 C ↘)
Sexuality										
Liberals (L)	0.21	!0.04	2 !100	!0.25	Diverging	0.25	!0.29	2 –100	!0.55	Converging
Conservatives (C)	0.52	0.34	!34	!0.18	(L ↘ 2 C ↘)	1.53	0.91	!41	!0.62	(C ↘ 2 L ↘)
Age										
Younger (Y)	0.45	0.41	!7	!0.03	Parallel	0.76	0.63	!16	!0.12	Converging
Older (O)	0.47	0.43	!8	!0.04	(Y → 1 O →)	0.03	!0.04	2 –100	!0.07	(Y ↘ 2 O ↘)
Race										
Younger (Y)	0.34	0.24	!29	!0.10	Diverging	0.34	!0.04	2 –100	!0.38	Diverging
Older (O)	0.32	0.27	!15	!0.05	(Y ↘ 2 O ↘)	0.31	0.07	!77	!0.24	(Y ↘ 2 O ↘)
Sexuality										
Younger (Y)	0.34	0.03	!90	!0.30	Diverging	0.76	!0.03	2 –100	!0.79	Diverging
Older (O)	0.35	0.24	!31	!0.11	(Y ↘ 2 O ↘)	0.66	0.28	!58	!0.39	(Y ↘ 2 O ↘)
Age										
White American (W)	0.45	0.41	!8	!0.03	Parallel	0.55	0.42	!23	!0.13	Converging
Black American (B)	0.51	0.46	!11	!0.06	(W → 1 B →)	0.37	0.39	2 4	2 0.02	(W ↘ & B →)
Race										
White American (W)	0.40	0.31	!25	!0.10	Converging	0.57	0.22	!61	!0.35	Converging
Black American (B)	!0.06	!0.04	!29	!0.02	(W ↘ & B →)	!1.10	!1.16	2 5	2 0.05	(W ↘ & B →)
Sexuality										
White American (W)	0.33	0.10	!70	!0.23	Parallel	0.67	0.12	!83	!0.55	Converging
Black American (B)	0.48	0.25	!48	!0.23	(W ↘ 1 B ↘)	1.03	0.42	!60	!0.61	(B ↘ 2 W ↘)
Age										
Female (F)	0.42	0.39	!7	!0.03	Parallel	0.43	0.35	!19	!0.08	Converging
Male (M)	0.51	0.47	!8	!0.04	(M → 1 F →)	0.72	0.56	!23	!0.17	(M ↘ 2 F ↘)
Race										
Female (F)	0.32	0.23	!26	!0.08	Parallel	0.25	!0.02	2	!0.27	Converging
					(M ↘ 1 F ↘)			!100		(M ↘ 2 F ↘)
Male (M)	0.35	0.27	!24	!0.08		0.42	0.06	!85	!0.36	
Sexuality										
Female (F)	0.32	0.09	!72	!0.23	Diverging	0.55	0.11	!80	!0.44	Converging
Male (M)	0.39	0.20	!50	!0.19	(F ↘ 2 M ↘)	1.01	0.37	!63	!0.63	(M ↘ 2 F ↘)
Age										
White American (W)	0.45	0.41	!8	!0.03	Parallel	0.55	0.42	!23	!0.13	Parallel
Asian American (A)	0.41	0.39	!6	!0.03	(W → 1 A →)	0.58	0.41	!30	!0.17	(W ↘ 1 A ↘)
Race										
White American (W)	0.40	0.31	!25	!0.10	Parallel	0.57	0.22	!61	!0.35	Parallel
Asian American (A)	0.35	0.27	!24	!0.08	(W ↘ 1 A ↘)	0.49	0.15	!70	!0.34	(W ↘ 1 A ↘)
Sexuality										
White American (W)	0.33	0.10	!70	!0.23	Parallel	0.67	0.12	!83	!0.55	Parallel
Asian American (A)	0.35	0.15	!58	!0.20	(W ↘ 1 A ↘)	0.80	0.36	!55	!0.44	(W ↘ 1 A ↘)
Age										
No college (NC)	0.38	0.43	2 14	2 0.05	Converging	0.40	0.38	!5	!0.02	Converging
College (C)	0.45	0.43	!5	!0.02	(NC ↗ & C →)	0.52	0.36	!31	!0.16	(NC → & C ↘)

(continued)

Table 5. (continued)

Attitude and group	Implicit attitudes					Explicit attitudes				
	Start	End	% 8	Raw 8	Interpretation (2007–2020)	Start	End	% 8	Raw 8	Interpretation (2007–2020)
Race										
No college (NC)	0.32	0.25	!21	!0.07	Parallel	0.29	!0.01	2! 100	!0.30	Parallel
College (C)	0.34	0.24	!28	!0.09	(NC ↘ 1 C ↘)	0.33	0.00	!99	!0.32	(NC ↘ 1 C ↘)
Sexuality										
No college (NC)	0.29	0.11	!60	!0.18	Parallel	0.63	0.06	!91	!0.57	Parallel
College (C)	0.34	0.13	!62	!0.21	(NC ↘ 1 C ↘)	0.71	0.13	!82	!0.58	(NC ↘ 1 C ↘)
Age										
Christian (Ch)	0.45	0.42	!8	!0.03	Parallel	0.49	0.40	!19	!0.10	Parallel
Nonreligious (Nr)	0.45	0.41	!8	!0.04	(Ch → 1 Nr →)	0.54	0.42	!23	!0.12	(Ch ↘ 1 Nr ↘)
Race										
Christian (Ch)	0.33	0.25	!25	!0.08	Parallel	0.30	0.00	!99	!0.30	Parallel
Nonreligious (Nr)	0.34	0.24	!28	!0.09	(Ch ↘ 1 Nr ↘)	0.38	0.02	!95	!0.36	(Ch ↘ 1 Nr ↘)
Sexuality										
Christian (Ch)	0.39	0.16	!60	!0.23	Parallel	0.88	0.29	!67	!0.59	Converging
Nonreligious (Nr)	0.29	0.05	!81	!0.23	(Ch ↘ 1 Nr ↘)	0.47	!0.03	2! 100	!0.50	(Ch ↘ 2 Nr ↘)
Sexuality										
Straight (S)	0.41	0.21	!48	!0.20	Parallel	0.93	0.45	!51	!0.47	Converging
Lesbian/gay (L/G)	!0.06	!0.23	2! 100	!0.17	(S ↘ 1 L/G ↘ 1 Bi ↘)	!0.62	!0.90	!45	!0.28	(S ↘ 2 L/G ↘)
Bisexual (Bi)	0.12	!0.09	2! 100	!0.21		!0.06	!0.40	2! 100	!0.34	(S ↘ 2 Bi ↘)

Note: Start and end values, as well as percentage of change (% **8**) and raw change (**8**), are calculated from the start and end points of the decomposed trend line (removing seasonality and noise). Using these decomposed trend values rather than raw monthly estimates eliminates results that may emerge from an outlier month that was unusually high or low. The interpretation column indicates whether the individual subgroup trends have moved in parallel or nonparallel (diverging or converging) directions on the basis of criteria outlined by Charlesworth and Banaji (2021); “=” indicates that the two groups moved at similar rates, “**2**” indicates that the first listed group showed a faster trend than the second listed group, and “&” indicates that the two groups moved in opposite directions. Arrows indicate the direction of the trend: downward to neutral (↘), no trend (→), or upward to neutral (↗).

Demographic differences in sexuality attitudes. Across all 14 years, implicit sexuality attitudes decreased in bias at similar rates (i.e., moved in parallel) regardless of respondents’ race (White, Black, and Asian Americans), education (college, no college), religion (Christian, Jewish, other, or nonreligious), or sexual orientation (straight, gay/lesbian, bisexual; Table 5). However, demographic differences persisted across respondent age and politics. Younger respondents decreased in implicit anti-gay/pro-straight bias by 90% over 2007 to 2020, whereas older respondents decreased at a slower rate of 31% (Fig. 2; Table 5). Similarly, liberal respondents in 2020 had reached implicit attitude neutrality (decreasing by more than 100%), whereas conservative respondents had decreased by a slower 34%; if past trends continue, conservatives would require at least another 16 years to reach the status of liberals’ now-neutral attitudes (Fig. 3; Table 5). In short, for both age and politics, the gaps between the implicit sexuality attitudes of young and old respondents and between liberal and conservative respondents have continued to grow. We also newly observed a difference between male and female respondents: Female

respondents decreased faster in implicit sexuality attitudes (72% drop) than male respondents (50% drop), a smaller difference than that across age and politics but nonetheless worth tracking in the future.

Relatively more demographic differences were observed for explicit sexuality attitudes. Specifically, although explicit sexuality attitudes decreased in parallel across respondent education, religion, and race (White and Asian Americans), demographic differences were present across age, politics, gender, race (White and Black Americans), and sexual orientation. As above, we found that younger respondents have decreased in explicit anti-gay/pro-straight bias faster than older respondents: Younger respondents have already reached neutrality (dropping by more than 100%), whereas older respondents dropped by a slower 60% (Fig. 2). Unlike implicit sexuality attitudes, however, the difference by politics indicated that the more biased conservative respondents had moved faster than the already neutral liberal respondents, meaning that the gap between the two groups (although still large) has been converging over time (Fig. 3). Nevertheless, it

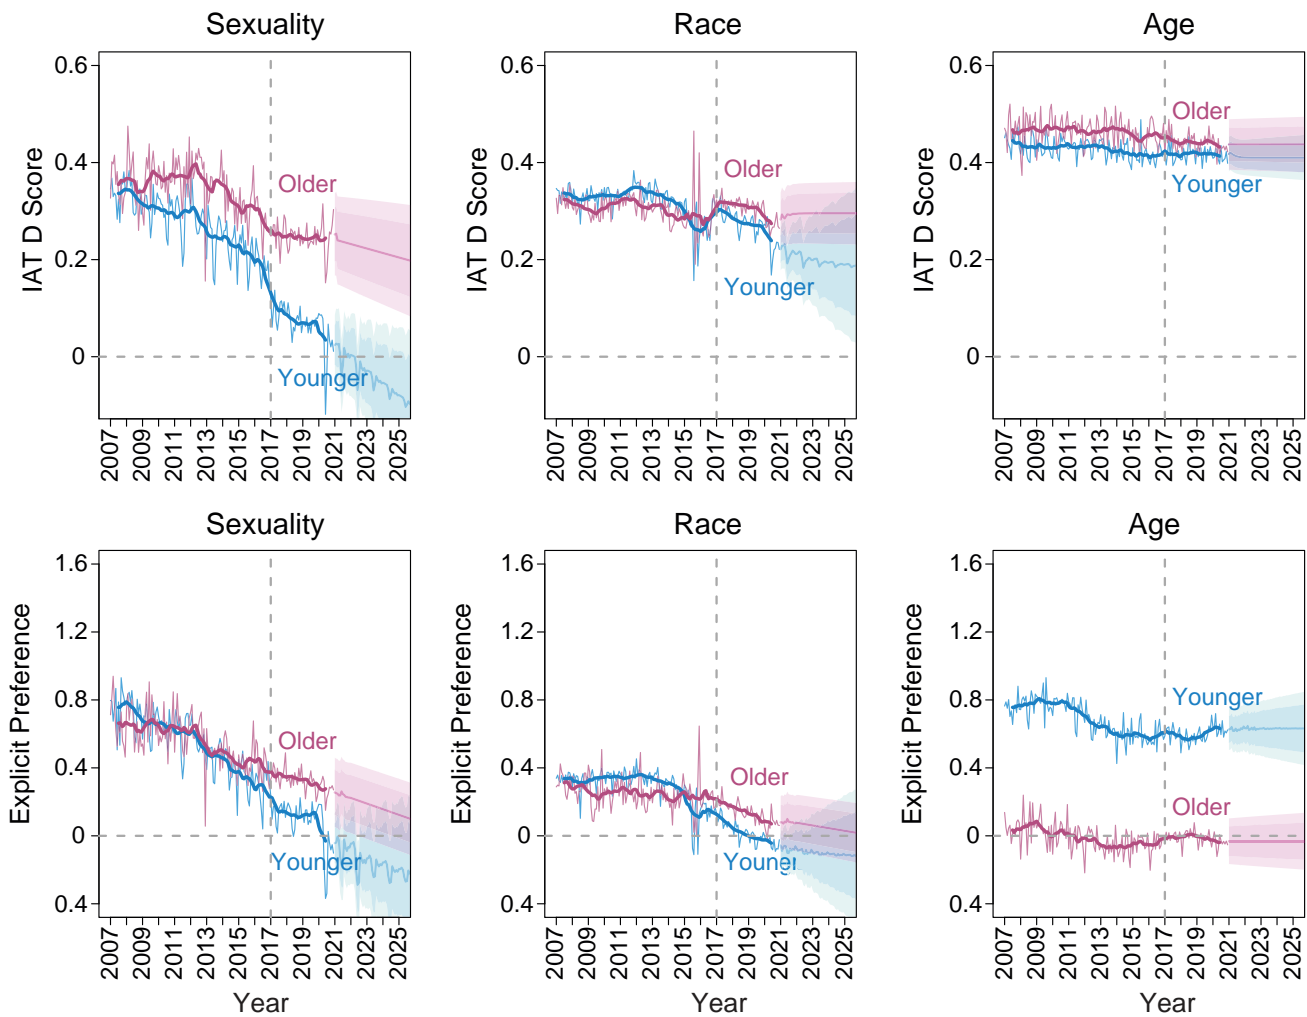


Fig. 2. Age differences on implicit (top row) and explicit (bottom row) measures of sexuality, race, and age attitudes from 2007 to 2025. Thin blue and red lines indicate observed monthly weighted means, and thick blue and red lines indicate decomposed trend lines of the observed monthly data (removing seasonality and noise), separately for younger adults (20 years) and older adults (40 years), respectively. Dark-shaded areas indicate 80% confidence intervals, and light-shaded areas indicate 95% confidence intervals of the autoregressive-integrated-moving-average model forecasts (for 2020–2025). The vertical gray line indicates the onset of the additional data (2017–2020) newly reported in this article. IAT 1 Implicit Association Test.

would take more than 200 years for conservative respondents to catch up to the low biases of liberal respondents if past trends continue (see the Supplemental Material). Finally, the differences by gender, race, and sexual orientation indicate that the (more biased) male respondents, Black American respondents, and straight respondents had moved faster than the (less biased) female respondents, White American respondents, and lesbian/gay or bisexual respondents, respectively, again resulting in converging gaps over time, because the more biased groups are catching up (Table 5).

Turning to the data from 2017 to 2020, demographic differences observed in the previous analysis across

politics, gender, and sexual orientation do not show such apparent differences in change. That is, from 2017 to 2020, comparisons of liberals and conservatives, men and women, White and Black Americans, and even straight and lesbian/gay/bisexual respondents revealed parallel rates of change for both implicit and explicit attitudes (see the Supplemental Material). The recent consistency across groups, although surprising, has also been documented on representative polls (e.g., Gallup recently found that Republicans’ approval of gay marriage increased by 9 percentage points, from 40% to 49% between 2016 and 2020, whereas Democrats similarly moved by 4 percentage points, from 79% to 83%; McCarthy, 2020).

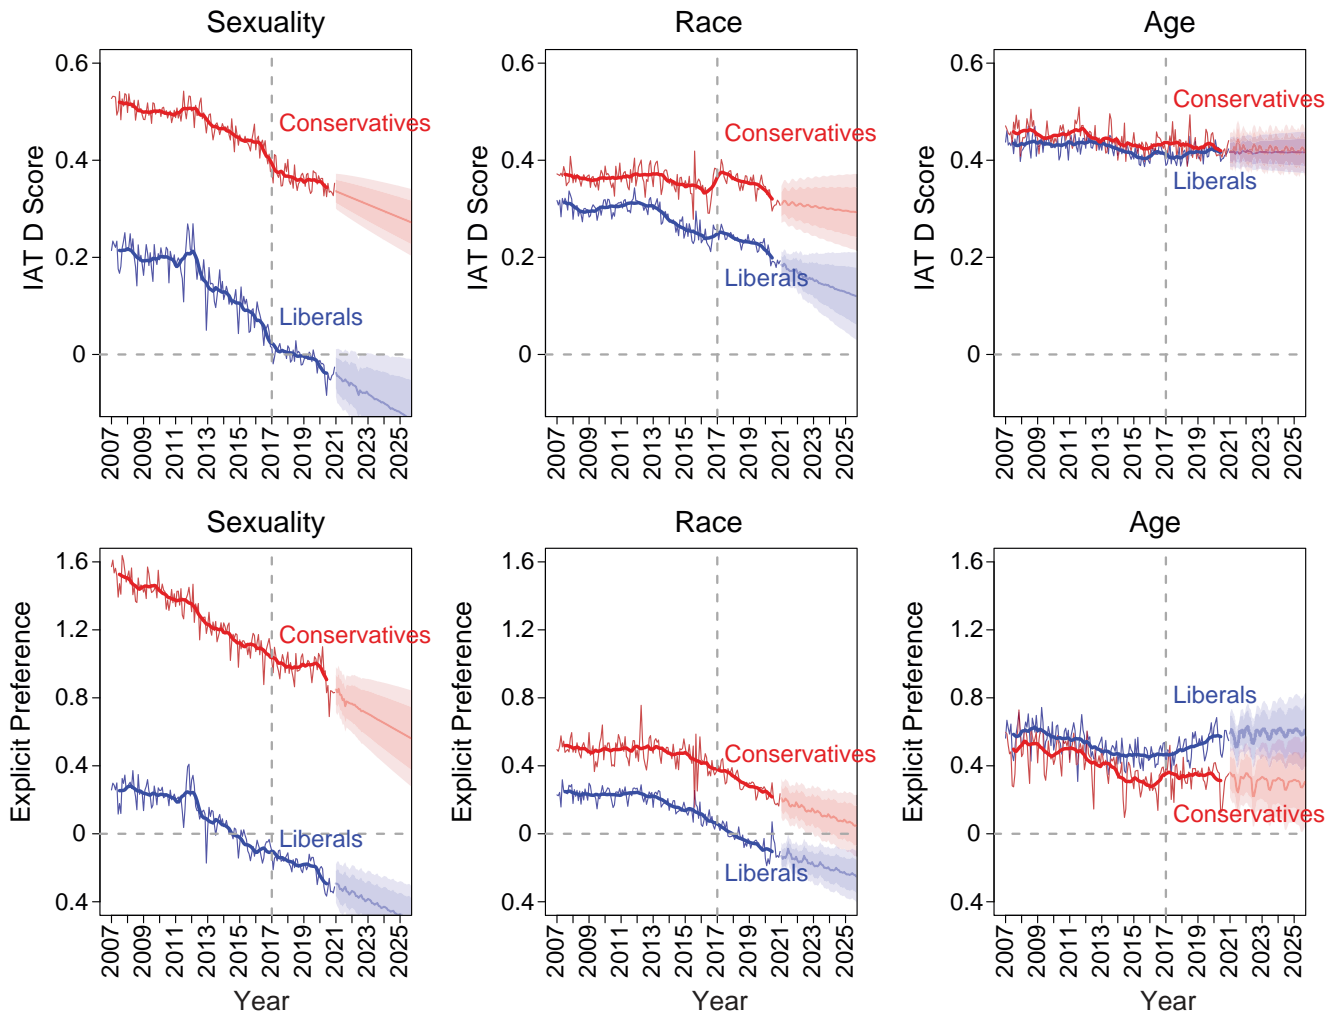


Fig. 3. Political differences on implicit (top row) and explicit (bottom row) measures of sexuality, race, and age attitudes from 2007 to 2025. Thin blue and red lines indicate observed monthly weighted means, and thick blue and red lines indicate decomposed trend lines of the observed monthly data (removing seasonality and noise), separately for liberals and conservatives, respectively. Dark-shaded areas indicate 80% confidence intervals, and light-shaded areas indicate 95% confidence intervals of the autoregressive-integrated-moving-average model forecasts (for 2020–2025). The vertical gray line indicates the onset of the additional data (2017–2020) newly reported in this article. IAT **1** Implicit Association Test.

Demographic differences in race attitudes. Across time, implicit race attitudes have decreased in anti-Black/pro-White bias in parallel regardless of respondents' gender, education, religion, and racial group (White versus Asian American identity; Table 5). However, as with sexuality attitudes, differences remained across respondent age and politics. Faster decreases in bias were observed among younger respondents (a 29% drop) than older respondents (a 15% drop; Fig. 2), as well as among liberal respondents (a 37% drop) than conservative respondents (a 14% drop; Fig. 3), both revealing growing gaps between these two groups' attitude trends. Again, although we emphasize that conservative respondents are also changing, the relatively slower rate of change means that, if past trends continue, conservatives will still

take at least another 31 years to reach implicit attitude neutrality, whereas liberals could reach that point in as little as 6 years. White, Asian, and Black Americans also continued to show differences in trends: Black Americans remained at relatively stable and weak pro-Black/anti-White implicit attitudes, whereas both White and Asian Americans decreased in their anti-Black/pro-White bias. Thus, the gap between Black Americans and White or Asian Americans has continued to decrease over time (Fig. 4).

Similar patterns of demographic similarities (and differences) were observed for explicit race attitudes: Decreases in anti-Black/pro-White bias progressed in parallel between 2007 and 2020 regardless of respondent education, religion, or White versus Asian American

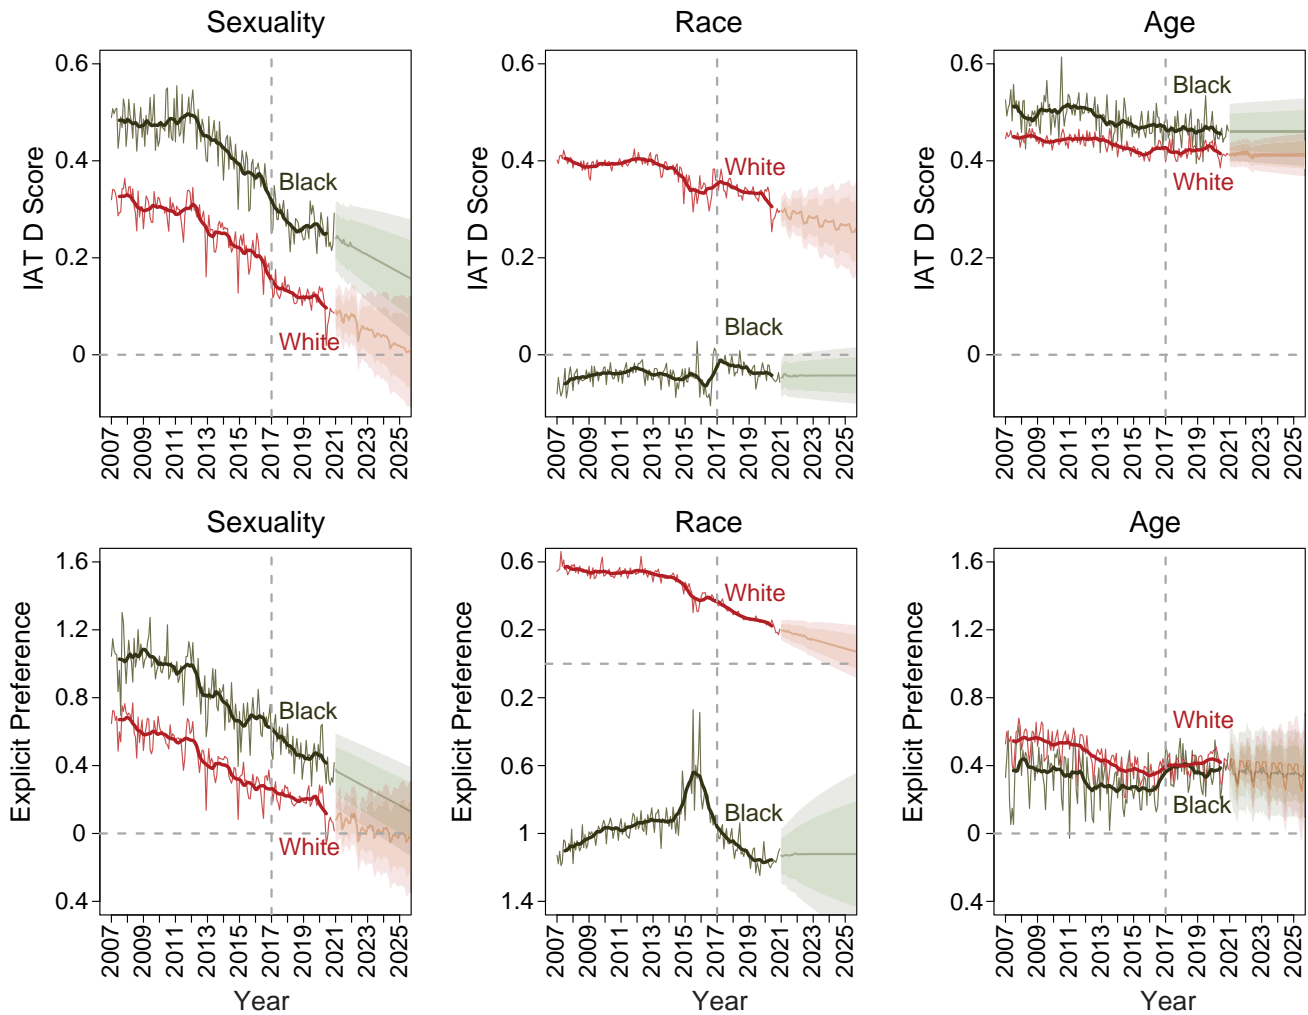


Fig. 4. Race differences on implicit (top row) and explicit (bottom row) measures of sexuality, race, and age attitudes from 2007 to 2025. Thin red and green lines indicate observed monthly weighted means, and thick red and green lines indicate decomposed trend lines of the observed monthly data (removing seasonality and noise), separately for White Americans and Black Americans, respectively. Dark-shaded areas indicate 80% confidence intervals, and light-shaded areas indicate 95% confidence intervals of the autoregressive-integrated-moving-average model forecasts (for 2020–2025). The vertical gray line indicates the onset of the additional data (2017–2020) newly reported in this article. Note that explicit race attitudes are visualized on a different y-axis scale than the other attitudes because of the uniquely lower explicit attitudes among Black Americans. IAT **1** Implicit Association Test.

identity. However, differences in explicit attitude trends persisted across respondent age, politics, White or Asian versus Black American identity, and (unlike implicit race attitudes) respondent gender. That is, as with implicit race attitudes, younger respondents decreased in bias faster (a drop of more than 100%; now passed explicit attitude neutrality) than older respondents (a drop of 77%; Fig. 2), and liberal respondents decreased in bias faster (a drop of more than 100%; now passed explicit attitude neutrality) than conservative respondents (a drop of 58%; Fig. 3). Next, for both gender and race differences, the more biased groups (i.e., men, White, Asian Americans) decreased in bias faster than the less biased groups (i.e., women, Black Americans), resulting in

converging gaps between these groups’ attitudes over time, perhaps because of the pressures on more biased groups to catch up to the new cultural norms of neutral explicit race attitudes.

For just 2017 to 2020 trends, we found that, as with sexuality attitudes, the differences in race attitude trends by gender, politics, age, and even race (White/Black) were largely eliminated. For instance, although Black Americans were moving to explicit attitude neutrality from below the zero line (i.e., decreasing in pro-Black/anti-White explicit attitudes) between 2007 and 2016, since 2017, Black Americans have switched directions and moved away from explicit attitude neutrality (started increasing in pro-Black/anti-White

explicit attitudes; Fig. 4). The recent widespread pull toward greater explicit pro-Black attitudes is therefore observed across every demographic subgroup, regardless of where the demographic subgroup's attitudes began (either above or below the zero line). Again, although there are long-term trends of differences across demographic groups by age, politics, and race, the recent data provide hints that the current change toward neutrality is occurring across the board. This result may seem surprising given that cultural divides in the United States on political issues, including race, feel particularly intransigent. However, the data point to the possibility that large-scale public engagement and widespread macrolevel events (e.g., BLM) may, in fact, have motivated change more broadly than assumed.

As a final note, we emphasize that the finding of explicit race attitudes crossing attitude neutrality by the end of 2020 was observed for nearly every demographic subgroup—the trajectories for female, younger, liberal, Christian, nonreligious, college-educated, and non-college-educated respondents all have ARIMA models that have already reached neutrality (Table 5). For the remaining subgroups (i.e., White, Asian, male, older, and conservative respondents), the forecasts include neutrality within the next few years (see the Supplemental Material). Even White respondents, who remain the furthest from neutrality, have ARIMA model forecasts in which the lower bound could reach neutrality as early as 2024 if past trends continue (see the Supplemental Material).

Demographic differences in age attitudes. From 2007 to 2020, implicit age attitudes were stable and parallel across nearly every demographic comparison, except for a new difference between college- and non-college-educated respondents. Specifically, over the 14 years, non-college-educated respondents (who are generally younger) increased in implicit anti-old/pro-young attitudes by 14%, whereas college-educated respondents remained stable (a drop of 5%; Table 5). Nevertheless, we emphasize that the dominant pattern for implicit age attitudes is one of similar, parallel stability across most demographic subgroups.

Many more demographic differences emerged for explicit age attitudes between 2007 and 2020: Here, demographic differences appeared across respondent politics, age, race (White and Black Americans), gender, and education. Indeed, the only parallel change for explicit age attitudes was across religious subgroups and White versus Asian Americans. First, in terms of political differences in explicit age attitudes, we found that liberals have followed a curvilinear pattern, decreasing until 2016 and then shifting to begin increasing in bias by approximately 22% between 2017 and

2020. Conservatives, in contrast, have consistently decreased in anti-old/pro-young explicit biases, dropping by 37% over the 14-year span (and decreasing by 11% between 2017 and 2020 alone). Next, in terms of age, race, gender, and education differences, we found that the gaps between these groups' attitudes have been converging over time because the previously more biased group (i.e., younger, White, male, non-college-educated) decreased in anti-old/pro-young explicit bias faster than the previously less biased group (i.e., older, Black, female, college educated) over all 14 years. Thus, as with explicit race and sexuality attitudes, these differences may be driven by the broader societal push to bring all demographic groups in line with social norms against explicit expressions of any social-group prejudice.

Finally, looking at 2017 to 2020 trends, we note that explicit anti-old/pro-young attitudes uniquely increased among younger and liberal respondents: All other groups moved toward lower bias (albeit at variable rates). These differences by age and politics rule out a cross-cutting source of explicit attitude change (e.g., anti-elderly associations of the COVID-19 pandemic) and, instead, align with earlier speculation about a generational war.

That is, if the data are revealing a conflict between younger people demanding liberalizing social change and older people who are perceived to be blocking such change with conservative values, we should observe the largest increases in anti-elderly/pro-young bias among young liberals. To explore this idea, we performed an exploratory, nonpreregistered analysis examining trends across four intersectional groups: young liberals, young conservatives, older liberals, and older conservatives (see the Supplemental Material). As expected, we found that the increase in explicit anti-elderly/pro-young bias from 2017 to 2020 was strongest among young liberals, who increased by 0.15 explicit attitude points (or about 20% in 4 years). In contrast, the largest decrease in explicit anti-elderly/pro-young bias was among older conservatives, who dropped by 0.13 explicit attitude points and ended in 2020 with neutral attitudes. Older liberals decreased by only 0.05 points, and young conservatives increased by 0.05 points. It therefore seems likely that the sources of explicit age-attitude increases were driven by the negative beliefs of young liberals about older (conservative) adults. This result is a notable departure from the pattern that we have seen across both implicit and explicit race and sexuality attitudes, where young respondents and liberals are the groups decreasing most rapidly in bias. Thus, young liberals are not ubiquitously the least biased and fastest progressing demographic group but, rather, are changing in some areas of intergroup

relations (race, sexuality) perhaps at the cost of others (anti-elderly bias).

Exploratory analyses of temporary increases in implicit attitudes

All earlier analyses of overall attitude trends and demographic differences in trends were preregistered. However, after inspecting the trends across all 14 years of data, we were unexpectedly confronted with a striking pattern of data: Four of the six implicit attitudes (race, skin tone, disability, and body weight) revealed visible, temporary increases around 2015 to 2017 that lasted approximately 1 year before returning to previous trends. What might help explain these short-term upticks in bias?

First, we ruled out deflationary explanations. We note that these increases are unlikely to be mere artifacts of sample change (as we controlled for such changes through weighting) or changes to the sample source (as we found no anomalies in the frequencies of participants coming from different sources). Moreover, the attitude increases are not likely to be errors in data archiving or other technical issues because they occurred at slightly different moments for each attitude and only in implicit attitudes, whereas archiving errors would be expected to affect all attitudes at a similar moment in time.

Finally, of the six attitude topics, the temporary increases were observed only when the group in question was targeted by particularly negative rhetoric from the new sociopolitical changes in 2015 to 2017 (i.e., race, skin tone, weight, and disability). For instance, a content analysis of 223 Trump tweets mentioning marginalized groups showed that 68% of those tweets discussed race/ethnicity and most often in a negative tone (Coe & Griffin, 2020). Additionally, Trump's highly publicized mocking of Serge Kovalski, a *New York Times* reporter with a disability, as well as his derision of celebrities and politicians as "fat" (including Rosie O'Donnell, Alicia Machado, and Kim Jong-un), created targeted negative media around the topics of disability and body weight. In contrast, Trump's negative comments were relatively rare toward the elderly and gay/lesbian people (Coe & Griffin, 2020). Although there were certainly harmful and biased policies targeted toward transgender people, the words and actions of Trump's campaign and early presidency appeared to recognize the massive shift that has occurred in public opinion about (cisgender) gay and lesbian people, even among conservatives (Pew Research Center, 2019). That the temporary attitude increases were observed most clearly in those Trump-targeted attitudes, but not in others, further suggests that the results are more likely

to be due to real social effects rather than to artifacts in the data.

Ultimately, the systematicity of such temporary increases is too remarkable not to offer both a detailed description of the empirical phenomenon as well as an initial exploratory case-study explanation of potential sources. To that end, we performed two post hoc non-preregistered exploratory analyses: (a) segmented regressions to describe the phenomenon and (b) a form of difference-in-differences regression to explore whether the increase is concentrated among conservative respondents and Republican states, an outcome that would be in line with the idea that the increase may arise, in part, from prejudice-emboldening actions and words surrounding Trump's campaign and early presidency.

Segmented regressions. We fitted segmented regressions to all implicit attitude trends to identify the timing, magnitude, and duration of the short-term increase in 2015 to 2017. Segmented regressions were fitted to the decomposed trend lines from 2007 to 2020 using the *segmented* package (Version 1.3; Muggeo, 2022) in the R programming environment, allowing the model to discover up to three break points empirically. In all cases, the three-break-point model provided very good fit to the trends (R^2 range 1 .88–.98), and adding further break points risked overfitting. In most cases, the model estimated break points that aligned with the true local minimum and maximum IAT score during the investigated period (e.g., if the true minimum was in April 2016, the model also estimated that the break point started at that month). However, in the one case of skin-tone attitudes, there was a discrepancy between the estimated break point (in July 2016) and the true empirical minimum (in March 2016); to maintain consistency, we report the model-estimated results below.

Results for the segmented regressions are summarized in Table 6 and Figure 5. Notably, neither implicit sexuality attitudes nor implicit age attitudes indicated any increases during the period of 2015 to 2017. Indeed, if anything, the segmented regression for implicit sexuality attitudes showed an inflection point in May 2016 that indicated more rapid decreases in bias until approximately April 2017, after which the trend slowed again to its previous rate of decrease. Additionally, implicit age attitudes had no significant short-term increases in 2015 to 2017. The identified break point in April 2015 indicated only a very slow long-term trend of weakly increasing attitudes that continued until January 2020.

In contrast, for implicit race and skin-tone attitudes, the segmented regressions estimated the temporary

Table 6. Segmented Regressions of Temporary Increases in Implicit Social-Group Attitudes (2014–2020)

Attitude	Estimated start		Estimated end		Slope of change		Raw change
	<i>M</i>	85% CI	<i>M</i>	85% CI	<i>b</i>	85% CI	
Sexuality	May 2016	[March 2016, July 2016]	April 2017	[January 2017, June 2017]	! 0.006	[! 0.007, ! 0.004]	! 0.06
Race	April 2016	[March 2016, June 2016]	April 2017	[February 2017, May 2017]	0.002	[0.002, 0.003]	Ž0.03
Skin tone	July 2016	[May 2016, August 2016]	February 2017	[December 2016, March 2017]	0.003	[0.001, 0.004]	Ž0.04
Age	April 2015	[February 2015, June 2015]	January 2020	[October 2019, March 2020]	0.00006	[0.00001, 0.0001]	Ž0.003
Disability	February 2016	[January 2016, March 2016]	March 2017	[February 2017, March 2017]	0.005	[0.004, 0.005]	Ž0.05
Body weight	August 2015	[July 2015, October 2015]	November 2016	[July 2016, October 2016]	0.003	[0.002, 0.004]	Ž0.03

Note: Estimated start and end mean month (as well as associated confidence intervals [CIs]) are estimated from segmented regression models with three empirical break points. Slope of change is the slope between the estimated start and end month. Raw change is the difference in Implicit Association Test D-score points between the estimated start and end month in the segmented regression.

increases to occur from approximately spring 2016 to spring 2017 (Table 6). Over this year, bias in race and skin-tone attitudes increased by approximately 0.03 and 0.04 IAT D-score points, respectively, or a change of approximately 10% in just 1 year. Indeed, the increase effectively reverted the attitude trends back to where they had been in mid-to-late 2014. For implicit disability

attitudes, the segmented regressions placed the temporary increase slightly earlier, spanning from approximately January 2016 to February 2017 and revealing an increase of Ž0.05 IAT D-score points in just 1 year (approximately an 11% increase). By the end of this increase, implicit disability attitudes were at their highest levels ever observed over all 14 years of data. Finally,

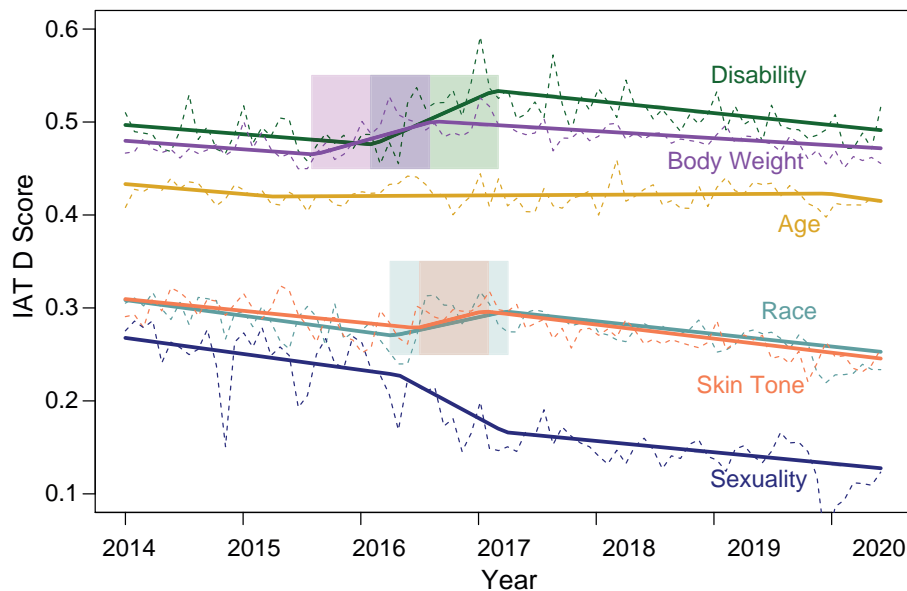


Fig. 5. Timing of temporary increases in six implicit attitude trends from 2014 to 2020. Thin dashed lines indicate raw monthly means, and thick solid lines indicate fitted segmented regressions. Shaded areas indicate the duration of the temporary increases identified for implicit disability, body-weight, race, and skin-tone attitudes (implicit age and sexuality attitudes had no short-term increases during the investigated period).

implicit body-weight attitudes revealed the earliest increase: Segmented regressions estimated that the uptick started in approximately August 2015 and ended in November 2016, ultimately increasing by $\bar{Z}0.03$ IAT D-score points and culminating in the highest levels of implicit body-weight bias observed over all 14 years.

Differences in temporary increases of implicit race attitudes by respondent politics and geography.

We next turn to an exploratory analysis that hints at a possible explanation for temporary increases in implicit race attitudes. First, we note that the temporary increase began in spring 2016 and, thus, any possible explanation would rest on a social change that also began around that time. Among the many changes occurring at this moment, one of the most notable, especially with respect to race attitudes, was the campaign and rise in popularity of Donald Trump. In particular, spring 2016 saw Trump secure the Republican presidential nomination in the primaries. This confirmation served as the first official endorsement of his views on race relations—views that had been accumulating through comments such as his stereotypes of Mexican immigrants (in June 2015) and his silence on condemning the Ku Klux Klan (in February 2016; see also Coe & Griffin, 2020). Political echo chambers being what they are, these opinions are likely to be concentrated and influential among conservative respondents and those living in strongly Republican states who were more exposed to Trump's messages through campaign rallies, Trump retweets, or Trump-favoring news media.

Here, we explore the potential role of Trump-related words and actions concentrated toward conservative and Republican respondents using a form of difference-in-differences regression, one of the strongest quasiexperimental approaches to understanding sources of change (e.g., Abadie & Cattaneo, 2018). Essentially, difference-in-differences analyses test whether differences in the degree of exposure to a certain event (e.g., difference between conservatives and liberals in the degree of exposure to Trump-related messages) correspond to differences in the degree of change in an outcome variable (e.g., conservatives' versus liberals' short-term attitude increases). In our case, the model approach was simple: Using the time-series trends computed above, we estimated a linear regression with an interaction between time (period of the short-term increase) and group (liberal vs. conservative respondents, or respondents living in Republican vs. Democrat states). If the interaction is significant, it tells us that the increase is concentrated among conservatives (or Republican states) and provides evidence in line with the notion that the increases could be at least partly attributed to social changes surrounding Trump's

campaign, candidacy, and early presidency, which had greater exposure among those groups.

We provide a case study using the race-attitude data. This attitude was chosen because (a) it is the largest sample, giving us precise estimates even after separating by demographics and geography; (b) of all the attitudes showing temporary increases, it is the one that has been most frequently and clearly targeted by Trump (Coe & Griffin, 2020); and (c) of all the attitudes showing temporary increases, it is the attitude that differs the most across political lines (i.e., liberals and conservatives have different magnitudes of implicit race attitudes). This political difference makes it more likely that we will be able to tease out the possible role of Trump-related words and actions in affecting one political group more than another (e.g., conservatives more than liberals).

In line with the role of Trump in the temporary increases, results from the difference-in-differences model showed a significant interaction between time and respondent politics, $b \mathbf{1} \neq 0.0028$, $SE \mathbf{1} 0.0002$, $p \mathbf{0} .001$ (Fig. 6; full results are reported in the Supplemental Material). Inspecting the fitted values is also illuminating: Although conservative respondents' attitudes increased by 0.04 IAT D-score points from April 2016 to April 2017, liberal respondents' attitudes increased by only 0.01 IAT D-score points over the same time span. Results were nearly identical for the complementary model performed at the state level comparing trends from the top-20 Republican states with the top-20 Democrat states on the basis of Trump's margin of votes. The state-level model showed a significant interaction, $b \mathbf{1} \neq 0.0021$, $SE \mathbf{1} 0.0002$, $p \mathbf{0} .001$, and again, the short-term increase was approximately 4 times larger in Republican states than in Democrat states (see the Supplemental Material).

Summary of key results

From among the complexity of data presented in the current article, we summarize four key conclusions. First, with the inclusion of the most recent data from 2017 to 2020, trends across all 14 years (2007–2020) have magnified the evidence of change in explicit and implicit attitudes. It is now incontrovertible that the societal attitudes of U.S. respondents have transformed in meaningful ways toward greater neutrality. Since 2007, explicit race, skin-tone, and sexuality attitudes have dropped by 98%, 79%, and 75%, respectively, and even the relatively slower-changing explicit attitudes toward disability, body weight, and age have dropped by 37%, 31%, and 22%, respectively. For implicit attitudes, race, skin-tone, and sexuality biases still have relatively far to go toward neutral but have nevertheless

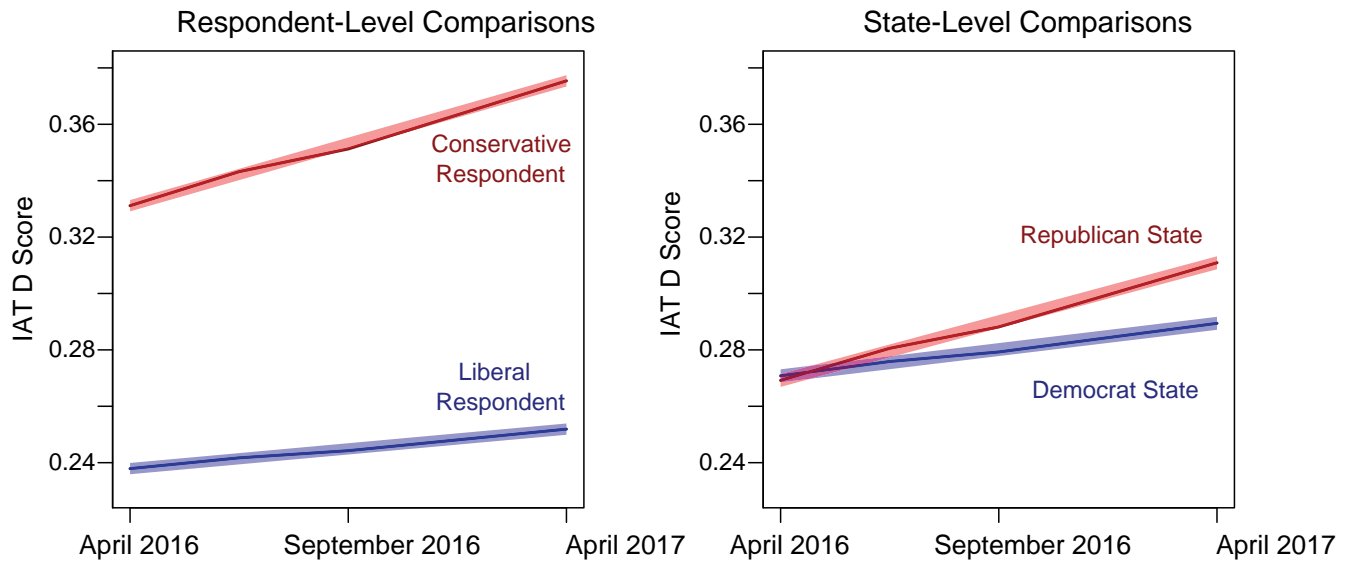


Fig. 6. Mean implicit race attitudes from April 2016 to April 2017 by respondent-level politics (left) and state-level politics (right). Solid lines indicate fitted values from interaction regression models; shaded areas indicate 95% confidence intervals. State-level comparisons used the top 20 Republican and top 20 Democrat states on the basis of Trump's margin of victory in the 2016 election. IAT 1 Implicit Association Test.

dropped by a noticeable 26%, 25%, and 65%, respectively. On the other hand, implicit age, disability, and body-weight attitudes, which had shown stability in a previous analysis, have continued to show little to no change (8%, 2%, and 1%, respectively, across 14 years).

Second, whether implicit attitudes were changing or remain stable, the observed trends were parallel across nearly all demographic groups, albeit with a few notable exceptions. For instance, the drops in implicit sexuality and race bias were observed across all demographic groups, although two demographic groups—those who are younger and those who self-identified as liberal—decreased in bias at a relatively faster rate than older or conservative respondents. Thus, implicit attitude change may be best understood as a consequence of macrolevel societal transformations that cut across many demographics in similar ways (e.g., sources of change such as social movements, ecological threats, or media; for a discussion, see Charlesworth & Banaji, 2021). In contrast, relatively more demographic differences emerged in the trends of change for explicit group attitudes, suggesting that these attitudes may be more tied to group-specific motivations (e.g., social-dominance orientation).

Third, both the overall trends and the demographic patterns in data from 2017 to 2020 were generally well predicted by ARIMA models fitted to past trends (2007–2016), with reasonable accuracy statistics. Indeed, if anything, the forecasts from past trends generally underestimated the true observed change from 2017 to 2020. The accurate predictions reveal both the

methodological strength of an ARIMA approach and, more substantively, the perhaps surprising long-term persistence in societal implicit and explicit attitude trends across a particularly tumultuous 4-year period.

Fourth and finally, despite the long-term persistence, some implicit (but no explicit) attitudes nevertheless showed temporary disruptions during the period of 2015 to 2017: Implicit race, skin-tone, disability, and body-weight attitudes each revealed year-long increases in bias that eventually returned to previous rates of change or stability after about a year. Post hoc explorations of these increases in implicit race attitudes showed that the increases from spring 2016 to spring 2017 were concentrated among conservative respondents and respondents in Republican states, in line with the idea that the increases may be due, in part, to the words and actions surrounding Trump's campaign and early presidency.

General Discussion

Using the largest record of continuously collected tests of implicit and explicit attitudes across 14 years (2007–2020), we found that, despite much tumultuous socio-political change from 2017 to 2020, attitude trends have generally persisted in long-term patterns from the past, whether decreasing or remaining stable in bias. Attitudes previously decreasing in bias—implicit sexuality, race, and skin-tone attitudes, and all explicit attitudes—continued along the path of decreasing bias since 2017, whereas attitudes that were previously stable—implicit age, disability, and body-weight attitudes—have also

continued to remain so. There are, of course, limits to the generalizability of such findings because of the nature of the data (i.e., a U.S. online convenience sample; see the Limitations section). Nevertheless, it now seems incontrovertible that, at least under some conditions, societal explicit and implicit attitudes *can* change over the long term in clearly profound ways: Across 14 years, the drops in implicit sexuality, race, and skin-tone bias have magnified to a decrease of 65%, 26%, and 25%, respectively, whereas the drops in explicit sexuality, race, and skin-tone bias now stand at 75%, 98%, and 79%, respectively, all since 2007.

Such clear evidence of change in implicit attitudes contrasts sharply with the stability of age, disability, and body-weight attitudes, reminding us that change, although possible, is not inevitable. Differing patterns also give greater credence to the change data, as they rule out the possibility that an unknown artifact of the method or data source may be detecting arbitrary change; rather, both change and stability seem robust. Renewed efforts for interventions are necessary to promote decreases in the particularly intransigent biases of age, disability, and body weight. Along these lines, the current data—which reveal widespread parallel change across most demographic groups—point to the interpretation that the most successful efforts for attitude change are likely to be macrolevel societal events that cut across demographic groups in similar ways (Charlesworth & Banaji, 2021). Efforts such as federal legislation and social movements have already shown an impact on implicit attitudes (Ofosu et al., 2019; Sawyer & Gampa, 2018); the current data newly reinforce the role of significant events and media (e.g., words and actions surrounding a president's campaign) as tools to, at least temporarily, spur population-level attitude change, although not always in the direction of decreasing bias.

Short-term increases in implicit attitudes

Exploratory post hoc segmented regressions identified that implicit race, skin-tone, disability, and body-weight attitudes each experienced temporary increases in bias at slightly different moments between 2015 and 2017. In contrast, implicit age and sexuality attitudes showed no increases, and neither did any explicit attitudes. That the increases were observed only in implicit attitudes adds new empirical evidence to the distinction between implicit and explicit attitude measures. We interpret this result in line with theorizing that explicit attitudes are more interwoven with concerns around self-consistency and stable values, whereas implicit attitudes are more attuned to the associations currently accessible in one's environment (Payne et al., 2017). Thus, when significant

events change the conversations and associations in the environment, implicit attitudes are more likely than explicit attitudes to update in response, even if only temporarily (Payne & Hannay, 2021).

It is equally notable that the increases in implicit attitudes were only temporary: Attitudes returned to their previous rates of change or to stable bias after approximately 1 year. This consistent duration of the upticks is an intriguing first hint at the impact and decay of social events on implicit attitudes (i.e., the “half-life” of events on attitudes; Page & Shapiro, 1992) and presents an exciting empirical discovery for future work to unpack.

More generally, the result of both temporary responsiveness and long-term persistence suggest new ways of thinking about principles of long-term change in implicit attitudes: At baseline, societal implicit attitudes appear to have homeostasis—they continue in slow but steady trends of change or stability that reflect the generally consistent social environment. However, like any adaptive system, implicit attitudes can also respond to shocks in the environment (e.g., a presidential campaign that emboldened prejudice). Following this shock, the attitudes may either adopt a durable new state if enough people are convinced (i.e., a tipping point) and/or if the event widely reshapes daily norms (e.g., the attitude inflection point following same-sex marriage legislation; Ofosu et al., 2019). Alternatively, if the new state is adopted only by a minority, or if the event is more transient (e.g., a presidential campaign), implicit attitudes may eventually return to their previous homeostasis as the shock fades from public memory.

Finally, although exploratory and speculative, initial post hoc analyses hint at possible sources for such temporary increases. We note that (a) the increases occurred only in those attitudes toward groups that were most strongly targeted by words and actions in Trump's campaign and early presidency, (b) the increases (at least in implicit race attitudes) appeared concentrated among conservative respondents and Republican states, and (c) the increases generally coincided with significant events related to Trump (e.g., Republican primaries). Each of these points provides early, speculative evidence in line with the idea that Trump-related words and actions may have, in part, contributed to temporary increases in implicit attitudes.

Limitations

The Project Implicit data set is a convenience sample from the United States and thus has inherent limitations regarding generalizability beyond the given population (e.g., non-U.S. settings, as well as the representative U.S. population). After all, as a convenience sample,

the data set can suffer from selection bias (Bethlehem, 2010). In particular, extreme ideological groups, such as the alt-right or White nationalists, are unlikely to be included in the current data. Although these groups represent only a fraction of the U.S. population, their views reflect explicit hostility, support for group-based dominance, and even violent discrimination (Forscher & Kteily, 2020). Future studies that proactively survey these demographic groups and include them in final population estimates of attitudes may find higher biases (and slower change) than the current data suggest. Although the database may be somewhat more reflective of the country than expected (because Project Implicit also receives volunteers assigned to participate for educational purposes), it is not designed to capture a fully representative sample of the population.

A second concern is the cross-sectional nature of the data. Cross-sectional data leaves open the possibility that the sample composition has changed over time. Although we attempt to control for this possibility using weighting and raking, it is still concerning that unobserved features of the sample may change (e.g., who identifies as liberal may be changing). Moreover, cross-sectional archival data means that causal conclusions of why change occurred is difficult to identify because of the co-occurrence of multiple significant events at the national level. For this reason, although we provide a case-study exploration of implicit race attitudes and suggest a possible source of change may be Trump-related words and actions concentrated among his base, we reiterate that these are only speculative initial explorations that were performed after seeing a striking pattern of data. The post hoc nature of such analyses means that we should not be too confident in the conclusions until further convergent evidence from other sources is provided (e.g., from public-opinion polls over the same time period). Future research using similar difference-in-differences methods but with preregistered hypotheses and systematic event analyses will be better equipped to elucidate potential causal sources of both long-term attitude trends and short-term responsiveness.

Transparency

Action Editor: Paul Jose

Editor: Patricia J. Bauer

Author Contributions

T. E. S. Charlesworth and M. R. Banaji developed the study concept. T. E. S. Charlesworth analyzed the data. Both of the authors interpreted the data, drafted the manuscript, and approved the final manuscript for submission.

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding

This research was supported by the Hao Family Inequality in America Support Grant and by grants from the Foundations of Human Behavior Initiative and the Mind Brain Behavior Interfaculty Initiative, all awarded to M. R. Banaji and T. E. S. Charlesworth.

Open Practices

Cleaned data and analysis scripts have been made publicly available via OSF and can be accessed at <https://osf.io/qywh4/>. Raw data are from the Project Implicit Demonstration Website database, archived at <https://osf.io/y9hiq/>. The design and analysis plans for the present study were preregistered at <https://aspredicted.org/mh958.pdf>. Deviations from the preregistered analysis plan are noted in the main text and/or the Supplemental Material available online. This article has received the badges for Open Data, Open Materials, and Preregistration. More information about the Open Practices badges can be found at <http://www.psychologicalscience.org/publications/badges>.



ORCID iD

Tessa E. S. Charlesworth  <https://orcid.org/0000-0001-5048-3088>

Supplemental Material

Additional supporting information can be found at <http://journals.sagepub.com/doi/suppl/10.1177/09567976221084257>

References

- Abadie, A., & Cattaneo, M. D. (2018). Econometric methods for program evaluation. *Annual Review of Economics*, *10*, 465–503. <https://doi.org/10.1146/annurev-economics-080217-053402>
- Banaji, M. R., & Greenwald, A. G. (2013). *Blindspot: Hidden biases of good people*. Delacorte Press.
- Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 361–382). Guilford Press.
- Bethlehem, J. (2010). Selection bias in web surveys. *International Statistical Review*, *78*(2), 161–188. <https://doi.org/10.1111/j.1751-5823.2010.00112.x>
- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, *6*(3), 242–261. https://doi.org/10.1207/S15327957PSPR0603_8
- Charlesworth, T. E. S., & Banaji, M. R. (2019). Patterns of implicit and explicit attitudes: I. Long-term change and stability from 2007 to 2016. *Psychological Science*, *30*(2), 174–192. <https://doi.org/10.1177/0956797618813087>
- Charlesworth, T. E. S., & Banaji, M. R. (2021). Patterns of implicit and explicit attitudes II. Long-term change and stability, regardless of group membership. *American Psychologist*, *76*(6), 851–869. <https://doi.org/10.1037/amp0000810>
- Coe, K., & Griffin, R. A. (2020). Marginalized identity invocation online: The case of President Donald Trump on

- Twitter. *Social Media and Society*, 6(1). <https://doi.org/10.1177/2056305120913979>
- Cohn, N., & Quealy, K. (2020, June 10). How public opinion has moved on Black Lives Matter. *The New York Times*. <https://www.nytimes.com/interactive/2020/06/10/upshot/black-lives-matter-attitudes.html>
- Cryer, J. D., & Chan, K.-S. (2008). *Time series analysis: With applications in R* (2nd ed.). Springer. <https://doi.org/10.1007/978-0-387-75959-3>
- Dimock, M., & Gramlich, J. (2021, January 29). *How America changed during Trump's Presidency*. Pew Research Center. <https://www.pewresearch.org/2021/01/29/how-america-changed-during-donald-trumps-presidency/>
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*, 54(1), 297–327. <https://doi.org/10.1146/annurev.psych.54.101601.145225>
- Forscher, P. S., & Kteily, N. S. (2020). A psychological profile of the alt-right. *Perspectives on Psychological Science*, 15(1), 90–116. <https://doi.org/10.1177/1745691619868208>
- Fraser, S., Lagacé, M., Bongué, B., Ndeye, N., Guyot, J., Bechard, L., Garcia, L., Taler, V., Adam, S., Beaulieu, M., Bergeron, C. D., Boudjemadi, V., Desmette, D., Donizzetti, A. R., Éthier, S., Garon, S., Gillis, M., Levasseur, M., Lortie-Lussier, M., . . . Wittich, W. (2020). Ageism and COVID-19: What does our society's response say about us? *Age and Ageing*, 49(5), 692–695. <https://doi.org/10.1093/ageing/afaa097>
- Gawronski, B. (2019). Six lessons for a cogent science of implicit bias and its criticism. *Perspectives on Psychological Science*, 14(4), 574–595. <https://doi.org/10.1177/1745691619826015>
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132(5), 692–731. <https://doi.org/10.1037/0033-2909.132.5.692>
- Hopkins, D. J., & Washington, S. (2020). The rise of Trump, the fall of prejudice? Tracking White Americans' racial attitudes via a panel survey, 2008–2018. *Public Opinion Quarterly*, 84(1), 119–140. <https://doi.org/10.1093/poq/nfaa004>
- Hyndman, R. J., & Athanasopoulos, G. (2013). *Forecasting: Principles and practice*. OTexts. <https://www.otexts.org/fpp>
- Hyndman, R. J., & Khandakar, Y. (2008). Automatic time series forecasting: The forecast Package for R. *Journal of Statistical Software*, 27(3). <https://doi.org/10.18637/jss.v027.i03>
- Jost, J. T. (2019). The IAT is dead, long live the IAT: Context-sensitive measures of implicit attitudes are indispensable to social and political psychology. *Current Directions in Psychological Science*, 28(1), 10–19. <https://doi.org/10.1177/0963721418797309>
- Marsden, P. V., Smith, T. W., & Hout, M. (2020). Tracking US social change over a half-century: The General Social Survey at fifty. *Annual Review of Sociology*, 46, 109–134. <https://doi.org/10.1146/annurev-soc-121919-054838>
- McCarthy, J. (2020, June 1). *U.S. support for same-sex marriage matches record high*. Gallup News. <https://news.gallup.com/poll/311672/support-sex-marriage-matches-record-high.aspx>
- Mueller, J. C., & McCollum, J. (2022). A sociological analysis of “OK Boomer.” *Critical Sociology*, 48(2), 265–281. <https://doi.org/10.1177/08969205211025724>
- Muggeo, V. M. R. (2022). *Package ‘segmented’* [Computer software]. Comprehensive R Archive Network. <https://cran.r-project.org/web/packages/segmented/segmented.pdf>
- Newman, B., Merolla, J. L., Shah, S., Casarez Lemi, D., Collingwood, L., & Karthick Ramakrishnan, S. (2020). The Trump effect: An experimental investigation of the emboldening effect of racially inflammatory elite communication. *British Journal of Political Science*, 51(3), 1138–1159. <https://doi.org/10.1017/S0007123419000590>
- Ofori, E. K., Chambers, M. K., Chen, J. M., & Hehman, E. (2019). Same-sex marriage legalization associated with reduced implicit and explicit antigay bias. *Proceedings of the National Academy of Sciences, USA*, 116(18), 8846–8851. <https://doi.org/10.1073/pnas.1806000116>
- Page, B., & Shapiro, R. Y. (1992). *The rational public: Fifty years of trends in Americans' policy preferences*. University of Chicago Press.
- Pasek, J. (2018). *Package ‘anesrake’* [Computer software]. Comprehensive R Archive Network. <https://cran.r-project.org/web/packages/anesrake/anesrake.pdf>
- Payne, B. K., & Hannay, J. W. (2021). Implicit bias reflects systemic racism. *Trends in Cognitive Sciences*, 25(11), 927–936. <https://doi.org/10.1016/j.tics.2021.08.001>
- Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2017). The bias of crowds: How implicit bias bridges personal and systemic prejudice. *Psychological Inquiry*, 28(4), 233–248. <https://doi.org/10.1080/1047840X.2017.1335568>
- Pew Research Center. (2017, October 5). *The partisan divide on political values grows even wider*. <https://www.people-press.org/2017/10/05/the-partisan-divide-on-political-values-grows-even-wider/>
- Pew Research Center. (2019, May 14). *Attitudes on same-sex marriage*. <https://www.pewforum.org/fact-sheet/changing-attitudes-on-gay-marriage/>
- R Core Team. (2020). *R: A language and environment for statistical computing* (Version 4.0.3) [Computer software]. <http://www.R-project.org>
- Sawyer, J., & Gampa, A. (2018). Implicit and explicit racial attitudes changed during Black Lives Matter. *Personality and Social Psychology Bulletin*, 44(7), 1039–1059.
- Schmidt, K., & Axt, J. R. (2016). Implicit and explicit attitudes toward African Americans and Barack Obama did not substantively change during Obama's presidency. *Social Cognition*, 34(6), 559–588. <https://doi.org/10.1521/soco.2016.34.6.559>
- Schmidt, K., & Nosek, B. A. (2010). Implicit (and explicit) racial attitudes barely changed during Barack Obama's presidential campaign and early presidency. *Journal of Experimental Social Psychology*, 46(2), 308–314. <https://doi.org/10.1016/j.jesp.2009.12.003>
- Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition*, 25(5), 638–656. <https://doi.org/10.1521/soco.2007.25.5.638>
- Smith, T. W. (1994). Is there real opinion change? *International Journal of Public Opinion Research*, 6(2), 187–203. <https://doi.org/10.1093/ijpor/6.2.187>
- Stimson, J. A. (2015). *Tides of consent: How public opinion shapes American politics* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9781316257791>