

Do American Voters Really Not Punish Overt Undemocratic Behavior at the Polls? Natural Experimental Evidence from the 2021 Insurrection of the U.S. Capitol

Sam van Noort*

First version: September 10, 2021

This version: September 25, 2021

Abstract

Existing research suggests that too few American voters hold politicians electorally accountable for overt undemocratic behavior to reasonably deter democratic backsliding. Evidence for this proposition comes primarily from hypothetical survey experiments with relatively modest treatments. I test this hypothesis using a natural experiment with a powerful real-world treatment: Donald Trump's incitement of the insurrection of the U.S. Capitol on January 6, 2021. The insurrection was unexpected to the general public, did not coincide with other events that could plausibly affect public opinion, and occurred while Gallup was conducting a nationally representative survey using random digit dialing. Comparing Republican Party support among respondents that were interviewed just before, and just after, the insurrection occurred suggests that the insurrection caused a 10.8% decline in support for the Republican Party. Voters predominantly moved to the Democratic Party, rather than Independent. Back-of-the-envelope calculations suggests that this electoral penalty is sufficient to decide presidential elections.

Keywords: Democratic Backsliding, Accountability, Voter Behavior

Words: 10,405 (excl. appendix)

*PhD Candidate at the University of Cambridge and Fellow at the Niehaus Center for Globalization and Governance of Princeton University. E-mail address: samvannoort@princeton.edu. Website: sites.google.com/view/samvannoort. I would like to thank Tanushree Goyal and Daniel Lyng for helpful comments and suggestions. I thank Jeremy Darrington, Bobray Bordelon, and Duy Trinh for excellent research assistance. This study was pre-registered at the EGAP registry.

“In general Americans do not prioritize democratic principles in our [sic] vote choices, and we are alarmingly willing to tolerate antidemocratic ideas and actions by co-partisans. Polarization seems to make this worse. If American democracy is at risk, citizens will not save it.” (Political Science Professor Robert Blair, as quoted in the New York Times (2021a))

1 Introduction

Politicians that cannot win through free and fair elections self-evidently have an incentive to rig the electoral system for political gain. While democratic norms among politicians could in principle self-constrain politicians from doing so, this mechanism, even if such norms exist, becomes more and more brittle ones the stakes of elections increases. Political scientists have therefore long understood that for democracy to be truly self-enforcing it is necessary that politicians anticipate that, were they to behave undemocratically, their own supporters would punish them by voting for a competitor in sufficiently large numbers to bring about their defeat. Anticipating this reaction, political leaders would rarely attempt to violate key democratic norms in the first place (Weingast, 1997).

It is extremely disconcerting therefore that a wide-range of recent studies find that the American electorate may no longer provide this shield of protection against democratic backsliding.¹ Foa and Mounk (2016) analyze descriptive survey data and find that support for democracy as a political system has been steadily declining in the United States in recent decades, particularly among younger generations. Graham and Svolik (2020) show, using a randomized survey experiment, that in regular state legislative elections between a Democrat and a Republican, and where the candidates differ substantively on economic and social policy (i.e., virtually every election in the United States), a mere 3.5% of the electorate

¹This concern is of course hardly limited to the United States alone. Indeed, a large comparative politics literature exists on the issue of democratic backsliding in general (e.g., Bermeo (2016), Cleary and Öztürk (2020), Diamond (2021), Ding and Slater (2021), Grillo and Prato (Forthcoming), Grossman et al. (2021), Haggard and Kaufman (2021), Lührmann (2021), and Waldner and Lust (2018)) and on the role of voters in tolerating democratic backsliding in particular (e.g., Acemoglu, Robinson and Torvik (2013), Cho and Hwang (2021), Magaloni (2006), Svolik (2019, 2020), and Wuttke, Gavras and Schoen (2020)).

is willing to punish clear undemocratic behavior by politicians if this means that they are forced to vote for a politician that is further away from their own policy preferences.² McCoy, Simonovits and Littvay (2020) and Albertus and Grossman (2021) meanwhile find survey experimental evidence that Americans are remarkably hypocritical when it comes to democratic backsliding, supporting the same anti-democratic state executive behavior when their own party is in power, while condemning it if the opposing party is in power.

This important existing evidence notwithstanding significant uncertainty remains as to whether Americans today do indeed not punish undemocratic behavior by elected officials sufficiently enough to reasonably influence elections. First, given that all existing evidence comes from hypothetical survey experiments and descriptive survey responses on abstract norm questions, it arguably remains unclear whether voters would not sufficiently react to *real-life* democratic transgressions that would actually affect their own life.³ Second, all existing studies employ important but nevertheless not truly fundamental and decisive forms of democratic backsliding (e.g., whether a candidate supported a proposal to reduce the number of polling places in areas that largely support the opposing party). It is arguably unclear therefore whether American voters will not be an effective electoral check to more extreme forms of democratic backsliding that would truly change the rules of the game (e.g., a state executive refusing to leave office after losing an election).

To address these issues I analyze the results of a novel natural experiment that occurred because Donald Trump’s incitement of the insurrection of the U.S. Capitol on January 6, 2021, unexpectedly occurred while Gallup was conducting a nationally representative public opinion survey among 1,023 adult Americans. Given that Gallup samples respondents using

²Carey et al. (2020) do find survey experimental evidence that American voters punish undemocratic behavior by elected officials *holding everything else constant* (i.e., not forcing respondent to give up some of their own policy preferences in order to punish undemocratic behavior). While noteworthy, this is arguably not the typical situation in contemporary American elections.

³Graham and Svobik (2020) are an exception to this. They also analyze the causal effect of Republican Greg Gianforte hitting a journalist in the face (which they define as antidemocratic behavior) during Montana’s 2017 special election for the U.S. House of Representatives. This paper differs from Graham and Svobik (2020) by focusing on a more extreme form of undemocratic behavior that had a more structural and institutional effect on democracy in America.

random digit dialing the probability of any particular individual to be interviewed before or after the insurrection is plausibly exogenous to any pre-treatment characteristics that may otherwise affect vote choice. Furthermore, given that there were arguably no other events surrounding January 6, 2021 that could plausibly explain major changes in political party alliance (see section 3.2.1 below) I can recover the average treatment effect of inciting the insurrection by simply comparing support for the Republican Party among respondents that happened to be interviewed just before, and just after, the insurrection occurred.

Using this natural experiment I find that Americans do in fact punish behavior as extreme as inciting an insurrection sufficiently strongly to reasonably influence elections. More specifically, I find that electoral support for the Republican Party went down by more than 10% as a direct result of Donald Trump’s incitement of the insurrection. Crucially, the majority of voters switched to supporting the Democratic Party, rather than becoming independents, causing the expected vote share of the Democratic Party to simultaneously increase by more than 8%. Similar effects are found on Donald Trump’s presidential approval rating more specifically. All these effects lasted for at least 1.5 weeks (when the survey ended) and plausibly for at least 3 more months after the fact. Analysis of treatment heterogeneity suggests that, if anything, these effects were slightly larger in the crucial swing states of the 2020 presidential election (i.e., Arizona, Colorado, Florida, Georgia, Iowa, Maine, Michigan, North Carolina, Ohio, Pennsylvania, Texas, and Wisconsin). Simple back-of-the-envelope calculations suggest that this electoral penalty for overt undemocratic behavior is more than sufficient to decisively influence U.S. presidential elections, even when taking possible idiosyncrasies generated by the Electoral College system into account.

These results suggests that the American electorate *does* form a credible electoral threat against undemocratic behavior by elected officials. This is at least the case when the undemocratic behavior in question entails something as extreme and overtly undemocratic as inciting an insurrection to overturn an election result. It remains an important question for future research whether forms of democratic backsliding that stop short of this high bar

(such as those studied by Graham and Svobik (2020)) are sufficiently punished at the polls in the real-world.

My results contrast relatively sharply with journalistic reporting regarding the January 6 insurrection of the U.S. Capitol. Most news reporting after the insurrection emphasized how radically different self-identifying Republicans and Democrats perceived and judged the events of January 6, 2021. The New York Times (2021*c*), Slate (2021), the Economist (2021), and many other well-respected news outlets, for example, reported extensively on a post-January 6 YouGov (2021) poll that found that approximately 45% of self-identifying Republicans actually supported the storming of the Capitol, that 52% of Republicans blamed Joe Biden, rather than Donald Trump, for inciting the unrest, and that while large majorities of Democrats and Independents saw the insurrection as a threat to democracy, only 27% of Republicans did. My results do not dispute these reports. Instead, my results show that there was a major shift in the *number* of people self-identifying as Republican and Democrat as a result of the insurrection. Within the now much smaller remaining group of Republican-leaning voters it could well be true that many trivialized or straight-out denied the nature and gravity of the events that occurred at the U.S. Capitol on January 6, 2021.

My results also contrast with a smaller number of journalistic reports that suggested that Republican voter registration went down in the weeks after the insurrection (both as compared to the Democratic Party, and as compared to the trend in the months leading up to January 6, 2021). My results differ from these reports by showing that the electoral penalty for the Republican Party was *much* (i.e., approximately two orders of magnitude) larger than the voter registration data suggested. The New York Times (2021*b*) reported that approximately 140,000 voters that were previously registered as Republicans dropped this affiliation in the four weeks after the insurrection in the 25 states for which voter registration data was readily available. After accounting for Democratic defections during this same period this implies a net gain for the Democrats of just 61,000 registered voters, or less than 0.1% of all registered voters, in those same 25 states (Agosta and Lightbourn, 2021). My

results suggest that this modest effect in the voter registration data is likely because most people, even if their party alliance has shifted in real-life, do not invest the effort to change their voter registration in the absence of an upcoming primary election.⁴

2 The insurrection of the U.S. Capitol

On November 3, 2020 Joe Biden (Democrat) won the U.S. presidential election of sitting U.S. president Donald Trump (Republican) with a little more than 7 million popular votes (and 74 Electoral College votes) difference. Despite this enormous difference in the number of votes, and the absence of any evidence of widespread voting fraud, Donald Trump claimed that the election was fraudulent and that he was the true winner of the election.

After several months of constant misinformation regarding the election, and a wide range of highly publicized court cases (which were all essentially dismissed for a lack of evidence), Donald Trump held a speech at the “Stop the Steal” rally on January 6, 2021. This rally was organized by his supporters on this specific day because the U.S. Constitution prescribes that the U.S. Congress has to official certify an election result on the January 6 following a presidential election.

In the speech Trump said, among other things:

“All of us here today do not want to see our election victory stolen by emboldened radical-left Democrats, which is what they’re doing. And stolen by the fake news media. That’s what they’ve done and what they’re doing. We will never give up, we will never concede. It doesn’t happen. You don’t concede when there’s theft involved.”

“Our country has had enough. We will not take it anymore and that’s what this is all about. And to use a favorite term that all of you people really came up with: We will Stop the Steal.”

⁴Analysis of treatment heterogeneity within the Gallup survey analyzed in this paper does not suggest that Donald Trump’s incitement of the insurrection of the U.S. Capitol had a different effect in the 25 states analyzed by the New York Times (2021*b*). The difference between my results and the voter registration data can thus not be explained by differences in sample composition.

“Because if Mike Pence does the right thing, we win the election. [...] All Vice President Pence has to do is send it back to the states to recertify and we become president and you are the happiest people.”

“[...] we’re going to walk down to the Capitol, and we’re going to cheer on our brave senators and congressmen and women, and we’re probably not going to be cheering so much for some of them.”

“Because you’ll never take back our country with weakness. You have to show strength and you have to be strong. ”

“We fight like hell. And if you don’t fight like hell, you’re not going to have a country anymore.”

“So let’s walk down Pennsylvania Avenue.”

Right after this speech, starting at 12:53 p.m. and ending around 7:30 p.m., a large mob of Trump supporters stormed the U.S. Capitol building, leading to the death of 5 people, the injury of at least 138 police officers, physical property damages in excess of 30 million dollars, the abrupt halt of an ongoing congressional debate, and the immediate evacuation of all members of Congress.

At 2:24 p.m., about one and a half hours after rioters had already breached the gated Capitol perimeter and at least half an hour after the first rioters had already breached the Capitol building, Donald Trump tweeted:

“Mike Pence didn’t have the courage to do what should have been done to protect our Country and our Constitution, giving States a chance to certify a corrected set of facts, not the fraudulent or inaccurate ones which they were asked to previously certify. USA demands the truth!”

While violent images of insurrectionists fighting with police officers inside and outside the Capitol building were aired across national and international television president Trump and his Defense Department took 3 hours and 19 minutes to approve a request by the Capitol Police for the D.C. National Guard to intervene and secure the Capitol building (Reuters, 2021b). Given that the Capitol Police force was ill-prepared for such a large and violent

group of rioters the late deployment of the National Guard gave the insurrectionists several hours to enter and vandalize the Capitol building largely unchallenged.

It was arguably only because members of Congress were evacuated just in time using underground tunnels that no parliamentarians were physically harmed. The weapons confirmed to have been carried by insurrectionists included, among other things: stun guns, pepper spray, zip-cuffs, Molotov cocktails, knives, and baseball bats.⁵ Furthermore, a car was found close to the Capitol building with a planted pipe bomb, an assault-style rifle equipped with a telescopic sight, and a handgun.

In the weeks following the insurrection a trial was held in the U.S. Congress to impeach president Trump. In this trial only 10 of the 207 Republican House of Representatives members and 7 of 50 Republican Senate members voted to impeach president Trump, leading to his acquittal (all Democrats in both houses voted for impeachment).⁶

I regard this event as a clear case of democratic backsliding that should be recognized and electorally sanctioned as such *if* the U.S. electorate is indeed to function as an effective check against overt undemocratic behavior by elected officials.

3 Identification strategy


To study whether the American electorate reacted to Donald Trump’s incitement of the insurrection of the U.S. Capitol in a way that is consistent with them checking overt undemocratic behavior by elected officials I exploit that the insurrection unexpectedly occurred while Gallup was conducting a nationally representative public opinion survey among 1,023 adult Americans using random digit dialing.

Figure 1 shows the number of respondents that were interviewed by Gallup before, on, and after January 6, 2021. As can be seen the survey was conducted from January 4 until

⁵Several other insurrectionist are currently charged with carrying handguns during the January 6 insurrection, but these court cases are still ongoing.

⁶Note that I regard the acquittal of President Trump by other Republican members of Congress as a part of the treatment – i.e., a clear and overt undemocratic act by the Republican Party at large. I therefore examine effects on the Republican Party as a whole, not only support for president Donald Trump himself.

Figure 1: Histogram of dates of interviews relative to treatment.



Note: Data refers to the year 2021. Data comes from Gallup’s January “Mood of the Nation” survey.

January 15, with 177 people having been interviewed before January 6, and 715 people having been interviewed after January 6.⁷

To identify the causal effect of Donald Trump’s incitement of the insurrection I employ a so-called “unexpected event during survey” design. In this case this design consists of a simple comparison of intended vote choice among respondents that happened to be interviewed just before, and just after, the insurrection occurred.

⁷As is standard in studies that use unexpectedly interrupted surveys for causal inference I drop the treatment day itself from the sample in my baseline results. This is because it is arguably unclear whether respondents that have been interviewed in the afternoon or evening of January 6 can already reasonably be expected to be aware of the occurrence and magnitude of the insurrection. It is therefore unclear whether these respondents should be considered as belonging to the treatment or control group. All results remain the same when assigning all the 118 respondents that were interviewed on January 6 to either the control or the treatment group. All results also remain the same when assigning all respondents that were interviewed after 6 p.m., 8 p.m., or 10 p.m. to treatment, and everyone else to control.

As rigorously explained by Muñoz, Falcó-Gimeno and Hernández (2020) identification of the average treatment effect (ATE) in this type of empirical design relies on three key assumptions. First, *compliance*, meaning that all respondents in the treatment group are actually aware of the occurrence of the insurrection. Second, *excludability*, meaning that the timing of interview should affect vote preferences only through the insurrection event, not through any other channel. Last, *temporal ignorability*, meaning that whether any particular individual is interviewed before or after January 6 should be orthogonal to any other individual-level characteristic that also affects vote preferences.

In the following three sections I explain why I believe that compliance, excludability, and temporal ignorability are likely to hold in this case.

3.1 Compliance

Identifying the ATE, as always, requires full compliance. Full compliance in this case necessitates that all respondents that were interviewed after January 6, 2021, are actually aware of Donald Trump’s incitement of the insurrection, and all respondents that were interviewed before January 6, 2021, are not.

Compliance is likely to hold in this case because the insurrection of the U.S. Capitol was a highly salient event that almost entirely dominated the national news from the afternoon of January 6, 2021 until weeks thereafter. Reflecting this, an Ipsos Poll that was conducted on January 6 and 7 (N = 536) found that only 1% of respondents indicated to not have heard about “[...] the recent (January 6, 2021) unrest at the Capitol during the certification of the (2020) election result.”⁸

It is highly likely therefore that (virtually) everyone that was interviewed on January 7 or thereafter was aware of what had happened and can thus legitimately be seen as “treated”. Everyone in the control group can meanwhile legitimately be seen as “untreated” because the insurrection of the U.S. Capitol was plausibly unexpected to the general public (see

⁸This data is available at: <https://ropercenter.cornell.edu/ipoll/study/31118152> (last accessed: September 12, 2021).

section 3.2.3) and because one can logically not be aware of an unexpected event that is yet to happen.⁹

3.2 Excludability

Excludability essentially requires that within the days surrounding the insurrection a dummy that takes the value 1 if a respondent is surveyed after January 6, and 0 if a respondent is surveyed before January 6, is not related to voter preferences through any other channel than Donald Trump’s incitement of the insurrection.

Muñoz, Falcó-Gimeno and Hernández (2020) highlight four potential ways through which excludability can fail in “unexpected event during survey” designs: (1) simultaneous events; (2) unrelated time trends; (3) endogenous timing of event; and (4) collateral events. In the following four subsections I discuss each of these potential threats to causal inference in turn.

3.2.1 Simultaneous events

Excludability would fail if another, unrelated, event occurred on January 6, 2021, that also influenced vote preferences.¹⁰ If this were to be the case any difference in outcomes post-treatment could be due to Donald Trump’s incitement of the insurrection, the other unrelated event, or a combination of both. In all cases this would mean that the estimated effect cannot be substantively interpreted.

To assess whether simultaneous events may be of concern here I read the entire editions of the New York Times, the Wall Street Journal, and the USA Today on each day from January

⁹More fundamentally one might argue that noncompliance is not in fact a threat to causal inference in this case. The problem with one-sided non-compliance (which is arguably the only possible non-compliance given that the insurrection was unexpected to the general public) is that the causal estimate may not be representative of the effect that would have occurred if everyone in the treatment group would actually have been treated (because one does not know how the respondents that did not take up treatment would have reacted if they in fact did). However, incorporating the effect among Americans that are not even aware of something as extreme as the insurrection of the U.S. Capitol is arguably hardly relevant for the purpose at hand because these type of people are in real-life most probably never sufficiently aware of potentially undemocratic behavior of politicians to form an effective check on democratic backsliding. In other words, the intension-to-treat effect is in that case arguably the estimand that is of actual substantive interest.

¹⁰Note that strictly speaking my results can only be spurious if another event occurred on January 6, 2021 that on caused a net shift in electoral support from the Republican Party towards the Democratic Party. Any event that caused a shift in the opposite direction would only lead me to underestimate the true effect.

4 to January 15.¹¹ For each day and each newspaper I code the title of all the articles on the front page, as well as the title of any other article that I believe could plausibly affect voter preferences. If necessary I provide a short summary of the article and for each article I code whether the news event described in the article may plausibly affect vote preferences, and if so in which direction (i.e., pro-Democratic or pro-Republican).

I report the results of this exercise in the online appendix. Taken together I do not find a compelling reason to believe that there were other events during this period that could plausibly explain a major shift in voter preferences from the Republican Party towards the Democratic Party.

To further reduce the potential influence of simultaneous events I run robustness checks whereby I narrow the amount of days surrounding January 6 in the sample.¹² As shown below this leaves the results unchanged.

3.2.2 Unrelated time trends

Another violation of the exclusion restriction can be generated by other variables that also affect voter preferences *and* changed continuously over the January 4 to January 15, 2021 period.

This is unlikely to be an issue in this case. First, linking Gallup’s December 2020 “Mood of the Nation” survey with the respondents that were interviewed on January 4 and 5 in my sample suggests that support for the Republican Party did not exert any downward trend leading up to January 6, 2021 (see online appendix). In the online appendix I also do not find a downward trend in Republican Party support in the post-election survey of the American National Election Study.

¹¹These newspapers are selected because they are the top three American newspapers by circulation.

¹²Note that there exists a trade-off between restricting the sample to days very close to January 6 to reduce the potential influence of simultaneous events, and precision (because standard errors increase when the number of observations goes down), generalizability (because a smaller sample, holding everything else constant, is less likely to accurately represent a diverse population), and potential imbalance-on-unobservables (because in a smaller sample there is a higher likelihood that the treatment and control group differ on independent determinants of voter preferences by random chance) (Muñoz, Falcó-Gimeno and Hernández, 2020, pp. 195–196). My solution to this trade-off is to report both the results with the full sample and with different degrees of sample restriction.

Second, if it is true that time during the January 4 to January 15 has a monotonic effect on vote preferences we would expect that many arbitrary partitions of time will yield statistically significant results. Using a specially designed randomization inference program that simulates the exact treatment assignment process within the full census of possible random assignments I find that this is not the case. Instead, it is exactly and only after January 6 that support for the Republican Party declines, and support for the Democratic Party increases (see online appendix).

3.2.3 Endogenous timing of event

A third potential source of exclusion restriction violations is if the insurrection itself occurred on January 6 for reasons that also influence voter preferences (confounding) and/or because expected post-January 6 changes in vote preferences caused the insurrection to occur on January 6 (reversed causality).

This issue is highly unlikely to be of concern in this case. The insurrection arguably only occurred on January 6, rather than some other day, because the U.S. Constitution prescribes that the final Electoral College votes have to be certified by Congress on the January 6 that follows a presidential election. It is for this reason that Donald Trump’s speech at the “Stop the Steal” rally, which led to the insurrection, occurred on this particular day. It is very unlikely therefore that the incitement of the insurrection of the U.S. Capitol was endogenously timed.¹³

3.2.4 Collateral events

A last issue with regard to excludability is the potential presence of collateral events. Collateral events would in this case be events that are also triggered by Donald Trump’s incitement

¹³According to the Federal Bureau of Investigation the storming of the U.S. Capitol building was itself not elaborately planned at all (although some very small extremist groups had been discussing options of this nature beforehand on social media). Rather, the insurrection appears to have been mostly the result of a large crowd getting out of control on the day itself (Reuters, 2021*a*). So not only was the insurrection plausibly unexpected for the general public, but the occurrence of the insurrection, and particularly its severity, was arguably unexpected even to the perpetrators themselves.

of the insurrection *and* which also cause changes in voter preferences *but* which in themselves have nothing to do with democratic backsliding. The presence of a collateral event would create a bundled treatment problem whereby any differences in vote preferences after January 6 may be due to Donald Trump’s incitement of the insurrection (the treatment of interest), the other non-democratic backsliding event triggered by the insurrection, or a combination of the two.

In the case of the insurrection of the U.S. Capitol there is arguably one important collateral event of potential relevance. The storming of the U.S. Capitol also exposed a major policy failure of the Washington, D.C. government to secure the Capitol building.¹⁴ It could be the case therefore that part of the change in voter preferences after January 6 were because of this public policy failure *rather* than due to Donald Trump’s and the Republican Party’ undemocratic behavior itself.

While relevant I believe this potential collateral event would only lead me to *underestimate* the degree to which Americans punish overt undemocratic behavior at the polls. This is because the government of Washington, D.C. was at the time of the insurrection fully controlled by Democrats (i.e., both the mayor – Muriel Elizabeth Bowser – was a Democrat and the Democratic Party held 11 out of the 13 seats in the Council of the District of Columbia, the other 2 seats being occupied by Independents, not Republicans). If anything, the police’s failure to secure the U.S. Capitol should thus weight against finding that Trump’s incitement of the insurrection led to a major electoral shift from the Republican Party *towards* the Democratic Party.

3.3 Temporal ignorability

Temporal ignorability essentially requires that people that had a higher/lower likelihood to prefer Democrats/Republicans/Independents in the absence of treatment, were not, for whatever reason, systematically more likely to be interviewed before/after January 6, 2021.

¹⁴This was arguably in no small part because president Trump and his Department of Defense did not employ the National Guard until several hours after the insurrection had started. As mentioned above I regard this later fact as part of the treatment of interest.

If temporal ignorability is violated it becomes impossible to tell whether any difference in vote preferences after January 6 are due to the treatment, due to differences in sample composition, or a combination of the two.

Muñoz, Falcó-Gimeno and Hernández (2020) discuss three ways how temporal ignorability can be violated in “unexpected event during survey” designs: (1) non-random fieldwork organization; (2) non-random reachability; and (3) non-random attrition. In the following three subsections I discuss each of these three potential threats to causal inference in turn.

3.3.1 Non-random fieldwork organization


The way Gallup recruits respondents makes it very unlikely that temporal ignorability is violated because of the way how the survey itself is constructed. Gallup interviews over the phone and samples respondents using random digit dialing. This ensures that any individual, regardless of where he/she lives, and regardless of his/her own background characteristics, has an equal probability to be contacted for an interview on any particular day.¹⁵

To test the assumption that whether respondents are interviewed before or after January 6 is as good as random I simply regress a dummy that takes the value 1 if a respondent is interviewed after January 6, and 0 if a respondent is interviewed before January 6, on a range of common determinants of vote preferences in the United States that can be regarded as pre-treatment. I include gender (being male, or not), religion (being Christian, or not), race (being white, or not), education (having a college education, or not), rural residency (living in a rural area, or not), and age (measured in years).

As expected the control and treatment group are balanced on these observables (see figure 2). This applies both in the full sample (left column) and when restricting the sample to include only days in close proximity to January 6, 2021 (middle and right column).

¹⁵Note that Gallup required for this survey that the sample consisted of at least 700 cellphone respondents, at least 300 landline respondents, and at least 30 Hispanic respondents. During the sampling process it turned out to be unnecessary to oversample these particular groups later in the survey to fulfill these quotas (see distribution of phone types and Hispanics across survey days in the online appendix). It is therefore not the case that Hispanics and people that use a particular type of phone are significantly overrepresented in the treatment group (see balance test in online appendix).

Figure 2: Balance on observables.



Note: Each column represents an individual OLS regression. Dependent variable is treatment status. Point estimates and 95% confidence intervals are reported.

3.3.2 Non-random reachability

Even though the structure of the survey makes it likely that temporal ignorability holds it may still be the case that temporal ignorability is violated due to the behavior of (prospective) respondents themselves. Most important in this regard is the issue of reachability. Gallup naturally seeks to interview respondents immediately after the random digit dialing algorithm provides a phone number for enumerators to call. Importantly, however, if a person cannot be reached or can be reached but only wants to be interviewed at some later time, Gallup is forced to call back later.

This implies that on the whole subjects that are more elusive and harder to reach have a higher probability to be interviewed later in the survey period (and as a result are more likely to end up in the treatment group in this case). This would be problematic for causal identification if those that are difficult to reach were already more likely to be Republicans/Democrats/Independents in the absence of treatment.

The data suggests that endogenous reachability differences are very unlikely to be of concern here. First, the number of tries before an interview was successfully completed is uncorrelated with supporting the Republican Party in my sample (OLS coefficient: -0.022; P-value: 0.199). Second, all results hold when dropping all respondents that were not successfully interviewed on the first try. Last, all results hold when simply controlling for the number of tries it took to interview any particular individual respondent.

3.3.3 Non-random attrition

A last channel through which temporal ignorability can be violated is if Donald Trump’s incitement of the insurrection made respondents that were interviewed by Gallup after January 6 more likely to lie about their support for the Republican Party (as compared to the two days before January 6).

While I cannot exclude this possibility for certain two empirical facts seem to suggest that social desirability bias is very unlikely to drive the results. If it were to be true that Republicans were disproportionally more likely to be interested in concealing their vote preference as a result of January 6, 2021 one would reasonably expect: (1) that non-responses on the vote preference question went up after January 6; and (2) that support for Independents, rather than Democrats, went up after January 6 (i.e., assuming of course that selecting Independents would still be socially “acceptable” while preferable over Democrats for “Republicans in disguise”).

This is not what the data suggests. Instead, there is no economically or statistically significant difference in non-response before and after January 6 (OLS coefficient: 0.004;

P-value: 0.661), and support tends to predominantly shift to the Democratic Party, rather than Independent, after January 6 (see results below).

4 Generalizability

The previous section has described how my empirical design is likely to recover a credible causal estimate of the ATE of Donald Trump’s incitement of the insurrection of the U.S. Capitol *within* the sample of respondents interviewed by Gallup between January 4 to January 15, 2021. An important, and entirely distinct, issue is whether this estimated effect also generalizes to the entire population of interest (i.e., all Americans that are eligible to vote).

Very unfortunately the total population of eligible voters is *not* the population that Gallup targets. Instead, Gallup samples by randomly selecting respondents from a list of all household telephone numbers that have been used at least ones in the past month within the United States.¹⁶ Strictly speaking, and implicitly assuming that nonresponse is orthogonal to vote preferences, the Gallup survey thus only generalizes to all adults that live in the United States and use a telephone at least ones per month (which is not necessarily the same population as those with voting rights).

Generalization to the population of all eligible American voters requires two conditions to hold. First, the treatment effect should not vary across any characteristic that differs across the population of adult Americans with a phone and the population of eligible American voters (ensuring that any random sample out of either population produces on average the same effect). Second, the standard condition that non-response should be orthogonal to vote preferences (ensuring that any random sample is not biased by those unwilling to be interviewed being systematically different in terms of vote preferences than those that did not accept to be interviewed) (Franco et al., 2017).

¹⁶This procedure is more reliable than using a simple phone book (where about 30% of phone numbers go unlisted).

Although these conditions ultimately concern counterfactual statements that cannot be directly tested using my data there is reason to believe that these conditions are likely to hold in this case.

First, the sample of Gallup’s January 2021 “Mood of the Nations” poll turns out to be relatively similar in terms of known determinants of Republican Party support, as compared to the best available data on the characteristics of the American electorate as a whole. As can be seen in table 1 the Gallup sample has approximately the same share of male, rural, and unemployed individuals as the American National Election Study (ANES) indicated existed in the American electorate in November 2020.¹⁷ Importantly, however, the Gallup sample is on average slightly more Christian, White, college educated, and of age than the U.S. electorate as a whole. To deal with this issue I create probability weights to correct for this (see last two columns in table 1).¹⁸ As shown below the results remain unchanged when estimating the effect in this weighted sample.¹⁹ In addition, I examine whether the treatment effect varies by any of the variables included in table 1 (which is necessary to cause external validity concerns because if all individuals tend to react the same to the treatment, any sample, even a non-random sample, would generate an estimate that generalizes to the entire population – Franco et al. (2017)). I here find that all these interaction terms are highly statistically and economically insignificant (P-value on OLS interaction terms always > 0.207).

Second, there is little reason to believe that in the United States non-response to telephone surveys is endogenous to vote preferences. Keeter et al. (2017) have shown that telephone survey estimates of party affiliation and political ideologically track well with estimates

¹⁷I use the ANES to benchmark my sample averages because other surveys with larger and more rigorous samples, such as the Current Population Survey, the Census, and the American Community Survey, do not include questions regarding voting eligibility. These alternative surveys can thus only be used to study the U.S. population of voting age, which is *not* the same as all Americans with voting rights (e.g., many (ex-)felons do not have voting rights in the United States).

¹⁸I create these weights using Deville and Särndal’s (1992) distance function. In terms of calibration I employ the principle of minimizing the distance between the smallest and largest weight. This leads to a weight range of 0.4 to 2.8 in this case (see online appendix for distribution of weights).

¹⁹Regrettably, this weighting exercise was not included in the pre-analysis plan.

Table 1: Congruence with American National Election Survey (ANES) on observables.

	Unweighted Gallup		Weighted ANES		Weighted Gallup	
	N	Mean	N	Mean	N	Mean
Male	892	0.520 (0.500)	7,372	0.479 (0.009)	849	0.479 (0.023)
Christian	869	0.719 (0.450)	7,354	0.469 (0.008)	849	0.469 (0.023)
White	886	0.790 (0.407)	7,342	0.662 (0.009)	849	0.663 (0.023)
College degree	889	0.481 (0.500)	7,308	0.355 (0.007)	849	0.355 (0.021)
Rural residency	876	0.349 (0.477)	7,353	0.398 (0.010)	849	0.398 (0.023)
Unemployed	890	0.055 (0.228)	7,375	0.049 (0.004)	849	0.049 (0.010)
Age	818	52.770 (17.181)	6,771	46.141 (0.351)	787	46.266 (0.772)

Notes: ANES includes both the pre- and post-election data. ANES is weighted using the probability weights provided in the data itself. Inverse probability weights for the Gallup survey are generated using the “sreweight” Stata package. For the age variable the sample is limited to people below 80 years of age because ANES does not measure age on the interval level above 79, while Gallup does. Gallup weights are generated using four age categories that can be distinguished in both datasets: 18–29, 30–49, 50–64, and 65+.

from high response rate surveys conducted in-person, such as the General Social Survey. Furthermore, analysis of national voter registration files suggest that voter registration with a particular party does not correlate with whether a person is willing to participate in a telephone survey.

5 Average treatment effects

5.1 Nonparametric estimates

In table 2 I present the number and percentage of respondents that indicated that, as of today, they support or lean towards supporting the Republican Party, the Democratic Party, or none of the above, before and after January 6, 2021.

As can be seen, before Donald Trump’s incitement of the insurrection of the U.S. Capitol the Republican Party had 31.6% of certain supporters, and 18.1% of likely supporters, leading to a maximum of approximately 49.7% of the popular vote. After the insurrection the

numbers are 24.6%, 14.3%, and 38.9%, respectively. In total, the Republican Party thus lost approximately 10.8% of its support as a result of January 6, 2021, and most of this electoral loss came from the group of voters that were solid Republican supporters a priori.

Where did the Republican supporters go as a result of January 6, 2021? As can be seen in table 2 most shifted towards the Democratic Party, which received an increase in total support of 8.4%. The remaining 2.4% of respondents came to reclassify themselves as Independents.

Two conclusions can be drawn from the results in table 2.

First, the Republican Party received a “double” electoral penalty as a result of Donald Trump’s incitement of the insurrection: not only did the Republican Party heavily lost many of its own supporters, but most of its own supporters moved directly to the only other party that can reasonably win contemporary U.S. presidential elections: the Democratic Party. This led to a net electoral penalty of 19.2% of the popular vote.

Second, the main effect is not driven by extremist Republican voters turning their back on the Republican Party because they felt that some other Republican leaders (e.g., Mitt Romney, Liz Cheney) were not sufficiently supportive of Donald Trump after he incited the insurrection (i.e., assuming of course that those that felt that the Republican establishment did not do enough to support Donald Trump in the wake of the insurrection would not reclassify themselves as Democrats to give expression to this belief). Rather, the effect appears mostly driven by moderate Republicans willing to move across party lines as a response to the overt undemocratic behavior of a Republican president.

Table 2: Party support before and after the January 6 insurrection of the U.S. Capitol.

	Before insurrection (January 4–5)		After insurrection (January 7–15)	
	Frequency	Percentage	Frequency	Percentage
Republican	56	31.64	176	24.62
Leaning Republican	32	18.08	102	14.27
Independent	13	7.34	70	9.79
Leaning Democrat	31	17.51	151	21.12
Democrat	45	25.42	216	30.21
Total	177	100.00	715	100.00

Notes: Data comes from Gallup’s January “Mood of the Nation” survey.

5.2 Parametric estimates

In order to test whether the results of table 2 are robust to controls, probability weights, geographical fixed effects, and the narrowing of the bandwidth of days around January 6 I turn to estimating the following set of (independent) equations with OLS:²⁰

$$R_i = \beta_0 + \beta_1 X_i + \epsilon_i \quad (1)$$

$$I_i = \alpha_0 + \alpha_1 X_i + \epsilon_i \quad (2)$$

$$D_i = \phi_0 + \phi_1 X_i + \xi_i \quad (3)$$

where R, I, and D are dummies that take the value 1 if a respondent identifies with Republican, Independent, or Democrat, and 0 otherwise. X is a dummy that takes the value 1 if respondent i was surveyed after January 6, 2021, and 0 if he/she was surveyed before this date. β , α , and ϕ are regression coefficients, and ϵ , ϵ , and ξ are errors terms, which I assume to be independently distributed from X given excludability and temporal ignorability. I report Huber–White heteroscedasticity-consistent standard errors.

Table 3 reports the results. Panel A columns (1), (2), and (3) replicate table 2, and find that Republican support went down with 10.8% as a result of Donald Trump’s incitement of the insurrection, and that 8.4% of this loss ended up at the Democratic Party.

²⁰I choose to use separate OLS regressions for simplicity. All results hold with binary and multinomial logit/probit as well (see online appendix).

Table 3: Effect January 6 insurrection on vote preference.

	Republican		Independent		Democrat	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: sample average treatment effect (SATE)						
6 January 2021	-0.108*** (0.042)	-0.095** (0.039)	0.024 (0.023)	0.020 (0.022)	0.084** (0.042)	0.075* (0.039)
Controls	No	Yes	No	Yes	No	Yes
Observations	892	849	892	849	892	849
Panel B: population average treatment effect (PATE)						
6 January 2021	-0.130** (0.057)	-0.125** (0.054)	0.056* (0.033)	0.050 (0.034)	0.074 (0.059)	0.076 (0.056)
Controls	No	Yes	No	Yes	No	Yes
Observations	849	849	849	849	849	849

Notes: Each column within panel A and B is a separate OLS regression. The dependent variable takes the value 1 if the respondent indicates that he/she supports or leans towards Republican/Independent/Democrat, and 0 otherwise. Huber–White heteroscedasticity-consistent standard errors in parentheses. Weights for the PATE estimates are generated using American National Election Study data on the percentage of male, Christian, White, college educated, rural, unemployed, and 18–29, 30–49, 50–64, and 65+ year old individuals that existed in the U.S. electorate in November 2020.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Even in the relatively small sample available these effects are statistically significant on at least the 95% level, and in the case of the main effect on Republican Party support also on the 99% level. This means that one would only find point estimates this large less than 1% of the times if the true effect of Donald Trump’s incitement of the insurrection within the entire population of all American adults with a phone was in fact null.²¹ This suggests that the results are very unlikely to be due to chance and very likely to generalize to Gallup’s target population of individuals that live in the United States and use a telephone at least ones per month.

Gallup’s target population is, however, not exactly the population of interest (i.e., all Americans with voting rights). I therefore probability weight the results using American National Election Study data on the percentage of male, Christian, White, college educated,

²¹This does assume that nonresponse is orthogonal to vote preferences in Gallup’s January 2021 telephone survey (which is shown for U.S. phone surveys in general by Keeter et al. (2017)).

rural, unemployed, and 18–29, 30–49, 50–64, and 65+ year old individuals that existed in the U.S. electorate in November 2020 (see section 4 for how these weights are generated).

Panel B columns (1), (3), and (5) report the weighted estimates. As can be seen using weights to make the sample more closely resemble the population of all eligible American voters, if anything, *strengthens* the results slightly. More specifically, in the weighted sample Donald Trump’s incitement of the insurrection of the U.S. Capitol is estimated to have led to an average 13% decrease in support for the Republican Party. 7.4% of this decrease can be attributed to migration towards the Democratic Party, and 5.6% to Independent.

As I generated the weights based on ten distinct covariates in a sample of only 849 the weighting exercise has somewhat increased the standard errors. The main effect is therefore, even while being economically larger, only statistically significant on the 95% level. This suggests that one would only find such a large effect, assuming that the true effect in the entire population of American voters is null, less than 5% of the time.²²

Columns (2), (4), and (6) show that both the weighted and unweighted results remain essentially unchanged when including the covariates used for the balance-on-observables test in figure 2.²³ All results also hold with geographical fixed effects for state, zipcode, or census region, and when including the following additional individual-level control variables: income, marital status, Muslim, Jewish, Atheist/Agnostic, 4-category Christian classification, self-identified intensity of religiosity, regularity of church attendance, 8-category education classification, Black, Hispanic, Asian, Native Hawaiian, and American Indian or Alaska Native ethnicity/race.

In figure 3 I examine whether the results change when limiting the sample to only respondents that are interviewed in the day(s) directly following January 6, 2021.²⁴ To do so I regress a dummy that takes the value 1 if a respondent supports or leans towards support-

²²Here again I am assuming non-response to be orthogonal to vote preferences.

²³Note that this analysis as well as the weighting exercise itself drops 43 observations due to missing data for some of the covariates. All results hold when adding the covariates one-by-one.

²⁴As mentioned before, limiting the sample to include only days in close proximity to January 6, 2021 limits the likelihood of confounding by collateral events at the cost of precision, generalizability, and a higher likelihood of confounding by imbalance-on-unobservables due to random chance.

ing the Republican Party, and 0 otherwise, on a treatment dummy that always assigns all respondents that were interviewed on January 4 and 5 to the control group, but gradually increases the number of days included in the treatment group with one day at a time. I report both the sample (circles) and population (diamonds) average treatment effects.

As can be seen in figure 3 the estimated effect is highly stable when including/excluding more days after the incitement of the insurrection. The point estimate in all cases hovers around a 10% reduction in vote share for the Republican Party (note that the estimate is not statistically significant on the 95% level in the first two days because of a lack of statistical power, *not* because the treatment effect is economically smaller).


Last, I examine whether, and how long, the negative effect of Donald Trump’s incitement of the insurrection of the U.S. Capitol lasted. Here I am relatively constrained by both the data and the empirical design because the January 2021 Gallup poll only ran until January 15, and, more fundamentally, because increasing the sample to several weeks or even months after January 6 is likely to compromise the identification strategy because it would significantly increase the probability that any observed differences in vote preferences are (also) driven by other events/trends that happened after January 6, 2021.

There is nonetheless at least some suggestive evidence that the effect was long-lasting.

First, there is no sign that the treatment effect diminished within the sample itself. Limiting the sample to the post-January 6 period I find that the number of days a respondent is interviewed after January 6 is uncorrelated with support for the Republican Party (OLS coefficient: 0.002; P-value: 0.803).

Second, I pool data from polls that Gallup conducted in the three months before and after January 6, and run an interrupted time series analysis. I here find that the incitement of the insurrection was correlated with a lasting 3-month decline in support for the Republican Party. Furthermore, an unstructured change point analysis identifies January 6 as the exact turning point in the day-level support for the Republican Party time-series (see online appendix).

Figure 3: SATE and PATE estimates with less days after January 6 included in the sample.



Note: Point estimates and 95% confidence intervals of OLS regressions of supporting the Republican party on treatment status. Dependent variable takes the value 1 if the respondent indicates that he/she supports or leans towards supporting the Republican Party, and 0 otherwise. The horizontal axis plots how many days after January 6, 2021 are included in the treatment group. The control group always consists of all respondents that were interviewed on January 4 and 5. Circles indicate SATE estimates. Diamonds indicate PATE estimates. Weights for the PATE estimates are generated using American National Election Study data on the percentage of male, Christian, White, college educated, rural, unemployed, and 18–29, 30–49, 50–64, and 65+ year old individuals that existed in the U.S. electorate in November 2020.

6 Back-of-the-envelope calculations of electoral penalty

A more than 10% reduction in the popular vote, and an approximately 8% increase in support for the only real competitor party, sounds like a sufficiently large electoral penalty to cause politicians that display overt undemocratic behavior to loose at upcoming elections (i.e.,

assuming of course that the overt undemocratic behavior does not in itself retard reasonably free and fair elections from taking place in the future).

Whether this electoral penalty will, however, be large enough in real world presidential elections is not necessarily clear because of the Electoral College system that blurs the relationship between popular vote shares and winning in presidential elections. Two simple stylized facts nonetheless suggest that the estimated electoral penalty should be more than sufficient to decisively influence presidential elections even when considering potential idiosyncrasies generated by the Electoral College.

First, recall from table 2 that support for the Republican Party, in the most generous counting possible, went from 49.7% before Donald Trump’s incitement of the insurrection, to 38.9% thereafter.²⁵ Even when taking into account that the Electoral College has tended to structurally favor the Republican Party in recent times it appears virtually impossible for any political party in the American political system to win the presidency with only 38.9% of the popular vote.²⁶ This is certainly the case when assuming that an effective two-party system remains in place. Note, for example, that no two-party presidential election in U.S. history was ever won with less than 46.1% of the popular vote (Donald Trump in 2016). But even if an effective two-party system does not remain in place, a 38.9% popular vote share seems to exclude almost certainly the possibility of receiving a plurality of Electoral College votes. Note, for example, that even Abraham Lincoln in the 1860 presidential election (which had four highly competitive contenders and was the election that was won with the lowest popular vote share in U.S. history) still received approximately 39.8% of the popular vote.

Second, holding the voter preferences of the past few presidential elections constant, it appears clear that if the Republican Party would find a way to recover from the estimated electoral penalty it is necessary that the electoral penalty is not quite 10% in the crucial

²⁵When using the probability weights devised in section 4 support for the Republican Party went from 44.4% before the insurrection to 31.4% thereafter.

²⁶The Electoral College has tended to favor the Republican Party in recent decades because: (1) the Republican Party has on average tended to win states with smaller pluralities in recent decades; and (2) the Electoral College gives disproportionately greater weight to small states, which have tended to lean Republican in recent decades.

swing states which tend to be decided by vote margins of (far) less than 10%. This is however not the case. When using the probability weights devised in section 4 I find that the effect does not differ substantively between swing and non-swing states (see the slope of the dotted lines in the right column of figure 4).²⁷ When analyzing the unweighted results I find that the electoral penalty is, if anything, slightly *larger* in the swing states (see left column of figure 4). This is not because the effect of Donald Trump’s incitement of the insurrection of the U.S. Capitol on Republican Party support is more negative in the swing states (the slope of the dotted line in the top left corner of figure 4) but because the migration towards the Democratic Party (as opposed to Independent) is even larger in the swing states (the slope of the dotted line in the bottom left corner of figure 4).

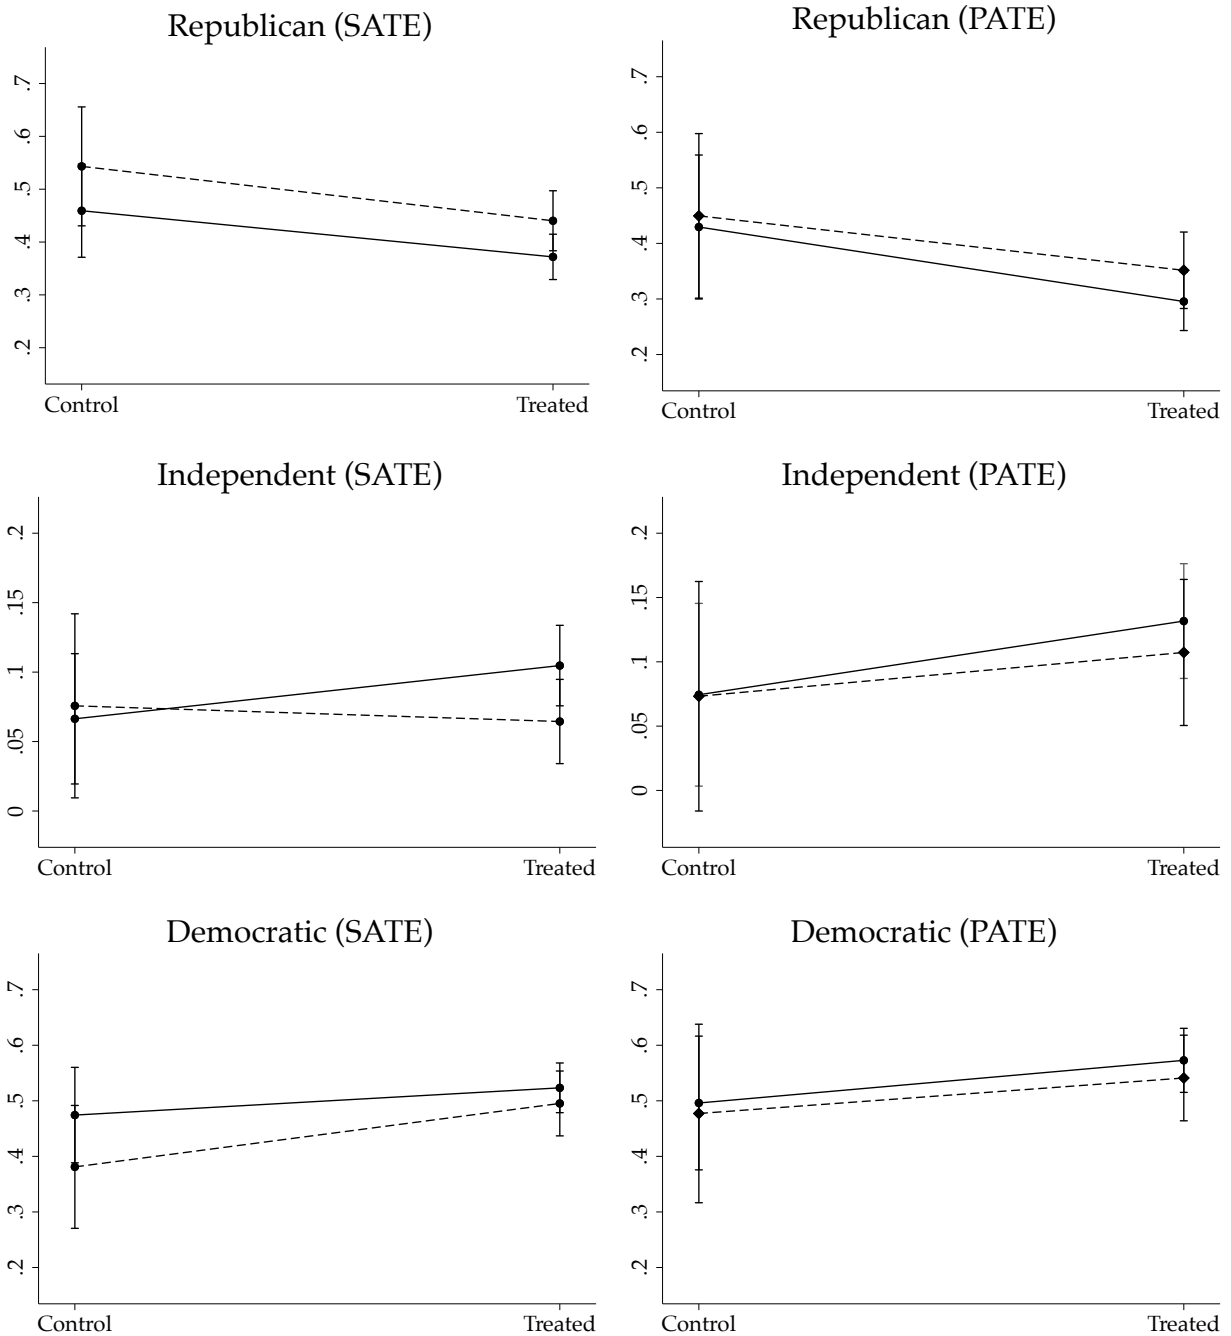
7 Effects on presidential approval rating

The analysis until now has focused on the effect of Donald Trump’s incitement of the insurrection of the U.S. Capitol on party-level voting preferences. This makes sense given that I conceptualize the incitement of the insurrection and the subsequent acquittal and trivialization of president Trump’s behavior by Republican members of Congress as one treatment, and, more importantly, because the U.S. political space can best be understood as a competition between two major political parties, as opposed to a competition between individual politicians. It is nonetheless also of interest how the incitement of the insurrection affected support for Donald Trump himself. Particularly because a significant part of Trump supporters were relatively non-traditional Republicans (Tucker et al., 2019).

To assess the effect of inciting the insurrection on support for Donald Trump as president I regress a dummy that takes the value 1 if a respondent indicates that he/she approves of Donald Trump’s job as president, and 0 if he/she disapproves, on a treatment dummy coding whether the respondent was interviewed before or after January 6, 2021.

²⁷I here code the swing states based on the 2020 presidential election: Arizona, Colorado, Florida, Georgia, Iowa, Maine, Michigan, North Carolina, Ohio, Pennsylvania, Texas, and Wisconsin.


Figure 4: Treatment heterogeneity across swing and non-swing states.



Note: The dashed lines plot the effect within all swing states of the 2020 Presidential election: Arizona, Colorado, Florida, Georgia, Iowa, Maine, Michigan, North Carolina, Ohio, Pennsylvania, Texas, and Wisconsin. Solid lines plot the effect within all other states. Weights for the PATE estimates are generated using American National Election Study data on the percentage of male, Christian, White, college educated, rural, unemployed, and 18–29, 30–49, 50–64, and 65+ year old individuals that existed in the U.S. electorate in November 2020.

Figure 5 reports the results. As can be seen the incitement of the insurrection also caused an approximately 10% decrease in Trump’s presidential approval rating. This result holds when including the controls of figure 2, when including geographical fixed effects, when using the probability weights devised in section 4, and when using binary logit/probit rather than OLS.²⁸

Figure 5: SATE and PATE estimates on Donald Trump’s presidential approval rating.



Note: Point estimates and 95% confidence intervals of OLS regressions of presidential approval rating on treatment status. Dependent variable takes the value 1 if the respondent approves of the job that president Trump is doing, and 0 if the respondent disapproves of the job that president Trump is doing. The horizontal axis plots how many days after January 6, 2021 are included in the treatment group. The control group always consists of all respondents that were interviewed on January 4 and 5. Circles indicate SATE estimates. Diamonds indicate PATE estimates. Weights for the PATE estimates are generated using American National Election Study data on the percentage of male, Christian, White, college educated, rural, unemployed, and 18–29, 30–49, 50–64, and 65+ year old individuals that existed in the U.S. electorate in November 2020.

²⁸Similar results are found when analyzing the effect on how a respondent thinks that “history will judge Donald Trump”, a question that Gallup asked with regard to nine recent presidents (see online appendix).

8 Conclusion

This paper has shown that the American electorate *does* act as an effective electoral check on democratic backsliding, at least when the undemocratic behavior by a politician is as extreme as inciting an insurrection to overturn the result of a free and fair election. This is in contrast to much recent survey experimental evidence that has generally found that only a small fraction of Americans tends to prioritize democratic principles in their electoral choices, particularly in the current highly polarized climate.

This is a hopeful message and reinforces Weingast’s (1997) claim, made long before the extreme partisan polarization dominating American politics today, that: “The U.S. Constitution has proved binding in practice partly because citizens are willing to defend it by reacting against proposed violations.” (p. 254)

At the same time it is also clear that the American electorate is very far from a *perfect* check on undemocratic behavior by elected officials. Indeed, if the American electorate was a perfect check on undemocratic behavior by politicians we would presumably not see the many instances of smaller, but nonetheless very serious, forms of democratic norm transgressions (e.g., gerrymandering). Further research is necessary with regard to the important question of how extreme undemocratic behavior has to become before American voters react in sufficient severity, and, crucially, how the electoral check on democratic backsliding can be strengthened by lowering the bar with regard to the type of undemocratic behavior that is electorally sanctioned at the polls.

References

- Acemoglu, Daron, James A Robinson and Ragnar Torvik. 2013. “Why Do Voters Dismantle Checks and Balances?” *Review of Economic Studies* 80(3):845–875.
- Agosta, Michael and Christy Lightbourn. 2021. “Are Republicans Leaving Their Party?” Available at: https://vvnstates.org/static/landing/media/are_republicans_leaving_the_party.pdf (last accessed: September 15, 2021).
- Albertus, Michael and Guy Grossman. 2021. “The Americas: When Do Voters Support Power Grabs?” *Journal of Democracy* 32(2):116–131.
- Bermeo, Nancy. 2016. “On Democratic Backsliding.” *Journal of Democracy* 27(1):5–19.
- Carey, John, Katherine Clayton, Gretchen Helmke, Brendan Nyhan, Mitchell Sanders and Susan Stokes. 2020. “Who Will Defend Democracy? Evaluating Tradeoffs in Candidate Support Among Partisan Donors and Voters.” *Journal of Elections, Public Opinion and Parties* 0(0):1–16.
- Cho, Young Ho and Injeong Hwang. 2021. “Who Defends Democracy and Why? Explaining the Participation in the 2016–2017 Candlelight Protest in South Korea.” *Democratization* 28(3):625–644.
- Cleary, Matthew R and Aykut Öztürk. 2020. “When Does Backsliding Lead to Breakdown? Uncertainty and Opposition Strategies in Democracies at Risk.” *Perspectives on Politics* pp. 1–17.
- Deville, Jean-Claude and Carl-Erik Särndal. 1992. “Calibration Estimators in Survey Sampling.” *Journal of the American statistical Association* 87(418):376–382.
- Diamond, Larry. 2021. “Democratic Regression in Comparative Perspective: Scope, Methods, and Causes.” *Democratization* 28(1):22–42.
- Ding, Iza and Dan Slater. 2021. “Democratic Decoupling.” *Democratization* 28(1):63–80.
- Economist. 2021. “Nearly half of Republicans support the invasion of the US Capitol.” January 7, 2021. Available at: <https://www.economist.com/graphic-detail/2021/01/07/nearly-half-of-republicans-support-the-invasion-of-the-us-capitol> (last accessed: September 15, 2021).
- Foa, Roberto Stefan and Yascha Mounk. 2016. “The Danger of Deconsolidation: The Democratic Disconnect.” *Journal of Democracy* 27(3):5–17.
- Franco, Annie, Neil Malhotra, Gabor Simonovits and LJ Zigerell. 2017. “Developing Standards for Post-Hoc Weighting in Population-Based Survey Experiments.” *Journal of Experimental Political Science* 4(2):161–172.
- Graham, Matthew H. and Milan W. Svobik. 2020. “Democracy in America? Partisanship, Polarization, and the Robustness of Support for Democracy in the United States.” *Amer-*

- ican Political Science Review* 114(2):392–409.
- Grillo, Edoardo and Carlo Prato. Forthcoming. “Reference Points and Democratic Backsliding.” *American Journal of Political Science* .
- Grossman, Guy, Dorothy Kronick, Matthew Levendusky and Marc Meredith. 2021. “The Majoritarian Threat to Liberal Democracy.” *Journal of Experimental Political Science* pp. 1–10.
- Haggard, Stephan and Robert Kaufman. 2021. *Backsliding: Democratic Regress in the Contemporary World*. Cambridge University Press.
- Keeter, Scott, Nick Hatley, Courtney Kennedy and Arnold Lau. 2017. “What Low Response Rates Mean for Telephone Surveys.” Pew Research Center. May 15, 2017. Available at: <https://www.pewresearch.org/methods/2017/05/15/what-low-response-rates-mean-for-telephone-surveys/> (last accessed: September 15, 2021).
- Lührmann, Anna. 2021. “Disrupting the Autocratization Sequence: Towards Democratic Resilience.” *Democratization* 28(5):1017–1039.
- Magaloni, Beatriz. 2006. *Voting for Autocracy: Hegemonic Party Survival and its Demise in Mexico*. Vol. 296 Cambridge University Press.
- McCoy, Jennifer, Gabor Simonovits and Levente Littvay. 2020. “Democratic Hypocrisy: Polarized Citizens Support Democracy-Eroding Behavior When Their Own Party is in Power.” *APSA Preprints*. doi: 10.33774/apsa-2020-sk4k9.
- Muñoz, Jordi, Albert Falcó-Gimeno and Enrique Hernández. 2020. “Unexpected Event during Survey Design: Promise and Pitfalls for Causal Inference.” *Political Analysis* 28(2):186–206.
- New York Times. 2021a. “How Far Are Republicans Willing to Go? They’re Already Gone.” June 9, 2021. Available at: <https://www.nytimes.com/2021/06/09/opinion/us-democracy-republicans.html> (last accessed: September 15, 2021).
- New York Times. 2021b. “‘There’s Nothing Left’: Why Thousands of Republicans Are Leaving the Party.” February 10, 2021. Available at: <https://www.nytimes.com/2021/02/10/us/politics/republicans-leaving-party.html> (last accessed: September 15, 2021).
- New York Times. 2021c. “Trump Has Always Been a Wolf in Wolf’s Clothing.” June 9, 2021. Available at: <https://www.nytimes.com/2021/01/07/opinion/trump-capitol-protests.html> (last accessed: September 15, 2021).
- Reuters. 2021a. “Exclusive: FBI Finds Scant Evidence U.S. Capitol Attack was Coordinated.” August 20, 2021. Available at: <https://www.reuters.com/world/us/exclusive-fbi-finds-scant-evidence-us-capitol-attack-was-coordinated-sources-2021-08-20/>

- (last accessed: September 23, 2021).
- Reuters. 2021b. “Pentagon Moved Slowly to Approve Jan. 6 Capitol Police Request for National Guard.” March 3, 2021. Available at: <https://www.reuters.com/world/us/pentagon-moved-slowly-approve-jan-6-capitol-police-request-national-guard-2021-03-03/> (last accessed: September 23, 2021).
- Slate. 2021. “Republicans Still Sympathize With the Insurrection.” April 15, 2021. Available at: <https://slate.com/news-and-politics/2021/04/republican-party-sympathize-capitol-insurrection.html> (last accessed: September 15, 2021).
- Svolik, Milan W. 2019. “Polarization versus Democracy.” *Journal of Democracy* 30(3):20–32.
- Svolik, Milan W. 2020. “When Polarization Trumps Civic Virtue: Partisan Conflict and the Subversion of Democracy by Incumbents.” *Quarterly Journal of Political Science* 15(1):3–31.
- Tucker, Patrick D, Michelle Torres, Betsy Sinclair and Steven S Smith. 2019. “Pathways to Trump: Republican Voters in 2016.” *Electoral Studies* 61:1–14.
- Waldner, David and Ellen Lust. 2018. “Unwelcome Change: Coming to Terms with Democratic Backsliding.” *Annual Review of Political Science* 21:93–113.
- Weingast, Barry R. 1997. “The Political Foundations of Democracy and the Rule of the Law.” *American Political Science Review* 91(2):245–263.
- Wuttke, Alexander, Konstantin Gavras and Harald Schoen. 2020. “Have Europeans Grown Tired of Democracy? New Evidence from Eighteen Consolidated Democracies, 1981–2018.” *British Journal of Political Science* pp. 1–13.
- YouGov. 2021. “Most Voters Say the Events at the US Capitol Are a Threat to Democracy.” January 6, 2021. Available at: <https://today.yougov.com/topics/politics/articles-reports/2021/01/06/US-capitol-trump-poll> (last accessed: September 15, 2021).