Research Article

ASSOCIATION FOR PSYCHOLOGICAL SCIENCE

Psychological Science
2019, Vol. 30(2) 174–192

DOI: 10.1177/0956797618813087
www.psychologicalscience.org/PS
SAGE

Patterns of Implicit and Explicit Attitudes: I. Long-Term Change and Stability From 2007 to 2016

Tessa E. S. Charlesworth and Mahzarin R. Banaji
Department of Psychology, Harvard University

Abstract

Using 4.4 million tests of implicit and explicit attitudes measured continuously from an Internet population of U.S. respondents over 13 years, we conducted the first comparative analysis using time-series models to examine patterns of long-term change in six social-group attitudes: sexual orientation, race, skin tone, age, disability, and body weight. Even within just a decade, all explicit responses showed change toward attitude neutrality. Parallel implicit responses also showed change toward neutrality for sexual orientation, race, and skin-tone attitudes but revealed stability over time for age and disability attitudes and change away from neutrality for body-weight attitudes. These data provide previously unavailable evidence for long-term implicit attitude change and stability across multiple social groups; the data can be used to generate and test theoretical predictions as well as construct forecasts of future attitudes.

Keywords

implicit attitude change, implicit association test, long-term change, time-series analysis, open data, open materials

Received 2/9/18; Revision accepted 8/27/18

The structure of the human brain has remained unchanged over thousands of years, but the products of the human brain—thoughts (beliefs) and feelings (attitudes)—are continually changing. Even within the relatively short history of the United States, examples of change on societally significant attitudes are easily found. For instance, from Puritan America through the 19th century, same-sex relations were punishable by death; today, same-sex marriage is legal across the United States. In 1958, only 4% of White Americans approved of Black–White marriages; today, 87% of White Americans approve (Newport, 2013).

The study of attitudes and attitude change is so fundamental to social psychology that every *Handbook of Social Psychology* since 1935 has included at least one chapter on the topic (Banaji & Heiphetz, 2010). As a result, much is known about the conditions that induce change in explicit, self-reported attitudes (Albarracin & Vargas, 2010). Specifically, theories of attitude strength and change (Petty & Krosnick, 1995) predict that resistance to explicit attitude change will cooccur with features such as high overall bias, strong intra-attitudinal linkages (indicated by high correlations among attitude measures), and low perceived societal priority (resulting in low discussion and elaboration).

Recently, attention has turned to the conditions for implicit attitude change. Initially hypothesized to be largely immutable (Bargh, 1999), implicit attitudes have since revealed short-term malleability following targeted interventions (Dasgupta, 2013; Lai et al., 2014). Yet the question of long-term implicit attitude change remains. With some exceptions (Devine, Forscher, Austin, & Cox, 2012; Gawronski, Morrison, Phills, & Galdi, 2017; McNulty, Baker, & Olson, 2014; McNulty, Olson, Jones, & Acosta, 2017), attempts to demonstrate long-term implicit attitude change have been unsuccessful (Lai et al., 2016) or failed to replicate (Forscher, Mitamura, Dix, Cox, & Devine, 2017). Moreover, investigations have largely focused on measuring single attitudes within an individual; comparisons of long-term implicit attitude change

Corresponding Author:
Tessa E. S. Charlesworth, Harvard University, Department of Psychology, William James Hall, 33 Kirkland St., Cambridge, MA 02138
E-mail: tet371@g.harvard.edu

across multiple attitudes at the population level remain unexplored.

Advantages of the Approach

This project examines the possibility of long-term implicit attitude change using data from the Project Implicit demonstration website (http://implicit.harvard.edu), which has collected two decades of continuous data from volunteers worldwide, yielding more than 20 million tests across 14 attitudes/stereotypes. We used a subset of these data involving 4 million tests that were collected continuously for 10 years across six social-group attitudes (sexuality, race, skin tone, age, disability, body weight), measured at the population level and analyzed with time-series models. This approach can newly address whether, and if so how, long-term change emerges in explicit and implicit attitudes.

Previous studies of attitude change have used within-persons repeated measures designs, ensuring strong internal validity but typically involving small samples, measurement of single attitude categories, and two discrete measures obtained within a brief period of time. An alternative approach, following social-survey methods, measures attitudes from different respondents across time, enabling large amounts of data to be collected continuously over years without concern for participant fatigue or practice effects. In giving up traditional within-persons measurement, we gain a singular opportunity to observe population-level change with continuous measurement across a decade.

Recently, such a population-level focus has been used to predict consequential outcomes (e.g., racial disparities in lethal force; Hehman, Flake, & Calanchini, 2017) as well as to reconsider the theoretical nature of implicit attitudes and their capacity for change (Payne, Vuletich, & Lundberg, 2017). Indeed, under this new perspective, population-level implicit attitudes are argued to be relatively more stable than individual-level attitudes because of greater stability in a culture than in an individual's daily experiences. Moreover, it should be preceded by situational change, including demographic, legislative, or explicit attitude change.

However, explicit-preceding-implicit attitude change is just one possible implicit–explicit relationship that has been documented in individual-level attitudes (Gawronski & Bodenhausen, 2006). Other patterns, including implicit-preceding-explicit, bidirectional, and unrelated change, may emerge and may differ across attitude targets. These data offer the first opportunity to empirically examine patterns of implicit–explicit change at the population level.

Additionally, this project compared change across social-group attitudes of sexuality, race, skin tone, age, disability, and body weight. Including six attitudes overcomes past limitations of almost exclusive study of Black–White racial attitudes and provides information on generalizability of change. Differences in rates of change across attitudes can also help rule out alternative explanations that would equally affect all attitudes, such as increasing awareness or test practice. Furthermore, comparisons across attitudes can test predictions regarding features of implicit attitudes that cooccur with change. This project examines predictions from explicit attitude theories (Petty & Krosnick, 1995) that implicit attitude stability will cooccur with higher bias, higher implicit–explicit correlations, and lower perceived societal priority (indicated by lower frequency of online searches).

Finally, this project introduces analytic improvements over previous studies using the same database (Sawyer & Gampa, 2018; Schmidt & Axt, 2016; Schmidt & Nosek, 2010; Westgate, Riskind, & Nosek, 2015). Past studies relied on linear multiple regressions—a model class that assumes linearity and independence of observations (i.e., no autocorrelations)—applied to data that violate both assumptions, thus opening the possibility of spurious conclusions about change. Furthermore, multiple regression is not designed to produce forecasts, an unnecessary limitation given theoretical and practical interest in predicting attitude trends. The current project employs autoregressive-integrated-moving-average (ARIMA) time-series models (Cryer & Chan, 2008) to address these concerns. There is growing appreciation that psychological data will benefit from adopting predictive machine-learning analyses (Yarkoni & Westfall, 2017), particularly when researchers are investigating the mechanics of cultural change (Varnum & Grossmann, 2017).

Method

Data source

Data, including respondents' zip codes, were retrieved from the Project Implicit demonstration website (https://implicit.harvard.edu/). Cleaned data used in the present study (without zip codes) are publicly available at https://osf.io/px8h3/; raw data (without zip codes), as well as further details about the website and test materials, are publicly available at https://osf.io/t4bnj/. All respondents were visitors to the Project Implicit website who provided informed consent and selected an implicit association test (IAT) from among the following: sexuality, race, skin tone, age, disability, and body weight. For all tests, the demographic questions, explicit measures, and IAT were presented to respondents in random order. Data inclusion began January 1, 2004, and ended December 31, 2016, for a total of 13 years. The fully available data from 2004 through 2016 were used for analyses of means and correlations. However, because of changes in the recording of explicit attitudes and demographics prior to 2007, only data collected after January 1, 2007, were available for the time-series models.

Body-weight and skin-tone IATs included missing data because no demographics or primary measures (IATs, explicit measures) were collected in particular months. Missing monthly averages for these tests were imputed using seasonal decomposition with linear interpolation (see Section 8 in the Supplemental Material available online). Additionally, stimuli for the body-weight IAT changed from face images to body-figure silhouettes in April 2010, resulting in two subsets of data (figure-stimuli test and face-stimuli test). Thus, data from both the recent figure-stimuli test (April 2010–December 2016) and the early face-stimuli test (March 2004–October 2011) are reported to account for the loss of early data in the figure-stimuli test.

Scores on the IAT were computed using the revised scoring algorithm (Greenwald, Nosek, & Banaji, 2003). Respondents whose scores fell outside of the conditions specified in the scoring algorithm did not have a complete IAT *D* score and were therefore excluded from analyses. Restricting the analyses to only complete IAT *D* scores resulted in an average retention of 92% of the complete sessions across tests. The sample was further restricted to include only respondents from the United States to increase shared cultural understanding of attitude categories. Finally, the sample was restricted to include only respondents with complete explicit measures and demographic information on age, gender, race/ethnicity, political ideology, education, and attitude-specific variables of sexuality, weight, and disability status. After these additional restrictions, an average of 62% of complete sessions remained across tests (for test-specific retentions, see Table S1 in the Supplemental Material). Supplemental analyses indicated that means and correlations (see Table S5.1 in the Supplemental Material), as well as patterns of change over time (see Table S5.2 in the Supplemental Material), were consistent for data without any exclusions (beyond having a complete IAT score), indicating that the excluded participants do not substantively alter the conclusions.

Sample demographics

Across all attitudes, a total final sample of 4,393,362 completed tests was obtained after retaining the maximum possible number of completed tests from U.S. respondents between 2004 and 2016. The unprecedented size of this sample and continuous measurement ensures that the aggregated estimates for each month are derived from reliably large samples with a median of 3,760 respondents per month (see Table S2 in the Supplemental Material).

Overall, the sample had a mean age of 27.46 years (*SD* = 11.91); 65.88% identified as female, 72.24% identified as White, 10.53% identified as Black or African American, 0.74% identified as American Indian, 5.09% identified as Asian, 1.88% identified as Black/White biracial, 4.85% identified as other biracial, and 4.99% identified as other or unknown; and 91.26% reported having completed a high school education or higher and 86.26% reported having completed a college education or higher. In addition, 45.98% identified as slightly, moderately, or strongly liberal; 25.77% as slightly, moderately, or strongly conservative; and 28.25% as politically neutral. Table S3 in the Supplemental Material provides test-specific demographic distributions.

Materials

IAT. The IAT (Greenwald, McGhee, & Schwartz, 1998) is a computerized task comparing reaction times to categorize paired concepts (in this case, social groups, e.g., young vs. elderly) and attributes (in this case, valence categories, e.g., good vs. bad). To sample the test, visit https://implicit.harvard.edu/implicit/takeatest.html. Respondents were presented with target stimuli (e.g., images of young and old faces, as well as good words, such as *joyful* and *friend*, and bad words, such as *evil* and *poison*), which were categorized into one of four groups (e.g., young, old, good, or bad). Average response latencies in correct categorizations were compared across two paired blocks in which participants categorized concepts and attributes with the same response keys. For illustration, in the age IAT, response latencies are compared across blocks in which (a) *young* + *good* have the same key and *old* + *bad* have the same key and (b) *young* + *bad* and *old* + *good* have the same keys. Faster responses in the paired blocks are assumed to reflect a stronger association between those paired concepts and attributes. In all tests, positive IAT *D* scores indicate a relative preference for the typically preferred group (in this example, young people). All IATs on the Project Implicit demo website use a standard seven-block format as described by Greenwald and colleagues (2003), with the order of the two paired blocks randomized across respondents.

Explicit preference. Explicit attitudes before 2007 were assessed on a 5-point Likert-type scale from −2 to 2, with higher scores indicating bias in favor of the typically preferred group (e.g., "I strongly prefer young people to old people") and lower scores indicating the reverse bias (e.g., "I strongly prefer old people to young people"). Explicit attitudes from 2007 onward were assessed on a

7-point Likert-type scale ranging from −3 to 3, with higher and lower scores reflecting the same preferences as above. In both cases, the midpoint of 0 represented equal liking of both groups.

Demographic variables. Respondents indicated their age, sex, education, political ideology, and ethnicity/race. For some tests, additional attitude-specific covariates were included: self-reported weight (for body-weight attitudes), self-reported sexual orientation (for sexuality attitudes), and self-reported disability status (for disability attitudes).

Results

Analytic strategy

Advantages of ARIMA models for analyzing change over time. Data were analyzed using ARIMA time-series models (Cryer & Chan, 2008; for an accessible introduction to time series in psychology, see Jebb, Tay, Wang, & Huang, 2015). We first provide a justification for using ARIMA models and then describe how the models are implemented; additional details are provided in the Supplemental Material.

ARIMA models offer several advantages over the multiple regression strategies used in previous examinations of Project Implicit data over time. First, and most fundamentally, measures over time are generally subject to significant autocorrelations (or serial dependence), meaning that the value measured at time t is dependent on, and highly correlated with, the value immediately before it. In other words, measures close in time are more related to each other than measures far in time. Indeed, in the current data, the median correlation (r) between monthly averages at month t and month t – 1 was .67, implying large and significant temporal dependence in the data. The presence of temporal dependencies in the data violates the assumption of independence in standard regression frameworks, leading to inefficient or biased model estimates. Indeed, Varnum and Grossmann (2017) provided an instructive summary of cases in which the failure to account for autocorrelations has led to spurious conclusions. Time-series models such as ARIMA are explicitly designed to accommodate the autocorrelated nature of time-series data and therefore address concerns raised by previous analytic strategies.

Second, visual inspection of the Project Implicit data shows that changes in implicit and explicit attitudes are characterized by substantial variability, seasonality (i.e., systematic within-year variability), and most importantly, nonlinearity. Thus, attempting to fit a single linear slope captures very little of the true nonlinear variation across time, resulting in poor model fit. Indeed, the median R^2 from multiple regressions with predictors analogous to those of previous studies (i.e., time, demographic covariates, and time-by-covariate interactions; see Schmidt & Axt, 2016; Schmidt & Nosek, 2010; Westgate et al., 2015) was .065 (see Section 13 in the Supplemental Material), implying that very little variance is accounted for when using these typical methods of analysis. ARIMA models account for nonlinearity by continuously updating their predictions on the basis of previous values and are therefore able to capture evolving nonlinear trends in the data.

Third, unlike regression models, ARIMA models can use information about nonlinearity and autocorrelation structures to provide optimal forecasts. In this case, they are designed to suggest upper and lower bounds of 95% confidence intervals (CIs) of attitude change into the future, that is, the time it will take for a specific attitude to reach neutrality or double in intensity. Forecasts provide intuitive estimates, akin to effect sizes with standard regression models, that are more easily interpreted than the ARIMA parameters alone. The inclusion of forecasts also aligns with recent calls to motivate psychology toward a predictive science, thereby improving scientific and applied understanding of the generalizability of trends into the future (Jebb et al., 2015; Varnum & Grossmann, 2017; Yarkoni & Westfall, 2017). Additionally, the current forecasts can be used to directly assess the impact of social events on the patterns of attitude change. Because Project Implicit data have the unique feature of being continuously updated, researchers can compare predictions generated by ARIMA models with future observed trajectories to quantitatively test whether social or political events substantively altered the patterns of attitudes outside of the provided CIs. At the least, these predictions are worth offering as a way of testing the viability of time-series models to account for attitude change.

Specifying and implementing ARIMA models. Time-series data are characterized, first, by differences in values at two time points; for example, the magnitude of an attitude measured on Saturday is different from the magnitude of an attitude measured on Sunday. Second, time series are characterized by the aforementioned autocorrelations; for example, attitudes measured on Saturday and Sunday are more similar to each other than attitudes measured on Saturday and Wednesday. Third, time series are characterized by lagged forecast errors; for example, the random errors in predicting attitudes on Saturday and Sunday are more similar to each other than the random errors in predicting attitudes on Saturday and Wednesday. ARIMA models describe these three features of time-series data using three parameters (p, d, q): d specifies the number of differencing parameters necessary to explain the differences between values, p specifies the number of autoregressive parameters used to explain the autocorrelations in the data, and q specifies the number of moving-average parameters used to explain the lagged forecast errors.

For example, to predict preferences on Sunday from preferences earlier in the week, the differencing (d) parameter would first be applied to ensure that the mean and variance across the week were stable, leaving only the daily fluctuations (i.e., "peaks and valleys"). These remaining daily variations would then be modeled by a combination of autoregressive (p) and moving-average (q) parameters. Autoregressive parameters would be used to model the consistent correlations between, for example, Sunday and Saturday (if p = 1) or Sunday and Friday (if p = 2). Moving-average parameters would be used to model the similarity in the random noise (i.e., error) in, for example, Sunday and Saturday (if q = 1) or Sunday and Friday (if q = 2).

These three parameters are first used to explain the nonseasonal component of the time-series data (i.e., the trends over the full time span investigated) but can also be extended to explain the systematic within-calendar-year variations, or *seasonality*, using seasonal ARIMA models (formally, SARIMA models). The same three parameters (p, d, q) are used to explain the seasonality component of the time-series data, with the same definitions as above. Finally, when time-series data include a clear and consistent slope over time, ARIMA models can include a drift parameter, which is analogous to a slope estimate in regression.

Thus, the final ARIMA models can include seven parameters, (p, d, q) (p, d, q) + drift, with the first three values specifying the order of nonseasonal parameters, the second three values specifying the order of seasonal parameters, and the inclusion of drift specifying any consistent slope. The order of the d parameter can be informative in revealing whether the time series is already stable (when d = 0) or is changing over time (when $d \neq 0$). The order of the autoregressive (p) and moving-average (q) parameters can reveal how many lags backward are necessary to predict the current measurement at time t. However, the order and values of autoregressive and moving-average parameters are generally not intuitively interpretable; we therefore focus on interpretation of the ARIMA model forecasts (see below).

In this article, ARIMA models were estimated using the automated algorithm implemented in the *forecast* package in the R programming environment (Hyndman & Khandakar, 2008; R Core Team, 2017). The algorithm explores the model space stepwise and chooses the best-fitting combination of nonseasonal, seasonal, and drift parameters, given the model's Akaike information criterion (AIC, a relative measure of model parsimony), as well as tests to ensure that the data are stationary (i.e., have been adequately differenced by the d parameter). Thus, all reported ARIMA models were entirely data driven.

Reporting and interpreting ARIMA forecasts. Mindful of the limits of any forecast, we offer three considerations in interpreting the results. First, in most cases, attitudes have shown a high degree of variability over time, visually manifesting in peaks and valleys. ARIMA models incorporate this variability and provide 95% CIs for their forecasts that include stability, change toward attitude neutrality, and change away from attitude neutrality. Although all directions of future change are therefore possible, we draw attention to the direction of the forecast nearest in time. For instance, we describe an attitude as forecast to move toward attitude neutrality if one bound of the forecast CI passes neutrality before the other bound of the CI passes double the initial attitude level. Neutrality and doubling are used as standards of change because they represent the same absolute value of change from the first measurement, but in opposite directions.

Second, we provide a supplemental internal analysis based on past data to allow greater confidence in the application of time-series models to offer predictions of future attitudes using the Project Implicit data set (see Section 12 in the Supplemental Material). To assess forecast accuracy for ARIMA models, we evaluated how well forecasts built from the first 8 years of data predicted the observed data of the last 2 years. Accuracy statistics (i.e., mean square error and mean percentage error) indicated that the ARIMA model approach had appropriate out-of-sample forecast accuracy when applied to implicit and explicit attitude change, thereby reinforcing confidence in the method and the results of forecasts into the future.

Third, the predictions offered are the best estimates based on currently available data of past trends. Thus, if data in the future change course because of unexpected shocks in the social or political climate, the predicted estimates would also be expected to systematically change in response. Notably, by offering these predictions, we gain the unique opportunity to quantitatively assess the impact of future social events in substantively altering the trajectories of implicit and explicit attitudes.

Addressing sample changes over time. Observed change may be an artifact of changes in the demographic composition of the sample over time. That is, observed attitude change toward neutrality could come not from true attitude change but from increasing numbers of female, liberal, young, non-White, or less-educated respondents, all of whom have been documented to have lower implicit bias (Nosek et al., 2007). Correlations between time (month) and each of these demographic variables indicate that, since 2007, the Project Implicit sample has become more liberal (see Table S4.1 in the Supplemental Material), which could create spurious movement toward neutrality across all attitudes. However, the sample has simultaneously become less female, less young, more White, and more educated (see Table S4.1), each of which could push the data away from neutrality over time. Evidently, it is necessary to separate out these spurious causes of change and identify the unique effect of time beyond changes to sample demographics.

To control for sampling changes, we calculated weights for all participants on the demographic variables of age, race, gender, education, and political orientation. Each individual participant was given a weight corresponding to his or her representativeness of the demographic frequencies from 2007 through 2016. For example, in the race IAT, the gender frequencies across the whole time span (2007–2016) were 61% female and 39% male, but in 2007, the gender frequencies were 63% female and 37% male. In the analysis, the data in 2007 would therefore be weighted such that women were "down-weighted" (given weights less than 1, because the data for female participants in 2007 were overrepresented by 2 percentage points relative to the frequencies for the whole time span), whereas men would be "up-weighted" (given weights greater than 1). These weighting procedures were performed for all demographic variables (not only for gender).

With the weighting values for each participant, we computed weighted monthly means for each attitude test, fitting the ARIMA models to the univariate time series of weighted monthly means. Weighting according to the sample demographics across the whole time span controls for changes in sample demographics between 2007 and 2016, thereby reducing the likelihood that any observed changes toward neutrality were due to increased participation from less-biased demographic groups and any observed changes away from neutrality were due to increased participation from more-biased demographic groups (for further details, see Section 6 in the Supplemental Material).

Addressing sample unrepresentativeness. The use of web-based data raises the concern of whether observed attitude change can be generalized from the Project Implicit sample to U.S. society. Sample demographics of race and education (based on high school attainment) approximated the population values for the United States, but the sample was younger, more liberal, and more female than the U.S. population (U.S. Census Bureau, 2016). To account for these demographic differences, we employed the same weighting procedures as above but weighted to the demographics of the 2007–2016 U.S. census. These procedures ensured that weighted monthly averages approximated the demographics in the U.S. population. Descriptions of overall patterns of change for census-weighted data largely replicated the results from within-sample weighted data, thus supporting the generalizability of our findings (see Section 7 in the Supplemental Material). However, the discrepancies between the sample and population demographics meant that, despite having stable weights, the census weighting was unable to converge. The results from census-weighted samples are therefore provided for illustration and should be interpreted with caution.

An additional argument for generalizability comes from convergent evidence in patterns of explicit attitude change using representative probability samples that are not considered to be affected by self-selection. Patterns of explicit attitude change observed in the present data are consistent with the patterns of change observed in many other surveys, including the rate and direction of change in beliefs about gay rights and race relations (e.g., General Social Survey, 2017).

Furthermore, a new study has revealed concordance between the magnitudes of implicit associations in the Project Implicit data and the magnitudes of linguistic associations in natural human language using the largest linguistic corpus of representative communication on the Internet, with more than 840 billion word tokens (Caliskan-Islam, Bryson, & Narayanan, 2016). Such consistency between, for example, the IAT magnitude of the *elderly* + *bad*/*young* + *good* association from Project Implicit and the magnitude of association of *elderly* + *bad*/*young* + *good* in Internet text provides confidence that the present database does not reflect the attitudes of a narrowly self-selected group. Finally, the increasingly wide usage of the Project Implicit website by institutions such as schools, governmental agencies, nonprofits, and for-profit corporations for organization-wide education has further reduced concern about self-selection and sample unrepresentativeness.

Addressing repeat test takers and regression to the mean. As the popularity of the Project Implicit website has increased over time, new respondents are added daily, but previous respondents have also returned to take additional tests. The initially extreme attitudes of novel test takers may gradually regress to the mean as they become experienced with the IAT. Thus, observed change toward neutrality could arise from increasing numbers of repeat test takers who are regressing to a more neutral attitude, rather than to genuine population attitude change. We addressed this concern by repeating the ARIMA models on the subset of data from respondents who reported having never taken an IAT before. Descriptions of overall patterns of change for these novel test takers were consistent with the patterns of change in the whole sample (see Section 11 in the Supplemental Material), implying that the cause of change cannot be attributed to repeat test takers and regression to the mean.

Addressing change across generational cohorts and demographic groups. Observed change toward neutrality could be due to (a) changes to the age of the sample, in which the average age of the sample decreases and younger respondents have lower bias; (b) cohort replacement, in which older generations are replaced by younger generations who have lower bias; or (c) period effects, in which external social or political forces affect all ages and cohorts in a society (Winship & Harding, 2008). These three factors (age, cohort, period) are linearly dependent, such that a period effect is equivalent to the combination of an age and cohort effect. For example, the extent of influence from external social forces (e.g., experiences in wartime) is entirely determined by the year in which one was born (e.g., 1935 vs. 1995) and the age at which one was measured (e.g., 10 years old vs. 30 years old). Thus, these three causes of change cannot be isolated from one another with the present data, a problem known as the identification problem. Nevertheless, we emphasize that changes to the age of the sample (i.e., age effects) are unlikely to account for the observed change because of the within-sample weighting procedures described above, which ensure that the average age of the sample remains consistent across the investigated time span.

To begin to examine the influence of cohort and period effects, we examined differences in patterns of change across four generational cohorts, defined by birth years: baby boomers (1945–1963), Generation X (1964–1975), millennials (1976–1995), and Generation Z (1996–2009). Data from earlier generations ("the silent generation") were not sufficient for monthly analyses (fewer than 100 observations per month). Additionally, data from Generation Z were included only after 2011, at which point monthly frequencies were sufficient.

To the extent that attitude change is observed only in younger generational cohorts, we can conclude that there is a cohort-by-period interaction, in which the social forces causing change (e.g., media campaigns or federal legislation) are predominantly affecting younger cohorts. In contrast, to the extent that attitude change is observed across all generational cohorts, we can conclude that the likely source of change is a period effect, in which the causes of change are widespread and affect respondents regardless of their age or cohort. Future research would benefit from modeling strategies that can further disambiguate age, cohort, and period effects with univariate time-series data.

In addition to investigations across generational cohorts, we examined patterns of change in demographic subgroups to address questions of generalizability across the disadvantaged and advantaged groups for each attitude test. To maintain the focus on illuminating overall patterns of change, we examined only directly relevant demographic subgroups for each attitude (e.g., respondent racial group for race and skin-tone attitudes). Research currently in preparation explores the full complexity of change by demographic subgroups and geographic location.

For both relevant demographic subgroups and generational cohorts, separate ARIMA models were fit to each subgroup; comparisons are reported for both past patterns of change (percentage change statistics) and future patterns of change from ARIMA model forecasts. Full models by relevant subgroups and cohorts are reported below for implicit attitudes and in Tables S9 and S10 in the Supplemental Material for explicit attitudes.

Addressing the relationship between explicit and implicit attitude change. The cross-temporal relationship between implicit and explicit attitude change in individual-level attitudes is a topic of substantial theoretical interest (Gawronski & Bodenhausen, 2006; Rydell & McConnell, 2006). The present data can be used to extend this discussion to long-term population-level attitude change. Specifically, we examined whether long-term implicit attitude change predicts explicit attitude change, or vice versa. This can be addressed by Granger causality models, which have been used in psychology to examine causes of cultural change (Grossmann & Varnum, 2015). Granger models of predictive causality examine whether the change in one variable (e.g., explicit attitudes) provides significantly more explanatory value about the change in a second variable (e.g., implicit attitudes), beyond using only the previous values of that second variable. In other words, does knowing past values of both explicit and implicit attitudes tell us significantly more about the patterns of implicit attitudes than merely knowing the past values of implicit attitudes alone?

The present study used Granger models to (a) predict implicit attitudes from explicit attitudes at lags of 1 month and 6 months and (b) predict explicit attitudes from implicit attitudes at lags of 1 month and 6 months. If only one direction of prediction is significant (e.g., only implicit to explicit), then one can conclude that change likely follows in that direction (e.g., implicit attitude change precedes explicit attitude change). If both directions of prediction are significant, then it is more likely that an exogenous third variable is causing change in both implicit and explicit attitudes. Finally, if neither direction is significant, then implicit change tells us little about explicit change, and vice versa; in this case, implicit and explicit attitudes may be changing for dissociable and independent reasons.

Addressing change as a function of societal priority of the attitude. At different points in history, particular social issues have received societal priority. For instance, during the American Civil War, the issue of slavery was of utmost social and political interest; during women's suffrage, women's right to vote became the focus of debate. In the United States today, race and sexuality attitudes appear to be societally prioritized (e.g., through the Black Lives Matter movement or legislation about same-sex marriage) and therefore are more frequently discussed than other attitudes, such as age or disability. Societal priority corresponds to more frequent and repeated exposure to debate or counterarguments that may, in turn, induce greater attitude change (Petty & Krosnick, 1995).

To determine whether change is indeed occurring faster on these prioritized attitudes, we examined the relative frequency of Google searches from January 2007 to December 2016 for three prejudice- and activism-related terms for each attitude—age: "ageism," "anti-old," "elder rights"; disability: "ableism," "anti-disabled," "disability rights"; race and skin tone: "racism," "anti-Black," "Black rights"; sexual orientation: "homophobia", "anti-gay," "gay rights"; and body weight: "sizeism," "anti-fat," "fat acceptance." Three commonly used terms were included for each social-group attitude to ensure that the relative rates of searches were not an artifact of any one search term comparison. Google searches have previously been validated as unobtrusive measures of social attitudes (Stephens-Davidowitz, 2014) and may be interpreted as proxies for the relative level of societal priority of an attitude.

Comparing the relative rates of group-related searches, we found that searches for "racism" were always more common than any other "-ism" term and that searches for "anti-gay" and "gay rights" were always more common than any other "anti-" or "rights" term, respectively (see Table S16 in the Supplemental Material). Out of a possible value of 100 (reflecting peak popularity), sexuality-related terms averaged a popularity of 18.53, race-related terms averaged 18.16, disability-related terms averaged 2.43, and neither age-related nor body-weight-related terms reached a score of 1. Thus, to the extent that societal priority and cultural-level frequency of discussion cooccur with relative rates of long-term implicit attitude change, it would be expected that race/skin-tone and sexuality attitudes will change faster than age, disability, or body-weight attitudes.

Implicit and explicit attitude means and correlations

All explicit and implicit attitudes showed significant preferences for the typically preferred group (i.e., pro-straight, pro-White, pro-light skin, pro-young, pro-abled, and pro-thin; see Table 1). The strongest implicit preferences were observed for disability, body-weight, and age attitudes, with relatively weaker implicit preferences observed for race, skin-tone, and sexuality

Table 1. Means and Correlations for Six Implicit and Explicit Social-Group Attitudes

		Implicit attitudes				Explicit attitudes				Correlation between explicit and implicit attitudes		
Attitude	*N*	*M* (*SD*)	95% CI	*t*	*d*	*M* (*SD*)	95% CI	*t*	*d*	*r*	95% CI	*t*
Sexuality	692,425	0.29 (0.48)	[0.29, 0.29]	495.24	0.60	0.54 (1.27)	[0.54, 0.54]	354.92	0.43	.42	[.41, .42]	381.61
Race	1,851,445	0.32 (0.44)	[0.32, 0.32]	972.40	0.71	0.32 (1.13)	[0.32, 0.32]	385.05	0.28	.32	[.32, .32]	459.66
Skin tone	488,330	0.31 (0.43)	[0.31, 0.31]	510.57	0.73	0.28 (1.00)	[0.28, 0.29]	197.60	0.28	.21	[.21, .22]	153.34
Age	588,230	0.44 (0.39)	[0.44, 0.44]	857.74	1.12	0.50 (1.23)	[0.50, 0.51]	314.24	0.41	.12	[.12, .13]	96.20
Disability	191,499	0.49 (0.44)	[0.49, 0.49]	493.63	1.13	0.49 (0.96)	[0.48, 0.49]	221.55	0.51	.14	[.14, .15]	63.17
Body weight (figure stimuli)	275,321	0.48 (0.41)	[0.48, 0.48]	609.65	1.16	0.92 (1.08)	[0.92, 0.93]	448.56	0.85	.22	[.22, .23]	119.88
Body weight (face stimuli)	306,112	0.40 (0.42)	[0.40, 0.40]	525.19	0.95	1.05 (1.14)	[1.05, 1.05]	508.80	0.92	.19	[.19, .19]	107.41

Note: All means and correlations are significantly different from zero ($p < .001$). CI = confidence interval.

attitudes. Similarly, strong explicit preferences were observed for disability, body-weight, and age attitudes, although strong preferences also emerged for explicit sexuality attitudes; relatively weaker explicit preferences were observed for race and skin-tone attitudes. Thus, if higher overall bias cooccurs with slower change, in line with the aforementioned predictions from attitude strength (Petty & Krosnick, 1995), then implicit age, disability, and body-weight attitudes should reveal slower change than implicit race, skin-tone, and sexuality attitudes. The relevance of differences in overall bias to rates of long-term change is discussed below.

Significant positive correlations between implicit and explicit attitudes were observed for all six attitudes, with the strongest correlations for sexuality and race attitudes and the weakest correlations for disability and age attitudes. The relevance of differences in correlation strength to rates of attitude change is discussed below. Notably, means and correlations for all attitudes closely replicated the results from previous large-scale analyses (Nosek et al., 2007), providing convergent evidence for the present sample.

Patterns of change

For each attitude, the reporting of results answered five questions. First, how fast have explicit attitudes changed over the past decade (see Table 2 and Fig. 1)? Rate of past change is indexed by the percentage change from the first to last months of the decomposed time-series trend, removing seasonality and noise. Second, how fast have implicit attitudes changed (by percentage change) and are they predicted to change in the future (see Table 2 and Fig. 1)? Predictions were derived from the number of months for the bounds of the 95% CIs of ARIMA forecasts of implicit attitudes to pass neutrality or double from the first month's value. Third, has the correlation between implicit and explicit attitudes changed over time (see Table 2)? Fourth, does implicit attitude change precede or follow explicit attitude change, as indexed by Granger causality models (see Table 3)? Fifth, does implicit attitude change generalize across generational cohorts (see Table 4) and across relevant demographic groups (see Table 5)?

Sexuality attitudes. Explicit sexuality attitudes revealed the largest overall change of any explicit attitude, moving toward neutrality by approximately 49% since 2007. Implicit sexuality attitudes also showed the most substantial overall change of all implicit attitudes, moving in the same direction but at a slower rate than explicit attitudes (changing toward neutrality by 33%). For implicit attitudes, the lower and upper bounds of the 95% CI of the ARIMA forecasts were predicted to pass neutrality by January 2025 and September 2045, respectively.

The correlation between implicit and explicit attitudes did not change substantively over time (decreased by ~3%). Granger causality models were inconclusive and revealed significant relationships in both directions, suggesting that an exogenous third variable may be causing simultaneous change in both implicit and explicit sexuality attitudes.

All respondents, regardless of sexual orientation, showed consistent past and future movement toward implicit pro-gay preference. Additionally, all cohorts changed toward attitude neutrality. This suggests that change in implicit sexuality attitudes may be caused by period effects. That is, the likely causes of changes in sexuality attitude are widespread shifts in the sociocultural climate that affect all ages, generational cohorts, and demographics. Such findings extend the conclusions from representative social surveys on change in explicit sexuality attitudes, which document movement toward neutrality across all birth cohorts (Rosenfeld, 2017).

Race attitudes. Over the past decade, explicit race attitudes have moved toward neutrality by approximately 37%. Implicit race attitudes have moved in the same direction but at a slower rate than explicit attitudes (changing toward neutrality by 17%), with a larger difference in rate of change than seen for sexuality attitudes. The pattern of change in implicit race attitudes revealed nonlinearity, with stability in early years followed by notable change since approximately 2012. The lower bound of the 95% CI of the ARIMA forecasts was predicted to pass attitude neutrality in August 2073, which is approximately half the time for the upper bound to reach double the level of initial bias.

Notably, the implicit–explicit correlation for race attitudes decreased by approximately 18% over the past decade, in contrast to the stability observed in implicit–explicit correlations for all other attitudes. The decrease in implicit–explicit correlations implies that, over time, the people holding strong implicit race attitudes are less likely to be the same people who hold strong explicit race attitudes (and, conversely, those holding weak implicit race attitudes are less likely to be those holding weak explicit race attitudes). In other words, implicit and explicit race attitudes show reduced correspondence over time, perhaps because of changing social desirability for this particular attitude.

Granger models reveal that implicit attitude change significantly predicted explicit change at lags of both 1 and 6 months, whereas the reverse direction (explicit to implicit) was not significant, suggesting that change in race attitudes likely flows from implicit to explicit attitudes.

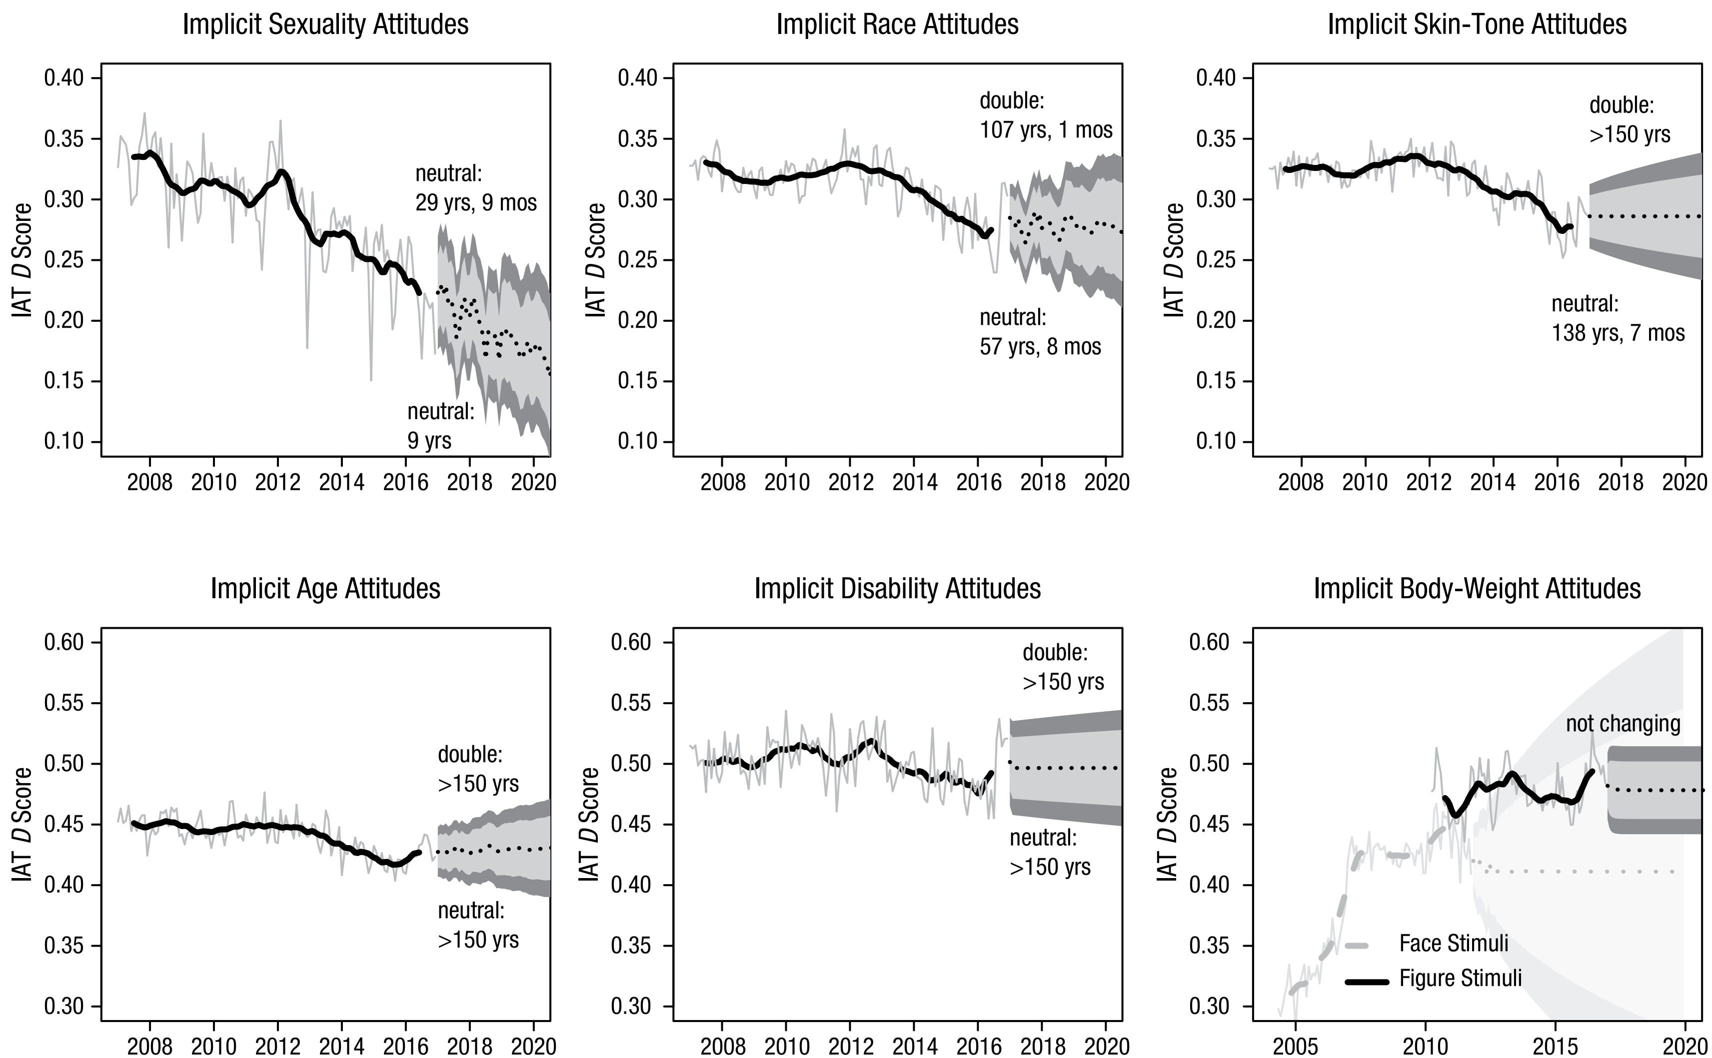

Fig. 1. *(continued on next page)*

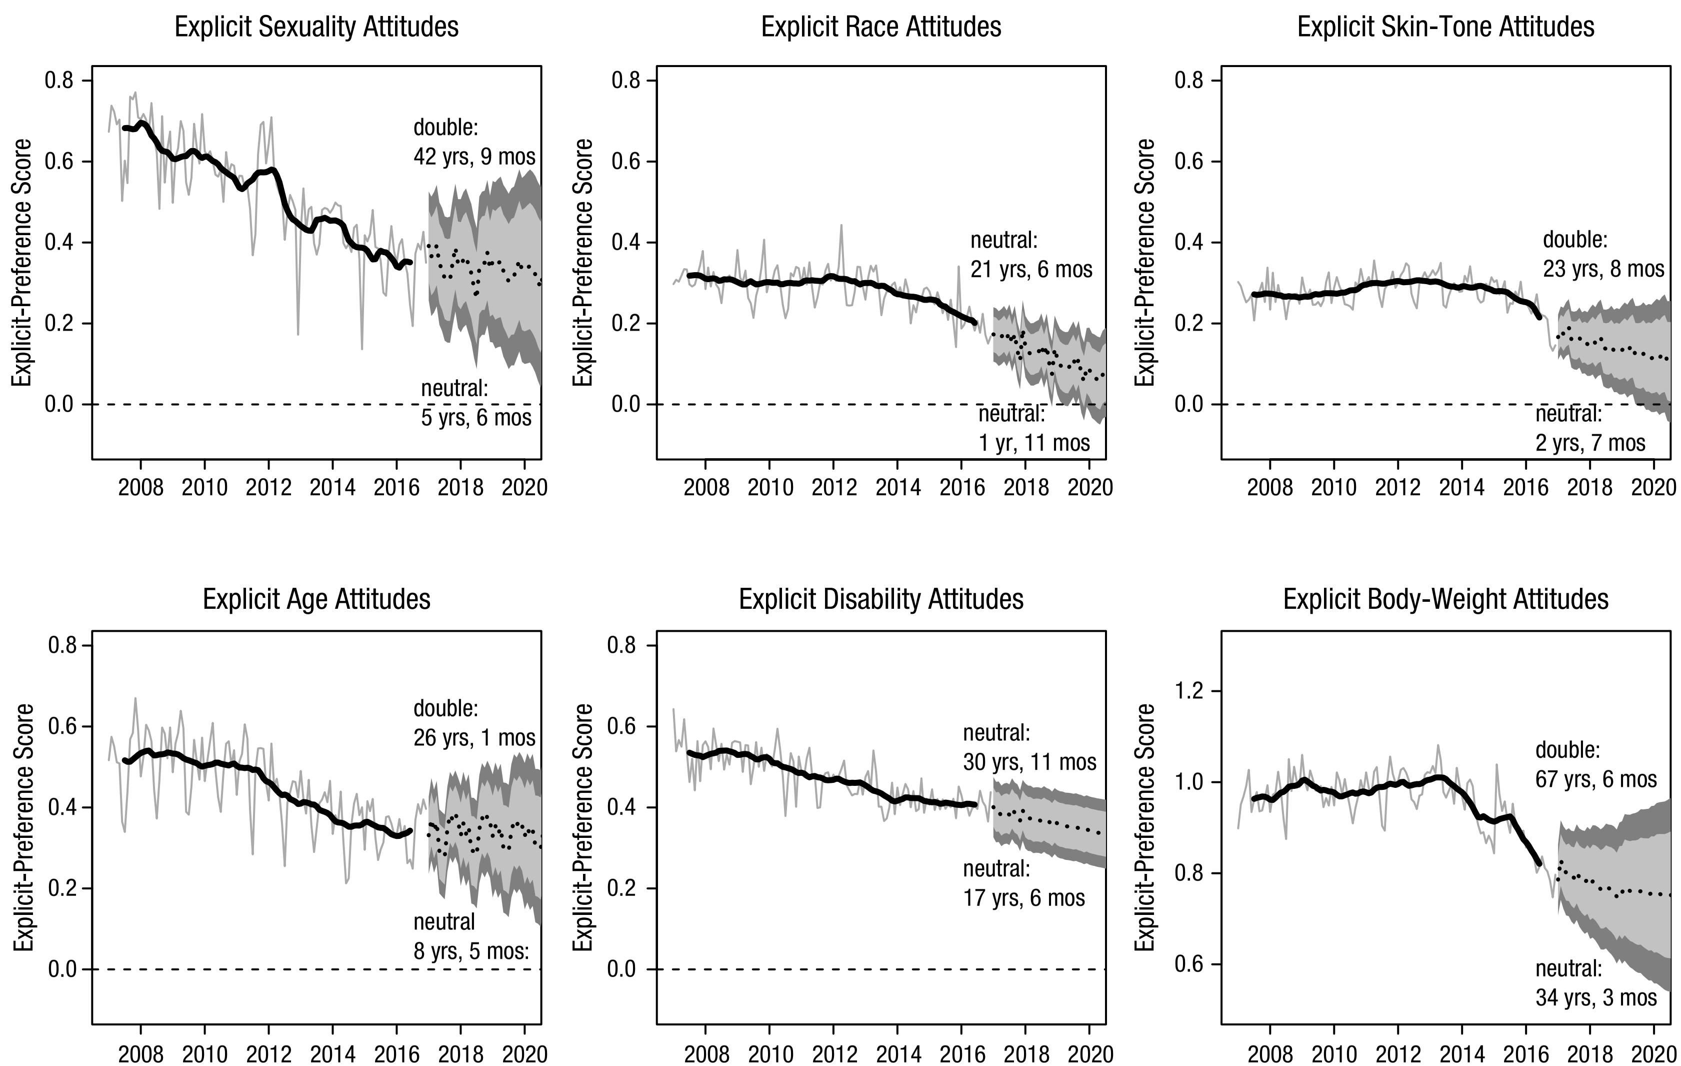

Fig. 1. Change and predicted change in implicit and explicit attitudes from 2007 to 2020: observed monthly weighted averages (2007–2016) of implicit association test (IAT) *D* scores (implicit attitudes; top two rows) and explicit-preference scores (explicit attitudes; bottom two rows), as well as forecasts of the autoregressive-integrated-moving-average (ARIMA) model (2017–2020). Solid black lines indicate decomposed trends of observed data (removing seasonality and noise), solid light-gray lines indicate the weighted monthly means from observed data, dotted black lines within the light-gray areas indicate the means of the ARIMA forecasts, light-gray areas indicate 80% confidence intervals (CIs), and dark-gray areas indicate 95% CIs of the ARIMA forecasts. Reported months to double indicate the number of months for the bound of the 95% CI to pass twice the level of initial bias; reported months to neutral indicate the number of months for the bound of the 95% CI to reach attitude neutrality. Implicit body-weight attitudes include data from both the face-stimuli test (dotted line and light-gray forecasts) and the figure-stimuli test (solid line and dark forecasts). Implicit age, disability, and body weight are on the same *y*-axis dimensions, whereas implicit race, skin tone, and sexuality are on the same *y*-axis dimensions. Explicit body-weight attitudes have different *y*-axis dimensions from all other explicit attitudes.

Table 2. Patterns of Change in Six Implicit and Explicit Social-Group Attitudes and Implicit–Explicit Correlations From 2007 to 2016

	Implicit attitudes				Explicit attitudes				Correlation between explicit and implicit attitudes		
Attitude	Starting raw value[a]	Ending raw value[a]	Percentage change in decomposed trend values[b]	ARIMA model parameters (p, d, q) (p, d, q)[c]	Starting raw value[a]	Ending raw value[a]	Percentage change in decomposed trend values[b]	ARIMA model parameters (p, d, q) (p, d, q)[c]	Starting r[a]	Ending r[a]	Percentage change in r
Sexuality	0.33	0.17	−33.46	(0, 1, 2) (2, 0, 0) + drift	0.67	0.35	−48.59	(0, 1, 2) (2, 0, 0)	.45	.43	−2.78
Race	0.33	0.30	−16.81	(1, 1, 1) (2, 0, 0)	0.30	0.17	−36.74	(0, 1, 1) (2, 0, 2) + drift	.35	.29	−18.20
Skin tone	0.33	0.29	−14.59	(0, 1, 1)	0.30	0.15	−21.04	(0, 1, 3) (1, 0, 1)	.22	.25	0.82
Age	0.45	0.42	−5.36	(0, 1, 2) (2, 0, 0)	0.52	0.40	−33.65	(0, 1, 2) (2, 0, 0)	.14	.12	1.84
Disability	0.51	0.52	−1.66	(0, 1, 2)	0.64	0.44	−23.97	(1, 1, 1) (2, 0, 0) + drift	.16	.19	−0.63
Body weight (figure stimuli)	0.48[d]	0.48	4.67	(1, 0, 0)	0.90[e]	0.80[e]	−14.81[e]	(0, 1, 2) (2, 0, 0)[e]	.24[f]	.21[g]	9.15
Body weight (face stimuli)	0.30[d]	0.41[d]	40.10	(0, 1, 1) (1, 0, 0)							

[a]Unless otherwise noted, starting values are from January 2007 and ending values are from December 2016. Starting and ending values are from the implicit association test (IAT; D scores) and 7-point explicit-preference scales. [b]Percentage change is between the first and last values of the decomposed time-series trend (removing seasonality and noise). Negative values indicate change toward neutrality (i.e., decreasing prejudice); positive values indicate change away from neutrality (i.e., increasing prejudice). [c]The first three parameters of the autoregressive-integrated-moving-average (ARIMA) model are nonseasonal, and the second three values are seasonal; drift is also included. In each set of parameters, d specifies the number of differencing parameters necessary to explain the differences between values, p specifies the number of autoregressive parameters used to explain the autocorrelations in the data, and q specifies the number of moving-average parameters used to explain the lagged forecast errors. [d]Implicit body-weight attitudes were measured in two tests: Figure stimuli were used from April 2010 to December 2016, and face stimuli were used from May 2004 to October 2011. [e]Explicit body-weight attitudes are from the combined data across the two tests, since explicit attitudes were not affected by the change in IAT stimuli. Thus, explicit body-weight attitudes are from January 2007 to December 2016, as with all other attitudes. [f]The starting implicit–explicit correlation for body-weight attitudes is from January 2007 and is therefore the correlation between explicit attitudes and implicit attitudes measured with the face-stimuli test. [g]The ending implicit–explicit correlation for body-weight attitudes is from December 2016 and is therefore the correlation between explicit attitudes and implicit attitudes measured with the figure-stimuli test.

Table 3. Results From Granger Tests of Predictive Causality on Implicit and Explicit Attitudes (With Trends)

Attitude	Implicit precedes explicit (1 month)	Implicit precedes explicit (6 months)	Explicit precedes implicit (1 month)	Explicit precedes implicit (6 months)
Sexuality	$F(1, 117) = 0.35$	$F(6, 107) = 3.18^{**}$	$F(1, 117) = 4.21^{*}$	$F(6, 107) = 2.24^{*}$
Race	$F(1, 117) = 11.33^{**}$	$F(6, 107) = 3.58^{**}$	$F(1, 117) = 0.35$	$F(6, 107) = 0.78$
Skin tone	$F(1, 117) = 0.41$	$F(6, 107) = 5.16^{***}$	$F(1, 117) = 1.91$	$F(6, 107) = 0.47$
Age	$F(1, 117) = 9.94^{**}$	$F(6, 107) = 2.71^{*}$	$F(1, 117) = 0.14$	$F(6, 107) = 5.85^{***}$
Disability	$F(1, 117) = 0.47$	$F(6, 107) = 1.99$	$F(1, 117) = 1.30$	$F(6, 107) = 0.49$
Body weight (figure stimuli)	$F(1, 78) = 0.20$	$F(6, 68) = 0.71$	$F(1, 78) = 1.29$	$F(6, 68) = 0.82$
Body weight (face stimuli)	$F(1, 55) = 0.12$	$F(6, 45) = 0.34$	$F(1, 55) = 0.79$	$F(6, 45) = 0.077$

$^{*}p < .05$. $^{**}p < .01$. $^{***}p < .001$.

Respondents of all racial groups (except Blacks/African Americans) moved toward implicit attitude neutrality; Black/African American respondents, however, had stable implicit pro-Black preferences over the past decade and were forecast to remain stable. Change toward neutrality was largest in millennials, with relative stability predicted for baby boomers and Generation Xers (as indicated by the absence of a differencing parameter in the ARIMA models). Because millennials are changing faster than older generational cohorts, the observed change in implicit race attitudes may be driven by a cohort-by-period interaction, in which the social forces driving attitude change are predominantly affecting younger cohorts.

Skin-tone attitudes. Explicit skin-tone attitudes have changed toward attitude neutrality by 21%, slower than sexuality and race attitudes. Implicit skin-tone attitudes revealed nonlinear change in the same direction but at a slower rate than explicit attitudes (changing by ~15%). The ARIMA forecasts for implicit skin-tone attitudes also indicate slower change than for implicit sexuality or race attitudes, with the lower bound of the 95% CI predicted to pass neutrality in July 2154 and the upper bound not predicted to pass doubling within the next 150 years.

The implicit–explicit correlation for skin-tone attitudes remained stable over the past decade, unlike the implicit–explicit correlation for race attitudes. Granger models indicated that implicit attitudes at a lag of 6 months significantly predicted explicit attitudes, whereas the reverse direction was not significant. Thus, as with race attitudes, this suggests that the direction of attitude change may flow from implicit to explicit attitudes.

Change toward implicit attitude neutrality was most notable among Black American respondents, whereas relative stability was observed among both White American and Asian American respondents. Change toward neutrality in implicit attitudes was also observed most strongly in millennials and Generation Zers, with relative stability (and even slight change away from neutrality) in baby boomers and Generation Xers. As with race attitudes, these generational patterns imply that the observed change toward neutrality could be attributed to a cohort-by-period interaction wherein the causes of change are largely focused on the attitudes of younger generations.

Age attitudes. Explicit age attitudes changed linearly toward attitude neutrality by approximately 34%. In contrast, implicit age attitudes revealed only slight change toward attitude neutrality over the past decade (changing by approximately 5%), moving in a parallel direction but at a slower rate than explicit attitudes. Indeed, age attitudes revealed the largest difference between rates of explicit and implicit change for any attitude. The upper and lower 95% CIs of ARIMA forecasts for implicit attitudes were not predicted to pass attitude neutrality or doubling within the next 150 years. Given the inherent uncertainty in forecasting over such long periods, forecasts beyond 150 years are best interpreted as attitude stability.

The implicit–explicit correlation for age attitudes also revealed stability over the past decade (increasing by less than 2%). Granger models were inconclusive regarding the direction of change, with significance in both directions, suggesting that an exogenous third variable may be causing the variability in both implicit and explicit attitudes. Finally, the stability in implicit attitudes was observed across all age groups and generational cohorts, with even the oldest cohorts and most elderly respondents showing stable implicit pro-young preferences.

Disability attitudes. Explicit disability attitudes changed toward neutrality by approximately 24%. However, no change was observed in implicit disability attitudes (changing by approximately 2%), and visual inspection shows

Table 4. Cohort Differences in Change for Six Implicit Social-Group Attitudes From 2007 to 2016

Attitude and generational cohort	*N*	Starting raw value[a]	Ending raw value[a]	Percentage change in decomposed trend values[b]	ARIMA model parameters (*p, d, q*) (*p, d, q*)[c]
Sexuality					
Baby boomers	41,203	0.28	0.21	−11.27	(0, 1, 1)
Generation Xers	65,168	0.30	0.21	−12.57	(0, 1, 1)
Millennials	400,882	0.34	0.16	−31.57	(1, 1, 1) (2, 0, 0)
Generation Zers	113,525	0.25[d]	0.26	−21.18	(1, 1, 3) (1, 0, 0)
Race					
Baby boomers	145,673	0.33	0.33	−1.47	(0, 0, 4)
Generation Xers	206,923	0.29	0.28	−13.83	(0, 0, 2)
Millennials	950,145	0.34	0.29	−19.98	(1, 1, 1) (1, 0, 1)
Generation Zers	173,932	0.29[d]	0.31	−5.62	(0, 1, 1)
Skin tone					
Baby boomers	36,963	0.35	0.42	4.20	(0, 0, 2)
Generation Xers	63,897	0.32	0.35	3.10	(3, 0, 1)
Millennials	259,142	0.32	0.27	−12.40	(2, 1, 3)
Generation Zers	47,556	0.33[d]	0.26	−14.12	(0, 1, 2)
Age					
Baby boomers	41,056	0.46	0.45	1.24	(0, 0, 2)
Generation Xers	45,845	0.46	0.45	−4.63	(2, 1, 1)
Millennials	289,760	0.45	0.43	−3.92	(0, 1, 2) (2, 0, 0)
Generation Zers	63,236	0.38[d]	0.41	10.59	(0, 1, 1)
Disability					
Baby boomers	17,783	0.59	0.73	9.18	(0, 1, 1)
Generation Xers	22,291	0.50	0.57	12.40	(0, 1, 3)
Millennials	107,650	0.50	0.49	0.67	(0, 0, 0)
Generation Zers	20,018	0.42[d]	0.47	27.55	(0, 0, 1)
Body weight (figure stimuli)					
Baby boomers	18,075	0.50[e]	0.53	4.34	(2, 1, 1)
Generation Xers	27,375	0.55[e]	0.51	7.69	(1, 0, 0)
Millennials	173,166	0.46[e]	0.47	6.50	(1, 0, 0) (1, 0, 2)
Generation Zers	55,820	0.35[d]	0.48	23.15	(4, 1, 0)
Body weight (face stimuli)					
Baby boomers	30,353	0.27[f]	0.41[g]	60.91	(0, 1, 1) + drift
Generation Xers	41,757	0.31[f]	0.45[g]	47.40	(0, 1, 4) + drift
Millennials	230,267	0.31[f]	0.42[g]	38.13	(0, 1, 1) + drift

Note: Baby boomers were born between 1945 and 1963, Generation Xers were born between 1964 and 1975, millennials were born between 1976 and 1994, and Generation Zers were born between 1995 and 2009.
[a]Unless otherwise noted, starting values are from January 2007 and ending values are from December 2016. Starting and ending values are from the implicit association test (*D* scores) and 7-point explicit-preference scales. [b]Percentage change is between first and last values of the decomposed time-series trend (removing seasonality and noise).
[c]The first three parameters of the autoregressive-integrated-moving-average (ARIMA) model are nonseasonal, and the second three values are seasonal; drift is also included. In each set of parameters, *d* specifies the number of differencing parameters necessary to explain the differences between values, *p* specifies the number of autoregressive parameters used to explain the autocorrelations in the data, and *q* specifies the number of moving-average parameters used to explain the lagged forecast errors. [d]Generation Z starting values were from January 2011 to ensure adequate sample size for each month. Generation Z body-weight (face-stimuli) values are not available because of the end date of the test. [e]Body-weight (figure-stimuli) tests started in April 2010. [f]Body-weight (face-stimuli) tests started in May 2004. [g]Body-weight (face-stimuli) tests ended in October 2011.

stability with a slight curvilinear trend of less neutral attitudes before approximately 2013 and slightly more neutral attitudes since 2013. As with age attitudes, rates of change for implicit and explicit disability attitudes show a large divergence, relative to sexuality and race attitudes. Neither the upper nor lower bounds of the 95% CIs of ARIMA forecasts for implicit disability attitudes were predicted to pass neutrality or doubling within the next 150 years.

Table 5. Relevant Demographic Differences in Change for Six Implicit Social-Group Attitudes From 2007 to 2016

Attitude and demographic group	N	Starting raw value[a]	Ending raw value[a]	Percentage change in decomposed trend values[b]	ARIMA model parameters (p, d, q) (p, d, q)[c]
Sexuality					
Straight	489,783	0.43	0.27	−24.91	(0, 1, 1) (2, 0, 0) + drift
Gay/lesbian	62,927	−0.15	−0.27	−80.73	(2, 1, 3) + drift
Race					
White American	1,059,974	0.40	0.36	−15.27	(0, 1, 2)
Black American	181,157	−0.089	0.0093	−3.61	(0, 0, 1) (1, 0, 0)
Asian American	73,051	0.32	0.31	−14.83	(0, 1, 1)
Skin tone					
White American	250,159	0.40	0.36	−11.45	(0, 1, 1)
Black American	74,991	0.095	0.10	−40.17	(0, 1, 3)
Asian American	20,973	0.30	0.27	−10.07	(0, 0, 1)
Age					
10–25 years	283,749	0.44	0.42	−6.54	(0, 1, 1) (2, 0, 0)
25–35 years	68,692	0.49	0.39	−6.79	(0, 1, 1) (1, 0, 0)
35–45 years	38,889	0.45	0.47	−3.17	(0, 0, 0)
45–55 years	31,945	0.45	0.43	−1.76	(0, 1, 1) (1, 0, 1)
55+ years	19,466	0.47	0.44	+2.53	(0, 0, 0)
Disability					
Disabled	24,127	0.42	0.43	−9.08	(0, 1, 2)
Not disabled	144,546	0.53	0.54	−0.69	(1, 1, 1)
Body weight (figure stimuli)					
Overweight	98,532	0.42[d]	0.44	+5.17	(1, 0, 0) (2, 0, 0)
Average weight	112,663	0.52[d]	0.52	+4.87	(3, 0, 0)
Underweight	24,184	0.50[d]	0.50	+3.20	(3, 0, 0) (1, 1, 1)
Body weight (face stimuli)					
Overweight	122,551	0.26[e]	0.38[f]	+48.25	(0, 1, 1) + drift
Average weight	135,468	0.34[e]	0.46[f]	+41.19	(0, 1, 1) + drift
Underweight	28,186	0.29[e]	0.43[f]	+26.47	(3, 1, 0)

[a]Unless otherwise noted, starting values are from January 2007 and ending values are from December 2016. Starting and ending values are from the implicit association test (*D* scores) and 7-point explicit-preference scales. [b]Percentage change is between first and last values of the decomposed time-series trend (removing seasonality and noise from the data). [c]The first three parameters of the autoregressive-integrated-moving-average (ARIMA) model are nonseasonal, and the second three values are seasonal; drift is also included. In each set of parameters, *d* specifies the number of differencing parameters necessary to explain the differences between values, *p* specifies the number of autoregressive parameters used to explain the autocorrelations in the data, and *q* specifies the number of moving-average parameters used to explain the lagged forecast errors. [d]Body-weight (figure-stimuli) tests started in April 2010. [e]Body-weight (face-stimuli) tests started in May 2004. [f]Body-weight (face-stimuli) tests ended in October 2011.

The implicit–explicit correlation for disability attitudes revealed stability, with change of less than 1% over the past decade. Granger models revealed no significant prediction in either direction of implicit or explicit change, implying that implicit and explicit attitudes are likely affected by dissociable processes. Finally, stability in implicit disability attitudes was observed for all respondents, regardless of disability status or generational cohort.

Body-weight attitudes. Explicit body-weight attitudes showed the slowest change toward attitude neutrality of all explicit attitudes, moving by approximately 15% over the past decade. Moreover, whereas all other implicit and explicit attitudes revealed trends toward neutrality, albeit at varying rates, implicit body-weight attitudes revealed movement away from neutrality over time, with slight changes in the figure-stimuli test (increasing by 5%) and large changes in the face-stimuli test (increasing by 40%).

In line with the small percentage change in the figure-stimuli test, the ARIMA model for the figure-stimuli test implied stability of implicit body-weight attitudes, with neither bound of the CIs predicted to pass neutrality or doubling. However, for the face-stimuli test, the best-fitting ARIMA model predicted change away from neutrality, with the upper bound of the 95% CI predicted to pass double the level of initial bias by June 2018, approximately 25 years before the lower bound would pass neutrality. Crucially, the recent data from the figure-stimuli test were included within the forecasted CIs of the face-stimuli test, suggesting that even the most recent stable data from the figure-stimuli test may still conform to a long-term pattern of change away from neutrality.

As with most other attitudes, implicit–explicit correlations for body-weight attitudes did not change over the past decade. Additionally, Granger causality models indicated no significant prediction of change in either direction, implying that implicit and explicit body-weight attitudes are changing independently.

For the figure-stimuli test, stability was observed across all respondents regardless of self-reported weight. Additionally, this stability was observed across baby boomers, Generation Xers, and millennials; Generation Zers revealed more substantial change away from neutrality than other generations, suggesting that the slight change away from neutrality may be attributable to a cohort-by-period interaction focused on the youngest cohort. For the face-stimuli test, movement away from neutrality was observed across all respondents, regardless of self-reported weight or generational cohort. This suggests that the early pattern of change away from neutrality could be attributed to a widespread period effect that changed the body-weight attitudes of all respondents in the years before 2010.

Summary of results. Even within just a decade, all explicit attitudes revealed change toward neutrality, implying that conscious and self-reported prejudice has decreased over time across attitudes. Crucially, long-term change was also observed across multiple implicit attitudes, with trends toward neutrality for sexuality, race, and skin-tone attitudes. Forecasts indicated vastly different rates of change among these three attitudes: Implicit sexuality attitudes were predicted to pass attitude neutrality as early as 9 years from 2016, whereas implicit skin-tone attitudes were predicted to take as long as 138 years. Moreover, the notable change in implicit sexuality attitudes was observed across all generations and demographic groups, whereas change in implicit race and skin-tone attitudes was observed most strongly in millennials and revealed demographic differences across racial groups.

Implicit disability and age attitudes revealed stability over time, regardless of generational cohort or demographics. Implicit body-weight attitudes, on both face-stimuli and figure-stimuli tests, revealed trends away from neutrality over time across all generations and demographics, in contrast to the direction of all other implicit and explicit attitudes.

General Discussion

Evidence for long-term change in population-level implicit attitudes counters the assumption that implicit attitudes, being less conscious and controllable than explicit attitudes, are necessarily immutable (Bargh, 1999). Instead, clear evidence of change across three attitudes suggests that implicit attitudes can be gradually and durably changed at the population level, in the direction of decreasing prejudice. This result complements the currently limited evidence of long-term individual-level implicit attitude change (e.g., Gawronski et al., 2017).

Notably, uncovering long-term attitude change requires statistical models that account for autocorrelations and nonlinearity. This article offers the first example of applying time-series (ARIMA) models to the study of implicit attitudes and the Project Implicit data, and the results challenge findings from linear multiple regressions. Past conclusions of stability in implicit race attitudes (Schmidt & Axt, 2016; Schmidt & Nosek, 2010) may result from fitting single linear slopes to nonlinear trends, thereby underestimating recent change. As in other areas of psychology, research on attitude change would benefit from embracing time-series models (Varnum & Grossmann, 2017).

Variability in rate and direction of change

Cross-attitude variability in the rate and direction of change warrants discussion. Examining overarching features that cooccur with changing attitudes reveals that, relative to stable attitudes (age, disability, body weight), changing attitudes (race, skin tone, sexuality) have lower overall bias, higher implicit–explicit correlations, and higher perceived societal priority (indexed by Google searches). That lower overall bias and higher societal priority cooccur with faster change is predicted from theories of attitude strength (Petty & Krosnick, 1995), newly showing that such predictions extend to implicit attitudes. However, that high implicit–explicit correlations correspond to faster change is, at first, unexpected: Under attitude-strength perspectives, high correlations are interpreted to reflect strong intra-attitudinal structure, which should correspond to slower change. An alternative interpretation, however, is that high implicit–explicit correlations reflect attitudes that are frequently discussed (Nosek, 2007). From this perspective, both implicit–explicit correlations and Google search prevalence imply that frequency of discussion is a consequential determinant of long-term change in implicit attitudes.

Whereas three theoretical explanations are offered for cross-attitude differences in rates of change, numerous additional distinctions could be drawn. For instance, rapid change in sexuality attitudes may arise from the unique concealability of sexual orientation (Pachankis et al., 2018), enabling positive contact before the stigma is revealed. In contrast, unique trends away from neutrality in implicit body-weight attitudes may arise from factors such as an increasing focus on health and the obesity epidemic, the increasing numbers of overweight individuals in the United States, and the perceived controllability of the stigma. Furthermore, we note that age, disability, and body-weight attitudes involve a perceived but measurable decline of the body and may therefore be seen to have an objective basis. In contrast, race, skin-tone, and sexuality attitudes are not rooted in objective evidence but have emerged for arbitrary and historic reasons (e.g., Sidanius & Pratto, 1999). Such differences in perceived objectivity may contribute to the relative stability of age, disability, and body-weight attitudes.

Variability in the implicit–explicit relationship over time

Implicit and explicit change between two out of six attitudes (disability and body weight) showed no significant relationship, supporting dual-process predictions of dissociable implicit and explicit change (Rydell & McConnell, 2006). However, a relationship was found between implicit and explicit change in four attitudes, although the direction was conclusive only for race and skin-tone attitudes. The presence of implicit–explicit associations speaks against complete dissociation, suggesting that the results may be better interpreted through frameworks allowing for interactive implicit–explicit change (e.g., the associative-propositional-evaluation, or APE, model; Gawronski & Bodenhausen, 2006).

Specifically, the current results imply that implicit attitude change precedes explicit attitude change for race and skin-tone attitudes (corresponding to APE Case 1), whereas exogenous influences likely cause change in both implicit and explicit sexuality attitudes (APE Case 6). In contrast, explicit age, disability, and body-weight attitude change may be negated before affecting implicit attitudes (APE Case 3). Although existing theories can inform these insights into the processes of attitude change, their predictions are derived from short-term individual-level data; processes of long-term population-level change may require theoretical revisions. Moreover, given variability across attitudes, researchers modeling implicit–explicit change would benefit from considering how the attitude target moderates processes of change.

Variability in change across generations and demographics

Most implicit attitudes showed generalizable trends across generational cohorts and across target and non-target demographic groups, despite the expectation that in-group preferences should dominate motivations to change (Tajfel, 1982). Indeed, implicit sexuality attitudes showed change across all cohorts and sexualities, newly suggesting that widespread and rapid attitude change is not limited to self-report (Rosenfeld, 2017). Nevertheless, implicit race and skin-tone attitudes revealed idiosyncratic differences by racial group, implying differential exposure to the causes of change, perhaps due to greater segregation by race than by other demographics (e.g., age, sexuality). Additionally, race and skin-tone attitudes showed the fastest change among millennials, suggesting that causes of change are predominantly affecting young generations.

Sample generalizability and source of change

Two limitations for interpretation are raised by the use of cross-sectional web-based data. First, can the observed change be generalized from this sample to U.S. society? The present sample is neither random nor representative, yet the results may be cautiously generalized for at least four reasons, elaborated in the Method section: (a) Weighting data to U.S. census demographics did not substantively alter descriptions of change, (b) similar change in explicit attitudes is observed in representative social surveys (General Social Survey, 2017), (c) similar magnitudes of implicit attitudes are documented in natural representative human language (Caliskan-Islam et al., 2016), and (d) the sample size and diversity offer improvement over small samples in much of psychology and poor response rates in probability samples.

The second potential limitation concerns the cause of change. Potential artifactual causes, including regression to the mean and increased representation of certain demographics, are shown not to account for the observed change. Additionally, although strategies to identify an independent age, period, or cohort effect remain beyond the scope of this article (Winship & Harding, 2008), initial explorations generally suggest cohort-by-period interactions in long-term population-level attitude change. Moreover, comparisons across attitudes provide insight into features that cooccur with, and may cause, implicit attitude change: low overall bias, high implicit–explicit correlation, and high societal priority. Addressing these limitations provides confidence in the project's contribution for generating and testing novel predictions about the patterns of long-term population-level implicit and explicit attitude change.

Action Editor

James K. McNulty served as action editor for this article.

Author Contributions

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

ORCID iD

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

Acknowledgments

We thank Patrick Mair for statistical guidance and Anthony Greenwald, Calvin Lai, Benedek Kurdi, Steven Pinker, Erin Westgate, and Christopher Winship for their comments on the manuscript.

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 Harvard University's Dean's Competitive Fund for Promising Scholarship awarded to M. R. Banaji.

Supplemental Material

Additional supporting information can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797618813087

Open Practices

Deidentified and cleaned data for this study, along with data-analysis scripts, have been made publicly available via the Open Science Framework and can be accessed at https://osf.io/px8h3/. The raw deidentified data and the materials from the Project Implicit demonstration website database are archived at https://osf.io/t4bnj/. The design and analysis plans for this study were not formally preregistered. The complete Open Practices Disclosure for this article can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797618813087. This article has received the badges for Open Data and Open Materials. More information about the Open Practices badges can be found at http://www.psychologicalscience.org/publications/badges.

References

Albarracin, D., & Vargas, P. (2010). Attitudes and persuasion: From biology to social responses to persuasive intent. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (Vol. 1, 5th ed., pp. 394–427). Hoboken, NJ: John Wiley & Sons.

Banaji, M. R., & Heiphetz, L. (2010). Attitudes. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (Vol. 1, 5th ed., pp. 353–393). Hoboken, NJ: John Wiley & Sons.

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). New York, NY: Guilford Press.

Caliskan-Islam, A., Bryson, J. J., & Narayanan, A. (2016). Semantics derived automatically from language corpora necessarily contain human biases. *Science*, *356*, 183–186. doi:10.1126/science.aal4230

Cryer, J. D., & Chan, K.-S. (2008). *Time series analysis: With applications in R* (2nd ed.). New York, NY: Springer-Verlag. doi:10.1007/978-0-387-75959-3

Dasgupta, N. (2013). Implicit attitudes and beliefs adapt to situations: A decade of research on the malleability of implicit prejudice, stereotypes, and the self-concept. In P. Devine & A. Plant (Eds.), *Advances in experimental social psychology* (Vol. 47, pp. 233–279). San Diego, CA: Academic Press. doi:10.1016/B978-0-12-407236-7.00005-X

Devine, P. G., Forscher, P. S., Austin, A. J., & Cox, W. T. L. (2012). Long-term reduction in implicit race bias: A prejudice habit-breaking intervention. *Journal of Experimental Social Psychology*, *48*, 1267–1278. doi:10.1016/j.jesp.2012.06.003

Forscher, P. S., Mitamura, C., Dix, E. L., Cox, W. T. L., & Devine, P. G. (2017). Breaking the prejudice habit: Mechanisms, timecourse, and longevity. *Journal of Experimental Social Psychology*, *72*, 133–146. doi:10.1016/j.jesp.2017.04.009

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, *132*, 692–731. doi:10.1037/0033-2909.132.5.692

Gawronski, B., Morrison, M., Phills, C. E., & Galdi, S. (2017). Temporal stability of implicit and explicit measures: A longitudinal analysis. *Personality and Social Psychology Bulletin*, *43*, 300–312. doi:10.1177/0146167216684131

General Social Survey. (2017). *GSS data explorer: Key trends*. Retrieved from https://gssdataexplorer.norc.org/trends

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, *74*, 1464–1480. doi:10.1037/0022-3514.74.6.1464

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I.

An improved scoring algorithm. *Journal of Personality and Social Psychology*, *85*, 197–216. doi:10.1037/0022-3514.85.2.197

Grossmann, I., & Varnum, M. E. W. (2015). Social structure, infectious diseases, disasters, secularism, and cultural change in America. *Psychological Science*, *26*, 311–324. doi:10.1177/0956797614563765

Hehman, E., Flake, J. K., & Calanchini, J. (2017). Disproportionate use of lethal force in policing is associated with regional racial biases of residents. *Social Psychological & Personality Science*, *9*, 393–401. doi:10.1177/1948550617711229

Hyndman, R. J., & Khandakar, Y. (2008). Automatic time series forecasting: The *forecast* package for R. *Journal of Statistical Software*, *27*(3). doi:10.18637/jss.v027.i03

Jebb, A. T., Tay, L., Wang, W., & Huang, Q. (2015). Time series analysis for psychological research: Examining and forecasting change. *Frontiers in Psychology*, *6*, Article 727. doi:10.3389/fpsyg.2015.00727

Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J. E. L., Joy-Gaba, J. A., . . . Nosek, B. A. (2014). Reducing implicit racial preferences: I. A comparative investigation of 17 interventions. *Journal of Experimental Psychology: General*, *143*, 1765–1785. doi:10.1037/a0036260

Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., . . . Nosek, B. A. (2016). Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General*, *145*, 1001–1016. doi:10.1037/xge0000179

McNulty, J. K., Baker, L. R., & Olson, M. A. (2014). Implicit self-evaluations predict changes in implicit partner evaluations. *Psychological Science*, *25*, 1649–1657. doi:10.1177/0956797614537833

McNulty, J. K., Olson, M. A., Jones, R. E., & Acosta, L. M. (2017). Automatic associations between one's partner and one's affect as the proximal mechanism of change in relationship satisfaction: Evidence from evaluative conditioning. *Psychological Science*, *28*, 1031–1040. doi:10.1177/0956797617702014

Newport, F. (2013, July 25). In U.S., 87% approve of Black-White marriage, vs. 4% in 1958. *Gallup*. Retrieved from http://www.gallup.com/poll/163697/approve-marriage-blacks-whites.aspx

Nosek, B. A. (2007). Implicit–explicit relations. *Current Directions in Psychological Science*, *16*, 65–69. doi:10.1111/j.1467-8721.2007.00477.x

Nosek, B. A., Smyth, F. L., Hansen, J. J., Devos, T., Lindner, N. M., Ranganath, K. A., . . . Banaji, M. R. (2007). Pervasiveness and correlates of implicit attitudes and stereotypes. *European Review of Social Psychology*, *18*, 36–88. doi:10.1080/10463280701489053

Pachankis, J. E., Hatzenbuehler, M. L., Wang, K., Burton, C. L., Crawford, F. W., Phelan, J. C., & Link, B. G. (2018). The burden of stigma on health and well-being: A taxonomy of concealment, course, disruptiveness, aesthetics, origin, and peril across 93 stigmas. *Personality and Social Psychology Bulletin*, *44*, 451–474. doi:10.1177/0146167217741313

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*, 233–248. doi:10.1080/1047840X.2017.1335568

Petty, R. E., & Krosnick, J. A. (Eds.). (1995). *Attitude strength: Antecedents and consequences*. Mahwah, NJ: Erlbaum.

R Core Team. (2017). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.

Rosenfeld, M. J. (2017). Moving a mountain: The extraordinary trajectory of same-sex marriage approval in the United States. *Socius*, *3*. doi:10.1177/2378023117727658

Rydell, R. J., & McConnell, A. R. (2006). Understanding implicit and explicit attitude change: A systems of reasoning analysis. *Journal of Personality and Social Psychology*, *91*, 995–1008. doi:10.1037/0022-3514.91.6.995

Sawyer, J., & Gampa, A. (2018). Implicit and explicit racial attitudes changed during Black Lives Matter. *Personality and Social Psychology Bulletin*, *44*, 1039–1059. doi:10.1177/0146167218757454

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*, 559–588. doi: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*, 308–314. doi:10.1016/j.jesp.2009.12.003

Sidanius, J., & Pratto, F. (1999). *Social dominance: An intergroup theory of social hierarchy and oppression*. New York, NY: Cambridge University Press. doi:10.2307/2655372

Stephens-Davidowitz, S. (2014). The cost of racial animus on a Black candidate: Evidence using Google search data. *Journal of Public Economics*, *118*, 26–40. doi:10.1016/j.jpubeco.2014.04.010

Tajfel, H. (Ed.). (1982). *Social identity and intergroup relations*. Cambridge, England: Cambridge University Press.

U.S. Census Bureau. (2016). *QuickFacts: UNITED STATES*. Retrieved from https://www.census.gov/quickfacts/table/PST045216/00

Varnum, M. E. W., & Grossmann, I. (2017). Cultural change: The how and the why. *Perspectives on Psychological Science*, *12*, 956–972. doi:10.1177/1745691617699971

Westgate, E. C., Riskind, R. G., & Nosek, B. A. (2015). Implicit preferences for straight people over lesbian women and gay men weakened from 2006 to 2013. *Collabra*, *1*(1), Article 1. doi:10.1525/collabra.18

Winship, C., & Harding, D. J. (2008). A mechanism-based approach to the identification of age–period–cohort models. *Sociological Methods & Research*, *36*, 362–401. doi:10.1177/0049124107310635

Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, *12*, 1100–1122. doi:10.1177/1745691617693393