



The Inseparability of Race and Partisanship in the United States

Sean J. Westwood¹ · Erik Peterson²

Accepted: 29 September 2020

© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

Many recent studies consider the overlapping nature of major political identities. Drawing on this research, we posit that partisanship and race are so enmeshed in the public mind that events which independently trigger one of these identities can also activate the other. We find support for this in three behavioral game experiments with 5496 respondents. These studies reveal what we refer to as the “parallel updating” of out-group affect. Shifts in racial affect are accompanied by simultaneous movement in attitudes and behavior towards members of the other political party. Conversely, changes in partisan affect co-occur with movement in views of racial out-groups. Our results speak to the inseparability of racial and partisan affect in the United States and suggest an important link between studies of racial animus and partisan affective polarization.

Keywords Polarization · Affect · Social identity

Foundational explanations for mass political behavior emphasize social group attachments. The large portion of the public that does not engage with politics on ideological terms can instead rely on others that share their race, class, or religion

Replication data and scripts are on the Political Behavior Dataverse: <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/IXOMN4>.

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s11109-020-09648-9>) contains supplementary material, which is available to authorized users.

✉ Sean J. Westwood
sean.j.westwood@dartmouth.edu

Erik Peterson
erik.peterson@tamu.edu

¹ Department of Government at Dartmouth College, 211 Silsby Hall 3 Tuck Mall, Hanover, NH 03755, USA

² Department of Political Science at Texas A & M University, 2010 Allen Building, 4348 TAMU, College Station, TX 77843, USA

(Berelson et al. 1954; Lipset 1960; Converse 1964; Conover 1988; Dawson 1994; Achen et al. 2016; Kinder and Kalmoe 2017). From this perspective even political parties draw their relevance from the set of social groups they represent (Campbell et al. 1960; Greene 1999; Green et al. 2002).

In an era of elevated out-party hostility, attention has turned towards how the intersection of multiple identities influences partisanship (Mason 2018; Klar 2018; Levendusky 2018). Evidence shows that social sorting—the alignment of an individual’s social group memberships with their party—amplifies partisan affective polarization (Mason 2015; Mason and Wronski 2018; Robison and Moskowitz 2019) and that those who strongly associate the opposing political party with other social out-groups hold more negative views of the other party (Ahler and Sood 2018).

Here we extend this research on the consequences of overlapping political identities with a focus on race and partisanship. Many studies establish the historical connection between these identities in the United States (e.g., White and Laird 2020; Acharya et al. 2020), and show that, in response to notable changes in elite-level politics, the alignment between race and partisanship among the mass public can shift (e.g., Carmines and Stimson 1989; Valentino and Sears 2005; Tesler 2013; Kuziemko and Washington 2018). Building on this important research, and further drawing on insights from psychology on the “spreading activation” of political concepts in memory (e.g., Lodge and Taber 2013), we argue for the need to also consider whether more rapid and interlinked shifts in out-group affect are possible when events raise the salience of racial or partisan identity. More specifically, based on this research we posit that partisanship and race are so enmeshed in the public mind that experiences which involve only one of the two groups may affect evaluations and behavior towards both, a connection with implications for the study of racial animus and partisan affective polarization.

While the rapid, interlinked shifts in out-group affect we propose are certainly not ruled out by previous work, existing research designs are unable to provide evidence of them. For this reason, a key contribution is to develop a new experimental framework that allows us to trigger negative and positive experiences with out-group members. We use this to measure what we refer to as the “parallel updating” of out-group affect. This is to say, we consider whether changes in both racial and partisan attitudes can be caused by a treatment based on only one of the two identities.¹ Using this approach in three behavioral game experiments with over 5,000 respondents, we find evidence that views of partisan and racial out-groups are inextricably connected and can rapidly update. Shifts in racial affect are accompanied by simultaneous movement in partisan affect and change behavior towards out-partisans. Conversely, shifts in partisan affect coincide with movement in views of racial out-groups.

Our findings extend current research in important ways. The rapid and interlinked shifts in out-group affect revealed in our experiments support an instinctive view of group-based political engagement that research in American politics has hinted

¹ We limit the scope of this paper to whites and African Americans. Historically these two racial groups are at the core of partisan and racial tensions and they therefore represent the cleanest test of our model.

at for decades (e.g., Converse 1964; Green et al. 2002; Huddy et al. 2015; Achen et al. 2016) and provides further evidence of the distinctive connection between race and partisanship in contemporary American politics. Moreover, while most previous work focuses on how race influences partisanship, we find evidence for a bidirectional relationship—one where partisan experiences can change both racial and partisan attitudes and where racial experiences also shape both racial and partisan attitudes (see also Engelhardt 2018). Finally, the influence of partisan affect on racial attitudes we identify substantially expands the implications of recent work on affective partisan polarization (for a review see Iyengar et al. 2019). Our results suggest that partisan affective polarization is partially a consequence of racial experiences triggered in daily interactions. Simultaneously, these findings suggest that partisan hostility reinforces racial hostility.

Social Identities in Politics

The dominant model of partisanship—derived from social identity theory (Tajfel 1970; Greene 1999; Huddy 2001; Mason 2018)—suggests partisan identity exists as a distinct construct from other social identities. According to this, individuals instinctively form attachments which produce in-group favoritism and out-group antipathy. This pattern occurs based on “minimal groups,” arbitrary researcher-generated divisions, and due to real-world social cleavages such as race and class. These attachments enable a public with limited sophistication and ideological understanding to make sense of politics (Campbell et al. 1960; Conover 1988; Dawson 1994; Green et al. 2002; Tesler 2016; Achen et al. 2016). The influence of partisan and other social identities extends into how people interact with out-group members in both political and non-political situations (Tajfel 1970; Iyengar and Westwood 2015).

Recent applications of social identity theory in politics extend in two directions. One line of research demonstrates that partisan identities have become increasingly salient in contemporary politics. Iyengar et al. (2012) document a substantial rise in affective polarization, the divide between the positive feelings of partisans toward the political party they identify with and their negative feelings toward the party they do not, over the past several decades (see also, Hetherington and Rudolph 2015; Mason 2015; Abramowitz and Webster 2016; Westwood et al. 2019). Whether assessed using self-reported, implicit, or behavioral indicators of partisan sentiment, negative affect toward the opposing party has reached unprecedented levels in contemporary politics (McConnell et al. 2018; Carlin and Love 2018; Theodoridis 2017; Iyengar and Westwood 2015; Huddy et al. 2015).

Another strand of research brings new attention to the relationships between social identities. “Socially sorted” individuals with multiple identities that align with their partisanship dislike their political opponents more than others with cross-cutting identities (Mason 2018, 2016; Levendusky 2018; Robison and Moskowitz 2019). This pattern does not occur in all circumstances as Klar (2018) isolates conditions in which a shared identity (gender) magnifies differences between members of different political parties. Margolis (2018), Egan (2020) and Davenport (2016)

reverse the canonical ordering of the social group-political party relationship, showing that political views may influence the adoption of some social identities. These studies show connections between social identities and partisanship, but also highlight the need for further consideration of exactly how various social identities relate.

Historical Alignments Between Race and Party

We build on an extensive literature establishing the central importance of racial and partisan identity in American politics (Campbell et al. 1960; Carmines and Stimson 1989; Dawson 1994; Kinder and Sanders 1996; Green et al. 2002; Hutchings and Valentino 2004; Valentino and Sears 2005; Tesler 2013; White et al. 2014; Kuo et al. 2017; Kuziemko and Washington 2018). Three aspects of this work are relevant to our present focus.

First, partisanship and race are closely linked, both in actuality and perceptually. Racial and partisan identities have become more closely associated since the 1970s (Carmines and Stimson 1989; Green et al. 2002; Tesler 2016; Kuziemko and Washington 2018). Perceptually, Ahler and Sood (2018) show that African Americans are regarded as prototypical Democrats, so much so that the public dramatically overestimates the share of the Democratic party that is black. Whites are strongly associated with the Republican party in open-ended survey responses (Rothschild et al. 2019; Halpern and Rodriguez 2018).

Second, the causal relationship linking attitudes related to race with partisanship runs primarily in one direction, with racial identity and attitudes toward racial groups influencing the development of partisanship and views of the political parties. In discussing the formation of partisanship, (Green et al. 2002, 109) note that “based on their understanding of which groups support each party and their own affinity for these groups, many citizens come to see themselves as members of partisan groups” (see also, Campbell et al. 1960). This perspective conceives of partisan identity as the product of affect toward a political party’s constituent groups, with racial groups among the most important (Converse 1964). Indeed, the association between partisanship and race arises during political socialization (Valentino and Sears 2005). The centrality of racial attitudes in the formation of partisanship (Sears and Funk 1999) and their subsequent stability (e.g., Schuman et al. 1997; Goldman and Hopkins 2019) is generally thought to limit the opportunity for a recursive process in which partisanship influences racial attitudes (but see Engelhardt 2018).

Third, after an understanding of the relationship between race and partisanship is established, citizens can update their views in response to notable changes in elite-level politics, particularly when aided by rhetoric emphasizing changes in political party reputations (Carmines and Stimson 1989; Kinder and Sanders 1996). Several examples illustrate this point. The passage of high-profile civil rights legislation prompted the public to associate the Democratic Party with African-Americans (Carmines and Stimson 1989; Green et al. 2002; Kuziemko and Washington 2018). While this led white southerners to shift to

the Republican party, the crystallized nature of partisan attachments meant this process took decades (Valentino and Sears 2005; Schickler 2016; Mason 2018). More recently, Barack Obama's emergence as the Democratic presidential nominee racialized partisan attitudes over the course of his presidency (Tesler 2013, 2016).

Psychological Links Between Race and Party

The second set of studies we build on consider the psychological nature of racial and partisan identities. The development of new measurement techniques, such as the implicit association test, has allowed researchers to show that political attitudes and identities can operate at a sub-conscious level. Studies that take this approach uncover an array of politically important implicit attitudes (Kam 2007; Mo 2015; Pérez 2016; Ryan 2017). Relevant to our present focus, these studies establish the existence of implicit attitudes regarding racial groups (Greenwald and Banaji 1995; Baron and Banaji 2006; Kinder and Ryan 2017; Pérez 2010) and political parties (e.g., Iyengar and Westwood 2015; Theodoridis 2017).

This psychological orientation towards the study of partisan and racial identities has implications for evaluations that involve these groups. It shows that, upon encountering a group-related stimulus, individuals arrive at an affective response prior to the point at which any cognitive assessments occur (Lodge and Taber 2013; Ryan 2017). This suggests the potential for rapid shifts in attitudes associated with racial and partisan identities.

Psychological treatments of identity also have implications for how racial and partisan identities relate to one another. Of particular relevance, studies in this tradition show that responses to group stimuli can be interlinked and simultaneously involve multiple identities, rather than occurring in isolation. The “spreading activation” literature in cognitive psychology finds that related concepts in memory are activated as a single unit (Collins and Loftus 1975; Quillian 1967). In this model groups are defined in memory as nodes connected by varying distances. When the distances between two nodes are large, activating one node is unlikely to trigger another. However, as the distance between two nodes shrinks, activating one cascades activation to the other, proximate node. For instance, evoking the concept “fruit” activates nodes related to “apple” and “pear” (Collins and Loftus 1975).

In an important application of these concepts, Lodge and Taber (2013) establish that social groups trigger these sub-conscious processes and show this goes on to shape political information processing. Building on this work, here we consider its implications for the structure of affect towards social out-groups. Given the systematic association between race and partisanship in the public mind, we argue that the proximity of these two concepts in memory is small enough to link racial and partisan identities together. This implies that affect towards out-groups on these dimensions can behave as a single unit (Wyer et al. 2014) and that activating racial identity (partisan identity) will also activate partisan identity (racial

identity). In the next section, we further develop this argument and its implications for out-group affect.

Considering the Parallel Updating of Racial and Partisan Affect

Studies of the alignment between race and partisanship in American politics show the central role of these two identities in the public mind and provide evidence on the relationship between racial and partisan attitudes. Psychological treatments of these identities suggest the possibility of rapid and interlinked shifts in views of partisan and racial groups due to spreading activation. We now synthesize the implications of these accounts for our key dependent variables: affect towards members of the opposing political party, a central element of studies of partisan affective polarization, and affect towards out-race groups, a focus in studies of racial animus.

We argue that rapid and interlinked shifts in affect towards racial and partisan out-groups may be possible when events raise the salience of either racial or partisan identity. We refer to this process as the *parallel updating* of racial and partisan affect. This perspective implies that partisanship and race are sufficiently enmeshed in public opinion that experiences which involve only one of the two groups may affect evaluations and behavior towards both. This is because, while there may be variability in the extent to which race and partisanship are linked in the minds of individuals, we suspect this association is sufficiently strong in American politics to affect most Americans. This perspective also means that parallel updating should occur in both directions, as research shows discussions of partisanship implicitly invoke race, as well as the reverse (e.g., Ahler and Sood 2018; Rothschild et al. 2019). Our expectation, then, is that primes that only invoke racial identity can have implications for both racial and partisan affect while primes that only expressly focus on partisan identity can also have implications for both partisan and racial affect.

Our central empirical focus in the rest of the paper will be using a series of experiments to create situations in which the parallel updating of out-party and out-race affect might occur and measuring whether this process happens. Before proceeding, a few points merit further attention. First, we again note our focus on the parallel updating of racial and partisan affect among Black and white respondents in the United States. This decision creates a “most likely” case for the parallel updating of out-group affect to occur. This is because these two racial groups have been at the core of partisan and racial tension. Focusing our attention in this manner has a key advantage. If parallel updating of out-group affect does not occur in these studies, it is also unlikely to happen in other situations that stray further from these conditions. However, this also means that while we anticipate this process to be weaker for identities that are less clearly associated and salient in the public mind, we do not start with firm expectations about how and whether it generalizes to other group pairings (e.g., those involving religion, class, gender). While we lack such expectations at the outset, aspects of our study are designed to provide some insight into the scope conditions around this process (e.g., we look at whether parallel updating can occur for out-groups that are less clearly connected to race or partisanship).

Second, our theory hinges on the notion that people connect race and partisanship as key political identities and change their views of out-groups on these dimensions in parallel. For socially sorted individuals (i.e., white Republicans and Black Democrats) this has straightforward implications for the structure of parallel updating. For instance we expect a Black Democrat who begins to view Republicans (their partisan out-group) in a more negative light to simultaneously adjust their views of whites (their racial out-group) in a negative direction. However, this theoretical logic also extends to those with cross-cutting identities (i.e., white Democrats and Black Republicans) where the “us against them” conception of the structure of out-group affect we use produces a less intuitive expectation. For white Democrats, for instance, it suggests that a negative shift in views of Republicans (their partisan out-group) will lead to a negative shift in views of Blacks (their racial out-group), a pattern of parallel updating clearly at odds with the real-world alignments of party and race.

We think it is important to state this expectation, and the reasoning behind it, clearly before we proceed in order to establish all the implications generated by this theoretical perspective. We also emphasize that, given the theoretical work on implicit attitudes we draw on, such parallel updating would reflect an automatic and pre-cognitive response, not a more reasoned assessment based on careful consideration of the partisan composition of various racial groups. In our analysis, we will separately consider our results among those with sorted and cross-cutting out-group identities to see if this expectation is born out. It may not occur if those with cross-cutting out-group identities do not engage in parallel updating or instead update their views by considering the empirical alignment between race and party.

Experiments to Study the Parallel Updating of Group Affect

We use a new experimental design to assess the parallel updating of out-group affect on party and race. While the next section fills in the specifics of this design, here we introduce the overall concept and establish its links to our theoretical motivation. We randomize the quality (i.e., positive or negative) of an individual’s interpersonal interaction with out-group members. The first quantity of interest to emerge from this design is the effect of the encounter on an individual’s views of the treated out-group, which we measure with survey or behavioral measures. We also measure their views or behavior toward other social groups (i.e., out-groups not emphasized in the interaction) to test for parallel updating. This allows us to estimate the effect of this interpersonal encounter on the participant’s behavior or affect toward these other social groups. In this conception, parallel updating is not mediated by the effect of the encounter on an individual’s views of the target group. Instead affect towards both the target and other groups occurs simultaneously due to the proximity of the two groups in memory.

Figure 1 shows how this design can be used to study the dynamic structure of partisan and racial affect. In each case, an initial shift in affect occurs due to an interpersonal interaction that is randomly assigned to be a positive or negative experience.

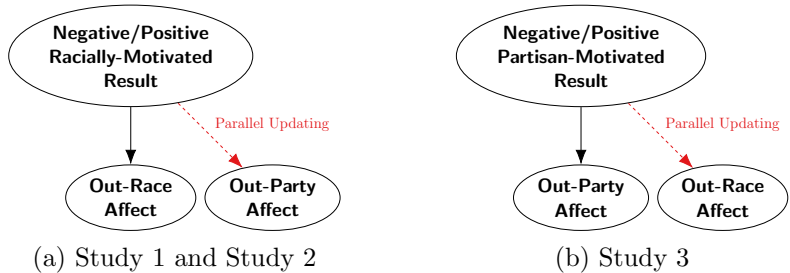


Fig. 1 Experimental design to study parallel updating of out-group affect

In the left panel, this interaction is with a member of a different race. In the right panel, the interaction is with a member of a different political party. The solid lines indicate the effect of the treatment on the corresponding outcome measure (i.e., racial affect for the out-race treatment and partisan affect for the out-party treatment). The dashed lines indicate the change expected due to parallel updating (i.e., partisan affect for the out-race treatment, racial affect for the out-party treatment).

This design is uniquely suited to measure whether the parallel updating of out-group affect occurs. We should only observe changes in partisan affect from the racial treatment in Panel A, or the reverse pattern in Panel B, if participants simultaneously update out-partisan affect when updating out-racial affect.

Delivering Treatments in a Behavioral Game

This section introduces the elements of our experimental design. We manipulate the quality of an individual’s interaction with an out-group member in a modified trust game (Berg et al. 1995). These games are commonly used to examine preferences for out-group cooperation by measuring the amount of money players allocate to others different from themselves on traits such as race, gender, or partisanship (Berg et al. 1995; Fershtman and Gneezy 2001; Habyarimana et al. 2007; Whitt and Wilson 2007; Carlin and Love 2018).

We depart from prior work by using the trust game setup to *administer* our treatment by manipulating the valence of a participant’s interaction with a member of an opposing social group in the game. We inform participants they are playing with an out-group member (i.e., someone from a different race or political party) and randomize whether they have a positive experience with the opposing player (i.e., receive a generous monetary allocation from them) or a negative experience (i.e., receive no money).

Typically, Player 1 in the trust game receives a cash allocation and is instructed to give “some, all, or none” of the money to a second player. The player is also told that the researchers will *triple* any amount given to Player 2 and that this other player could, at his or her discretion, return some, all, or none of the money back to Player 1. The implication is this: the more Player 1 expects reciprocity from Player 2, the more money they should allocate to them.

Political Behavior

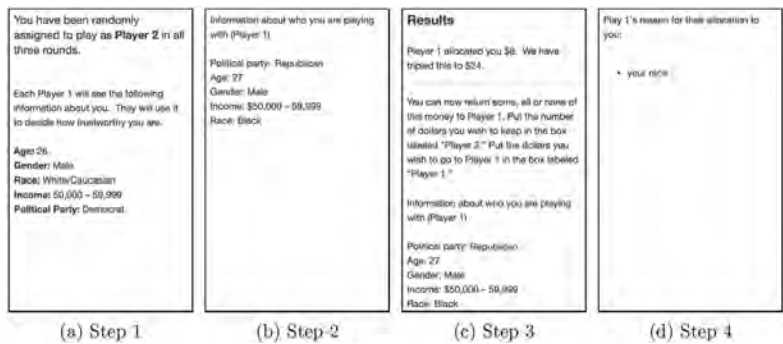


Fig. 2 Steps in the modified trust game

Our games use this basic framework, but alter it in several ways. First, participants—when being treated—take the role of Player 2. This means they observe the allocation another player makes to them. Second, participants interact with a contrived opponent who offers allocations based on a pre-determined script. Third, we assign demographic traits to this opponent, ensuring they represent an out-group (i.e., differ from the participant with respect to either race or partisanship).

During each round, participants were told that the other player saw a demographic profile that revealed their race, age, gender, income and partisanship (Fig. 2). They also saw a demographic profile of the other player displaying these same variables (Fig. 2b). In Study 1, participants played against an out-race opponent, but were provided with information about this player’s partisanship, which was randomly assigned at the participant-level as either a Democrat or Republican (i.e., participants played against someone of the same partisanship across all three rounds). The opponent’s other profile attributes were randomly assigned in each round of play.² In Study 2, participants always played against a member of the opposing race and saw no partisan information. Finally, in Study 3, participants always played against a member of the opposing political party, and the race of the opposing player was randomly assigned at the participant level (i.e., either Black or White) and remained the same in every round.

With these modifications, the game unfolded like a standard trust game. Participants read instructions, saw three example rounds, and completed two comprehension questions.³ In an essential piece of the design, participants were randomly assigned to receive either a positive or negative result in the game. Those in the positive result condition received allocations of \$8, \$7 and \$8 (tripled to \$24, \$21 and \$24). This greatly exceeds the typical allocation of around \$4.50 per round made in previous studies (Iyengar and Westwood 2015; Carlin and Love 2018). However,

² Gender was either ‘Male’ or ‘Female’, age was drawn between 25 and 35, and income was drawn from four brackets: ‘\$30,000–\$39,999’, ‘\$40,000–\$49,999’, ‘\$50,000–\$59,999’, and ‘\$60,000–\$69,999.’

³ If they missed a comprehension question, participants were given the answer and asked the questions again. Those failing the questions three times were removed from the survey.

in a dose-response pre-test discussed in the supporting materials, we found that this average value is perceived by participants as a “negative” result. We therefore use larger values in the positive conditions to ensure participants perceive favorable intent, a pattern confirmed in manipulation checks embedded in these studies. In contrast to this positive group, those in the negative result condition were not allocated anything in each of the three rounds. Following each round, participants were told why the other player made their allocation, and for both the positive and negative result conditions, the explanation emphasized the social identity focused on in that study. This is designed to ensure the treatment is perceived as intended. For both the positive and negative result conditions we reported the same reasoning. For Study 1 and Study 2, this was race (round 1), race and income (round 2), and race (round 3). For Study 3, it was partisanship (round 1), partisanship and income (round 2), and partisanship (round 3).⁴ At the end of the game, participants saw a summary of the results for each round.

Finally, we measured out-group affect as our primary outcome. Studies 1 and 3 used attitudinal measures and participants reported their feelings toward several different groups in a battery of feeling thermometers. In Study 2, which instead used a behavioral measure, participants continued in the game as Player 1 and made allocations to four other players, allowing us to see how they interacted with out-group members when they could make the initial allocation. At the end of all three studies, participants received the value of their wallet as a cash bonus after applying a .05 multiplier.

The design offers precise, randomized control over the quality of an individual’s out-group interaction. It also enables participants to clearly understand the identity and motivation of the out-group member they interact with. Finally, although the modified trust game provides a new delivery mechanism, this approach to shifting out-group affect via interpersonal interactions appears elsewhere in the literature on affective polarization (e.g., Kuo et al. 2017; Suhay et al. 2017).

Evidence for the Parallel Updating of Racial and Partisan Affect

We implemented three versions of this experimental design. Study 1 (n = 1697; recruited via Qualtrics Panels)⁵ considers whether the parallel updating of racial and partisan affect occurs due to explicitly racial encounters, in this case delivered in the trust game via a positive or negative interaction with an out-race member. We measure the effects of this encounter with feeling thermometers in which respondents assessed racial and partisan out-groups. In Study 1 we also randomly assigned partisanship to the out-race players. We do so because, absent this information,

⁴ We added income as a reason to make our treatment less obvious.

⁵ The sample included quotas to approximate age and gender benchmarks from the American Community Survey.

participants might infer the party of the other player based on the racial cue, complicating the interpretation of our results.⁶

Study 2 further examines the implications of out-race interactions, but uses a behavioral outcome measure ($n = 1975$; recruited via Qualtrics Panels with the same criteria as Study 1). We anticipate that interactions with a member of a different race should also influence behavior toward members of a different political party (Iyengar and Westwood 2015; Carlin and Love 2018). This addresses concern that measures of out-party affect based on survey responses reflect partisan cheer-leading (Bullock et al. 2015) or a lack of constraining social norms when assessing out-partisans in surveys (Iyengar and Westwood 2015). As in Study 1, participants learned the race of the other player, always assigned as the out-race, as well as a randomly assigned gender, income and age for each profile. Unlike in Study 1, the profile did not include information about the partisanship of the other player, allowing us to assess whether similar findings occur when party is not mentioned.

To capture behavioral responses, we observe allocations in later rounds of the game. Recall that in the treatment rounds, participants received allocations from other players in the Player 1 role. In the additional rounds their role is reversed and they make the initial allocation as Player 1 to another Player 2. Participants in each of the rounds were given the full allocation of \$10 and a demographic profile of the contrived Player 2. In the first and second round, we offered a racial cue (white or black) and omitted partisanship. In the third and fourth round, we provided a partisan cue (Democrat or Republican) and no racial cue. As with the treatment rounds, we reported a randomly assigned gender, income, and age for each profile. Our outcome measure, the amount the participants offered to other players in these subsequent rounds, captures their incentivized, behavioral preferences for out-group cooperation.

Finally, Study 3 tests for the reverse of the pattern explored in Study 1 and Study 2: parallel updating based on partisan and not racial treatments ($n = 1824$; again recruited via Qualtrics Panels with the same criteria as the first two studies). The outcome in this study was once again a series of feeling thermometers. Unlike in the first two studies, individuals now played against a member of a different political party. For example, Democratic participants encountered three Republican players. Like in Study 1, we filled in information about the ancillary characteristics of the other player through random draws, in this case by randomly assigning the player's race at the participant level to be either Black or white.

Manipulation Checks

Before examining the main results, we first assess whether the experiments worked as intended. Across all three experiments, a crucial requirement is that

⁶ In our robustness tests we show that this additional information has no effect on our results, meaning that participants are making an implicit internal connection between race and party and are neither imputing or relying on explicitly provided information.

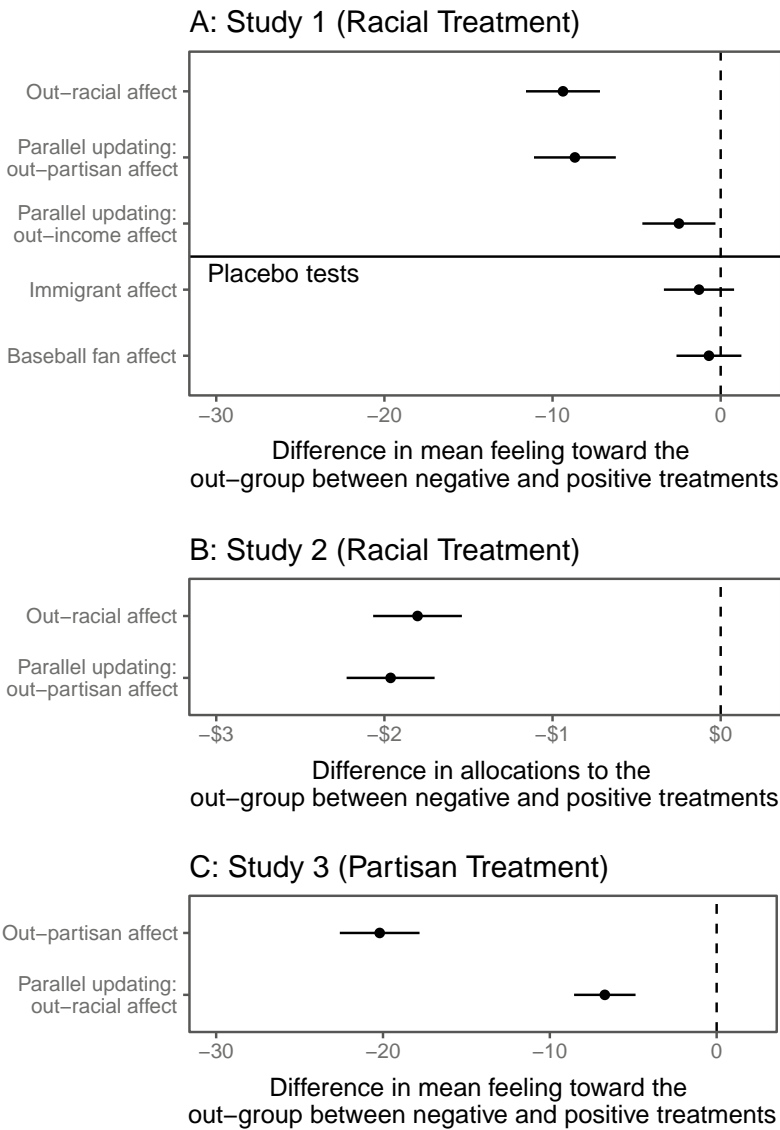


Fig. 3 Parallel updating across studies

participants perceived the treatments in the intended manner. To assess this, after each study, participants were asked, “Would you say that you were treated fairly or unfairly when playing the game?” and responded on a four-point scale that ran

from “Very Fairly” (1) to “Very Unfairly” (4). In Study 1, those in the positive results condition were significantly less likely to report they were treated unfairly ($M=1.72$) than their counterparts in the negative results condition ($M=3.18$; mean difference = 1.46 scale points, 95% Confidence Interval [1.38, 1.54]). Results were similar in Study 2, with a mean difference of 1.00 (95% Confidence Interval [.93, 1.07]). Likewise in Study 3 the mean difference in fairness was = 1.65 (95% Confidence Interval [1.58, 1.72]). Altogether this indicates the trust game interactions worked as intended.

Main Experimental Results

Study 1

We now move on to the substantive results of Study 1 and examine the effect of an interaction with an out-race member on subsequent allocations to other players of a different race. The top row of Fig. 3a shows that the interpersonal interactions in the game altered feeling thermometer assessments of the out-race in the expected manner. On average, those in the negative condition assessed the out-race at 52.48 points, while those in the positive condition placed the out-race at 61.86 points, a difference of -9.38 points (95% confidence interval [-11.58, -7.18]).

The second row of Fig. 3a tests for evidence of parallel updating of attitudes toward out-partisans when participants experience out-race interactions. When participants were treated positively by a member of the out-race in the initial rounds of the game, they subsequently reported more favorable views toward out-party members than when they were treated poorly by an out-race member. Those in the positive out-race interaction condition evaluated the out-party at 43.8 points on the feeling thermometer, while those in the negative out-race interaction condition evaluated the out-party at 35.1 points. The second row of Fig. 3a displays the difference between these two groups. Relative to those with a positive out-race interaction, those in the negative condition evaluated the out-party -8.67 points (95% confidence interval [-11.10, -6.24]) lower. Once again, an out-race interaction has roughly as large an effect on partisan affect as it does on racial affect, even though the game only focused around racial identity.

Our model posits a distinctive association between race and partisanship in American politics. We test for this by looking for parallel updating of another prominent social group: economic class (income). Income and race are less connected in the public mind (Jackman and Jackman 1983), so we expect substantially smaller amounts of parallel updating of attitudes toward income groups based on a racial treatment. Here, we code the in-group and out-group by family income split at the median household income in 2018. The third row of Fig. 3a shows the effect of the racial treatment on income, where the mean difference was -2.48 (95% confidence interval [-4.66, -0.31]). As expected, the effect was stronger for tightly connected identities (race and partisanship) and weaker for the less clearly connected cleavage of income.

Beyond this theoretically relevant group, it is possible those in the negative result treatment arm exited the game with a hostile attitude toward all groups, even those not referenced or connected to the attributes in the game. We allay this concern by including two placebo groups into the set of feeling thermometers: immigrants and baseball fans. When analyzing these outcomes we find no effect of the treatment onto affect toward immigrants (mean difference = -1.29 , 95% Confidence Interval $[-3.38, 0.80]$) between those in the positive result treatment arm ($M = 60.00$) and the negative treatment arm ($M = 58.71$). For baseball fans, there were also no differences (mean difference = -0.70 , 95% Confidence Interval $[-2.63, 1.23]$) between those in the positive result treatment arm ($M = 58.65$) and the negative treatment arm ($M = 57.95$).

Study 2

We extend the attitudinal results from Study 1 with a behavioral measure of parallel updating of race and partisan affect. The top row of Fig. 3b shows that in Study 2 those in the negative condition allocated \$3.17 to the out-race player. Those in the positive condition allocated \$4.97, a mean difference of -1.80 (95% confidence interval $[-2.07, -1.54]$). There is evidence of parallel updating of behavior towards the other party based on racial treatments. Figure 3b shows that average allocations to the out-party player in the negative results condition were \$2.72, while in the positive results condition, the allocation was \$4.68. The mean difference between treatments is $-\$1.96$ (95% confidence interval $[-2.22, -1.70]$).⁷

Study 3

Finally, Study 3 shows that parallel updating occurs both from race to partisanship and from partisanship to race. Figure 3c shows the expected movement on the treated group. The average out-party affect among those assigned to have a negative interaction with an out-partisan was 32.01 points. Meanwhile, in the positive interaction condition, mean out-party affect crossed the midpoint of the scale and reached 52.21 points, a mean difference of -20.20 points (95% confidence interval $[-22.59, -17.81]$).

In Studies 1 and 2, we found the effect of the racial treatment on partisan behavior to be approximately the same size as the effect of the racial treatment on racial affect. Here, we find a substantially larger effect of a partisan treatment on the partisan outcome measure (nearly twice as large) and a smaller effect on the racial outcome measure (77% of the size of the effect from race to partisanship), although we utilize a different outcome measure. Nevertheless, the effect of partisan treatments on race is still large and significant. Mean out-race affect in the negative (partisan)

⁷ We also collected feeling thermometers after the additional game grounds, though we believe they are not appropriate for comparison to Study 1 because of the additional time and interaction between treatment and measurement. Nonetheless, they show results that are substantively and significantly consistent with Study 1

treatment was 55.64 points, while mean out-race affect in the positive (partisan) treatment was 62.34 points, for a difference of -6.70 points (95% confidence interval $[-8.54, -4.86]$).

Parallel updating of racial attitudes when faced with partisan stimuli is smaller in magnitude than parallel updating of partisan attitudes when presented racial stimuli. We suspect this more modest result stems in part from the social norms constraining self-reports of hostile feeling towards out-race members (Iyengar and Westwood 2015), something that does not apply to feelings towards members of the other political party in Study 1.

Summary

Altogether, these three studies provide strong support for the parallel updating of racial and partisan affect. We show that interactions based only on race (Study 1 and Study 2) or only on partisanship cause citizens to update their attitudes towards both partisan and racial out-groups. This occurs for both attitudes toward out-groups (Study 1 and Study 3) and behavior towards them (Study 2). Moreover, this bidirectional relationship is not explained by various forms of expressive responding as there is no movement on unrelated placebo groups. Finally, we show the special nature of race and partisanship by demonstrating that parallel updating is conditional on how close groups are perceived, with significantly weaker results for the parallel updating between out-groups based on race and income. We now turn to a variety of tests that consider the robustness of these relationships.

Robustness

We now turn to a variety of tests that consider the robustness of this parallel updating of partisan and racial affect. In the sections that follow we first show our findings do not hinge on participants working to infer the unmentioned characteristics about the out-group member they interacted with, a pattern that would be inconsistent with our model of parallel updating. Second, we show that we are not merely observing a halo effect where all groups are assessed positively or negatively based on the assigned treatment, updating, like the interactions in the behavioral game, is restricted to out-groups. Third, we show that parallel updating occurs for both sorted and unsorted participants. Fourth, we show that parallel updating is not moderated by racial bias. Finally, we show that parallel updating occurs consistently across partisan and racial subgroups.

Parallel Updating Persists when Partisan Information is Available

In an attempt to break the link between racial treatments and out-party affect, we randomized the partisanship of each Player 1 assigned to each participant. The purpose of this analysis is to ensure the observed changes in partisan affect stem from the racial treatment alone and do not depend on the reported partisanship of the set

of Player 1s. Put another way, do partisans only update out-party attitudes when harmed/helped by an out-party member? Or, as we argue, are changes in partisan attitudes driven by reactions to the race of the opposing player in a process of parallel updating?

Our results show that participants are inattentive to the reported partisanship of Player 1 in the games. Instead, out-party attitudes are updated based on the result of the racial interaction. Specifically, there is never a main effect of Player 1 partisanship on out-group affect (see Appendix Table A7 for full regression results). We also find no interaction between the treatment and the randomly assigned partisanship of Player 1 for out-race affect ($\beta = -1.51$, 95% confidence interval $[-6.36, 3.33]$), or out-party affect ($\beta = -1.11$, 95% confidence interval $[-5.49, 3.28]$). This shows that any updates to out-race and out-party affect are not an artifact of the reported partisanship of the opposing player.

We further probe the lack of an effect of the partisanship of Player 1 on out-party affect by regressing out-group affect on co-partisanship. There is no significant effect of shared partisanship between the participant and each Player 1 on out-group affect (see Appendix Table A8 for full regression results). We find no interaction between co-partisanship and the treatments on out-race affect ($\beta = 0.27$, 95% confidence interval $[-4.11, 4.64]$) or out-party affect ($\beta = -4.52$, 95% confidence interval $[-9.34, 0.31]$). Moreover, the effect of the racial treatment on out-race affect remains ($\beta = 9.51$, 95% confidence interval $[6.45, 12.58]$), as does the parallel updating effect on out-party affect ($\beta = 11.13$, 95% confidence interval $[7.75, 14.51]$).

Results showing parallel updating of both racial and partisan attitudes from a racial treatment are robust to providing direct partisan cues. This is evidence that the effects we identify are due to links between race and partisanship in the minds of participants and not the provided partisan information.

Parallel Updating Occurs on Out-Group Attitudes

One concern with interpreting these findings is the potential presence of a “halo” effect, in which those who had positive interactions were subsequently more positive towards all types of groups. To address this alternative interpretation, we examine a different outcome variable, the difference in feeling thermometer ratings of different political parties (in-party minus out-party) and the difference in feeling thermometer ratings of different racial groups (in-race minus out-race). This difference out any baseline increase in affect towards all groups. More specifically, if participants in the positive condition are merely more positive towards all types of groups, we should see no significant difference on these two outcome measures.

In study 1, racial polarization decreased in the positive condition relative to the negative condition (mean difference = 9.20, 95% confidence interval $[6.79, 11.61]$). The mean difference in affective polarization between treatments was smaller, at 2.55 (95% confidence interval $[-0.62, 5.72]$). While in the anticipated direction where these interactions matter more for out-groups than for in-groups, this last difference does not reach statistical significance.⁸

⁸ We find less evidence of this concern in Study 2 and 3.

In study 2 participants did penalize both in-race and out-race players in the negative condition, but, the penalty was larger for the out-race (mean difference = \$0.60, 95% confidence interval [\$0.37, \$0.83]). Similarly, participants offered penalty both in-party and out-party players in the negative condition, but they were again much harsher toward the out-partisan (mean difference = \$0.56, 95% confidence interval [0.32, 0.81]). The similar magnitude of these two treatment effects once again indicates a strong linkage between race and party.

In study 3 we find that relative to the negative treatment, the positive treatment significantly reduced racial affective polarization (mean difference = 3.64, 95% confidence interval [1.43, 5.86]) and partisan affective polarization (mean difference = 13.92, 95% confidence interval [10.86, 16.98]). Our treatment moved both in-group and out-group measures, but had a much larger effect on out-group measures.

Parallel Updating Occurs Among Sorted and Unsorted

As we have previously discussed, the psychological framework we draw on suggests that interactions with one out-group can also shape affect towards another out-group. If this model holds, we anticipate similar degrees of parallel updating for those with sorted (i.e., white Republicans and Black Democrats) and those with cross-cutting (i.e., white Democrats and Black Republicans) identities on race and partisanship. However, we were clear in noting that this framework can imply an unintuitive pattern of shifts in affect among some individuals with cross-cutting identities in this area. It implies, for example, that improving a white Democrat's views of the out-race (Blacks) should lead to improved views of the out-party (Republicans). Here we consider whether this pattern holds or if, instead, there are attenuated or reversed patterns of parallel updating on racial and partisan affect among those with cross-cutting identities relative to their sorted counterparts.

For this test we look at the parallel updating effect (i.e., the consequences of the treatment for the identity that was not emphasized during the trust game interactions) across our three experiments. We do so separately for individuals who possess sorted social identities and those possessing cross-cutting identities. Figure 4a shows that in Study 1 parallel updating was similar for those with sorted out-group identities (mean difference = -8.09, 95% confidence interval [-11.24, -4.94]) and those with cross-cutting out-group identities (mean difference = -9.44, 95% confidence interval [-13.28, -5.61]). An interaction testing for a difference between those with sorted out-group identities and those with cross-cutting out-group identities is not significant ($\beta = -1.35$, 95% confidence interval [-6.40, 3.69]).

Similar to Study 1, Fig. 4b shows that in Study 2 the out-race treatment affected out-party allocations on the behavioral outcome to a similar degree for those with sorted out-group identities (mean difference = \$-1.95, 95% confidence interval [\$-2.24, \$-1.66]) and those with cross-cutting out-group identities (mean difference = \$-2.02, 95% confidence interval [\$-2.60, \$-1.43]). An interaction testing for a difference between those with and without sorted out-group identities shows no significant difference ($\beta = -0.06$, 95% confidence interval [-0.71, \$0.58]).

Finally, in Study 3 we again find that sorted out-groups are not necessary for parallel updating to occur (Fig. 4c). Indeed, although the parallel updating effect was

slightly larger for sorted out-groups relative to cross-cutting out-groups, the difference is not significant in an interaction model (mean difference = -4.62 , 95% confidence interval $[-9.39, 0.14]$). The positive treatment continued to increase out-race affect relative to the negative condition for those with sorted out-groups (mean difference = -8.94 , 95% confidence interval $[-11.92, -5.96]$) and those with cross-cutting out-groups (mean difference = -4.57 , 95% confidence interval $[-6.77, -2.36]$).

Placing these three tests together, we observe consistent evidence for the predictions based on conceiving of parallel updating as an “us against them” process where shifts in the views of one out-group generate simultaneous change in views towards the related out-group. We fail to see evidence of a process in which individuals with cross-cutting identities do not engage in parallel updating or take account of the empirical distributions of race and party when updating their views.

This pattern supports the idea that individuals are responding to their own internalized mappings between the out-race and the out-party and are not merely responding to identities that are sorted or aligned. But why, for example, would a white Democrat who is treated negatively (positively) by a black Player 1 update her affect toward Republicans? Consistent with past work (see Lodge and Taber 2013), this is because the updating process is peripheral and not deeply considered. Participants are reacting to someone who differs on one dimension (race), then reacting as if they differ on other dimensions (partisanship), and updating assessments of both out-groups simultaneously. As we show throughout our early analysis, this updating has real consequences on behavior and attitudes toward out-groups.

Parallel Updating Is Not Moderated by Racial Resentment

One concern is that parallel Updating of race to partisanship may be less general than we theorize and, in particular, might be conditional on the racial resentment of the respondent. To account for this possibility, we embedded four items from the racial resentment scale in our survey.⁹ These items were measured pre-treatment and were separated from the main component of the survey with a brand recognition distraction task. We bin a participant’s level of racial resentment into terciles for the interaction (Hainmueller et al. 2019).¹⁰ Using this measure, we show that racial resentment does not significantly moderate the relationship between our treatment and either out-party or out-race affect (see Appendix Table A9 for full regression results).

Parallel Updating Occurs Across Partisan and Racial Subgroups

Finally, we break out the effect by race and partisan subgroup to consider whether our effects are driven by a single racial or partisan subgroup in Study 1. Table 1 shows that we observe significant parallel updating in the expected direction for all values of race and party.¹¹ The effects are strongest for Blacks and Democrats, but are consistently large across all sub-groups in our sample.

⁹ See the supporting materials for details on the questions included in this scale.

¹⁰ Results are similar when using the raw measure.

¹¹ We show that this is also the case for Study 2 and Study 3 in the supporting materials.

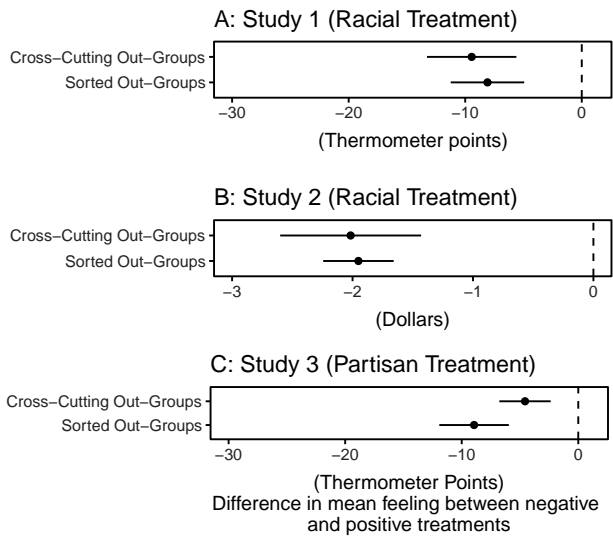


Fig. 4 Sorting, cross-cutting and parallel updating

Discussion and Conclusion

Using three experiments, we find evidence for the parallel updating of racial and partisan affect. This shows that partisanship and race are inseparably linked in American political behavior. These findings expand on previous considerations of the connection between race and partisanship in several ways. First, we find dynamic and interlinked shifts in out-group affect. This is an element missing from prior studies that, due to their focus on the consequences of shifts in elite-level politics, consider shifts in the relationship between race and party over a longer window time window (Carmines and Stimson 1989; Green et al. 2002; Valentino and Sears 2005; Tesler 2013; Kuziemko and Washington 2018). Second, we show that parallel updating is bidirectional—with movement happening from partisanship to race and race to partisanship. This is a departure from previous studies that have generally focused on the influence of racial attitudes on partisanship (see also, Engelhardt 2018)

A key contribution of our study is to show a link between examinations of racial animus and partisan affective polarization. Our evidence demonstrates that partisan affective polarization is not merely the consequence of a growing sense of partisanship as a social identity (Iyengar and Westwood 2015) or a greater alignment between other social identities and partisanship (Mason 2018). Although both contribute to this growing affective divide, we show that partisan affective polarization is also a consequence of the non-political racial interactions of Americans. Out-race interactions rapidly spill into assessments of the other political party even for those who are not sorted. Moreover, in a finding that broadens the implications of the growing body of research on affective partisan polarization (Iyengar et al. 2019), we find that partisan hostility also contributes to racial animosity. To understand

Table 1 Parallel updating by subgroup

| | <i>Dependent variable: Out-Party Affect</i> | | | |
|--------------------|---|-------------------------------|-------------------------------|-------------------------------|
| | White respondents | Black respondents | Democratic respondents | Republican respondents |
| | (1) | (2) | (3) | (4) |
| Positive treatment | 7.769*** (5.175, 10.364) | 15.563*** (8.931, 22.195) | 11.138*** (7.740, 14.535) | 6.275*** (2.834, 9.715) |
| Constant | 36.665*** (34.868, 38.461) | 24.531*** (19.698, 29.364) | 31.316*** (28.952, 33.681) | 38.721*** (36.315, 41.127) |
| Observations | 1,456 | 241 | 828 | 869 |

*Note**p<0.05; **p<0.01; ***p<0.001

partisan affective polarization, our results suggest, requires us to consider partisanship and race not only as related groups, but as inseparable in the minds of Americans.

Finally, our findings have implications for the study of political identity. First, they can potentially explain why partisan bias manifests in apolitical contexts—people have politics at mind and are unable to sever it from apolitical assessments. Second, they suggest a need to expand models of partisan affect to consider rapid parallel updating of affect from identities that are connected to partisanship. Theorizing of partisan identity as merely a consequence of sorting among other social identities assumes a slow-moving and unresponsive relationship that is not compatible with our results. Finally, this examinations reinforces how race and partisanship relate in American politics. At best, conceptualizing these two identities as separate is an incomplete account of how they operate. At worst, viewing these identities in a siloed manner understates the extent to which racial attitudes influence partisan attitudes and vice versa.

Our work should be assessed within a set of limitations. Methodologically, our experiments provide high internal validity, but in settings that are somewhat artificial. Further consideration of the real-world circumstances that can generate shifts in out-group affect of the type we study here is needed. Substantively, there are additional forms of identity and potential moderators of this relationship that are needed in future studies. Our initial investigation examines two of the most substantial cleavages in American society, race and partisanship, and focuses on racial and partisan groups that possess clear and defined in-groups and out-groups. Whereas Democrats and Republicans stand in opposition, it is less clear what the appropriate out-group is for pure independents (partisans, members of the parties but not leaners, etc.). need for additional work also extends to ethnic categories that overlap racial categories such as Latinos and to those who are multiracial. It is also possible the connection between race and partisanship may vary by geography (for example see Acharya et al. 2020), individually or by partisan strength. This paper lays a foundation for this future research into the nature of group affect.

Acknowledgements The authors thank Dan Butler, Justin Grimmer, Daniel Hopkins, Shanto Iyengar, Yphatach Lelkes, Jonathan Mummolo, and conference participants at UCSD, University of Pennsylvania and Facebook for helpful comments.

References

- Abramowitz, A. I., & Webster, S. (2016). The rise of negative partisanship and the nationalization of US elections in the 21st century. *Electoral Studies*, 41, 12–22.
- Acharya, A., Blackwell, M., & Sen, M. (2020). *Deep roots*. Princeton: Princeton University Press.
- Achen, C. H., & Bartels, L. M. (2016). *Democracy for realists: Why elections do not produce responsive government*. Princeton: Princeton University Press.
- Ahler, D. J., & Sood, G. (2018). The parties in our heads: Misperceptions about party composition and their consequences. *Journal of Politics*, 80(3), 964–981.
- Baron, A. S., & Banaji, M. R. (2006). The development of implicit attitudes: Evidence of race evaluations from ages 6 and 10 and adulthood. *Psychological Science*, 17(1), 53–58.
- Berelson, B., Lazarsfeld, P., & McPhee, W. (1954). *Voting: A study of opinion formation in a presidential campaign*. Chicago: University of Chicago Press.
- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1), 122–142.
- Bullock, J. G., Gerber, A. S., Hill, S. J., & Huber, G. A. (2015). Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, 10, 519–578.
- Campbell, A., Converse, P. E., Miller, W. E., & Stokes, D. E. (1960). *The American Voter*. New York: Wiley.
- Carlin, R. E., & Love, G. J. (2018). Political competition, partisanship and interpersonal trust in electoral democracies. *British Journal of Political Science*, 48(1), 115–139.
- Carmines, E. G., & Stimson, J. A. (1989). *Issue evolution: Race and the transformation of American politics*. Princeton: Princeton University Press.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407.
- Conover, P. J. (1988). The role of social groups in political thinking. *British Journal of Political Science*, 18(1), 51–76.
- Converse, P. E. (1964). The nature of belief systems in mass publics. In E. David (Ed.), *Ideology and its discontents* (pp. 206–261). New York: The Free Press.
- Davenport, L. (2016). Beyond Black and White: Biracial attitudes in contemporary US politics. *American Political Science Review*, 110(1), 52–67.
- Dawson, M. (1994). *Behind the mule: Race and class in African American politics*. Princeton: Princeton University Press.
- Egan, P. (2020). Identity as a dependent variable: How Americans shift their identities to better align with their politics. *American Journal of Political Science*, 64(3), 699–716.
- Engelhardt, A. M. (2018). Racial attitudes through a partisan lens. Working Paper.
- Fershtman, C., & Gneezy, U. (2001). Discrimination in a segmented society: An experimental approach. *Quarterly Journal of Economics*, 161, 351–377.
- Goldman, S. K., & Hopkins, D. (2019). Past place, present prejudice: The impact of adolescent racial context on white racial attitudes. *Journal of Politics*, <https://doi.org/10.1086/706461>.
- Green, D., Palmquist, B., & Schickler, E. (2002). *Partisan hearts and minds*. London: Yale University Press.
- Greene, S. (1999). Understanding party identification: A social identity approach. *Political Psychology*, 20(2), 393–403.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4.
- Habyarimana, J., Humphreys, M., Posner, D. N., & Weinstein, J. M. (2007). Why does ethnic diversity undermine public goods provision? *American Political Science Review*, 101(4), 709–725.

- Hainmueller, J., Mummolo, J., & Yiqing, X. (2019). How much should we trust estimates from multiplicative interaction models? Simple tools to improve empirical practice. *Political Analysis*, 27(2), 163–192.
- Halpern, D., & Rodriguez, P. (2018). Partisan representations: Partisan differences in semantic representations and their role in attitude judgments. In: Proceedings of the 40th Annual Conference of the Cognitive Science Society. pp. 445–450.
- Hetherington, M. J., & Rudolph, T. J. (2015). *Why Washington won't work: Polarization. Political trust and the governing crisis*. Chicago: University of Chicago Press.
- Huddy, L. (2001). From social to political identity: A critical examination of social identity theory. *Political Psychology*, 22(1), 127–156.
- Huddy, L., Mason, L., & Aarøe, L. (2015). Express partisanship: Campaign involvement, political emotion and partisan identity. *American Political Science Review*, 109(1), 1–17.
- Hutchings, V., & Valentino, N. (2004). The centrality of race in American politics. *Annual Review of Political Science*, 7(1), 383–408.
- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The origins and consequences of affective polarization in the United States. *Annual Review of Political Science*, 22, 129–146.
- Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*, 76(3), 405–431.
- Iyengar, S., & Westwood, S. J. (2015). Fear and loathing across party lines: New evidence on group polarization. *American Journal of Political Science*, 59(3), 690–707.
- Jackman, M. R., & Jackman, R. W. (1983). *Class awareness in the United States* (Vol. 343). Berkeley: University of California Press.
- Kam, C. D. (2007). Implicit attitudes, explicit choices: When subliminal priming predicts candidate preferences. *Political Behavior*, 29(3), 343–367.
- Kinder, D. R., & Sanders, L. M. (1996). *Divided by color*. Chicago: University of Chicago Press.
- Kinder, D. R., & Kalmoe, N. P. (2017). *Neither liberal nor conservative: Ideological Innocence in the American Public*. Chicago: University of Chicago Press.
- Kinder, D., & Ryan, T. (2017). Prejudice and politics re-examined: The political significance of implicit racial bias. *Political Science Research and Methods*, 5(2), 241–259.
- Klar, S. (2018). When common identities decrease trust: An experimental study of partisan women. *American Journal of Political Science*, 62(3), 610–622.
- Kuo, A., Malhotra, N., & Mo, C. H. (2017). Social exclusion and political identity: The case of Asian American partisanship. *Journal of Politics*, 79(1), 17–32.
- Kuziemko, L., & Washington, E. (2018). Why did the democrats lose the South? Bringing new data to an old debate. *American Economic Review*, 108(10), 2830–2867.
- Levendusky, M. (2018). Americans, not partisans: Can priming American national identity reduce affective polarization? *Journal of Politics*, 80(1), 59–70.
- Lipset, S. (1960). *Political man: Social bases of politics*. New York: Doubleday.
- Lodge, M., & Taber, C. S. (2013). *The rationalizing voter*. Cambridge: Cambridge University Press.
- Margolis, M. (2018). *From politics to the pews: How partisanship and the political environment shape religious identity*. Chicago: University of Chicago Press.
- Mason, L. (2015). 'I Disrespectfully agree': The differential effects of partisan sorting on social and issue polarization. *American Journal of Political Science*, 59(1), 128–145.
- Mason, L. (2016). A cross-cutting calm: How social sorting drives affective polarization. *Public Opinion Quarterly*, 80(S1), 351–377.
- Mason, L. (2018). *Uncivil agreement: How politics became our identity*. Chicago: University of Chicago Press.
- Mason, L., & Wronski, J. (2018). One tribe to bind them all: How our social group attachments strengthen partisanship. *Advances in Political Psychology*, 39(1), 257–277.
- McConnell, C., Margalit, Y., Malhotra, N., & Levendusky, M. (2018). The economic consequences of partisanship in a polarized Era. *American Journal of Political Science*, 62(1), 5–18.
- Mo, C. H. (2015). The consequences of explicit and implicit gender attitudes and candidate quality in the calculations of voters. *Political Behavior*, 37(2), 357–395.
- Pérez, E. (2016). *Unspoken politics: Implicit attitudes and political thinking*. Cambridge: Cambridge University Press.
- Pérez, E. (2010). Explicit evidence on the import of implicit attitudes: The IAT and immigration policy judgments. *Political Behavior*, 32(4), 517–545.

- Quillian, M. R. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral Science*, 12(5), 410–430.
- Robison, J., & Moskowitz, R. (2019). The group basis of partisan affective polarization. *Journal of Politics*, 81(3), 1075–1079.
- Rothschild, J. E., Howat, A. J., Shafraneck, R. M., & Busby, E. C. (2019). Pigeonholing partisans: Stereotypes of party supporters and partisan polarization. *Political Behavior*, 41(2), 423–443.
- Ryan, T. (2017). How do indifferent voters decide? The political importance of implicit attitudes. *American Journal of Political Science*, 61(4), 892–907.
- Schickler, E. (2016). *Racial realignment: The transformation of American liberalism, 1932–1965*. Princeton: Princeton University Press.
- Schuman, H., Steeh, C., Bobo, L., & Krysan, M. (1997). *Racial attitudes in America: Trends and interpretations*. Cambridge: Harvard University Press.
- Sears, D. O., & Funk, C. L. (1999). Evidence of the long-term persistence of adults political predispositions. *Journal of Politics*, 61(1), 1–28.
- Suhay, E., Bello-Pardo, E., & Maurer, B. (2017). The polarizing effects of online partisan criticism: Evidence from two experiments. *International Journal of Press*, 23(1), 95–115.
- Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American*, 223(5), 96–102.
- Tesler, M. (2013). The return of old-fashioned racism to White Americans' partisan preferences in the early Obama era. *Journal of Politics*, 75(1), 110–123.
- Tesler, M. (2016). *Post-racial or most-racial? Race and politics in the Obama Era*. Chicago: University of Chicago Press.
- Theodoridis, A. G. (2017). Me, myself and (I), (D), or (R)? Partisanship and Political cognition through the lens of implicit identity. *Journal of Politics*, 79(4), 1253–1267.
- Valentino, N. A., & Sears, D. O. (2005). Old times there are not forgotten: Race and partisan realignment in the contemporary South. *American Journal of Political Science*, 49(3), 672–688.
- Westwood, S. J., Peterson, E., & Lelkes, Y. (2019). Are there still limits on partisan prejudice? *Public Opinion Quarterly*, 83(3), 584–597.
- White, I. K., & Laird, C. N. (2020). *Steadfast democrats*. Princeton: Princeton University Press.
- White, I. K., Laird, C. N., & Allen, T. D. (2014). Selling out? The politics of navigating conflicts between racial group interest and self-interest. *American Political Science Review*, 108(4), 783–800.
- Whitt, S., & Wilson, R. K. (2007). The dictator game, fairness and ethnicity in postwar Bosnia. *American Journal of Political Science*, 51(3), 655–668.
- Wyer, J., Robert, S., & Srull, T. K. (2014). *Advances in social cognition: A dual process model of impression formation* (Vol. I). Hove: Psychology Press.

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