Unraveling a "Cancel Culture" Dynamic: When, Why, and Which Americans Sanction Offensive Speech

Nicholas C. Dias, University of Pennsylvania
James N. Druckman, University of Rochester
Matthew S. Levendusky, University of Pennsylvania

There is a growing belief that many Americans shun, ostracize, or "cancel" those they dislike or those who make disagreeable statements. Yet, no empirical work has explored the prevalence or motives of this type of sanctioning or how Americans perceive it. Using a nationally representative survey with an embedded conjoint experiment, we find that Americans vastly overestimate how likely other people—especially out-partisans—are to cancel others. Nevertheless, they accurately perceive what motivates others to cancel: disagreeable and offensive statements, not disliked speakers. Additionally, we find that Democrats and Republicans are similarly motivated to cancel, although canceling behavior out in the world may more commonly come from Democrats. Our findings highlight how "cancel culture" could limit harmful speech but encourage self-censorship and partisan animus. They also reveal the normative fault lines underlying debates about free speech in contemporary society.

On July 7, 2020, over 150 writers and academics published a letter in *Harper's Magazine*. It warned of a "vogue for public shaming and ostracism" that was "spreading more widely in our culture" and increasingly constricting the "free exchange of information and ideas, the lifeblood of a liberal society" (Ackerman et al. 2020). Less than a week after the letter circulated widely, a second letter with over 160 signatories criticized the first as a "caustic reaction to a diversifying industry . . . that's starting to challenge institutional norms that have protected bigotry." The second letter's authors asserted that the *Harper's* letter ignored "how marginalized voices have been silenced for generations" (*The Objective* 2020).

This exchange encapsulates a growing debate about the practice of socially or economically sanctioning those who make offensive statements, otherwise known as "canceling" (see Norris 2023; Romano 2020). In recent years, "cancel culture" has been the focus of intense media attention: prominent outlets such as *The New York Times* bemoan the phenomenon in op-eds, and they report on the canceling of public figures ranging from James Cordon to J. K. Rowling and even Mr. Potato Head (e.g., The Editorial Board 2022). Some have suggested that this behavior is not limited to a narrow swath of committed activists but is widespread among the American public (e.g., Ackerman et al. 2020).

For all this popular attention, little empirical evidence exists to inform the debate. We begin to compile such evidence by focusing on one dimension of cancel culture: the decision to "cancel" or sanction others.[1] Specifically, we explore (1) how Americans cancel others in reality, as best as can be assessed,

Nicholas C. Dias (niccdias@sas.upenn.edu) is a PhD candidate in the Department of Political Science and Annenberg School for Communication at the University of Pennsylvania, Philadelphia, PA 19104. James N. Druckman (jdruckma@UR.Rochester.edu) is a professor in the Department of Political Science, University of Rochester, Rochester, NY 14627. Matthew S. Levendusky (mleven@upenn.edu) is a professor in the Department of Political Science at the University of Pennsylvania and the Stephen and Mary Baran Chair in the Institutions of Democracy at the Annenberg Public Policy Center, Philadelphia, PA 19104.

Support for this research was provided by the Annenberg Public Policy Center. The preregistration for this article's experiment is available on AsPredicted (https://aspredicted.org/y3mn5.pdf). Replication files are available in the *JOP* Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An online appendix with supplementary material is available at https://doi.org/10.1086/733004.

1. Sanctioning has clear implications for another aspect of cancel culture, self-censoring, which stems from a fear of social isolation (Noelle-Neumann [1980] 1993). We use "cancel" to refer to the sanctioning part of cancel culture while noting there are additional elements to the "culture."

Published online March 10, 2025.

The Journal of Politics, volume 87, number 2, April 2025. © 2025 Southern Political Science Association. All rights reserved. Published by The University of Chicago Press for the Southern Political Science Association. https://doi.org/10.1086/733004

and (2) Americans' perceptions of others' canceling behavior. How frequently do Americans sanction others for offensive speech? Why do they do it? How do Americans perceive the prevalence and motives of canceling? We derive hypotheses about these questions and test them using a survey that contained a preregistered conjoint experiment.[2]

We educe four key findings. First, Americans overestimate—by at least a factor of two—how often their fellow citizens cancel others. In particular, partisans inflate how willing members of the other political party are to cancel speech that offends them. Second, source biases play a minor role in decisions to cancel others: generally, citizens do not care who makes offensive statements—be it a member of the other party, someone of another race/ethnicity, or an elected official. Instead, citizens cancel offensive statements that run counter to their ideological leanings, regardless of who says them. It is content, more so than identity, that drives decisions to cancel. Third, citizens accurately perceive that others cancel based on the content of what people say, even if they misperceive who cancels and how often. Fourth, we find limited evidence that these dynamics differ between Democrats and Republicans. Despite the popular perception that Democrats cancel more than Republicans, in many cases, members of both parties are similarly unlikely to cancel others.

These findings align with past work that suggests that media and psychological biases lead individuals to dramatically misperceive the behavior of others, especially members of the opposing political party (Ahler and Sood 2018; Druckman et al. 2022). More importantly, they provide a much-needed empirical foundation for the debate about what canceling behavior represents for America's political culture. Citizens' tendency to overestimate the prevalence of canceling behavior suggests that the threat of self-censorship is acute (Noelle-Neumann [1980] 1993) and may jeopardize America's traditional commitment to free speech. Moreover, the widespread misperception that those from the other party cancel frequently could exacerbate partisan animus (Iyengar et al. 2019) and limit contact between citizens who disagree (Carlson and Settle 2022).

However, instances of canceling cannot be reduced to disingenuous partisan attacks, and citizens recognize that they are unlikely to be canceled just because of who they are. Instead, canceling behaviors are intended to counter what citizens perceive as harmful or threatening speech. As such, this behavior is arguably laudable, especially when "canceled" statements attack those who are already threatened or marginalized (i.e., racial minorities, gender nonconforming people). Ultimately, whether individuals view canceling as salubrious or deleterious depends on how they weigh competing values: a classically liberal commitment to unfettered speech versus a desire to protect marginalized populations in the public sphere.

FREE SPEECH NORMS, CANCEL CULTURE, AND SOCIAL SANCTIONS

Political scientists have spent decades studying citizens' willingness to allow "the expression of ideas or interests that one opposes," also known as political tolerance (Sullivan, Pierson, and Marcus 1982, 2). This work typically focuses on whether individuals support the government's efforts to limit a "least liked" group's political and civil rights, such as the right to hold a public rally (for a review, see Gibson 2011).[3] A parallel body of scholarship explores the regulation of hate speech, focusing again on governmental efforts to ban or regulate it (Strossen 2018). The overarching finding of these studies is that Americans generally oppose government infringements on even the most offensive speech, particularly when compared to citizens of other democratic nations (Wike 2016).

Government regulation of speech clearly matters, as illustrated by a recent spate of laws banning public school teachers from discussing racism, gender identity, and sexual orientation (Greene 2022; Jones and Franklin 2022). Yet, focusing only on government action misses an equally crucial piece of the puzzle: how Americans constrain one another's speech informally via social and economic sanctions (Balkin 2018; Schauer 2010). Today, speech spreads with unparalleled speed and reach via social media platforms. Because these platforms often have lax rules surrounding what cannot be said (Gillespie 2018), citizens now set many, if not most, of the boundaries around acceptable speech (Waisbord 2020).

Indeed, some argue that the apparent rise in canceling behavior signals that Americans' preference for unfettered speech is shifting in favor of one that balances free expression with social justice (Chong, Citrin, and Levy 2024). On the one hand, defenders of canceling behavior argue that "canceled" statements are often prejudicial and deter members of marginalized groups from participating in civic discourse (e.g., McGowan 2022). On the other hand, critics argue that canceling behaviors belie faux outrage and reflect a desire to disqualify members of the opposing political party ("outpartisans") from participating in civic discourse (e.g., Trump 2020).

2. See https://aspredicted.org/y3mn5.pdf for the preregistration.

3. This approach cannot easily distinguish censorship of disliked groups from censorship of disagreeable ideas. Moreover, focusing on censorship of a "least liked" group may hide intolerance of more mainstream groups, such as political parties.

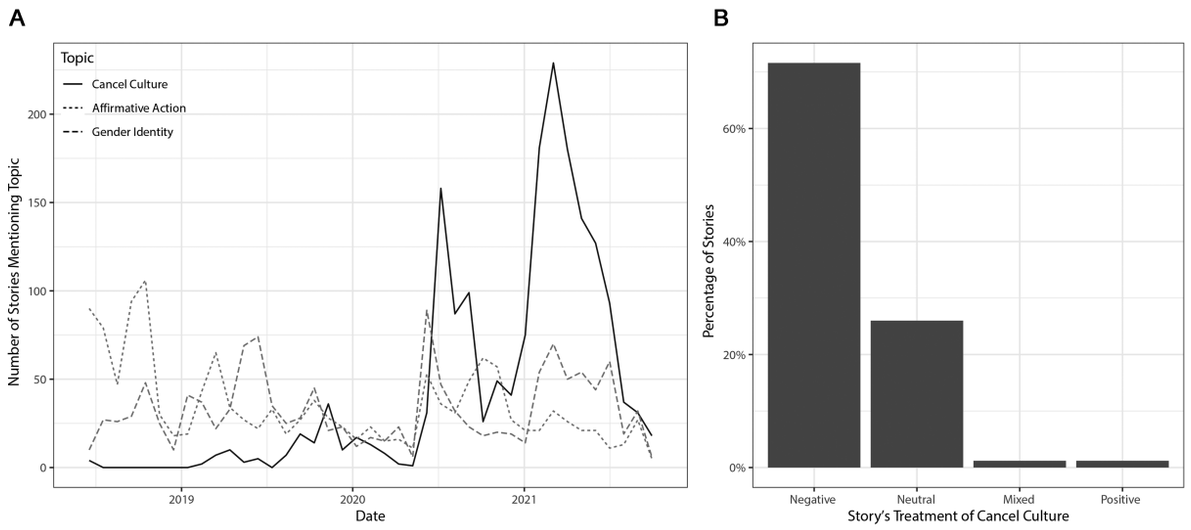

Figure 1. News coverage of cancel culture is voluminous and negative. Panel A depicts the number of stories or television programs published by *The New York Times*, *The Wall Street Journal*, CNN, MSNBC, and Fox News on each topic from June 15, 2018, to September 24, 2021. Data points represent 30 day windows. Panel B depicts the frequency with which these stories discussed cancel culture in different tones, as determined by a hand-coded random subset of 250 stories.

Motivations for canceling aside, many scholars argue that more-than-minimal constraints on speech jeopardize citizens' ability to openly evaluate all ideas to ensure that the best and truest guide society (Mill [1859] 2003; Scudder 2020). While "cancel culture" is a recent phenomenon (Romano 2020), scholars have long worried that the fear of social sanctions could cause citizens to self-censor.[4] For example, the spiral of silence theory posits that individuals censor themselves for fear of sanctions or losing social status, if they perceive themselves to hold an unpopular opinion (Noelle-Neumann [1980] 1993, 42–50). While self-censorship is a vitally important topic, we focus on the informal sanctions that can lead to self-censorship. We view our study as just one part of a broader research program about cancel culture, which should ultimately explore both sanctioning and self-censoring.

WHO CANCELS AND HOW OFTEN DO THEY DO SO?

As with many political phenomena, we expect the media to shape people's behaviors and beliefs about canceling (Levendusky and Malhotra 2016; Wilson, Parker, and Feinberg 2020). Along these lines, figure 1A shows that the two most widely circulated newspapers in the United States (*The Wall Street Journal* and *The New York Times*) and the three most-watched cable news channels (CNN, MSNBC, and Fox News) together published more than 1,760 stories about "cancel culture" between June 2018 and September 2021. By comparison, the widely discussed topics of "affirmative action" and "gender identity" were used in roughly 1,440 and 1,350 stories, respectively, during the same period. While coverage of cancel culture waxed and waned, often in response to high-profile canceling incidents, it had clearly become a ubiquitous topic by mid-2020: An average of at least 51 stories per month discussed cancel culture, overtaking coverage of other topics and signaling its prominent place in public discourse.

Fascinatingly, not only did the outlets frequently discuss cancel culture, but they also consistently framed it in very specific ways. Nearly all stories discussed cancel culture in the abstract, quoting politicians' complaints about canceling behavior (e.g., President Trump condemning it as far-left fascism; see Trump 2020). Fewer stories reported on a small number of prominent instances of canceling (e.g., the reaction to Mr. Potato Head becoming gender-neutral; see Viser 2021). And virtually no stories described instances of ordinary citizens canceling others. Moreover, canceling was typically discussed with a negative valence. In figure 1B, we present the results of hand-coding the tone of 250 randomly selected stories (see app. A for coding details). Approximately 72% of stories about cancel culture framed the issue in negative terms.

Two hypotheses follow from this media coverage. First, given the adverse media coverage of canceling behaviors, we suspect that most Americans perceive them negatively and are thus unlikely to cancel others. This aligns with the reality that most Americans dislike the political conflict inherent in canceling (Klar and Krupnikov 2016) and spend little time engaging with politics, even on social media (Wojcieszak et al. 2022).

4. Cancel-culture discussions can also be seen as an extension of debates over political correctness in the 1990s (Shapiro et al. 2021) and others dating back decades, if not centuries (Mishan 2020).

Canceling is not widespread (H1). A minority of Americans cancel others.[5]

Second, we expect that the deluge of media coverage will lead many citizens to assume that canceling behaviors are more prevalent than they are. When estimating how often an event occurs, people rely on the speed and ease with which the topic comes to mind (Tversky and Kahneman 1973). The news media's coverage of canceling, especially high-profile cases of it, therefore skew citizens' perceptions of its frequency. Consequently, there should be a gap between how much citizens actually cancel and how much they perceive that others cancel.

Canceling is overestimated (H2). Americans will perceive canceling to occur at significantly higher rates than it actually does, all else constant.

The overwhelmingly negative coverage of cancel culture has further implications. People like to believe that they are better than others and consequently perceive the world in biased, self-enhancing ways (Hoorens 1993). Given that the media portrays canceling behaviors negatively, Americans may be inclined to believe that others cancel more than they do, as this belief boosts their self-esteem. This disjuncture should be especially pronounced with regard to out-partisans: Because partisanship is increasingly central to Americans' political identities, Americans are incentivized to maximize the status of their own political party (the in-party) and minimize that of the other major party (the out-party; see Iyengar et al. 2019). Americans therefore are likely to think that others—and especially out-partisans—cancel more than they do.

Canceling is (mostly) for out-partisans (H3). Americans will perceive that in-partisans cancel more than themselves, all else constant. Americans will perceive that out-partisans cancel more than in-partisans, all else constant.

WHY DO AMERICANS CANCEL?

As important as how often others cancel is why they cancel at all. Commentators have offered competing narratives of what motivates canceling behavior. Defenders of canceling behavior argue that it is an appropriate response to statements that dehumanize and disempower marginalized groups, such as transphobic or racist remarks (McGowan 2022). Others claim that canceling behaviors represent an attempt by one partisan side to exclude the other from civic discourse, perhaps stemming from partisan animus (Trump 2020).

This debate has important implications for understanding the normative implications of canceling behavior. Regardless of its motivation, canceling may undermine the fundamental liberal values of free speech (Strossen 2018). Yet, with the exception of extremist groups such as neo-Nazis, canceling a speaker based on their identity additionally amounts to questionable discrimination or bias.[6] If instead canceling occurs in response to offensive statements, regardless of who says them, canceling would at least imply an effort to counter the perceived harm of offensive speech.[7] What, then, leads citizens to cancel others?

When deciding whether to cancel, citizens may indeed attend to a speaker's political identity—namely, their partisanship. Specifically, citizens may be more likely to sanction out-partisans, regardless of what they say. Partisans display less sensitivity to uncivil comments made by in-partisans (Druckman et al. 2019) and support news outlets that criticize the out-party but not their in-party (Lelkes and Westwood 2017). This type of partisan bias would align with a large literature showing that individuals are intolerant toward disliked groups, including the out-party (Iyengar et al. 2019). Additionally, this source dynamic may extend to groups stereotypically associated with the parties. Citizens possess strong stereotypes about the racial composition of the parties: Republicans are white, and Democrats are nonwhite (e.g., Ahler and Sood 2018). Unless told otherwise, individuals will likely infer that Black speakers are Democrats and white speakers are Republicans. Consequently, Democrats will be more likely to cancel Republicans and white speakers, whereas Republicans will be more likely to cancel Democrats and Black speakers.

Cancelers sanction disliked speakers (H4). Americans will cancel those from the out-party and racial groups stereotypically associated with the out-party more than those from the in-party and its associated racial groups, all else constant.

A speaker's occupation or social role may also influence whether they are canceled. Citizens have different expectations

5. This hypothesis was not explicitly preregistered, though it was strongly implied by our preregistration. We offer this hypothesis to structure the presentation of our findings.

6. Extremist identities are outside the scope of what we consider here. Instead, we focus on more conventional identities, most notably partisanship and race.

7. Of course, determining when statements may be justifiably canceled is normatively complicated. We set that issue aside for now and return to it in the conclusion.

of what someone says based on their visibility and ability to influence others (Chung and Moon 2016). For example, Americans may be especially likely to punish an elected official who says something offensive, given strong norms of elite civility in public discourse (Jamieson and Hardy 2012). Likewise, because celebrities and college faculty can powerfully influence their fans and students (Jackson and Darrow 2005), citizens may be more critical of what they say. Americans may therefore be more likely to cancel public figures than private ones.

Cancelers sanction public figures (H5). Americans will be more likely to cancel public figures (elected officials, professors, celebrities) than private figures, all else constant.

Perhaps most obviously, what a person says—not just who they are—likely influences how individuals react to others' speech. Citizens care about the content of others' speech and react negatively to speech that attacks or offends, even when it comes from in-partisans (Costa 2021). In separate pretest data (reported in app. B), we found that people were far more likely to cancel statements they perceived as offensive. Moreover, a statement's offensiveness largely reflected whether the statement was ideologically disagreeable to the respondent. As such, we predict that Democrats will cancel right-leaning statements, and Republicans will cancel left-leaning ones.

Cancelers sanction disagreeable speech (H6). Americans will be more likely to cancel statements that contradict their party's ideology, relative to statements that affirm this ideology, all else constant.

Beyond what actually motivates canceling behavior, it is important to ask how people perceive the motivations behind canceling. Different perceptions about what causes sanctions likely lead citizens to self-censor in distinct ways (Noelle-Neumann [1980] 1993). If citizens believe they will be canceled because of who they are, they may be likely to self-censor a broad swathe of political speech (offensive or not) or avoid political conversations entirely. However, if citizens perceive canceling behavior as triggered by a well-defined set of offensive statements, they may feel comfortable discussing most topics. Given the many suggestions that media have made about why citizens cancel, it is unclear how citizens will perceive these motives.

RQ1. How do Americans see canceling motives? Will Americans perceive canceling behavior to be motivated by source or content?

Finally, an obvious point of interest is whether Republicans and Democrats differ in how they perceive and engage in canceling behaviors. One might expect differences, given rhetoric from the political right that suggests Democrats are the primary perpetrators of canceling behavior (Fahey, Roberts, and Utych 2023). Indeed, in his infamous speech at Mount Rushmore, President Trump remarked, "[T]here is a new far-left fascism that demands absolute allegiance. If you do not speak its language . . . you will be censored, banished, blacklisted, persecuted and punished" (Trump 2020). Recent surveys suggest that such rhetoric affects how the public sees canceling behavior (e.g., Vogels et al. 2021). Thus, we will note partisan differences (or the lack thereof) as we discuss our various findings.[8]

SURVEY AND EXPERIMENTAL DESIGN

To examine our hypotheses and research questions, summarized in table 1, we fielded a survey containing a preregistered conjoint experiment. Surveys are crucial to understanding the prevalence of canceling. As mentioned, news media stories focus on high-profile incidents of canceling behavior that may not represent the general phenomenon. Furthermore, many types of canceling behavior (which we discuss below) cannot easily be observed outside of a survey context. And while canceling often occurs via social media, the unrepresentativeness of social media users (Mislove et al. 2021) and limits on data access would similarly skew conclusions.

Our data come from a sample of 1,752 American adults from Bovitz Forthright, a high-quality online panel.[9] The sample was quota-matched to represent American adults on age, gender, education, Census region, and race (see app. C for sample demographics). These respondents are recruited via mail campaigns based on address-based probability sampling, as well as via online ads, and their data have been used extensively in political science (e.g., Druckman et al. 2019). We collected our data between September 24 and October 4, 2021.

The survey began by asking respondents if they had ever engaged in six different canceling behaviors (described below) in the real world. In addition, respondents indicated which (if any) of the behaviors they had seen others do and who (i.e., Democrats or Republicans) they had seen do them. This allows us to assess how often individuals have canceled in the past, relative to their perceptions of what others have done. To avoid cueing respondents, we did not call any of these activities "canceling."

8. We preregistered hypotheses about partisan differences. Due to space constraints, we defer our discussion of them to app. G.

9. We drop pure Independents ($n = 296$) from analyses that examine partisanship as a predictor.

Table 1. Summary of Hypotheses and Research Questions

Canceling Prevalence	
Canceling is not widespread (hypothesis 1)	A minority of Americans cancel others.
Canceling is overestimated (hypothesis 2)	Americans will perceive canceling to occur at significantly higher rates than it actually does, all else constant.
Canceling is (mostly) for out-partisans (hypothesis 3)	Americans will perceive that in-partisans cancel more than themselves, all else constant. Americans will perceive that out-partisans cancel more than in-partisans, all else constant.
Canceling Motives	
Cancelers sanction disliked speakers (hypothesis 4)	Americans will cancel those from the out-party and racial groups stereotypically associated with the out-party more than those from the in-party and its associated racial groups, all else constant.
Cancelers sanction public figures (hypothesis 5)	Americans will be more likely to cancel public figures (elected officials, professors, celebrities) than private figures, all else constant.
Cancelers sanction disagreeable speech (hypothesis 6)	Americans will be more likely to cancel statements that contradict their party's ideology, relative to statements that affirm this ideology, all else constant.
How do Americans see canceling motives? (research question 1)	Will Americans perceive canceling behavior to be motivated by source or content?

Identifying canceling behaviors is a challenge, as there is no consensus on what constitutes canceling or sanctioning someone. We turned to two data sources to develop our list of canceling behaviors. First, we reviewed news stories that mentioned "cancel culture" to see what activities they describe (see app. D). Second, we conducted an open-ended pretest wherein respondents described how they had responded to, or had seen others respond to, disagreeable or offensive speech. From this, we formulated our list of six behaviors: criticizing the speaker on social media, complaining to the speaker's employer, boycotting the speaker's employer or merchandise, boycotting or protesting at an event where the speaker is participating, reporting the speaker to (or trying to ban the speaker from) a social media site, and publishing the speaker's personal information online (i.e., "doxing").[10] Question wordings are provided in appendix E.

Although not exhaustive, this list encompasses potentially anodyne activities (e.g., social media criticism) and more severe actions (e.g., doxing). These behaviors largely involve clear efforts to silence or ostracize others, consistent with our discussion above. The one possible exception is social media criticism, which could constitute ostracization or counter-speech, depending on its content and intention. This reflects the ambiguous reality of cancel culture. As Rauch (2021, 217–20) notes, the line between counter-speech and canceling is blurry, so we opted to study an inclusive list of behaviors. Our findings are consistent across canceling behaviors, as we discuss below, and dropping any given behavior does not qualitatively change our results (see app. F).

After respondents had indicated whether they or others had canceled anyone in the real world, they participated in a conjoint experiment. The experiment began by asking respondents to read four hypothetical scenarios wherein speakers made potentially offensive statements. For each scenario, respondents reported how likely they would be to engage in each of our six behaviors on a four-point scale ranging from "not at all likely" to "very likely." As summarized in table 2, each scenario varied the speaker's partisanship or race, the speaker's role, and what the speaker said. This allows us to test our hypotheses about what drives canceling behavior.

Our "cancelable" statements focused on race and gender identity, given their centrality to actual instances of canceling (Clark 2020). We identified these cancelable statements using Google News (see app. D), simplifying them for brevity

10. Although important, we excluded deplatforming from our list, as it is a sanction taken by social media companies and not individuals.

Table 2. Conjoint Experiment Attributes and Levels

Speaker identity	Democratic, Republican, Black, white
Speaker role	Elected official, professor, celebrity, voter, university student
Statement ideology	Left-leaning: • The Founding Fathers were racist. • America is a racist nation. • All police are bad. • Schools should require students to learn about transgender life and why it is normal. Right-leaning: • Confederate statues are about America's heritage and are not racist. • Some races are less intelligent than others. • Athletes who kneel during the National Anthem should be kicked off their teams. • There is no such thing as transgender, only male and female.

and ensuring an equal number of left-leaning and right-leaning statements. For example, one left-leaning statement read "The Founding Fathers were racist," while a similar right-leaning statement read "Confederate statues are about America's heritage and are not racist."[11] While each statement is offensive to some, most fall short of hate speech (that is, they do not contain slurs or threats of violence). Attitudes toward hate speech are important but distinct from our purpose here (see Rasmussen 2022). One likely exception is the statement "some races are less intelligent than others," which many (including the authors) would consider hate speech. Whether this statement is right-leaning is debatable. We included it in this classification here given its association with The Bell Curve (Herrnstein and Murray 1994).

We next presented respondents with four additional, identically randomized scenarios wherein speakers made potentially offensive statements. This time, we asked respondents how others would respond. In the first three of these scenarios, respondents predicted which canceling behaviors out-partisans would likely engage in (e.g., what a Democratic respondent thought Republicans would do) using the same response scale as above. For the fourth and final scenario, respondents predicted what in-partisans would do. We compare the results of these third-person conjoints (which gauge perceptions of why others cancel) to the first-person conjoints (which gauge why respondents cancel) to assess whether respondents misperceive the motives for canceling behavior.

AMERICANS OVERESTIMATE THE PREVALENCE OF CANCELING

We start by comparing the extent to which Americans actually engage in canceling behavior with their perceptions of the prevalence of these behaviors. In figure 2A, we compare the percentage who report having engaged in each canceling behavior in the real world to the percentage who report having seen others engage in that behavior in the real world.

These results show a dramatic gap between what citizens themselves do and what they think others do. Consistent with "canceling is not widespread" (hypothesis 1), a minority of respondents engage in any canceling behavior. Furthermore, as our "canceling is overestimated" hypothesis (hypothesis 2) predicts, respondents report that other citizens are more likely to engage in each canceling behavior than they themselves are. These misperceptions are quite dramatic: People overestimate canceling behavior by a factor of two or more. For example, while 7% of respondents had boycotted or protested an event, 37% had seen others do so. In the most striking example, respondents were 10 times more likely to have seen someone else dox someone's identity than to have done it themselves. Clearly, people's perceptions of canceling behaviors' prevalence far outpace the number who actually engage in these behaviors.

After asking respondents whether they had seen others engage in our six canceling behaviors in the real world, we asked respondents whether those they saw engaging in those

11. One could argue that these statements discourage members of marginalized groups from participating in civic discourse (Munger 2017) and thus constitute a kind of canceling in themselves. This is an important extension to consider in future work.

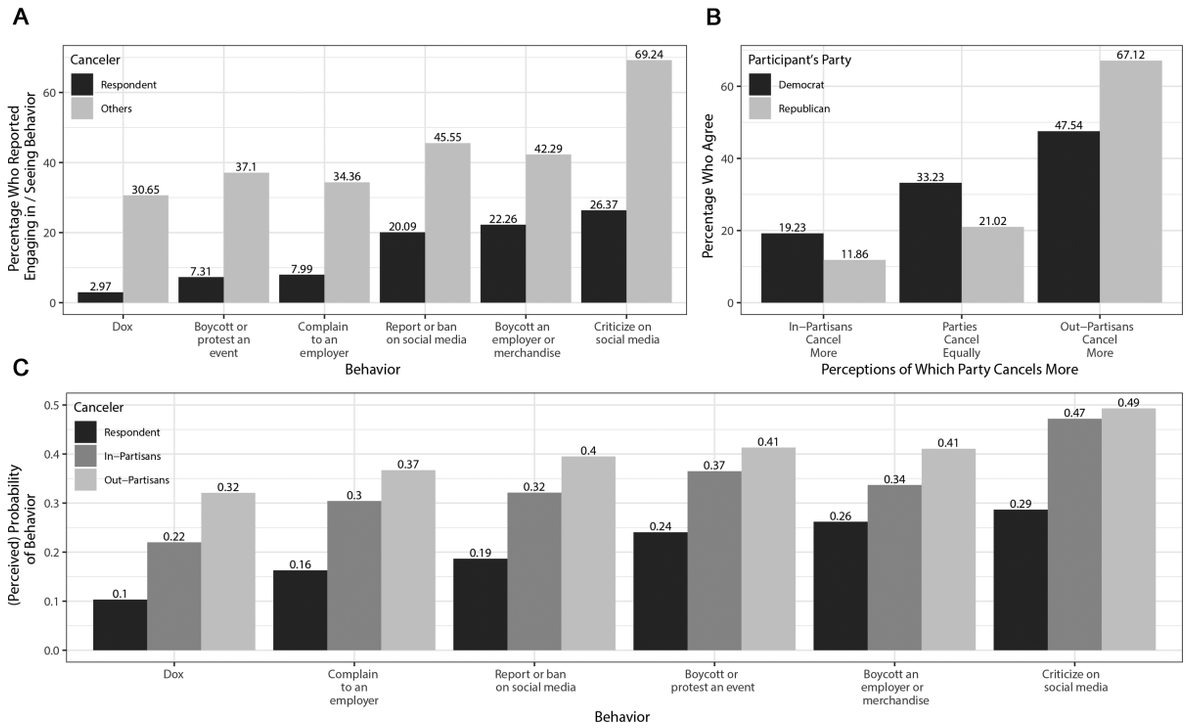

Figure 2. Americans overestimate how much others, especially out-partisans, engage in canceling. Panel A shows how often respondents saw others engage in each canceling behavior in the real world, relative to how much they themselves engaged in those behaviors. Panel B shows how often respondents thought real-world canceling behavior came from the out-party versus the respondent's in-party. Panel C shows how likely, in our conjoint experiment, respondents thought they themselves, in-partisans, and out-partisans would engage in each canceling behavior.

behaviors were more often Democrats or Republicans. Consistent with "canceling is (mostly) for out-partisans" (hypothesis 3), figure 2B shows that respondents are very likely to believe that out-partisans cancel more than in-partisans. This pattern holds across parties: 48% of Democratic respondents believe Republicans cancel more than Democrats, and 67% of Republican respondents believe Democrats cancel more than Republicans. Notably, then, while members of both parties perceive out-partisans to cancel more, this pattern is more pronounced for Republicans. This pattern coheres with stereotypes that Democrats cancel more than Republicans (Fahey et al. 2023). It also is consistent with the partisan breakdown in actual canceling behavior: as shown in appendix G, Democrats consistently report having canceled more than Republicans, with differences ranging from 1.9 to 13.6 percentage points.

One important limitation of the perceptual data used above is that they could reflect the acts of a few active "supercancelers" who may lean left politically. Our third-person conjoint data allow a further test of misperceptions. Recall that respondents predicted how Democrats and Republicans would respond to conjoint scenarios. We compare these estimates (i.e., perceptions) to respondents' own responses to the conjoint scenarios (i.e., reality). This enables us to overcome the potential biases of the questions used earlier, which might be skewed by a few highly visible cancelers or incidents. While media reports may still skew perceptions of how others would respond to particular scenarios, here, respondents would be inaccurately projecting from media to the likelihood of canceling in specific cases, rather than accurately reporting on skewed media coverage.

We present the results in figure 2C. Consistent with our "canceling is not widespread" hypothesis (hypothesis 1), respondents themselves are unlikely to cancel others.[12] Moreover, across scenarios, Democrats and Republicans cancel at similar rates (see app. G). This contrasts with differences in the frequencies presented above. We can only speculate about what drives this difference. However, we note that our conjoint scenarios describe an equal number of potentially offensive statements for the Left and Right. The partisan differences in real-world canceling behavior might stem from variations in the "supply" of statements that each

12. Some individuals, such as those high in partisan animosity, may be more likely to sanction. We test for such heterogeneous effects in app. F and find little evidence of them.

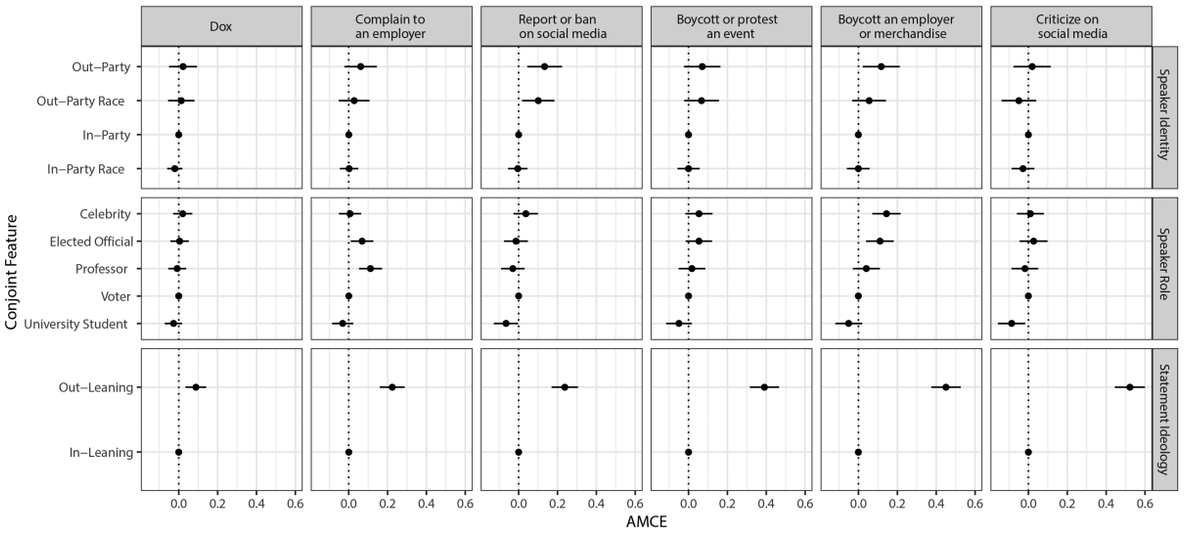

Figure 3. Americans cancel disagreeable statements, not disliked speakers. This figure shows which factors drove respondents to engage in each canceling behavior in our conjoint experiment. Points and bars represent AMCEs and 95% confidence intervals.

party finds offensive. That is, Democrats and Republicans may be equally likely to cancel, but Democrats might have seen more cancelable statements at the time we collected our data. We return to this point in our conclusion.

Consistent with "canceling is (mostly) for out-partisans" (hypothesis 3), in-partisans are perceived to be somewhat more likely to cancel than respondents are, and out-partisans are perceived to be the most likely to cancel. Respondents overestimate the likelihood that in-partisans will cancel by a factor of 1.3 to 2.2 and overestimate the likelihood that out-partisans will cancel by a factor of 1.6 to 3.2. Both parties perceived the other to be more likely to cancel (see app. G), underscoring the ubiquity of these misperceptions. This perceived partisan gap is slightly smaller than that for other negative perceptions of out-partisans. For instance, Druckman et al. (2022) find that partisans overestimate how many out-partisans are ideological or frequently discuss politics by a factor of 3.5. While these gaps may reflect some innumeracy among the public (Duffy 2019), the fact that the misestimates vary across behaviors and groups suggests that the misperceptions reflect respondents' real beliefs too. At the very least, innumeracy interacts with real beliefs to produce self-reported misperceptions.

AMERICANS CANCEL WHAT IS SAID, NOT WHO SAYS IT

Americans perceive that others cancel far more than they do, implying that citizens enter interactions with inaccurate expectations that could lead them to self-censor (Noelle-Neumann [1980] 1993) or avoid political conversations entirely (Carlson and Settle 2022). This may be considered a feature or a bug of canceling behavior, depending on one's view on what speech is (or should be) discouraged by this behavior. Yet, what speech is discouraged depends on how Americans view the rationale for canceling: if citizens misperceive why others cancel, they may self-censor more speech than is necessary to avoid social sanctions.

Why do Americans engage in canceling behavior? We turn to our first-person conjoint data for the answer. Following Hainmueller, Hopkins, and Yamamoto (2014), figure 3 presents the average marginal component effects (AMCE) of our conjoint features on each canceling behavior. We use the prefixes "in-" and "out-" to indicate how a speaker's identity corresponds to the respondent's partisanship. For example, "in-party" refers to a Democratic or Black speaker when the respondent is a Democrat. Similarly, we use the terms "in-leaning" and "out-leaning" to indicate statements associated with the Left or Right of the ideological spectrum.

The results in figure 3 are remarkably consistent across canceling behaviors. In contrast to "cancelers sanction disliked speakers" (hypothesis 4), respondents are not much more likely to cancel out-partisans or racial groups associated with the out-party compared to in-partisans or associated groups (AMCEs: –0.05 to 0.13). Moreover, in contrast with "cancelers sanction public speakers" (hypothesis 5), respondents are not much more likely to cancel public figures compared to private ones (AMCEs: –0.09 to 0.14).

Instead, canceling appears to be driven by statement content: as predicted by the "cancelers sanction disagreeable speech" hypothesis (hypothesis 6), respondents are considerably more

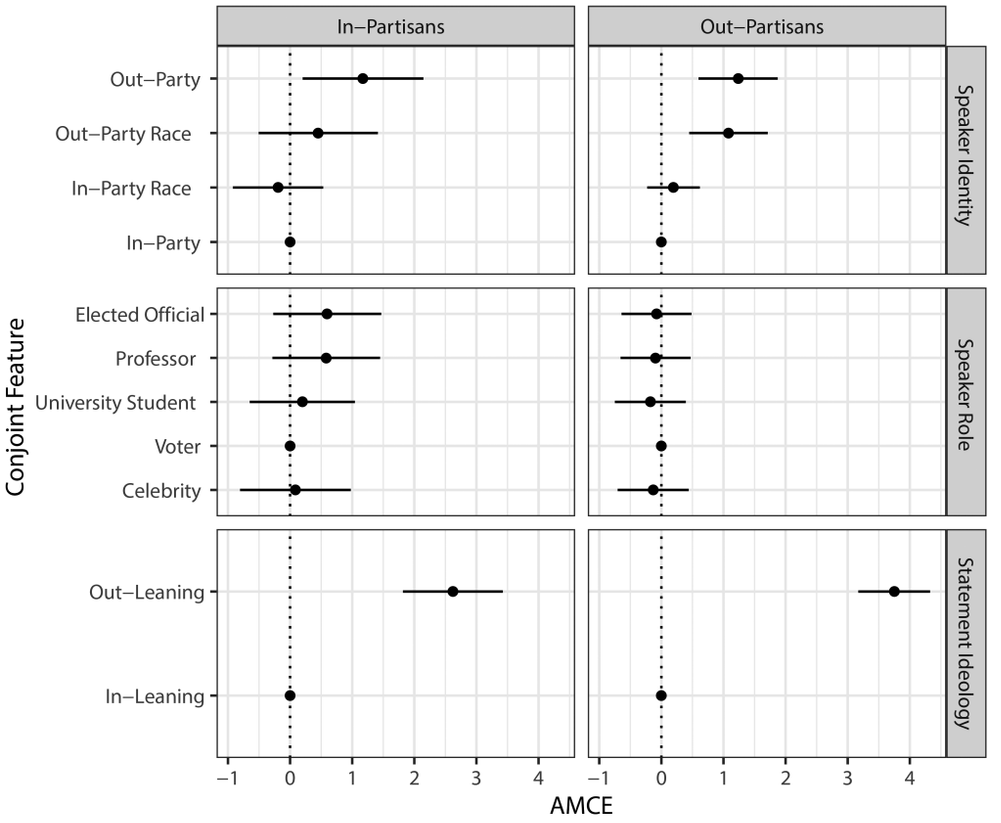

Figure 4. Americans accurately perceive others to cancel disagreeable statements. This figure shows which factors drove the perception that in-partisans and out-partisans would cancel, according to our conjoint experiment. Points are the AMCEs with bars representing 95% confidence intervals. Conjoint attributes (e.g., "In-Party") are defined relative to the supposed canceler (indicated in the facet label).

likely to cancel ideologically disagreeable statements compared to agreeable ones (AMCEs: 0.09–0.52).[13] Across canceling behaviors, the effect of statement ideology is consistently at least 1.8 times larger than the largest source-based effect. Furthermore, in the case of behaviors like social media criticism, the effect of an "out-leaning" statement is nearly 20 times that of source identity. In appendix F, we find little evidence that the effect of statement ideology stems from inferred party loyalty (i.e., a source effect; Dias and Lelkes 2022).

Thus, the clearest and most important factor that drives canceling is what speakers say, not who they are. People do not cancel chiefly on the basis of a speaker's identity, as one might expect in the era of partisan polarization (Iyengar et al. 2019). And while public officials, celebrities, and professors may get more attention when they make controversial statements, individuals do not hold them to a higher standard. In appendix G, we show that the dominance of statement ideology manifests for both Democrats and Republicans. Moreover, in appendix F, we find little other evidence for heterogeneous treatment effects. These findings suggest that canceling does not primarily involve punishing those we dislike but involves minimizing speech perceived as harmful (as dictated by one's political ideology).

AMERICANS ACCURATELY PERCEIVE WHY CANCELING OCCURS

Finally, how accurately do Americans perceive what drives canceling behavior? We use our third-person conjoints to answer research question 1: How do Americans see canceling motives? Figure 4 presents the AMCEs of our conjoint features on the perceived likelihood that in-partisans and out-partisans would engage in any canceling behavior. To simplify the presentation of our results, our dependent variable is a 0–18 score representing the sum of the likelihoods of engaging in each of our six canceling behaviors.[14] The facet labels at the top of figure 4 indicate whether respondents were judging the likely response of in-partisans or out-partisans. By contrast, conjoint attributes such as "in-party" are defined relative to

13. In app. B, we show that perceived statement offensiveness drives sanctioning.

14. As shown in app. F, our results in fig. 4 are consistent across canceling behaviors.

the supposed canceler (i.e., the facet labels). For example, when out-partisans are the cancelers in question (as denoted by the facet label), conjoint attributes referencing the "in-party" refer to the respondent's out-party.

As shown in figure 4—and in stark contrast to their misperceptions of who cancels and how often—Americans' perceptions of why others cancel are generally accurate. Consistent with figure 3, Americans perceive that both in-partisans and out-partisans cancel primarily in response to offensive statements that contradict their ideology (AMCEs: 2.62–3.75). People realize that canceling primarily occurs when content offends. That said, unlike actual canceling behavior, people anticipate both in-partisans and out-partisans to be somewhat more likely to cancel one another (AMCEs: 1.17–1.24). This appears to be the result of respondents thinking that others are more likely to cancel with malice. In appendix G, we show that respondents were more likely to view identical canceling behaviors as nefariously motivated (e.g., biased, over-sensitive, punishing) when said behaviors were performed by others, particularly out-partisans. Again, this is true for both Democrats and Republicans. Overall, these findings suggest that citizens accurately perceive canceling as triggered by a fairly well-defined set of offensive statements.

CONCLUSION

Cancel culture has generated a vigorous debate about whether Americans actually ostracize one another for what they say and, if so, what it implies for Americans' political culture. Defenders of canceling behavior argue that canceled statements are often prejudicial and, when left unchecked, deter members of marginalized groups from participating in the public sphere (e.g., McGowan 2022). Critics, however, argue that canceling belies faux outrage and a desire to disqualify political opponents from civic discourse (e.g., Trump 2020) or interferes with the marketplace of ideas (Gillespie 2022; Scudder 2020). Our work is among the first to empirically investigate the prevalence and motives of canceling among the American public, as well as Americans' perceptions of canceling behavior.

Our results complicate the narratives of both defenders and critics of canceling behavior. First, we find that Americans overestimate—by at least a factor of two—how often their fellow citizens cancel others. In particular, partisans vastly inflate how willing out-partisans are to cancel speech that offends them. The politicized nature of these misperceptions could exacerbate partisan animus (Iyengar et al. 2019) and discourage cross-party dialogue (Carlson and Settle 2022). Elites, in turn, may stoke this sentiment and galvanize witch-hunts against political opponents.

For critics of canceling, these misperceptions may also underscore that the threat of self-censorship is real and could generate a "spiral of silence" (Noelle-Neumann [1980] 1993), especially when a small number of cancelers can prompt institutional action (e.g., a formal investigation, firing). Alternatively, defenders of canceling might view these misperceptions as beneficial, insofar as they discourage speech that might render the public sphere inhospitable to marginalized communities. A crucial question for future research is whether individual-level traits (e.g., conflict aversion, innumeracy) that lead to self-censoring could also affect canceling behavior. One could study this experimentally by varying a speaker's perceived audience and asking participants to choose whether to endorse particular statements. Furthermore, it will be important to explore the extent of self-censorship, perceptions of its frequency, and how those have changed over time (Gibson and Sutherland 2023). Insights like these would go far to illuminate the crucial role of self-censorship in cancel culture.

Second, people are not canceled for who they are but, rather, for what they say. Although partisan biases characterize perceptions of who cancels and how often, canceling behaviors do not amount to disingenuous partisan attacks. Citizens do not care who makes offensive statements—be it a member of the other party, another race/ethnicity, or an elected official. Instead, Americans cancel disagreeable or offensive statements, even if those statements come from in-partisans. Democrats and Republicans were similarly motivated to cancel offensive statements in our conjoint experiment. This suggests that the parties agree that canceling behavior is sometimes appropriate, but they differ on what statements should be canceled, a point we return to below.

Third, citizens recognize that others cancel mostly in response to what is said, not who says it. This qualifies some, but not all, concerns about canceling behavior and the self-censorship that it likely inspires: if citizens believed they would be canceled because of who they are, they may be likely to self-censor a broad swathe of political speech or avoid political conversations entirely. Instead, the restriction here is more about a set of debates around what exactly is within the boundaries of acceptable speech.

Fourth, given a similar set of circumstances, Democrats and Republicans are similarly likely to engage in canceling behavior. They also cancel for similar reasons—i.e., offensive and ideologically disagreeable ideas. This defies the common narrative that Democrats cancel more than Republicans (Fahey et al. 2023; Trump 2020). However, judging by reports of real-world behavior, Democrats do cancel more than Republicans. Why this discrepancy? A plausible explanation, in our view, is differences in the "supply" of offensive statements: partisan biases in who makes offensive statements, or who is offended, may shift over time. At the time our data were collected, the supply of offensive statements may have

had a right-leaning bias. This balance may now have shifted, given recent Republican activism pushing back on discussions of race, gender, and sexuality in schools. Alternatively, Democrats could experience greater encouragement to cancel in the real world, consistent with a norm-based explanation of decreasing tolerance for prejudicial speech (Chong et al. 2024).

In closing, we note two points underemphasized in earlier commentary on cancel culture. The extant normative literature argues about whether the harm created by speech justifies restricting the marketplace of ideas (Mill [1859] 2003; Scudder 2020). In particular, classical liberals argue that speech should only be restricted under extremely rare circumstances, that citizens should be tolerant of offense, and that the proper response to offensive speech is counter-speech (e.g., Gillespie 2022). Yet, our results highlight that Americans agree that some speech should be restricted but differ on what speech should be restricted. Democrats seem more inclined to cancel speech that attacks marginalized communities (Chong et al. 2024), whereas Republicans are more inclined to cancel speech that challenges traditional social orders or American exceptionalism (see Jost 2021).

In another way, our results highlight that media and elite debate over "cancel culture" marks a shift in who determines the boundaries of acceptable expression. Although Americans do not disproportionately punish public figures for offensive comments, these cases have attracted outsized attention. This is perhaps because elites have rarely had to account for their actions and their speech is seen by many more people (and hence faces a higher likelihood of criticism). Historically, political elites—such as politicians, journalists, lawyers, and judges—largely determined what speech was acceptable. Now, in the age of social media, ordinary Americans increasingly set the limits of acceptable discourse, and those limits evolve rapidly.

ACKNOWLEDGMENTS

We thank Mira Dedhia, Ben Hempker, Kirsten Huh, Ami Ikuenobe, Maryarita Kobotis, Aaron Lazara, Jack McGovern, Evan Myers, Uday Tandon, and Anna Wang for research assistance. We also thank Jim Gibson, Dan Hopkins, Lizzie Martin, Jesper Rasmussen, Rune Slothuus, Alex Tolkin, Sean Westwood; seminar participants at the European University Institute, Stony Brook University, Northwestern University, Aarhus University, and the 2022 Conference on Bystanders and Online Hostility; our anonymous referees; and the editors for their invaluable comments.

REFERENCES

Ackerman, Elliot, Saladin Ambar, Martin Amis, Anne Applebaum, Marie Arana, Margaret Atwood, John Banville, et al. 2020. "A Letter on Justice and Open Debate." *Harper's Magazine*, Jul 7. https://harpers.org/a-letter-on-justice-and-open-debate/ (accessed March 14, 2024).

Ahler, Douglas J., and Gaurav Sood. 2018. "The Parties in Our Heads." *Journal of Politics* 80 (3): 964–81.

Balkin, Jack. 2018. "Free Speech Is a Triangle." *Columbia Law Review* 118 (7): 2011–56.

Carlson, Taylor N., and Jaime E. Settle. 2022. *What Goes Without Saying*. Cambridge, UK: Cambridge University Press.

Chong, Dennis, Jack Citrin, and Morris Levy. 2024. "The Realignment of Political Tolerance in the United States." *Perspectives on Politics* 22 (1): 131–52.

Chung, Sungeun, and Shin-Il Moon. 2016. "Is the Third-Person Effect Real?" *Human Communication Research* 42 (2): 312–37.

Clark, Meredith. 2020. "Drag Them: A Brief Etymology of So-Called 'Cancel Culture.'" *Communication and the Public* 5 (3–4): 88–92.

Costa, Mia. 2021. "Ideology, Not Affect." *American Journal of Political Science* 65 (2): 342–58.

Dias, Nicholas, and Yphtach Lelkes. 2022. "The Nature of Affective Polarization." *American Journal of Political Science* 66 (3): 775–90.

Druckman, James, S. R. Gubitz, Matthew S. Levendusky, and Ashley Lloyd. 2019. "How Incivility On Partisan Media (De-)Polarizes the Electorate." *Journal of Politics* 81 (1): 291–95.

Druckman, James N., Samara Klar, Yanna Krupnikov, Matthew Levendusky, and John B. Ryan. 2022 "(Mis)Estimating Affective Polarization." *Journal of Politics* 84 (2): 1106–17.

Duffy, Bobby. 2019. *Why We're Wrong About Nearly Everything*. New York: Basic Books.

The Editorial Board. 2022. "America Has a Free Speech Problem." *New York Times*, March 18. https://www.nytimes.com/2022/03/18/opinion/cancel-culture-free-speech-poll.html (accessed March 14, 2024).

Fahey, James J., Damon C. Roberts, and Stephen M. Utych. 2023. "Principled or Partisan?" *American Politics Research* 51 (1): 69–75.

Gibson, James L. 2011. "Political Intolerance in the Context of Democratic Theory." In Robert Goodin, ed., *The Oxford Handbook of Political Science*. Oxford: Oxford University Press, 409–27.

Gibson, James L., and Joseph L. Sutherland. 2023. "Keeping Your Mouth Shut." *Political Science Quarterly* 138 (3): 361–76.

Gillespie, Nick. 2022. "Greg Lukianoff: Saving the 'Culture of Free Speech.'" *Reason Podcast*, June 8. https://reason.com/podcast/2022/06/08/greg-lukianoff-saving-culture-free-speech/.

Gillespie, Tarleton. 2018. *Custodians of the Internet*. New Haven, CT: Yale University Press.

Greene, Peter. 2022. "Teacher Anti-CRT Bills Coast To Coast." *Forbes*, February 16. https://www.forbes.com/sites/petergreene/2022/02/16/teacher-anti-crt-bills-coast-to-coast-a-state-by-state-guide/ (accessed March 14, 2024).

Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis." *Political Analysis* 22 (1): 1–30.

Herrnstein, Richard J., and Charles Murray. 1994. *The Bell Curve*. New York: Free Press.

Hoorens, Vera. 1993. "Self-Enhancement and Superiority Biases in Social Comparison." *European Review of Social Psychology* 4 (1): 113–39.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22 (1): 1–18.

Jackson, David J., and Thomas I. A. Darrow. 2005. "The Influence of Celebrity Endorsements on Young Adults' Political Opinions." *Harvard International Journal of Press/Politics* 10 (3): 80–98.

Jamieson, Kathleen Hall, and Bruce Hardy. 2012. "What Is Civil Engaged Argument and Why Does Aspiring to It Matter?" *PS: Political Science and Politics* 45 (3): 412–15.

Jones, Dustin, and Jonathan Franklin. 2022. "Not Just Florida. More than a Dozen States Propose So-Called 'Don't Say Gay' Bills." *NPR*, April 10. https://www.npr.org/2022/04/10/1091543359/15-states-dont-say-gay-anti-transgender-bills (accessed March 14, 2024).

Jost, John T. 2021. *Left & Right*. Oxford: Oxford University Press.

Klar, Samara, and Yanna Krupnikov. 2016. *Independent Politics*. Cambridge, UK: Cambridge University Press.

Lelkes, Yphtach, and Sean J. Westwood. 2017. "The Limits of Partisan Prejudice." *Journal of Politics* 79 (2): 485–501.

Levendusky, Matthew, and Neil Malhotra. 2016. "Does Media Coverage of Partisan Polarization Affect Political Attitudes?" *Political Communication* 33 (2): 283–301.

McGowan, Mary Kate. 2022. "Just Words." *Australasian Philosophical Review* 5 (2): 129–49.

Mill, John Stuart. (1859) 2003. *On Liberty*. New Haven, CT: Yale University Press.

Mishan, Ligaya. 2020. "The Long and Tortured History of Cancel Culture." *New York Times*, December 3. https://www.nytimes.com/2020/12/03/t-magazine/cancel-culture-history.html (accessed March 14, 2024).

Mislove, Alan, Sune Lehmann, Yong-Yeol Ahn, Jukka-Pekka Onnela, and J. Rosenquist. 2021. "Understanding the Demographics of Twitter Users." In *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 5. Washington, DC: Association for the Advancement of Artificial Intelligence, 554–57.

Munger, Kevin. 2017. "Tweetment Effects on the Tweeted." *Political Behavior* 39 (3): 629–49.

Noelle-Neumann, Elisabeth. (1980) 1993. *The Spiral of Silence*. 2nd ed. Chicago: University of Chicago Press.

Norris, Pippa. 2023. "Cancel Culture: Myth or Reality?" *Political Studies* 71 (1): 145–74.

The Objective. 2020. "A More Specific Letter on Justice and Open Debate." *The Objective*, July 10. https://objectivejournalism.org/2020/07/a-more-specific-letter-on-justice-and-open-debate/ (accessed March 14, 2024).

Rasmussen, Jesper. 2022. "The (Limited) Effects of Target Characteristics on Public Opinion of Hate Speech Laws." *PsyArXiv*, June 8. https://osf.io/preprints/psyarxiv/j4nuc.

Rauch, Jonathan. 2021. *The Constitution of Knowledge*. Washington, DC: The Brookings Institution.

Romano, Aja. 2020. "Why We Can't Stop Fighting about Cancel Culture." *Vox*, August 25. https://www.vox.com/culture/2019/12/30/20879720/what-is-cancel-culture-explained-history-debate (accessed March 14, 2024).

Schauer, Frederick. 2010. "Facts and the First Amendment." *UCLA Law Review* 57 (4): 897–919.

Scudder, Mary. 2020. *Beyond Empathy and Inclusion*. Oxford: Oxford University Press.

Shapiro, Ari, Alejandra Marquez Janse, Mia Venkat, Noah Caldwell, and Patrick Jarenwattananon. 2021. "How Cancel Culture Became Politicized—Just Like Political Correctness." *NPR*, July 26. https://www.npr.org/2021/07/09/1014744289/cancel-culture-debate-has-early-90s-roots-political-correctness (accessed March 14, 2024).

Strossen, Nadine. 2018. *Hate: Why We Should Resist It through Free Speech, Not Censorship*. Oxford: Oxford University Press.

Sullivan, John, James Pierson, and George Marcus. 1982. *Political Tolerance and American Democracy*. Chicago: University of Chicago Press.

Trump, Donald. 2020. "Remarks by President Trump at South Dakota's 2020 Mount Rushmore Fireworks Celebration." Speech transcript, Keystone, SD, July 4. https://trumpwhitehouse.archives.gov/briefings-statements/remarks-president-trump-south-dakotas-2020-mount-rushmore-fireworks-celebration-keystone-south-dakota/.

Tversky, Amos, and Daniel Kahneman. 1973. "Availability." *Cognitive Psychology* 5 (2): 207–32.

Viser, Matt. 2021. "Early in Biden's Presidency, GOP Shows the Places They'll Go." *Washington Post*, March 5. https://www.washingtonpost.com/politics/gop-seuss-muppets-neanderthal/2021/03/05/9d4e92b6-7dd8-11eb-a976-c028a4215c78_story.html (accessed March 14, 2024).

Vogels, Emily A., Monica Anderson, Margaret Porteus, Chris Baronavski, Sara Atske, Colleen McClain, Brooke Auxier, Andrew Perrin, and Meera Ramshankar. 2021. "Americans and 'Cancel Culture': Where Some See Calls for Accountability, Others See Censorship, Punishment." *Pew Research Center*, May 19. https://www.pewresearch.org/internet/2021/05/19/americans-and-cancel-culture-where-some-see-calls-for-accountability-others-see-censorship-punishment/ (accessed March 14, 2024).

Waisbord, Silvio. 2020. "Mob Censorship." *Digital Journalism* 8 (8): 1030–46.

Wike, Richard. 2016. "Americans More Tolerant of Offensive Speech than Others in the World." *Pew Research Center*, October 12. https://www.pewresearch.org/short-reads/2016/10/12/americans-more-tolerant-of-offensive-speech-than-others-in-the-world/ (accessed March 14, 2024).

Wilson, Anne E., Victoria A. Parker, and Matthew Feinberg. 2020. "Polarization in the Contemporary Political and Media Landscape." *Current Opinion in Behavioral Sciences* 34:223–28.

Wojcieszak, Magdalena, Andreu Casas, Xudong Yu, Jonathan Nagler, and Joshua A. Tucker. 2022. "Most Users Do Not Follow Political Elites on Twitter; Those Who Do Show Overwhelming Preferences for Ideological Congruity." *Science Advances* 8 (39): eabn9418.