Quarterly Journal of Political Science, 2024, 19: 387–432

District Competitiveness Increases Voter Turnout: Evidence from Repeated Redistricting in North Carolina

Robert Ainsworth[1], Emanuel Garcia Munoz[2] and Andres Munoz Gomez[3]*

[1] *University of Florida, Gainesville, FL, USA; robert.ainsworth@ufl.edu*
[2] *Palm Beach Atlantic University, West Palm Beach, FL, USA; emanuel_garciamunoz@pba.edu*
[3] *University of Florida, Gainesville, FL, USA; amunozgomez@ufl.edu*

ABSTRACT

We study whether competitive legislative districts cause higher voter turnout. To do so, we employ rich data on the 2006 to 2020 elections in North Carolina. We make use of variation in district competitiveness due to repeated bouts of redistricting, a process in which district boundaries are redrawn. Specifically, we compare people who share the same districts in each legislative chamber (U.S. House, NC Senate, NC House) before redistricting but who differ in districts after redistricting. We match these people on demographics, party registration, and pre-redistricting turnout. We then track their turnout behavior in post-redistricting elections. For the U.S. House, switching from an uncompetitive "80–20" district to a competitive "55–45" district increases turnout by a rate of 1 percentage point per election of exposure. For the state chambers, the magnitude is 0.6. Effects are highly persistent

*For helpful comments, we thank Matt Backus, Michael Best, Damon Clark, Rajeev Dehejia, Allan Drazen, Carlos Estrada, Anthony Fowler, Jon Hamilton, Ethan Kaplan, Ilyana Kuziemko, Cameron LaPoint, Suresh Naidu, Andy Pham, Cristian Pop-Eleches, Richard Romano, Steve Slutsky, Miguel Urquiola, Stephane Wolton, and seminar participants at Columbia University, the University of Maryland, the Public Choice Society Annual Conference, the North American Summer Meeting of the Econometric Society, and the APPAM Fall Research Conference. All errors are our own.

Online Appendix available from:
http://dx.doi.org/10.1561/100.00022114_app
Supplementary Material available from:
http://dx.doi.org/10.1561/100.00022114_supp
MS submitted on 24 August 2022; final version received 22 November 2023
ISSN 1554-0626; DOI 10.1561/100.00022114

and sum across chambers. They appear to be explained in part by a learning channel, where living in a competitive district induces people to believe that races can be competitive.

Keywords: Voter turnout; competitiveness; gerrymandering

A central question in research on voting is whether the competitiveness of the electoral environment affects turnout. This question is of both scientific and policy interest. The answer has the potential to illuminate the mechanics of voting behavior. It also can help guide policymakers in designing institutions that promote turnout.

From a theoretical point of view, the relationship between turnout and competitiveness is uncertain. On the one hand, there are ample reasons why a person may be more likely to vote when subject to a competitive electoral environment. On the other hand, there are also stories in which competitiveness may not affect turnout. Consistent with the theoretical ambiguity, empirical evidence on the link between competitiveness and turnout is mixed. Papers using cross-sectional techniques often measure a strong association (Cancela and Geys, 2016). In addition, papers have identified non-zero causal impacts in lab experiments (Agranov *et al.*, 2017) and when studying natural experiments in Europe.[1] However, the few causal papers that study the United States find little effect. For instance, Enos and Fowler (2014) and Gerber *et al.* (2020) detect minimal responses in field experiments that randomly provide people with information on competitiveness. Also, Moskowitz and Schneer (2019) observe only modest impacts when they examine the short-run consequences of changes in the competitiveness of voters' congressional districts.[2]

In this paper, we expand on the work in Moskowitz and Schneer (2019). Like them, we study the turnout impacts of the competitiveness of American legislative districts. Unlike them, we provide results over both the short- and long-run and for districts associated with both federal and state legislatures. These additions are important due to three features of legislative districts. First, districts are meant to last for multiple elections (up to a decade). Thus, it is interesting to know how turnout depends on the length of time that a person spends in a competitive or uncompetitive district. Second, districts are eventually redrawn. Thus, we want to see whether effects persist or decay

[1]These relate to runoff elections (Fauvelle-Aymar and Francois, 2006; Indridason, 2008; Paola and Scoppa, 2014; Simonovits, 2012) and access to polling information (Bursztyn *et al.*, 2023; Morton *et al.*, 2015).

[2]Suggestively, Cantoni and Pons (2022) show that competitiveness is correlated with states' causal effects on turnout. However, they cannot disentangle whether this is due to competitiveness or other attributes of a state.

after a person's district is replaced. Third, people have multiple districts (one for each legislative chamber). Thus, we care about how effects aggregate over each of the districts that a person is in.

To obtain results, the paper uses rich longitudinal data from the state of North Carolina. The data includes each person who was registered to vote (a "registrant") in each of the 2006 to 2020 elections. We exploit variation in competitiveness due to "redistricting", the process in which district boundaries are redrawn. Specifically, we study people who lived in the same districts before redistricting but who are assigned to districts of differing competitiveness during redistricting. We match these people on covariates — including pre-redistricting turnout — and we run tests to confirm that the matching eliminates bias. We then examine how the redistricting-induced differences in competitiveness align with post-redistricting differences in turnout.

One point about the paper bears emphasizing. The primary treatment variable is the underlying competitiveness of a legislative district, not of a race in a particular election-year.[3] We focus on *district* competitiveness because it is a lever that can be manipulated by policymakers. Namely, the measures of district competitiveness that we use can be calculated at the time of redistricting. As a result, the causal effects that we recover can be employed during redistricting to forecast the impacts of different district configurations. Nonetheless, when we explore mechanisms, we also examine the effects of exposure to competitive races.

We use multiple measures of district competitiveness. The measures center on predicting a district's two-party vote share in a generic race in a post-redistricting election. The predictions make use of data on individual-level turnout and precinct-level votes from pre-redistricting elections. Two of the measures are new. A third is based on the Partisan Voting Index (PVI) from the Cook Political Report, which is relied on in prior literature. When exploring mechanisms, we also use two measures of a race's competitiveness. These are the amount of spending in the race and the closeness of the vote margin.

Our empirical strategy builds on a growing literature that uses redistricting to gain variation in the characteristics of legislative districts.[4] In North Carolina, there were repeated bouts of redistricting in the 2010s. In 2011, districts were redrawn both for the federal legislature (the U.S. House of Representatives) and for the two state legislatures (the NC Senate and NC

[3] A district's competitiveness is the district's propensity, across election-years, to have competitive races. Whether a district has a competitive race in a particular year depends in part on stochastic factors, such as the electoral swings that occur in the year and the quality of the candidates that run in the race.

[4] Examples include Ansolabehere *et al.* (2000) and Sekhon and Titiunik (2012) on the strength of the incumbency advantage, Henderson *et al.* (2016) on an incumbent's race/ethnicity, Fraga (2016) on a district's racial composition, Moskowitz and Schneer (2019) on district competitiveness, and Fraga *et al.* (2021) on a district's partisan makeup.

House). Depending on the chamber, districts were redrawn again — in response to court orders — in 2015, 2017, and 2019. We define a redistricting "episode" as an instance in which districts for a particular chamber are redrawn. We make use of all the episodes in the 2010s via a stacking approach.

The empirical strategy has two key components. First, we divide North Carolina into regions based on the districts that were used before and after a redistricting episode. In our main analysis, we define regions as areas that have all the same pre-redistricting districts and that differ in districts for only a single chamber after redistricting. This construction lets us identify the effect of changing competitiveness for just one of a person's three districts. In an additional analysis, we define regions based only on pre-redistricting districts, and we run comparisons for people who may differ in multiple post-redistricting districts. This lets us explore how the effects of competitiveness aggregate across legislative chambers.

The second component is to exact-match individuals within regions. Matching deals with the issue that there may be differences between people who are placed into more or less competitive districts, even among those who live in the same region (Henderson *et al.*, 2016). We match on a variety of demographic and political variables. In different specifications, these include race/ethnicity, gender, age, home value, neighborhood characteristics, additional districts, registration party, and turnout in up to three pre-redistricting elections.

We validate our strategy in multiple ways. Among other things, we show that redistricting-induced differences in competitiveness are not associated with turnout in pre-redistricting elections that are not used in matching. Also, results exist for all competitiveness measures and are robust to altering the construction of regions or changing the covariate sets on which we match. In addition, results are not distorted by pre-redistricting district experiences or by the fact that, in North Carolina, a district's competitiveness is correlated with both the share of its residents who are racial minorities and the share who are Democrats.

Our main finding is that district competitiveness influences turnout. Further, effects (i) scale with the number of elections that people spend in competitive districts, (ii) persist after districts are replaced, and (iii) are additive across legislative chambers. In order to understand magnitudes, consider the impact of switching between two hypothetical districts — an "80–20" district (in which the parties are expected to split the two-party vote 80% to 20%) and a "55–45" district (where the expected split is 55–45). For the U.S. House, we find that switching into the more competitive district would cause turnout to grow by 1 percentage point for each election spent in the district. For the state chambers, the magnitude is 0.6 percentage points. If people were to experience a change in districts for multiple chambers, the overall effect would be the sum of the chamber-specific effects. To understand the persistence finding, suppose that in a subsequent round of redistricting, people in the 80–20 and 55–45 districts are put into districts of equal competitiveness. We find that people who had been living in the more competitive district would continue to turn out more. In fact, we see no evidence of a decay in the effect for at least three elections after the districts are replaced, the longest period available in the data.[5]

We lend insight into the pattern of effects by showing a few additional results. First, district competitiveness operates by increasing the number of consistent voters, defined as those who turn out in a given election and all (observable) subsequent elections. By contrast, district competitiveness does not work by inducing one set of people to turn out in one election and a substantially different set to turn out in a different election. Second, effects depend on exposure to competitive races — and, specifically, to races with close vote margins. Third, there is limited heterogeneity in effects across registrant types. However, effects are larger for young people and largest for young people in well-educated neighborhoods. They are also larger for whites than for racial minorities, which we show is not due to the influence of majority–minority districts.

Given these findings, we interpret the results as suggestive of a learning story. Competitive districts give people a chance to realize that races can be competitive. This realization — together with other ways by which voting is habit-forming — cause the affected people to continue turning out for many future elections. Effects are magnified for people with more opportunity to learn: those who have less electoral experience (young people), those with more ability to discern competitiveness (young people in well-educated areas), and those who live in competitive districts for longer. That said, we cannot rule out that other explanations also play a role.

We assess external validity by benchmarking our results with Moskowitz and Schneer (2019). That paper has data from all 50 states; thus, the exercise lets us compare North Carolina with the country as a whole. When we use the same empirical strategy and time period, we match the paper's magnitudes. As such, our findings may have national relevance — though we note, as a caveat, that the comparison can only be conducted during a portion of our sample period.

The last part of the analysis is to investigate how the competitiveness of North Carolina's recent legislative districts affected election outcomes. To do this, we consider a counterfactual situation where each district used during 2012–2020 was a 55–45 district. We ask how turnout and votes would have differed if people had lived in these (moderately) competitive districts rather than their actual districts. We predict that turnout would have increased

[5]The discussion is for general-election turnout. Competitiveness does not seem to impact turnout in primaries.

for most registrants. By 2020, overall turnout would have been higher by 2.5 percentage points (a 4.1% change). Increases are similar across parties in percentage-point terms. However, they are larger for Democrats on aggregate. We translate the turnout impacts into effects on votes by using information on registrants' probabilities of preferring Democratic or Republican candidates. We find that competitive districts would have caused little change in statewide vote margins. This is because the boost in turnout among Democrats is driven by people with a non-trivial chance of voting for Republicans.

Our paper relates to a number of literatures. First, the paper contributes to the previously discussed literature on whether electoral competitiveness spurs turnout. Our results indicate that it can, even in a context where elections feature a variety of races on the ballot. Second, the paper adds to a voluminous literature on the determinants of voting behavior. It provides empirical evidence that is consistent with models in which voters learn from past experiences (Bendor *et al.*, 2003; Esponda and Pouzo, 2017; Kanazawa, 1998) and with research that voting is habit-forming (Coppock and Green, 2016; Fujiwara *et al.*, 2016; Gerber *et al.*, 2003; Meredith, 2009). Third, the paper contributes to research on majority–minority districts (Davis, 2019; Fraga, 2016; Hayes and McKee, 2012; Washington, 2012). It suggests that the uncompetitive nature of these districts reduces turnout. Fourth, the paper adds to a recent literature that uses causal methods to explore which electoral policies affect turnout (Bhatt *et al.*, 2020; Braconnier *et al.*, 2017; Cantoni, 2020; Cantoni and Pons, 2019; Kaplan and Yuan, 2020). It highlights that the drawing of competitive or uncompetitive legislative districts is one such policy.

Conceptual Overview

We next provide an overview of some key concepts that underlie the paper's results. We first consider the relationship between turnout and competitive races. We then discuss the impact of competitive districts. Finally, we review evidence on persistence in voting behavior.

From a theoretical perspective, there are stories both supporting and questioning whether competitive races affect turnout. On the one hand, there are many reasons why they may increase turnout. First, people may think they have a better chance of influencing a race's outcome when the race is competitive and thus likely to have a close vote margin (Downs, 1957). Second, the expressive benefits that people get from voting may grow with a race's competitiveness (Hamlin and Jennings, 2011; Kawai *et al.*, 2021; Riker and Ordeshook, 1968). Third, people in competitive races may feel a stronger civic duty to vote (Dellavigna *et al.*, 2016; Feddersen and Sandroni, 2006). Fourth, competitive races may generate more media attention or popular buzz

(Clarke and Evans, 1983). This could help to remind potential voters that an election is occurring or to reduce their uncertainty about the quality of the candidates (Degan and Merlo, 2011). Fifth, competitive races may attract better candidates, who are able to motivate additional turnout (Stephanopoulos and Warshaw, 2020). Sixth, national and local political parties may devote more resources to competitive races. This could lead to more spending on get-out-the-vote efforts, including voter registration drives, advertising campaigns, or door-knocking and canvassing operations (Cox and Munger, 1989; Enos and Fowler, 2018; Hill and McKee, 2005; Shachar and Nalebuff, 1999).

Nonetheless, there are also reasons why a race's competitiveness may not matter for turnout. First, there is evidence that people have little sense of whether a race is competitive (Gerber *et al.*, 2020; McDonald and Tolbert, 2012; Moskowitz and Schneer, 2019); given this, the stories that depend on voters being aware of competitiveness may not be applicable. Second, elections often feature multiple races, such as races for different political offices, as well as ballot initiatives and referenda. It is possible that the competitiveness of a single race on the ballot is not enough to sway a person's overall turnout decision. Third, the net effect of get-out-the-vote efforts is unclear. Namely, parties may work both to increase turnout for their supporters and to decrease turnout for other types of voters (Spenkuch and Toniatti, 2018). Aggregated over all parties, these contrasting efforts may cancel. Thus, it may be that competitive races have little effect on turnout.

The turnout effect of competitive districts depends closely on that of competitive races. The main manner by which competitive districts may matter is by increasing the probability that a person experiences a competitive race. That said, districts may also matter on their own, as a result of being used for multiple elections. In particular, living in a competitive district makes people likely to experience a succession of competitive races, including some that have close vote margins. This gives people more chance to realize that races in their district are competitive or that races can be competitive in general. It also may allow for the development of social processes that encourage turnout via peer effects, and it may lead to stronger local parties, which are better at engaging residents and being a presence in the community. Nonetheless, the impacts of the above mechanisms are uncertain. For instance, evidence is mixed as to whether learning about competitiveness affects turnout. Also, as with get-out-the-vote efforts, the net effect of stronger parties may be zero. Finally, peer effects are a downstream phenomenon and arise only if there are some individuals for whom competitiveness matters. Thus, it is unclear a priori whether competitive districts affect turnout.

There is now a large literature that studies persistence in voting behavior; it has a few findings. First, "some people always vote and others never vote" (Gerber *et al.*, 2020), and the share who vote increases with age (Fujiwara *et al.*, 2016). Second, voting exhibits habit formation, meaning that the very act of voting makes a person more likely to vote in the future.[6] Third, people vote "retrospectively", based on past experiences (Healy and Malhotra, 2013). In addition, they do so in a boundedly rational manner in which they often rely on rules of thumb (Bendor *et al.*, 2003; Esponda and Pouzo, 2017) and in which they are heavily influenced by experiences in young adulthood (Ghitza *et al.*, 2022).

Given the existing literature and the nature of legislative districts, we predict that if district competitiveness affects turnout, its effects will be persistent. This is for three reasons. First, habit formation implies that any shock to turnout will partially persist. Second, under some of the mechanisms by which district competitiveness may matter, the conditions that spur turnout will also partially persist. In particular, if it takes time for people to notice that their district is competitive, for local parties to strengthen, or for there to be an enhanced social norm regarding voting, then it will also take time for these factors to decay once districts are replaced.[7] Finally, if people rely on rules of thumb — e.g., "I vote because I remember that one close race" — rather than continuous updating, then turnout effects may never decay. This could help to explain why turnout increases with age.

Preliminaries

We now discuss the paper's setting, data, and competitiveness measures.

Setting

The setting is the 2006 to 2020 general elections in North Carolina. A few features of this setting are worth mentioning. First, in North Carolina, general elections involve multiple races. In midterms, there are legislative races, races for judicial offices, ballot initiatives, local races, and, possibly, races for the U.S. Senate. In presidential years, there are also races for president and for state offices, such as governor or attorney general.

Second, in North Carolina, there are three legislative chambers in which races vary by district: the U.S. House of Representatives, the NC Senate, and the NC House. During the sample period, the U.S. House had 13 districts, the NC Senate had 50, and the NC House had 120. In all chambers, representatives face reelection every two years.

[6] Fujiwara *et al.* (2016) show that a percentage point increase in turnout in a presidential election increases turnout in the next presidential election by 0.6–1 percentage points. Coppock and Green (2016) and Meredith (2009) find that being just age-eligible to vote in one election increases turnout in future elections for up to 20 years. Coppock and Green (2016) also find that habit formation is stronger for elections of the same type.

[7] In contrast, expenditures from non-local parties are likely to respond quickly to current competitiveness.

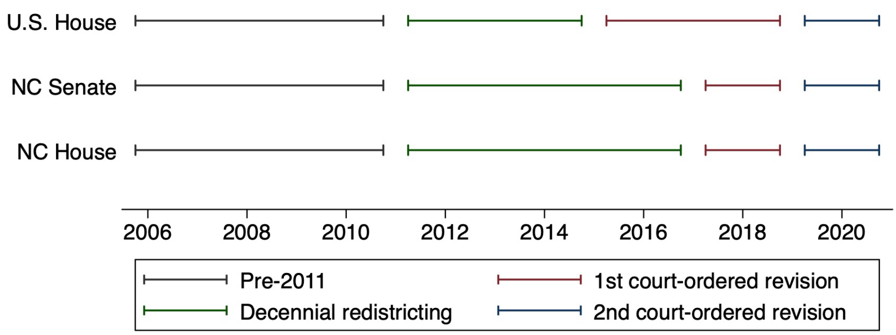

Figure 1: The districts used in recent North Carolina elections.

The figure summarizes the district configurations that were used in elections since 2006. Districts for elections prior to 2011 were drawn by the state legislature in 2002 for the U.S. House and in 2004 for the NC Senate and NC House. In addition, districts for the NC House were modified to a small degree between the 2008 and 2010 elections. Due to the minor nature of this change, we ignore it. See the text for details on the districts used after 2011.

Third, legislative districts in North Carolina were redrawn on multiple occasions during the 2010s. The timeline of these changes is summarized in Figure 1. In 2011, districts were redrawn for each of the three legislative chambers, as part of the U.S.'s decennial redistricting process. The new districts were drawn by the Republican-controlled state legislature and were immediately challenged in court. The districts for the U.S. House were used for the 2012 and 2014 elections. However, in advance of the 2016 election, they were deemed a "racial gerrymander" and were overturned.[8] The legislature then revised the districts, and the new configuration was used for the 2016 and 2018 elections. For the NC Senate and the NC House, the original 2011 districts were used during 2012 to 2016. Prior to the 2018 election, they too were ruled a racial gerrymander and were revised. Finally, in 2019, a state court overturned the revised districts for all of the chambers. The legislature then drew new districts, which were used for the 2020 election. In sum, in the 2010s, districts were redrawn on three occasions for each of the three chambers. This makes for a total of nine redistricting "episodes".

The North Carolina setting is advantageous for three reasons. First, the state government supplies publicly available data that is unusual in its richness. The data allows observing everyone who was registered in a given election and tracking their voting behavior and exact addresses over a long period of

[8]Specifically, in Harris v. McCrory (2016), a federal court said that the district configuration was excessive in the extent to which it packed racial minorities into the same districts.

time.[9] Second, the large number of redistricting episodes means that there is considerable variation in district competitiveness, which increases statistical precision. Third, the fact that districts are both implemented and revised lets us explore dynamics. Namely, we can see how effects develop when a district is in use and how they change when the district is replaced.

Despite the advantages, North Carolina also has traits that may limit its representativeness. First, due to repeated redistricting, North Carolinians experienced frequent changes in districts during the sample period. By contrast, in other states, districts often last an entire decade. The frequency of redistricting may affect the degree to which North Carolinians notice and respond to district attributes. Second, North Carolina commonly has competitive races at the statewide level. These give residents a strong incentive to turn out, even in uncompetitive districts. Meanwhile, there are other states where statewide races are one-sided and where a competitive legislative race would be one of the most salient contests on the ballot. Fortunately, we are able to provide evidence that effects in North Carolina are similar to those in other states, assuaging concerns.

Data

We use data on registrants, legislative races, legislative districts, and precinct-level vote shares.

The data on registrants comes chiefly from the state election authority, the North Carolina State Board of Elections (NC SBE). From this source, we obtain snapshots of North Carolina's voter registration database in each year from 2006 to 2020. The snapshots provide information on all individuals who were registered to vote in the state at a specified point in time. Importantly, they include a unique registrant ID, which allows us to link them longitudinally.

The first step in compiling the registrant data is to choose a population from which to draw registrants. We focus on people who were registered in the election prior to redistricting, which we call the "baseline" election.[10] For people in the baseline election, we use the linked registration snapshots to acquire a number of covariates. First, we obtain a registrant's year of birth, gender, self-identified race/ethnicity, date of first registration, and year of death. Second, we obtain the registrant's turnout behavior, party registration,

---

[9] Similar data can be obtained at a national level from private vendors. However, due to high costs, researchers generally make sacrifices when buying it. For instance, the data in Moskowitz and Schneer (2019) ends in 2014, a mere two elections after the decennial redistricting. Similarly, the data in Cantoni and Pons (2022) lists people's locations at the level of the county, which is too aggregate for our purposes.

[10] For the decennial redistricting episodes, the baseline election is 2010. For the first court-ordered revision, it is 2014 for the U.S. House and 2016 for the state chambers. For the second revision, it is 2018 for all chambers.

legislative districts, and exact address in each of the 2006 to 2020 elections.[11] We then merge covariates from two other sources. From U.S. Census data, we add the population density of the census block in which the registrant lived during the baseline. In addition, we add the median household income and the share college graduates in the registrant's baseline block-group. From the NC One Map project, we collect information on the value of the property parcel associated with the registrant's baseline address.

The second type of data concerns the legislative races that occurred in North Carolina during 2006 to 2020. This data is of two varieties. First, from the NC SBE, we gather data on race vote shares. We use these to calculate the closeness of a race, which we define as 1 minus the absolute two-party vote-share margin. Second, we assemble data on race spending. For U.S. House races, this comes from the FEC; for NC Senate and NC House races, it is from Follow the Money. We focus on the total spending in a race, which we compute by summing campaign contributions and independent expenditures for the two highest vote-getters. We convert to a per-capita measure by dividing by the district population.

The third category of data is on the districts that were used in North Carolina during 2006 to 2020. We obtain information on district boundaries from the NC SBE. Merging with the registrant data, we calculate the shares of registrants in a district who are racial minorities and Democrats. We also compute three measures of a district's competitiveness.

The last type of data is precinct-level vote shares. We use these in calculating the competitiveness measures, and we obtain them from the NC SBE. The vote shares are for the 2008–2020 elections. They cover U.S. House races and a variety of statewide races (U.S. President, U.S. Senate, NC Governor, etc.). Summary statistics for the data are presented in Online Appendix A1.

Measuring district competitiveness

We calculate district competitiveness measures in two steps. First, we predict district vote shares using only the information available at the time of redistricting. Second, we define competitiveness as 1 minus a district's absolute predicted two-party vote-share margin. Under this scaling, competitiveness equals 1 in a "50–50" district where Democrats and Republicans are predicted

[11] People do not appear in a registration snapshot if they have not yet registered or if their registration lapses. This occurs if they die, or if they do not vote in four consecutive North Carolina general elections and also do not respond to a mailed information card. When conducting the analysis, we drop people who die before the 2020 election. For other people with lapsed registrations, we set their address to the value from when they were last registered, and we set the turnout variable to zero. A related issue is that we lack data on registration or voting from states other than North Carolina. This means that we misclassify residential location and, possibly, turnout for people who leave the state. We discuss this in more depth when explaining our empirical strategy.

to be equally popular. It equals 0 in a "100–0" district where one party is predicted to gain all of the two-party vote.[12]

The competitiveness measures differ in the procedures that they use for predicting district vote shares. The procedure for our main measure has four steps. First, we regress precinct-level vote shares from the baseline election on precinct-level means of voter characteristics. Second, we use the regressions to generate individual-level predictions for a person's probability of preferring candidates from each party. We compute these for all baseline registrants, including non-voters. Third, we aggregate the preference probabilities to the district level by averaging over registrants whose baseline address falls within district boundaries. Fourth, in averaging, we weight the preference probabilities by registrants' pre-redistricting turnout histories. The weights tilt the averages toward the preferences of registrants who are more likely to turn out in post-redistricting elections.

The procedure for the second measure is similar but does not use a regression to calculate individual-level preference probabilities. Instead, it assigns each registrant a value equal to the vote share in the registrant's precinct. Thus, the measure is based on a weighted average of precinct-level vote shares. Finally, the procedure for the Cook measure is again similar, but with two distinctions. First, it uses only vote shares from the presidential race; second, it relies on data from two presidential elections prior to redistricting, not the baseline election.

For reasons that will be explained later, we calculate separate versions of the competitiveness measures for each baseline (or for each set of two presidential elections). Thus, for a given district, we have a version that uses the baseline for the redistricting episode that created the district (the "primary" version). However, we also have secondary versions that use the baselines for the other episodes. This way, we observe the competitiveness of each district during each baseline.[13]

We test the quality of the main competitiveness measure in Figure 2. The figure shows how the measure predicts outcomes in legislative races: specifically, it plots the closeness and spending in a race against the competitiveness of the race's district. It reveals that the measure has considerable predictive power, with a best-fit line generating an R-squared of 0.35 for closeness and 0.32 for spending. Quality tests for the other competitiveness measures are provided in Online Appendix A2. These measures are highly correlated with our main measure, but have worse predictive power.

[12] For an "80–20" district, competitiveness is 0.4; for a "55–45" district, it is 0.9.

[13] The only exception is that we do not have primary versions for the pre-2011 districts. This is because the baseline elections for these districts occurred before the start of our data. Full details on how we calculate competitiveness can be found in Online Appendix A2.

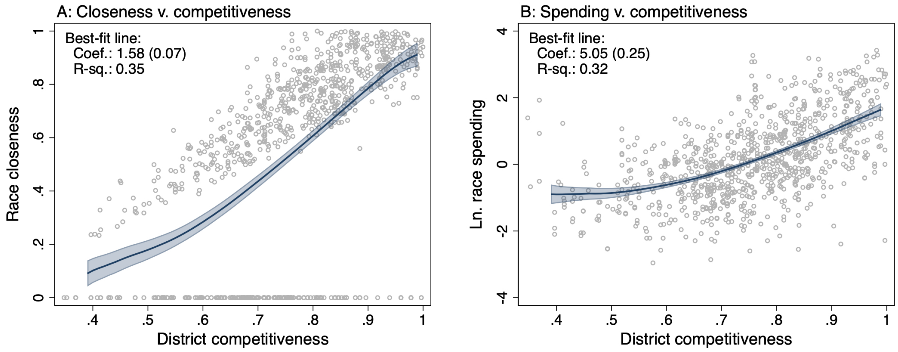

Figure 2: Predicting race outcomes using district competitiveness measures.

The figure plots outcomes in legislative races against the competitiveness of the races' districts. Competitiveness is the main measure and is calculated during the baseline election associated with a district's own redistricting episode (i.e., it is the "primary" version). "Race closeness" is 1 minus the absolute two-party vote-share margin. It equals 0 when a race is uncontested. "Ln. race spending" is the natural log of per-person spending in the race, measured in 2010 dollars. The sample is all legislative races that occurred in North Carolina during the 2012 to 2020 elections.

The Effects of Assigned Competitiveness

We now turn to the causal analysis, which we break into three parts. The first and second parts relate to the effects of increasing competitiveness in one chamber while holding it constant in other chambers. In the first part (this section), we use an intent-to-treat framework where we study the effects of the competitiveness to which a registrant is assigned during redistricting. In the second part (the next section), we build an IV model, and we use it to quantify the effects of the competitiveness that a registrant actually experiences. In the third part, we modify the IV model and present results related to the effects of increasing competitiveness in multiple chambers at once. These last results let us see how the impacts of experienced competitiveness aggregate across chambers.

Empirical Strategy

The empirical strategy for the first part of the causal analysis is the foundation for the later parts. The basic idea is to compare outcomes over time for people with the same covariates and who had the same districts before redistricting but who differ in districts for a single chamber after redistricting.[14]

[14]The strategy is similar to those in Henderson *et al.* (2016) and Moskowitz and Schneer (2019), but with some important differences. We provide a detailed comparison in the Online Appendix.

For simplicity, we explain the approach in relation to a single redistricting episode — i.e., a single instance in which districts for a legislative chamber are redrawn. We call this episode the "focal episode", and we call the chamber associated with this episode the "chamber of interest". We start by making two sample restrictions. First, we limit attention to people who were registered in the focal episode's baseline election. Second, we drop people who die before the 2020 election. Restricting to baseline registrants allows us to avoid any sample selection that may occur if competitiveness induces additional people to register and thus appear in our data. Nonetheless, it also means that we ignore one of the channels by which competitiveness could operate (the registration margin). Dropping individuals who die lets us avoid deflating the results by including observations for people who are unable to turn out.

We denote elections relative to the focal episode by τ, with the first post-redistricting election being $\tau = 0$ and the first pre-redistricting election being $\tau = -1$. A key variable for registrant i is her "assigned district". This district, labeled a_i, is the post-redistricting district in the chamber of interest that contains i's baseline address — it is the district to which i gets "assigned" by the episode. In post-redistricting elections, the assigned district is the same as i's actual district if i remains at her baseline address and if the district is still in use. However, it is not if i moves or if there has been a subsequent episode. We label the competitiveness of i's assigned district as c_{a_i}. Relatedly, we label the competitiveness of the district in which i resides in election τ as $\mathrm{c}_{i\tau}$. c_{a_i} is a fixed value for i, while $\mathrm{c}_{i\tau}$ may change over time for i if i moves among districts or if districts are replaced in later episodes. In constructing c_{a_i} and $\mathrm{c}_{i\tau}$, we measure competitiveness during the focal episode's baseline election. This way, the variables reflect only the information available at the time of redistricting.[15] Finally, we let $\mathrm{to}_{i\tau}$ be an indicator equal to 1 if i turns out in τ.[16]

Unfortunately, we cannot recover the causal effect of assigned competitiveness by simply regressing turnout on c_{a_i}. This is because there may be third factors that are associated with both variables. As an example, if districts in one redistricting episode are drawn in part based on districts in previous episodes, then assigned competitiveness may be correlated with a registrant's past district experiences. Similarly, if there is geographic clustering by political preferences, then assigned competitiveness may be correlated with the district characteristics that registrants experience in other legislative chambers.

[15] Borrowing the earlier language, c_{a_i} uses the "primary" version of competitiveness. $\mathrm{c}_{i\tau}$ uses the primary version if i's district in τ was created by the focal episode; it uses a "secondary" version (measured during the focal episode's baseline) if the district was created by a different episode.

[16] $\mathrm{to}_{i\tau}$ is 0 if either (i) i is registered in τ but does not turn out or (ii) i is not registered in τ. Also, "turn out" means that i submits a ballot in the election. It is not conditioned on i voting in any particular race on the ballot.

Finally, registrants assigned to more or less competitive districts may simply have different attributes, leading them to have different turnout propensities for reasons unrelated to assigned competitiveness.

We deal with these issues by using a matching strategy that exploits our rich, individual-level data. The strategy proceeds in two steps. First, for each episode, we divide North Carolina into regions. Regions are defined as areas that have the same pre-redistricting districts, as well as the same post-redistricting districts for chambers other than the chamber of interest. Second, we exact-match registrants according to their covariates and the region they lived in during the baseline election. Matching on baseline region accounts for differences in pre-redistricting experiences and in experiences in other chambers. This is because registrants with the same baseline region are likely to have shared the same districts before redistricting and to share the same districts in other chambers after redistricting. Matching on covariates accounts for differences in the types of people who live in parts of regions that get assigned to more or less competitive districts.

The matching procedure partitions registrants into distinct groups, which we call "match-groups". Specifically, the match-group of registrant i, g_i, is the set of registrants who lived in i's region during the baseline and who share the same covariates as i. To explore robustness, we construct match-groups using different covariates in different specifications. In our main specification, we include gender, three race/ethnicity groups, five age groups, three groups for the share college graduates in the registrant's baseline block-group, three groups for the registrant's party registration in the baseline election, the registrant's history of turnout in the three elections prior to redistricting, and the election in which the registrant first registered in North Carolina. In other specifications, we include the value of the registrant's baseline property parcel, the population density of the registrant's baseline census block, median household income in the registrant's baseline block-group, and additional district variables.[17]

After matching, we run regressions of the form:

$$\mathrm{Y}_{i\tau} = \theta_\tau \cdot \mathrm{c}_{a_i} + \theta_{g_i\tau} + \theta_{i\tau}. \tag{1}$$

Here, $\mathrm{Y}_{i\tau}$ is an outcome, such as turnout, $\theta_{g_i\tau}$ is a match-group fixed effect that varies by relative election, $\theta_{i\tau}$ is an error term, and θ_τ is the coefficient of interest. θ_τ represents the τ-specific, within-match-group association between the outcome and assigned competitiveness.[18] In order for θ_τ to have a causal

[17] The matching strategy is discussed in depth in Online Appendix A3.

[18] A point to emphasize is that Eq. (1) allows a different intercept for each match-group in each election. This is less parametric than using separate fixed effects for match-groups and elections. Also, Eq. (1) does not include registrant fixed effects. This is because these raise the computational burden and are mostly redundant after we control for match-groups. In robustness checks, we show that adding them does not change the results.

interpretation, it must be that, within match-groups, there are no third factors that are correlated with both assigned competitiveness and the election-τ outcome. This claim is plausible given the construction of match-groups. In addition, we provide evidence in support of it by showing that c_{a_i} does not predict various variables of concern.

Additional Details

Four additional details are worth mentioning regarding the empirical strategy.

First, we deal with the fact that our setting contains multiple redistricting episodes by stacking the episodes on top of each other. Specifically, we construct a dataset where a registrant's observations appear multiple times, once for each episode in which the registrant was registered in the episode's baseline election. In doing, we re-define the index i to be a registrant-episode combination (although, out of convenience, we sometimes refer to i as a registrant). Importantly, match-groups, g_i and assigned districts, a_i are episode-specific. Thus, they are constant for a registrant within an episode but differ across episodes.[19]

Second, in our strategy, not every match-group contributes to the estimation. Notably, when all registrants in a match-group are assigned to the same post-redistricting district, the group has no variation in c_{a_i} and does not influence the estimated value of θ_τ. (This is because the only controls in Eq. (1) are match-group fixed effects.) In order to be transparent about the effective sample size, we drop the irrelevant groups before obtaining results.[20]

Third, our analysis faces a limitation due to the fact that we lack data from states other than North Carolina. In particular, our turnout variable captures only within-state turnout: it is zero for all registrants who leave, even if they turn out elsewhere. This means that the results combine a potentially non-zero effect for registrants who have not left with a zero effect for registrants who have. This is likely a lower bound for impacts on turnout in any state.[21]

[19] An advantage of stacking is that it has been shown to resolve issues with regressions that use two-way fixed effects in settings with staggered treatments (Baker *et al.*, 2022; Cengiz *et al.*, 2019). We mostly avoid these issues, since our main model, Eq. (1), has only one set of fixed effects. Nonetheless, in robustness checks, we sometimes add additional sets. By stacking, we ensure that these specifications are valid.

[20] Summary statistics for the estimation sample are presented in Online Appendix A3. The sample includes almost 9 million combinations of registrants and episodes and over 500,000 match-groups. On average, there are 17.3 registrants per match-group, with a standard deviation of 38. The sample draws from across the state and is representative of all North Carolina registrants in terms of covariate means and standard deviations.

[21] Fortunately, we find evidence against selective attrition. In the Online Appendix, we show that being placed into a more competitive district does not affect the probability of moving within North Carolina; thus, it probably also does not influence the probability of leaving the state.

Fourth, we cluster standard errors by a registrant's baseline district in the chamber of interest. These districts are an aggregation of match-groups, which are the units within which we exploit variation in assigned competitiveness. We choose this more aggregate level because there may be correlation in errors among registrants who were previously in the same district. Nonetheless, we show that significance is similar under alternative ways of clustering.

Results

Results for the effects of assigned competitiveness are summarized in Figure 3. The figure provides event-study plots for estimates of the θ_τ coefficients

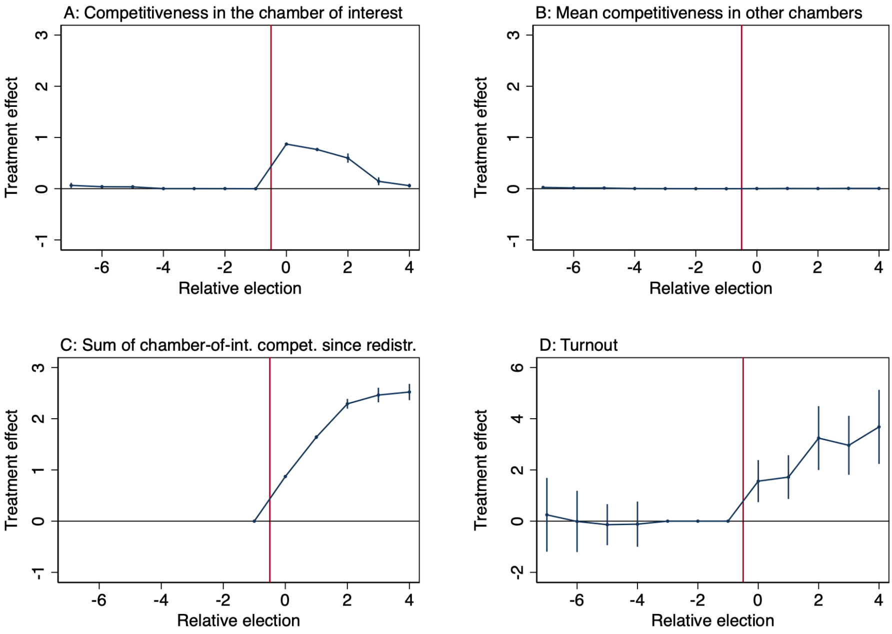

Figure 3: The effects of assigned competitiveness, c_{a_i}, on experienced competitiveness and turnout.

The figure plots coefficient estimates and 95% confidence intervals for θ_τ, the coefficients on assigned competitiveness, c_{a_i}, in Eq. (1). Results are obtained by running τ-specific regressions of the listed variables on c_{a_i} and match-group fixed effects. "Competitiveness in the chamber of interest" is $c_{i\tau}$. "Mean competitiveness in other chambers" is the average competitiveness that i experiences in chambers other than the chamber of interest in τ. "Sum of chamber-of-int. compet. since redistr." is $C_{i\tau}$. "Turnout" is $to_{i\tau}$ and is denominated in percentage points. Results are not provided for $\tau \leq -2$ in Panel C because $C_{i\tau}$ is 0 in all pre-redistricting elections. The regressions use all redistricting episodes. Competitiveness variables are constructed using the main competitiveness measure. Standard errors are clustered by baseline district in the chamber of interest.

in Eq. (1). The four panels are for four different outcome variables: the competitiveness that i experiences in election τ in the chamber of interest, $\mathrm{c}_{i\tau}$ (Panel A); the mean competitiveness that i experiences in τ in chambers other than the chamber of interest (Panel B); the sum of the chamber-of-interest competitiveness that i has experienced since redistricting, $\mathrm{C}_{i\tau} = \sum_{h=0}^{\tau} \mathrm{c}_{ih}$ (Panel C); and whether i turns out in τ, $\mathrm{to}_{i\tau}$ (Panel D). The values in the figure represent the predicted difference in the variables — within match-groups — between registrants assigned to a highly competitive 50–50 district versus a highly uncompetitive 100–0 district.

Figure 3 provides four takeaways. First, being assigned to a more competitive district causes a registrant to experience increased competitiveness in the chamber of interest in elections after redistricting (Panel A). The effects are close to one-for-one at first but fade over time as registrants move out of their assigned districts and as the districts are modified in subsequent episodes.

Second, our empirical strategy appears to deal with the major sources of concern. Namely, within match-groups, assigned competitiveness is not associated with: (i) the competitiveness that registrants experience before redistricting (Panels A and B), (ii) the competitiveness that registrants experience after redistricting in chambers other than the chamber of interest (Panel B), or (iii) pre-redistricting turnout (Panel D). Result (i) implies that the strategy is not distorted by lingering effects of pre-redistricting experiences. Result (ii) implies that it is also not influenced by spillover effects from other chambers. Finally, results (i) and (iii) together suggest that the strategy is not confounded by omitted registrant characteristics.

Third, assigned competitiveness affects turnout. People who are placed into a more competitive district turn out more in post-redistricting elections (Panel D). In combination with the prior results, this increase can credibly be understood as causal.

Fourth, the effects on turnout have an interesting shape. They increase steadily for elections $\tau = 0, \ldots, 2$ and then flatten out but remain elevated in $\tau = 3$ and 4. The first three elections are when the assigned districts are mostly still in use (depending on the episode) and when assigned competitiveness has predictive power for the competitiveness that a registrant currently experiences, $\mathrm{c}_{i\tau}$. By contrast, in the last two elections, the original districts have all been replaced and assigned competitiveness is unrelated to current competitiveness. In sum, turnout effects grow as people spend more time in competitive districts and persist after new redistricting episodes eliminate differences in current competitiveness. Another way to say this is that turnout appears to depend on the *cumulative* competitiveness that a person has experienced. Effects track those on $\mathrm{C}_{i\tau}$ (the sum of experienced competitiveness since redistricting; Panel C), not those on $\mathrm{c}_{i\tau}$ (currently experienced competitiveness; Panel A).

The shape of the turnout effects can be seen with more granularity in Online Appendix Figure A1. This figure presents event studies that are fit separately by “treatment group” — defined as episodes whose districts last for the same number of elections. Depending on the group, we can observe outcomes for up to three elections while districts are in use and up to three elections after districts are replaced. The figure shows again that turnout effects mirror gaps in cumulative competitiveness: effects grow while districts are in use and endure after districts are replaced. This is even though the relationship between assigned and current competitiveness becomes zero.

The findings in this section are not driven by the specification or by sample changes across elections. Online Appendix Figure A2 shows that turnout effects are similar under a variety of specifications — including matching on additional district variables and adding registrant-episode fixed effects — and when using a balanced panel. Another worry relates to experiences in legislative races. This is that effects on race experiences may last longer than those on current competitiveness — if, for instance, it takes time for the strength of local parties to readjust following a subsequent redistricting episode. Further, persistence in effects on race experiences could explain the persistence in effects on turnout. Online Appendix Figure A3 negates this concern: effects on race closeness and spending parallel those on current competitiveness, falling to zero in $\tau = 3$ and 4.

The Effects of Experienced Competitiveness

As seen previously, the competitiveness to which a registrant is assigned in redistricting has a positive but imperfect relationship with the competitiveness that the registrant experiences in elections after redistricting. This is because the registrant may move out of his assigned district and because the district will at some point be replaced in a subsequent episode.

We now explore the effects of experienced competitiveness. One can gain a rough sense of these by dividing the values in Panel D of Figure 3 by corresponding values in the other panels. We quantify the effects precisely by using an IV model.

IV Model

The IV model instruments for experienced competitiveness using assigned competitiveness. Given the prior results, we relate turnout with experienced competitiveness in its cumulative form, $\mathrm{C}_{i\tau}$, not its current form, $\mathrm{c}_{i\tau}$. We instrument for $\mathrm{C}_{i\tau}$ using the sum of assigned competitiveness since redistricting, which we label $\mathrm{C}_{a_i\tau}$. This variable has more predictive power for $\mathrm{C}_{i\tau}$ than does the level of assigned competitiveness, c_{a_i}. This is because it can incorporate the fact that, in a given relative election, there may be differences across episodes in whether there has been a subsequent episode. In particular, in constructing $\mathrm{C}_{a_i\tau}$, we sum over only the elections in which the assigned district remains in use.

For i's focal episode, let τ_i^l be the last election before the next episode. Then $\mathrm{C}_{a_i\tau}$ is:

$$\mathrm{C}_{a_i\tau} = \begin{cases} \sum_{h=0}^{\tau} \mathrm{c}_{a_i} = (\tau+1)\cdot \mathrm{c}_{a_i} & \text{if } \tau \leq \tau_i^l \\ \sum_{h=0}^{\tau_i^l} \mathrm{c}_{a_i} = (\tau_i^l+1)\cdot \mathrm{c}_{a_i} & \text{if } \tau > \tau_i^l. \end{cases}$$

In turn, the IV model is:

$$\begin{aligned} \mathrm{to}_{i\tau} &= \alpha_\tau \cdot \mathrm{C}_{i\tau} + \alpha_{g_i\tau} + \alpha_{i\tau} \\ \mathrm{C}_{i\tau} &= \beta_\tau \cdot \mathrm{C}_{a_i\tau} + \beta_{g_i\tau} + \beta_{i\tau}. \end{aligned} \tag{2}$$

The first line is the "structural equation", and the second is the "first stage". $\alpha_{g_i\tau}$ and $\beta_{g_i\tau}$ are τ-specific fixed effects for match-groups. β_τ is the τ-specific, within-match-group association between $\mathrm{C}_{a_i\tau}$ and $\mathrm{C}_{i\tau}$. The construction of $\mathrm{C}_{a_i\tau}$ means that this association is always close to 1. α_τ is the coefficient of interest: it is the effect on turnout of a one-unit increase in the cumulative competitiveness that i has experienced in the chamber of interest in elections between redistricting and τ.

In order for Model (2) to identify α_τ, two conditions must be met. First, β_τ must not be 0, which we demonstrate empirically in Online Appendix A4. Second, after controlling for the fixed effects, $\mathrm{C}_{a_i\tau}$ cannot be related to $\mathrm{to}_{i\tau}$ through any channel other than its correlation with $\mathrm{C}_{i\tau}$; we present evidence in support of this in Online Appendix A5.[22] Since every coefficient in Model (2) varies by τ, we estimate the model by running separate IV regressions for each relative election.[23]

Main Results

Coefficient estimates and standard errors for α_τ are displayed in the first five columns of Table 1. The estimates are all positive and similar in magnitude.

[22] Technically, the second condition applies only under an assumption of constant treatment effects. With heterogeneous treatment effects, we instead need independence, exclusion, and monotonicity. In addition, our results would recover a local average treatment effect (LATE). α_τ would not be the effect of $\mathrm{C}_{i\tau}$ for all registrants; rather, it would be an average effect for compliers, defined as those whose experienced competitiveness is influenced by their assigned competitiveness in the intended direction. We discuss this fully in Online Appendix A7.

[23] An alternative to using an IV strategy is to simply drop registrants who move out of their assigned districts. This is the approach in Moskowitz and Schneer (2019). Dropping movers causes assigned and experienced competitiveness to be the same while a redistricting episode's districts are in use. However, it ignores that these may differ once the districts are modified in a later episode. In Online Appendix A16, we show that the two strategies generate similar results in elections before new episodes.

Table 1: The turnout effects of district competitiveness.

	Election relative to redistricting, τ						Chamber	
	Zero	One	Two	Three	Four	All	U.S. House	NC legisl.
Sum of competitiveness in a registrant's districts, $C_{i\tau}$	1.79***	1.07***	1.38***	1.14***	1.42***	1.30***	2.03***	1.22***
	(0.480)	(0.252)	(0.274)	(0.231)	(0.282)	(0.234)	(0.738)	(0.243)
Turnout percentage	64.7	49.9	61.8	49.8	62.0	58.1	58.3	58.0
Clusters	338	255	163	151	151	338	36	302
Registrants	5,203,371	4,486,052	3,928,260	3,839,532	3,839,532	5,203,371	1,586,751	4,653,157
Registrant-episodes	8,769,574	6,883,563	5,442,632	5,135,610	5,135,610	8,769,574	1,700,565	7,069,009
Registrant-episode-elections	—	—	—	—	—	31,366,989	6,505,033	24,861,956

The table presents results from Models (2) and (3). The first five columns are for Model (2). They show coefficient estimates and standard errors for α_τ from τ-specific IV regressions of $\mathrm{to}_{i\tau}$ on $C_{i\tau}$ and match-group fixed effects. The last three columns are for Model (3). They show coefficient estimates and standard errors for α from 2SLS regressions of $\mathrm{to}_{i\tau}$ on $C_{i\tau}$ and match-group-by-τ fixed effects. The 2SLS regressions use all post-redistricting elections. The "Chamber" columns use just the redistricting episodes for the listed chambers, while the other columns use all episodes. Competitiveness variables are constructed using the main measure. Coefficient estimates and standard errors are denominated in percentage points. "Turnout percentage" is the percent of observations that turned out to vote. "Registrants" is the number of distinct registrants. "Registrant-episodes" is the number of registrant-episode combinations. For the first five columns, this is the number of observations. For the last three columns, the number of observations equals "Registrant-episode-elections". Standard errors are clustered by baseline district in the chamber of interest. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Depending on the relative election, a one unit increase in $\mathrm{C}_{i\tau}$ is estimated to increase turnout by between 1.1 and 1.8 percentage points.

The stability of the estimates across relative elections reinforces the earlier evidence that turnout depends on cumulative competitiveness. If turnout depended on some other form of competitiveness, then we would expect the α_τ coefficients to have a systematic trend. For instance, suppose turnout is a moving-average process where lagged competitiveness has a smaller impact than current competitiveness. Then the α_τ coefficients would decline with τ. This is because lagged competitiveness makes up a growing share of $\mathrm{C}_{i\tau}$ as τ increases. Moreover, in our setting, most of the instrument-induced variation in $\mathrm{C}_{i\tau}$ is due to variation in the competitiveness that registrants experience in the first few elections after redistricting. (As seen in Figure 3 and Online Appendix Figure A2, the relationship between assigned and current competitiveness starts out positive but becomes zero.) Thus, the estimates for the later α_τ coefficients are identified almost entirely off differences in lagged competitiveness. If this type of competitiveness mattered less, then those estimates would be small. Instead, they are comparable to the estimates for the earlier coefficients. Of course, there may be a certain number of lags after which the impact of lagged competitiveness does shrink and the α_τ coefficients do decline. But this number is larger than what we are able to study (see Online Appendix A8 for further discussion).

Given the observed stability of α_τ, we run a version of the IV model where we pool all post-redistricting elections and fit a single coefficient on $\mathrm{C}_{i\tau}$. This model uses the same first stage as (2) but has a different structural equation:

$$\mathrm{to}_{i\tau} = \alpha \cdot \mathrm{C}_{i\tau} + \alpha_{g_i\tau} + \alpha_{i\tau}. \tag{3}$$

The coefficient α is the average effect of $\mathrm{C}_{i\tau}$ across τ. In practice, it is similar to the election-specific effects. The "All" column in Table 1 shows that the estimate for α is 1.30, with a 95% confidence interval of 0.84–1.76. All but one of the estimates for α_τ fall within this interval.

The last columns of Table 1 present results by legislative chamber. For the U.S. House, a one unit increase in $\mathrm{C}_{i\tau}$ is estimated to cause a 2.03 percentage point increase in turnout. For the state chambers, the effect is 1.22 percentage points. The larger effect for the U.S. House is unsurprising, as races for this chamber generally have higher profiles.

The magnitudes in Table 1 can be understood in different ways. The simplest interpretation relates to a comparison between a highly competitive 50–50 district and a highly uncompetitive 100–0 district. This is that the effects capture the impact of one election of exposure to the former district versus the latter. Other interpretations relate to smaller gaps in competitiveness but for multiple elections. For instance, the effects can be understood as the impact of spending two elections in a moderately competitive 55–45 district versus an uncompetitive 80–20 district. Impacts for one election with this gap would be half as large: 1.02 percentage points for the U.S. House and 0.61 percentage points for the state chambers.[24]

Robustness

The results in Table 1 are robust to various concerns.

First, they are robust to different specifications and to using a balanced panel. Table 2 presents estimates of α for four alternative versions of Model (3). The first column accounts for possible bias due to registrants' experiences in pre-redistricting elections. It does this by adding linear controls for characteristics of registrants' pre-redistricting districts. The second column deals with post-redistricting experiences in other chambers. It redefines the treatment variable, $C_{i\tau}$, to be the sum of post-redistricting competitiveness over all chambers, not just the chamber of interest. The third column adds registrant-episode fixed effects. The last column creates a balanced panel by using just the decennial redistricting episodes. The coefficient estimates from these models range from 1.29 to 1.33, nearly identical to Table 1 value of 1.30.

Second, Online Appendix Tables A1 and A2 show that results are not affected by our choice of instruments or by which competitiveness measure we

Table 2: Alternative estimates of α.

	Controlling for pre-redistricting experiences	Summing over competitiveness in all chambers	Adding registr.-episode fixed effects	Using just the decennial episodes
Sum of competitiveness in a registrant's districts, $C_{i\tau}$	1.30***	1.29***	1.31***	1.33***
	(0.233)	(0.231)	(0.242)	(0.244)
Turnout percentage	58.1	58.1	52.6	56.4
Clusters	338	338	338	151
Registrants	5,203,371	5,203,371	5,203,371	3,839,532
Registrant-episode-elections	31,366,989	31,366,989	70,156,592	25,678,050

The table provides coefficient estimates and standard errors for α from modified versions of Model (3). "Controlling for pre-redistricting experiences" adds linear controls for sums of the district characteristics that registrants experienced in pre-redistricting elections. These include: district competitiveness, district share minority, district share Democratic, race closeness, and the natural log of race spending. "Summing over competitiveness in all chambers" redefines $C_{i\tau}$ to be the sum of competitiveness in a registrant's districts in all chambers, not just the chamber of interest. "Adding registr.-episode fixed effects" pools data from all relative elections, including pre-redistricting elections, and adds fixed effects for registrant-episode combinations. "Using just the decennial episodes" restricts the sample to that for the decennial redistricting episodes, which each have data for five post-redistricting elections. Other details are the same as in Table 1.

[24] The 55–45 v. 80–20 comparison is apt, as in practice there are no 100–0 districts and few 50–50 ones.

use.[25] Third, Online Appendix Tables A3 and A4 reveal that results vary by only a limited amount when we match in different ways. This is the case even when we expand the covariate set on which we match to include the value of the registrant's baseline property parcel, the population density in the registrant's baseline census block, the median household income in the registrant's baseline block-group, or the districts that the registrant lived in during pre-baseline elections. It is also the case when we drop regions and instead match on a set of the registrant's districts in both pre- and post-redistricting elections.

Fourth, results are not driven by district attributes other than competitiveness. Due to the existence of majority–minority districts, in North Carolina a district's competitiveness is negatively correlated with both its share minority and share Democratic. In Online Appendix A9, we test for bias due to these other treatments by fitting Model (3) on a trimmed sample on which the correlations are zero. The estimate for α is 1.48, quite close to the main value of 1.30.

Fifth, results are not an artifact of stacking. In Online Appendix A10, we fit models that do not rely on stacking and continue to recover similar estimates.

Another way that Model (3) may be misspecified is if turnout is nonlinear in the degree of a district's competitiveness. For instance, if a given increase in competitiveness matters more at a high versus a low level, then turnout would not depend on $\mathrm{C}_{i\mathcal{T}}$. Instead, it would depend on the sum of a nonlinear function of competitiveness. We examine this possibility in Online Appendix A11 and find considerable evidence in favor of linearity.

Finally, the degree of statistical significance is robust. In Online Appendix A12, we show that standard errors are similar (but slightly smaller) when we use different clustering strategies. We also show that significance survives a permutation test based on randomization inference. In the test, we simulate the distribution of coefficient estimates that would pertain if one were to randomly reassign competitiveness values to districts (rather than use districts' actual competitiveness). In none of 200 simulation runs do we get an estimate as large as 1.30.

Additional Results

To gain deeper insight into the relationship between turnout and district competitiveness, we present a few additional results.

We start by exploring the nature of the increase in turnout. Is it due to additional people being induced to become consistent voters, who vote in all subsequent elections? Or is one set of people spurred to turn out in one election and a substantially different set impacted in the next? In Table 3, we

[25] That said, effects are slightly smaller in standard deviation units when relying on the Cook measure.

Table 3: The effects of district competitiveness on consistent turnout.

	Turnout type		
	Any	Consistent	Inconsistent
Panel A: All registrants			
Sum of competitiveness in a registrant's districts, $C_{i\tau}$	1.30^{***}	1.29^{***}	0.015
	(0.234)	(0.232)	(0.078)
Turnout percentage	58.1	48.6	9.5
Clusters	338	338	338
Registrants	5,203,371	5,203,371	5,203,371
Registrant-episode-elections	31,366,989	31,366,989	31,366,989
Panel B: Voted in the baseline			
Sum of competitiveness in a registrant's districts, $C_{i\tau}$	1.00^{***}	1.24^{***}	-0.238^{**}
	(0.215)	(0.272)	(0.096)
Turnout percentage	85.4	76.6	8.7
Clusters	338	338	338
Registrants	2,712,457	2,712,457	2,712,457
Registrant-episode-elections	14,444,506	14,444,506	14,444,506
Panel C: Didn't vote in the baseline			
Sum of competitiveness in a registrant's districts, $C_{i\tau}$	1.54^{***}	1.33^{***}	0.209^{**}
	(0.290)	(0.251)	(0.095)
Turnout percentage	34.8	24.6	10.1
Clusters	338	338	338
Registrants	3,006,258	3,006,258	3,006,258
Registrant-episode-elections	16,922,483	16,922,483	16,922,483

The table presents results from versions of Model (3) that use alternative turnout variables. "Any" is $\text{to}_{i\tau}$, as in the "All" column of Table 1. "Consistent" is an indicator equal to one if i turns out in τ and in all later elections. "Inconsistent" equals 1 if i turns out in τ but not in all later elections. Consistent and inconsistent turnout sum to $\text{to}_{i\tau}$. Panels B and C are limited to registrants who did (did not) turn out in the baseline election. Other details are the same as in Table 1.

fit Model (3) for two alternative turnout variables. "Consistent" turnout is an indicator equal to one for i in τ if i turns out in τ and in all later elections that can be observed in the data. By contrast, "Inconsistent" turnout equals one if i turns out in τ but does not turn out in at least one later election. These variables sum to $\text{to}_{i\tau}$. Thus, by running the IV model on them, we can see how much of the increase in turnout is for turnout of each type. We present results for all registrants and by whether a registrant voted in the baseline election. Splitting by baseline turnout lets us see if effects differ for people who already vote versus those who do not.

Table 3 shows that district competitiveness operates largely by creating consistent voters. For people who voted in the baseline, a one unit increase in competitiveness causes a 1.00 percentage point increase in the probability of turning out in the given election, a 1.24 percentage point increase in the probability of turning out in the given election and all later elections, and a 0.24 percentage point decrease in the probability of turning out in the given election but not in a later election. Thus, for these registrants, competitiveness reduces the share who do not turn out, slightly reduces the share who inconsistently turn out, and sizably increases the share who always turn out. For people who did not vote in the baseline, impacts are 1.54 percentage points on any turnout, 1.33 percentage points on consistent turnout, and 0.21 percentage points on inconsistent turnout. Thus, even for people who do not already vote, competitiveness operates mainly by increasing the share of consistent voters — although, for these people, it also generates a small rise in the share of inconsistent voters.[26] The fact that district competitiveness mainly creates consistent voters helps to explain the persistence in turnout effects. Namely, additional registrants are induced to become people who "always vote"; since these people do not stop turning out, the turnout effects do not decay.

We next examine the role of competitive races. In Table 4, we run versions of Model (3) where we build $\mathrm{C}_{i\tau}$ and $\mathrm{C}_{a_i\tau}$ using measures of race competitiveness: the closeness of the vote margin and the natural log of spending. The α coefficients from these models reveal whether people turn out more when they live in districts that *end up* having competitive races (rather than in those we predict as *likely* to). For context, we also run a version of Model (3) that uses district competitiveness, as in Table 1. Further, we run a "horse race" where we jointly include all the variables in the same specification. In order to allow comparing magnitudes, we standardize the measures of race or district competitiveness before constructing $\mathrm{C}_{i\tau}$ and $\mathrm{C}_{a_i\tau}$.

The results reveal that the turnout effects of district competitiveness can be explained in large part by exposure to competitive races. Also, effects are especially dependent on exposure to races with close vote margins and less so to races with high spending. When standardized race closeness is included alone, a one unit increase in its sum causes an increase in turnout of 0.24 percentage points. This is considerably larger than the corresponding values for standardized district competitiveness or race spending of 0.18 and 0.17 percentage points. Similarly, in the horse race, the bulk of the weight is placed

[26]For the full sample, the effects on inconsistent turnout cancel between those who did or did not vote in the baseline, and the impact on any turnout is entirely explained by an increase in consistent turnout. Also, the results in Table 3 are not simply because there are some episode–election combinations for which there are few or no later elections that have turnout data. In Online Appendix Table A5, we repeat the analysis but restrict the sample to episode-election combinations with at least one, two, or three future elections with data. We continue to find that effects are almost entirely due to increases in consistent turnout.

Table 4: The turnout effects of race competitiveness.

	(1)	(2)	(3)	(4)
Sum of standardized district competitiveness in a registrant's districts	0.181*** (0.032)			0.086** (0.042)
Sum of standardized race closeness in a registrant's districts		0.238*** (0.044)		0.156*** (0.059)
Sum of standardized ln. race spending in a registrant's districts			0.173*** (0.043)	0.048 (0.057)
Turnout percentage	58.1	58.1	58.1	58.1
Clusters	338	338	338	338
Registrants	5,203,371	5,203,371	5,203,371	5,203,371
Registrant-episode-elections	31,366,989	31,366,989	31,366,989	31,366,989

The table provides results for versions of Model (3) where $C_{i\tau}$ and $C_{a_i\tau}$ are built using either district competitiveness, race closeness, or the natural log of race spending. The model in Column 4 includes each of the treatment variables in Columns 1–3. It instruments for them using the union of the instruments in those columns. Measures of district or race competitiveness are standardized before building $C_{i\tau}$ and $C_{a_i\tau}$. Standard deviations of the measures are shown in Online Appendix Tables A19 and A20. The only difference between Column 1 and the "All" column of Table 1 is that district competitiveness is now standardized. See Table 1 for additional details.

on race closeness, with some on district competitiveness and very little on race spending.

We next explore additional heterogeneity in the effects of district competitiveness. Key features of this heterogeneity are summarized in Figure 4, while full results are provided in Online Appendix Tables A6–A8 and A48. The figure indicates that district competitiveness matters for turnout for almost all types of registrants and in both midterm and presidential elections. Differences in coefficient estimates across registrant and election types are usually not statistically significant. Nonetheless, estimates are somewhat larger for young registrants (age 35 or below; 1.70) and are especially large for young registrants who live in well-educated neighborhoods (those where the share of college graduates is greater than 0.4; 2.40). Estimates are also larger for whites than for racial minorities. By contrast, estimates are similar by gender and for midterms versus presidential years.[27]

[27]In Online Appendix A9, we show that the heterogeneity by race is not due to the correlations between district competitiveness, share minority, and share Democratic. The similarity in midterm and presidential elections is somewhat surprising, given that legislative races are closer to the top of the ticket in midterms. A countervailing factor is that the two election types differ in their voting populations, with larger turnout in presidential years (Online Appendix Table A8). Thus, in these years, there may be more people who are on the margin of turning out.

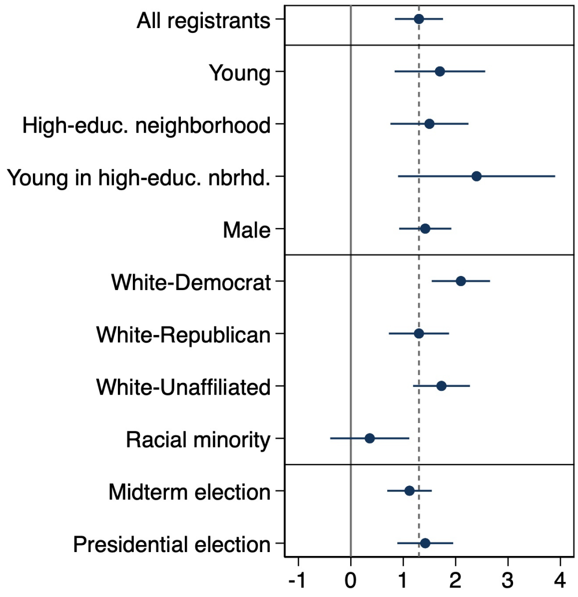

Figure 4: Heterogeneity in the turnout effects of district competitiveness.

The figure summarizes heterogeneity in the turnout effects of district competitiveness. It shows coefficient estimates and 95% confidence intervals for α in versions of Model (3) that are fit using only the specified types of registrants or elections. "All registrants" corresponds with the "All" column of Table 1. "Young" is registrants who were 35 or younger during the baseline election. "High-educ. neighborhood" means that the share college graduates in the baseline block-group is greater than 0.4. Party registration is measured during the baseline. Full results are provided in Online Appendix Tables A6–A8 and A48. Other details are as in Table 1.

Finally, in Online Appendix A13, we show that competitiveness does not impact turnout in primary elections. This may seem to conflict with the results about consistent turnout. In particular, if a person is induced to become a consistent general-election voter, one may expect the person to vote more in all elections. That said, primaries are quite different from general elections; for instance, they have only half the turnout of midterms (Online Appendix Tables A8 and A54). In addition, prior research suggests that people's voting habits differ by election type (Coppock and Green, 2016). Thus, it is possible to consistently vote in general elections but rarely vote in primaries.

Discussion

The visual evidence in Figure 3 and Online Appendix Figure A1 and the stability of α_τ in Table 1 imply a high degree of persistence in the turnout effects of district competitiveness. Notably, effects are never observed to decay, even up to three elections after differences in current competitiveness disappear. This persistence is greater than what is found in the literature on habit formation.[28]

A possible reason for the discrepancy is that the persistence in our setting may be due to more than just habit formation. As examples, there may be slow adjustment to district conditions — it may take time for people to realize that their new districts are not as competitive, for social norms to dissipate, or for local parties to weaken — and people may exhibit rule-of-thumb behavior where they are influenced by memories of close races for many years. With this in mind, we next try to better understand the mechanisms that drive our results.

We consider three broad channels. The first is a campaign-based story. This is that competitive districts increase turnout via elite mobilization. Namely, national parties and local stakeholders spend resources to activate voters and bring them to the polls, such as via advertising or door-knocking and canvassing.

The campaign-based story clashes with at least two of our results. First, it struggles to explain the high degree of persistence in effects. In Online Appendix Figure A3, we showed that the impacts of assigned competitiveness on race spending follow those on current competitiveness and become zero by a few elections after redistricting. Thus, people placed in more versus less competitive districts eventually see no boost in the intensity of campaigns, and yet continue turning out more. In the campaign-based story, the persistence in turnout gaps could only be due to habit formation (unless there is continued outreach from local parties that is not reflected in spending). However, as discussed, this would imply a larger amount of habit formation than is found in the literature. The campaign-based story also fails to rationalize Table 4. If our results were driven by campaigning, then we would expect turnout effects to depend most heavily on exposure to races with high spending. Instead, they depend more strongly on exposure to races with close vote margins. Finally, the story is challenged by prior research. Moskowitz and Schneer (2019) show that spending in competitive legislative races is allocated toward TV advertising, which has been found to minimally affect turnout. Thus, the campaign-based story seems unable to accommodate our results.

The second channel is a candidate-based story. This is that popular incumbents elicit turnout from their constituents, and that incumbents who are more popular are more willing to run in competitive districts after redistricting. We test this story by estimating Model (3) for registrants who are assigned to districts in which their pre-redistricting incumbents do not run. These

[28] While Coppock and Green (2016) find up to 20 years of persistence for young voters, they find less overall. Fujiwara *et al.* (2016) find decay in the effects of a rainfall shock in just the next election.

registrants have no attachment to the candidates in their assigned districts. Thus, the effects of competitiveness for them are purged of interactions with incumbency. The results (in Online Appendix Table A9) do not decline, casting doubt on the story.

The final channel is a learning story. This is that people in competitive districts become more interested in voting because they realize either that races in their district are competitive or that races can be competitive generally. The realization could occur due to viewing polls and vote margins, via peer communications, or as a result of media coverage and campaign spending. Once people are aware of competitiveness, their turnout may increase due to instrumental considerations, concerns about civic duty, a desire to gain expressive benefits, or compliance with social pressure. Finally, for various reasons, people may keep turning out even after their districts are replaced. These include slow adjustment to the new districts, general habit formation, or the development of rules of thumb regarding the importance of voting.

The results broadly support the learning story. First, the story implies a high degree of persistence, as we find. Second, it aligns with the fact that turnout depends on exposure to races with close vote margins. Third, it explains why effects are larger for young voters (they are more impressionable), for young people in well-educated neighborhoods (they may pay more attention), and for people who live in competitive districts for longer (they receive more signals).

The learning story is also consistent with evidence in laboratory experiments and in natural experiments in Europe that people respond to competitiveness in their voting decisions. By contrast, it conflicts with two other findings from the literature. These are the lack of effect in field experiments that provide Americans with information on competitiveness and the fact that people seem to know little about competitiveness.

We cannot fully reconcile this divergence, but we offer two possible resolutions. First, the learning opportunities from living in a competitive district may be of a higher dosage than the marginal increases in information that people gained in the experiments. Second, people appear to have some information on competitiveness. For Congressional districts, Moskowitz and Schneer (2019) show a positive association between the Cook measure of competitiveness and people's beliefs in the small amount of survey data that is available on the topic. For state-level races, Gerber *et al.* (2020) find that observed vote margins are a strong predictor of beliefs. In addition, they find that beliefs are more accurate for individuals with more education. Thus, it may be that, while people have little sense of competitiveness, they have enough for their behavior to be swayed by competitive districts. In sum, ample evidence backs the learning story — although we take pains to note that other explanations may also play a role.

External Validity

We evaluate the external validity of our results by benchmarking them with those of Moskowitz and Schneer (2019). Moskowitz and Schneer study only the decennial redistricting episode for the U.S. House. In addition, they use a slightly different empirical strategy than we do, and they examine outcomes in only the 2012 and 2014 elections. Nonetheless, they have data from all 50 states. Thus, by appropriately comparing our results with theirs, we can isolate the role of the North Carolina setting. We do this in Online Appendix A16 and closely match their effects. Thus, there seems to be nothing peculiar about North Carolina. Instead, our results may have broad relevance for other states. That said, the comparison is only for the U.S. House and only during the early part of our sample period. As a result, we cannot prove generalizability for the other chambers or for the later elections. The second limitation is important because the later elections are when North Carolinians saw unusually frequent changes in their districts.

Additivity Across Chambers

We conclude the causal analysis by providing evidence that the effects of district competitiveness are additive across legislative chambers.

To do this, we modify our previous empirical approach. In the prior two sections, we restricted attention to registrants who differ in assigned districts for only a single chamber. Now, we allow there to be within-match-group differences in assigned districts for multiple chambers. Implementing the new approach involves making two specific changes. First, we alter the definition of a redistricting episode: we now define an episode as an instance in which districts are redrawn for any chamber, not for a particular chamber. Second, we redefine the regions that we rely on in constructing match-groups: we have them depend only on pre-redistricting districts, not also on post-redistricting districts (see Online Appendix A3.3).[29]

We assess additivity in three ways. The first strategy is to test whether the combined effect of changing competitiveness for multiple chambers is equal to the sum of the effects for each constituent chamber. We implement the test by fitting a version of Model (3) that uses a special form of the treatment and instrument. The treatment is a weighted sum of competitiveness in a

[29] In the new approach, there are four redistricting episodes, which occurred in 2011, 2015, 2017, and 2019. Given that there is no longer a chamber of interest, we also must change how we cluster standard errors. We choose to cluster by the intersection of baseline districts. In the earlier analysis, these units yielded similar significance as the baseline chamber-of-interest district (Online Appendix A12).

registrant's districts for the U.S. House and state chambers. It is:

$$\mathrm{C}_{i\tau} = \hat{\alpha}^{\mathrm{USH}} \cdot \sum_{h=0}^{\tau} \mathrm{c}_{ih}^{\mathrm{USH}} + \hat{\alpha}^{\mathrm{NC}} \cdot \sum_{h=0}^{\tau} (\mathrm{c}_{ih}^{\mathrm{NCS}} + \mathrm{c}_{ih}^{\mathrm{NCH}}).$$

Here, $\hat{\alpha}^{\mathrm{USH}}$ and $\hat{\alpha}^{\mathrm{NC}}$ are the coefficient estimates from the "Chamber" columns in Table 1. The instrument, $\mathrm{C}_{a_i\tau}$, is an analogous weighted sum of assigned competitiveness. Thus, in this model, we weight competitiveness according to its estimated effect in each type of chamber. If the effects of competitiveness are additive across chambers, then the coefficient on $\mathrm{C}_{i\tau}$ should equal 1. That is, the impact of an increase for multiple chambers should be the sum of the impacts we obtained when studying changes for one chamber at a time. By contrast, if effects aggregate according to a different functional form, then the coefficient on $\mathrm{C}_{i\tau}$ may not equal 1.

The test supports additivity. As seen in Column 1 of Table 5, the estimate for the coefficient on $\mathrm{C}_{i\tau}$ is 1.01, insignificantly different from 1.

The second strategy is to re-estimate α^{USH} and α^{NC} using the new variation. In Column 2 of Table 5, we fit a version of Model (3) that has two treatment variables: the sums of competitiveness since redistricting in the U.S. House and the state chambers. The instruments are analogous sums of

Table 5: Effects calculated using differences in assigned districts for multiple chambers.

	(1)	(2)
Weighted sum of competitiveness in a registrant's districts: all chambers	1.01^{***} (0.185)	
Sum of competitiveness in a registrant's districts: U.S. House		2.57^{***} (0.774)
Sum of competitiveness in a registrant's districts: state chambers		1.12^{***} (0.233)
Turnout percentage	58.0	58.0
Clusters	540	540
Registrants	5,604,366	5,604,366
Registrant-episode-elections	27,919,274	27,919,274

The table presents results from versions of Model (3) that allow match-groups to have differences in assigned districts for multiple legislative chambers. The treatment variable in Column 1 is a weighted sum of competitiveness in a registrant's U.S. House and state legislative districts. In the sum, competitiveness for each type of chamber is weighted by the corresponding coefficient estimate in Table 1 (2.03 for the U.S. House and 1.22 for the NC Senate and NC House). The instrument is an analogously weighted sum of assigned competitiveness. The model in Column 2 includes two treatment variables: the sums of competitiveness in the U.S. House and the state chambers. The instruments are analogous sums of assigned competitiveness. Standard errors are clustered by the intersection of baseline districts.

assigned competitiveness. Under additivity, the coefficient estimates for this model should match the values in the "Chamber" columns of Table 1. That is, we should recover the same estimates regardless of whether we conduct comparisons among people who differ in assigned districts for multiple chambers or for just a single chamber.[30]

The results again endorse additivity. The coefficient estimates are close to those in Table 1, and none of the differences are statistically significant.

The third strategy is to test for interactions in the effects. This test reveals whether the impact of increasing competitiveness in one chamber depends on the degree of competitiveness in a person's districts in other chambers. In Online Appendix Table A10, we repeat the analysis in Columns 1 and 2 of Table 5 but add interaction terms. These are sums of the products of competitiveness in different chambers. We instrument for them using analogous sums of the products of assigned competitiveness. The results counter the existence of interactions. Allowing for them hardly changes the main estimates, and the terms lack statistical significance. (Only two of 14 are individually significant — and at only a 10% level; joint tests have marginal p-values.)

Thus, there is considerable evidence that effects are additive across chambers. First, changing competitiveness for multiple chambers generates an impact that is the sum of the chamber-specific effects. Second, we recover similar estimates for the chamber-specific effects whether we compare people who differ in multiple chambers or a single chamber. Third, interaction terms are mostly insignificant. The additivity finding — together with the prior evidence that turnout depends on cumulative competitiveness — suggests that the effects of district competitiveness can "add up" and may in some instances become sizable. We illustrate this hypothetically in Online Appendix A14 and investigate it empirically in the next section.

The Impacts of North Carolina's Legislative Districts

As a last exercise, we study the impacts of North Carolina's recent legislative districts. We consider a counterfactual in which all districts used during the 2012–2020 elections had been moderately competitive 55–45 districts, and we simulate changes in turnout and votes. The analysis involves combining the causal effects from before with information on registrants' district histories and party preferences. The results are specific to the North Carolina setting.

[30] Unfortunately, it is possible for the two sets of estimates to differ even if effects are additive. This is because the approach of using variation from multiple chambers may suffer from multicollinearity. Namely, if competitiveness is correlated across chambers, then the model may struggle to parse the distinct effects of each chamber type. As such, the test is only partially informative. It cannot disprove additivity; yet it can serve as evidence in favor of it.

Nonetheless, they show how effects add up, and they give insight into the types of magnitudes we might observe if other states were to make districts competitive.

Methodology

The simulation has four steps.

First, we limit attention to people who were registered at the start of the analysis period. Specifically, we consider 2010 registrants who do not die before the 2020 election. In this sample, everyone is subject to the counterfactual treatment for the same length of time.[31]

Second, we calculate impacts on registrants' turnout probabilities. To understand the calculation, let t index calendar-year elections (not elections relative to redistricting), and let $\hat{\alpha}_i^{\mathrm{USH}}$ and $\hat{\alpha}_i^{\mathrm{NC}}$ be estimates of the causal effects of competitiveness for i in the U.S. House and state chambers. Depending on the specification, $\hat{\alpha}_i^{\mathrm{USH}}$ and $\hat{\alpha}_i^{\mathrm{NC}}$ may be the same for all registrants or may vary by i's characteristics (in order to reflect the heterogeneity seen previously). Also, let $0.9 - \mathrm{c}_{ih}^{j}$ be the difference in competitiveness between that of a 55–45 district and that of i's actual chamber-j district in election h. Finally, let $\sum_{h=2012}^{t}(0.9 - \mathrm{c}_{ih}^{j})$ be the sum of these differences for elections between 2012 and t. Then, for each election $t = 2012, \ldots, 2020$, we compute:

$$\Delta_{it}^{\mathrm{to}} = \hat{\alpha}_i^{\mathrm{USH}} \cdot \sum_{h=2012}^{t} (0.9 - \mathrm{c}_{ih}^{\mathrm{USH}}) + \hat{\alpha}_i^{\mathrm{NC}} \cdot \sum_{h=2012}^{t} [(0.9 - \mathrm{c}_{ih}^{\mathrm{NCS}}) + (0.9 - \mathrm{c}_{ih}^{\mathrm{NCH}})]. \tag{4}$$

$\Delta_{it}^{\mathrm{to}}$ is the predicted change in i's turnout probability in election t due to living in competitive districts. Its functional form reflects the earlier findings that turnout depends on cumulative competitiveness and that effects are additive across chambers. Notably, Eq. (4) says that $\Delta_{it}^{\mathrm{to}}$ is a weighted sum of the change in i's competitiveness since redistricting in each chamber, with weights that reflect the chambers' causal effects of competitiveness.

Third, we quantify how competitive districts would have impacted registrants' vote probabilities. This calculation uses two inputs. First, we suppose that competitiveness has no effect on registrants' preferences over political parties; as a result, impacts on votes are due only to changes in turnout. Second, we generate individual-level party preference probabilities. These are election-specific predictions for a registrant's probability of preferring the

[31] As before, we do not adjust for the fact that people may move to other states. In this way, the exercise reveals the influence of competitive districts on the within-state voting behavior of still-alive (as of 2020) 2010 registrants.

Democratic or Republican candidate in a generic contest.[32] Let i's preference probability for party k in election t be written p_{it}^{k}. Then, the change in i's vote probability for the party in the election is:

$$\Delta_{it}^{k} = \mathrm{p}_{it}^{k} \cdot \Delta_{it}^{\mathrm{to}}.$$

It is the change in i's turnout probability scaled by the probability that i prefers the party.

In a last step, we summarize $\Delta_{it}^{\mathrm{to}}$ and Δ_{it}^{k} over the registrants in the sample. This lets us see the aggregate impacts of competitive districts.

The Competitiveness of Registrants' Districts

Before turning to the results, we appraise registrants' experiences with respect to competitiveness. For each registrant, we calculate the average competitiveness of the registrant's districts during the 2012 to 2020 elections:

$$\bar{\mathrm{c}}_i = \frac{1}{15} \sum_{j \in \{\text{USH, NCS, NCH}\}} \sum_{t=2012}^{2020} \mathrm{c}_{it}^{j}.$$

We then examine the distribution of $\bar{\mathrm{c}}_i$, separately by a registrant's race and party. The results are in Figure 5. For each group, the figure displays a histogram of $\bar{\mathrm{c}}_i$ and lists summary statistics.

The figure has five takeaways. First, among whites, the distribution of $\bar{\mathrm{c}}_i$ does not depend on a registrant's party registration. Second, racial minorities experienced somewhat less competitive districts during 2012–2020 than did whites. Third, registrants on average lived in modestly uncompetitive districts during this period. For whites, the mean of $\bar{\mathrm{c}}_i$ is 0.78, which corresponds to a 61–39 district; for minorities, it is 0.75, or a 62.5–37.5 district. Fourth, there is substantial variation in experiences across registrants. For whites, the 10th percentile of $\bar{\mathrm{c}}_i$ is 0.69 (a 65.5–34.5 district), and the 90th percentile is 0.88 (a 56–44 district); for minorities, values are 0.63 (a 68.5–31.5 district) and 0.85 (a 57.5–42.5 district).[33] Fifth, if all districts had been 55–45, most registrants

[32] The preference probabilities are similar to the values that were used in computing the competitiveness measures (Online Appendix A2). We derive them in three steps. First, in each election, we regress precinct-level vote shares on precinct-level means of voter characteristics. Second, for each registrant in the sample, we multiply the registrant's covariates by the regression coefficient estimates, obtaining contest-specific predictions. Third, we average the contest-specific predictions over the contests in the election, and we bound the averages between 0 and 1.

[33] Online Appendix Tables A11 and A12 show how registrants' experiences with respect to competitiveness differed by chamber and election. The tables reveal that the districts for the state chambers were less competitive than the districts for the U.S. House. In addition, the districts used in 2018 and 2020 (after the court-ordered revisions) were more competitive than those from earlier elections.

would have seen an increase in $\bar{c}_i$, although some would have seen a slight decrease.

Results

We calculate results for three specifications. These differ in how we model heterogeneity in the effects of district competitiveness across registrants. The main specification ("race-party") lets $\hat{\alpha}_i^{\mathrm{USH}}$ and $\hat{\alpha}_i^{\mathrm{NC}}$ vary by the same four race and party groups as in Figure 5. The second specification ("age-education") lets them vary by the four age and education groups used in Online Appendix Table A7. The last specification ("common") restricts $\hat{\alpha}_i^{\mathrm{USH}}$ and $\hat{\alpha}_i^{\mathrm{NC}}$ to be the same for all registrants. This specification is appropriate if the heterogeneity

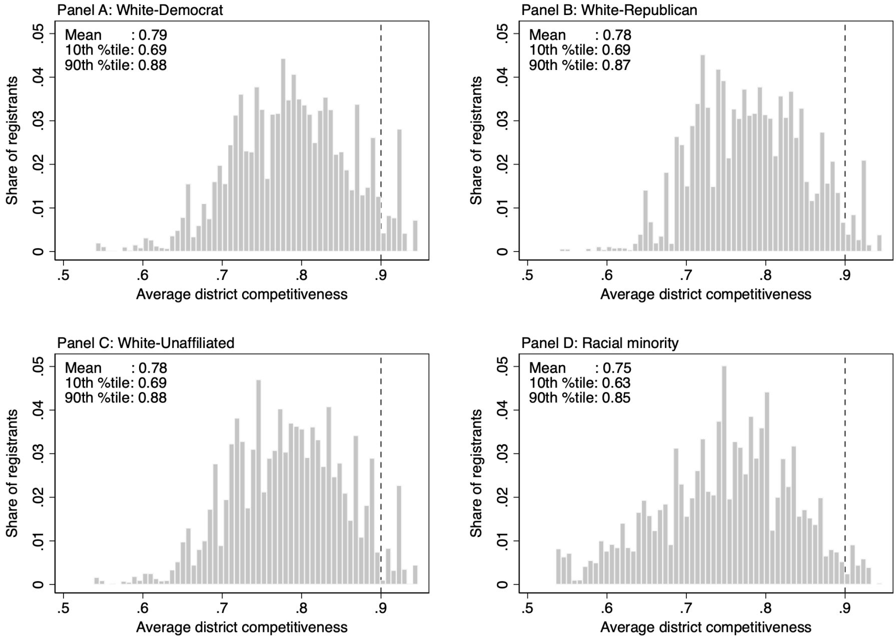

Figure 5: Histograms of $\bar{c}_i$ by a registrant's race and party.

The figure plots histograms of $\bar{c}_i$ for four groups of registrants. "Mean", "10th %tile", and "90th %tile" are the mean, the 10th percentile, and the 90th percentile in the given groups. Competitiveness is the main measure and is calculated in 2010. The sample includes 5,661,564 registrants. Of these, 1,262,551 are white-Democrats, 1,703,424 are white-Republicans, 1,138,183 are white-Unaffiliated, and 1,557,406 are minorities. Also, 2,450,388 are Democrats, 1,793,941 are Republicans, and 1,417,235 are Unaffiliated. The dashed vertical lines indicate the value of $\bar{c}_i$ under 55–45 districts.

Table 6: The change in registrants' turnout probabilities under 55–45 districts.

	All registrants			Means by race and party				Means by party		
Election	Mean	10th percentile	90th percentile	White-Dem.	White-Rep.	White-Unaffil.	Minority	Dem.	Rep.	Unaffil.
2012	0.60	0.11	1.23	0.91	0.60	0.76	0.24	0.59	0.58	0.66
2014	1.20	0.22	2.44	1.82	1.19	1.52	0.49	1.18	1.15	1.31
2016	1.70	0.29	3.41	2.58	1.68	2.16	0.67	1.66	1.62	1.86
2018	2.15	0.36	4.38	3.27	2.13	2.75	0.82	2.09	2.06	2.36
2020	2.53	0.39	5.23	3.85	2.55	3.25	0.92	2.44	2.46	2.79

The table provides summary statistics for Δ_{it}^{to}. That is, it shows how turnout would have differed in the 2012–2020 elections if registrants had lived in 55–45 districts during this period. The values in the "All registrants" columns are the election-specific mean, 10th percentile, and 90th percentile of Δ_{it}^{to} over everyone in the sample. The values in the "Means by race and party" and "Means by party" columns are election-specific means of Δ_{it}^{to} for people in the specified groups. All values are denominated in percentage points. Results are for the main specification of the simulation.

by registrant characteristics is merely noise.[34] In presenting results, we focus first on the main specification and then later provide a comparison with the alternative specifications.

Table 6 shows how competitive districts would have affected registrants' turnout probabilities. It presents summary statistics for Δ_{it}^{to} and reveals a few points. First, registrants would have been more likely to turn out during 2012–2020 if they had lived in competitive districts. For 2012, the average change in the probability of turning out is an increase of 0.6 percentage points. By 2020, this value grows to 2.5 percentage points. Second, there is substantial variation in Δ_{it}^{to} across registrants. For instance, in 2020, the 10th percentile of Δ_{it}^{to} is an increase of 0.4 percentage points, while the 90th percentile is an increase of 5.2 percentage points. Third, much of the variation in impacts can be explained by a registrant's race-party group. Across groups, values are largest for white-Democrats and smallest for racial minorities.[35] By contrast, there are more limited differences by party: in 2020, the mean of Δ_{it}^{to} is an increase of about 2.45 percentage points for Democrats and Republicans and 2.8 percentage points for Unaffiliated registrants.

Table 7 presents effects on aggregate turnout. It shows how the number of registrants who turn out would have changed if everyone had lived in competitive districts. For context, it also lists the number who actually did

[34] The values of $\hat{\alpha}_i^{\text{USH}}$ and $\hat{\alpha}_i^{\text{NC}}$ that we use are displayed in Online Appendix Table A13.

[35] The small impact for minorities is due to their small values of $\hat{\alpha}_i^{\text{USH}}$ and $\hat{\alpha}_i^{\text{NC}}$. As discussed earlier, minorities experienced the least competitive districts among the four race-party groups.

Table 7: The change in aggregate turnout under 55–45 districts.

		Change in turnout							
			By race and party				By party		
Election	Actual turnout	All	White-Dem.	White-Rep.	White-Unaffil.	Minority	Dem.	Rep.	Unaffil.
2012	3,644,162	34,090	11,471	10,148	8,670	3,802	14,436	10,333	9,321
2014	2,442,412	68,137	22,927	20,284	17,329	7,598	28,852	20,655	18,630
2016	3,445,209	96,165	32,555	28,578	24,578	10,454	40,706	29,092	26,368
2018	2,762,620	121,706	41,344	36,341	31,322	12,700	51,216	36,974	33,516
2020	3,467,293	143,393	48,651	43,445	37,033	14,264	59,711	44,168	39,514

The table summarizes the impact of 55–45 districts on aggregate turnout. "Actual turnout" is the number of registrants in the sample who turned out in the specified election. "Change in turnout" is how turnout would have differed if all districts used since 2012 were 55–45. Values in these columns are calculated by summing Δ_{it}^{to} over the specified group for the given election. They equal the product of the values in Table 6 and the number of registrants in the group. Results are for the main specification.

turn out in each election. The results differ from those in Table 6 in that they incorporate information on the size of each race-party or party group.

The results indicate that competitive districts would have generated a considerable increase in aggregate turnout, especially in later elections. For instance, in 2020, overall turnout would have been higher by 143,393 registrants. This is a 4.1% increase over the election's actual turnout among the sample. As in Table 6, the increase in aggregate turnout would have been largest for white-Democrats and smallest for minorities. By contrast, the results by party are now different. For aggregate turnout, Democrats would have seen the largest change. For instance, in 2020, Democratic turnout would have been higher by 59,711, or 42% of the total increase. This reflects that Democrats are the largest party by registration (see the notes to Figure 5).

Table 8 presents impacts on votes. It shows how the aggregate number of votes received by each party would have differed under competitive districts. In the table, the columns titled "Change in votes" display the sum of Δ_{it}^{k} over all registrants in the sample. The column titled "Net change for Dem." lists the difference between the change for Democrats and Republicans, revealing impacts on vote margins. For context, the "Predicted votes" columns provide predictions for how the sample's registrants voted in each election.[36]

[36] These are calculated by summing p_{it}^{k} over the in-sample registrants who turned out. Note that the sum of the "Predicted votes" columns in Table 8 is less than the "Actual turnout" column in Table 7. This is because people can vote for third-party candidates or abstain. The same story holds for the sum of the "Change in votes" columns in Table 8 and the "All" column in Table 7.

Table 8: The change in aggregate votes under 55–45 districts.

	Predicted votes		Change in votes		Net change
Election	Democrats	Republicans	Democrats	Republicans	for Dem.
2012	1,742,498	1,781,188	15,529	16,971	−1,442
2014	1,123,158	1,240,290	32,363	32,672	−309
2016	1,585,846	1,740,649	43,046	48,550	−5,505
2018	1,288,959	1,407,595	58,262	60,677	−2,415
2020	1,656,080	1,741,297	65,608	73,864	−8,256

The table summarizes the impact of 55–45 districts on aggregate votes. "Predicted votes" is the predicted number of votes for each major party among the sample. For party k and election t, it is the sum of p_{it}^{k} over in-sample registrants who turned out in the election. "Change in votes" is how votes would have differed if all districts used since 2012 were 55–45. For party k and election t, it is the sum of Δ_{it}^{k} over all in-sample registrants. "Net change for Dem." is the difference between the change in votes for Democrats and Republicans. Results are for the main specification.

The results suggest that competitive districts would have led to a higher number of votes; however, the impact on vote margins would have been negligible. For example, in 2020, vote totals would have increased by 65,608 for Democratic candidates and 73,864 for Republican ones. The change in the vote margin would have been a relative loss in votes for Democrats of 8,256.

The fact that competitive districts would have hardly affected vote margins is surprising. We found previously that aggregate turnout would have increased more for Democratic registrants than for other registrant groups. Thus, one might expect that Democratic candidates would similarly gain a larger increase in votes.

To understand this puzzle, we examine the party preferences of the people who would have been induced to turn out by competitive districts. For different registrant groups, we compute the percent of the group's increase in votes that accrues to each party. Specifically, for group s in election t and for party k, we calculate:

$$f_{st}^{k} = 100 \cdot \frac{\sum_{i \in s} \Delta_{it}^{k}}{\sum_{i \in s} \Delta_{it}^{\mathrm{to}}}.$$

f_{st}^{k} reveals the partisan composition of group s's increase in votes.

Values of f_{st}^{k} are exhibited in Table 9. They show that there is an asymmetry in how changes in a party's turnout translated into changes in votes for the party. Notably, for Republicans, the increase in turnout due to competitive districts generated votes that went almost entirely to Republican candidates ("By party — Rep." columns). By contrast, for Democrats, the increase in turnout yielded a more balanced allocation between the parties ("By party — Dem." columns). The reason for the latter finding is twofold. First, most

Table 9: The partisan composition of different groups' increases in votes.

	By race and party								By party					
	White-Dem.		White-Rep.		White-Una.		Minority		Dem.		Rep.		Unaffil.	
Election	Dem.	Rep.	Dem.	Rep.	Dem.	Rep.	Dem.	Rep.	Dem.	Rep.	Dem.	Rep.	Dem.	Rep.
2012	65	30	8	90	47	45	86	12	71	25	8	90	49	43
2014	60	34	12	86	56	38	87	12	66	29	12	86	58	36
2016	64	32	12	86	41	49	84	13	69	27	12	86	43	48
2018	59	37	15	83	56	43	86	13	65	31	15	83	58	41
2020	64	34	14	85	44	50	83	15	69	29	14	85	45	49

The table displays f^k_{st} for different groups of registrants. That is, for each group, it shows the percent of the group's increase in votes, due to 55–45 districts, that accrues to each major party. The table's second (third) row specifies the group s (party k). Values of f^k_{st} do not sum to 100 across the major parties because voters may choose other options, such as third parties or abstaining. Results are for the main specification.

of the increase in turnout for Democrats was due to increases among white-Democrats (Table 7). Second, the white-Democrats who were induced to turn out were moderates who had a decent chance of voting for Republican candidates ("White-Dem." columns of Table 9).

Results for the alternative specifications are provided in Online Appendix Tables A14–A17. As in the main specification, they show that competitive districts would have spurred higher turnout but would have caused only small changes in vote margins. One difference is that, in the alternative specifications, vote margins always shift in favor of Democrats. This is because, in these specifications, more of the increase in Democratic turnout is from registrants with a high chance of preferring Democratic candidates, such as minority registrants.

In sum, the exercise offers two conclusions about North Carolina's recent legislative districts. First, these districts reduced turnout in comparison with a counterfactual of competitive districts. However, second, they likely had little effect on parties' statewide vote shares or on who won statewide races. The only way they could have is via channels not captured in our analysis, such as by affecting whether people become registered.

Conclusion

In this paper, we asked a central question in research on voting — whether competitive electoral environments induce additional turnout.

We studied this question in the context of American legislative districts and found a number of results. First, competitiveness does indeed spur turnout. For the U.S. House, switching from an 80–20 district to a 55–45 district increases turnout by an average of 1 percentage point per election of exposure. For the NC Senate and the NC House, the magnitude is 0.6 percentage points. Second, effects are long-lasting, with no evidence of decay in the available data. Third, effects sum across legislative chambers. Fourth, effects may be due in part to a learning channel, whereby living in a competitive district induces people to believe that races can be competitive.

One question with respect to our results is how big are the effects that we recover? The answer can be understood in three ways. First, we confirm the earlier finding of Moskowitz and Schneer (2019) that most of the correlation between district competitiveness and turnout is due to selection bias. For instance, in their replication of the non-causal literature, they show that turnout is 5–6.75 percentage points higher in 55–45 v. 80–20 U.S. House districts. By contrast, we find that the causal effect of switching between these districts for one election is only a single percentage point. Second, the impact of a moderate dosage of exposure to district competitiveness is somewhat small; however, it is on par with the effects of other voting interventions and policies.

For instance, we find that spending one election in 55–45 v. 80–20 districts for all three legislative chambers would increase turnout by 2.24 percentage points. In comparison, the increases generated by in-person canvassing, all-mail ballots, a 1 s.d. decrease in distance to the polling site, 10 days of early voting, and same-day registration are 2.5, 2–4, 1–3, 2.2, and 5 percentage points (Cantoni, 2020; Gerber *et al.*, 2013; Green *et al.*, 2013; Highton, 2004; Kaplan and Yuan, 2020). Finally, the impact of a large dosage of exposure to district competitiveness is larger than the effects of these other treatments. For instance, spending three elections in all 55–45 v. 80–20 districts would raise turnout by 6.72 percentage points.

The results in our paper have practical implications for the drawing of legislative districts. First, they imply that competitiveness is one of the criteria along which districts should be evaluated. Historically, districts have often been drawn to be uncompetitive, so as to protect incumbents. Unfortunately, these "bipartisan" gerrymanders suppress turnout, one of the core means of participation in democracy. Second, the results suggest that majority–minority districts should be configured carefully. These districts should have a large enough minority share for minorities to hold sway over electoral outcomes; however, they also should be competitive enough to ensure that outcomes aren't a foregone conclusion. Third, the results highlight that policymakers may be able to influence elections by manipulating district competitiveness. In particular, policymakers can mobilize or demobilize voters by placing them in more or less competitive districts, which can then potentially alter statewide vote shares — although the districts that were used in North Carolina during the 2010s do not seem to have had this effect.

Importantly, the results in this paper are applicable during redistricting. Stakeholders can use them to forecast the turnout impacts of different district configurations. They can do this simply by measuring competitiveness and multiplying by our causal effects.

Despite the advances in this paper, many questions remain. One task for future work is to explore persistence in the context of other voting policies. Do these policies also have the feature that their effects can "add up"? A second task is to more fully disentangle the channels via which competitiveness drives turnout. Third is to examine effects on outcomes other than turnout. How does competitiveness influence legislator behavior, pork barrel spending, the incumbency advantage, voter preferences, etc.? Some research already exists in this area (Finnegan, 2021; Lindgren and Southwell, 2014; McCarty *et al.*, 2019), but there is much still to be done. Finally, competitiveness is only one of the attributes of districts that we might care about. Others include demographic and partisan composition, whether the districts are geographically compact, whether they preserve or divide communities, and whether they generate a legislative seat share that proportionally reflects the statewide vote share.

More work is needed to understand the importance of these other attributes so that we can gain a sense of the optimal district configuration.

References

Agranov, M., J. K. Goeree, J. Romero, and L. Yariv. 2017. "What Makes Voters Turn Out: The Effects of Polls and Beliefs". *Journal of the European Economic Association*. 16(3): 825–56.

Ansolabehere, S., J. M. Snyder, and C. Stewart. 2000. "Old Voters, New Voters, and the Personal Vote: Using Redistricting to Measure the Incumbency Advantage". *American Journal of Political Science*. 44(1): 17–34.

Baker, A. C., D. F. Larcker, and C. C. Y. Wang. 2022. "How Much Should We Trust Staggered Difference-in-differences Estimates?" *Journal of Financial Economics*. 144(2): 370–95.

Bendor, J., D. Diermeier, and M. Ting. 2003. "A Behavioral Model of Turnout". *The American Political Science Review*. 97(2): 261–80.

Bhatt, R., E. Dechter, and R. Holden. 2020. "Registration Costs and Voter Turnout". *Journal of Economic Behavior & Organization*. 176: 91–104.

Braconnier, C., J.-Y. Dormagen, and V. Pons. 2017. "Voter Registration Costs and Disenfranchisement: Experimental Evidence from France". *American Political Science Review*. 111(3): 584–604.

Bursztyn, L., D. Cantoni, P. Funk, F. Schonenberger, and N. Yuchtman. 2023. "Identifying the Effect of Election Closeness on Voter Turnout: Evidence from Swiss Referenda". *Journal of the European Economic Association*.

Cancela, J. and B. Geys. 2016. "Explaining Voter Turnout: A Meta-analysis of National and Subnational Elections". *Electoral Studies*. 42: 264–75.

Cantoni, E. 2020. "A Precinct Too Far: Turnout and Voting Costs". *American Economic Journal: Applied Economics*. 12(1): 61–85.

Cantoni, E. and V. Pons. 2019. "Strict ID Laws Don't Stop Voters: Evidence from a U.S. Nationwide Panel, 2008–2018". Working Paper No. 25522. National Bureau of Economic Research.

Cantoni, E. and V. Pons. 2022. "Does Context Outweigh Individual Characteristics in Driving Voting Behavior? Evidence from Relocations within the United States". *American Economic Review*. 112(4): 1226–72.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs". *The Quarterly Journal of Economics*. 134(3): 1405–54.

Clarke, P. and S. H. Evans. 1983. *Covering Campaigns: Journalism in Congressional Elections*. Stanford, CA: Stanford University Press.

Coppock, A. and D. P. Green. 2016. "Is Voting Habit Forming? New Evidence from Experiments and Regression Discontinuities". *American Journal of Political Science*. 60(4): 1044–62.

Cox, G. W. and M. C. Munger. 1989. "Closeness, Expenditures, and Turnout in the 1982 U.S. House Elections". *The American Political Science Review*. 83(1): 217–31.

Davis, B. R. 2019. "Are Majority Minority Districts Too Safe?: A Look at the Alabama State Legislature". *Du Bois Review: Social Science Research on Race*. 16(1): 157–75.

Degan, A. and A. Merlo. 2011. "A Structural Model of Turnout and Voting in Multiple Elections". *Journal of the European Economic Association*. 9(2): 209–45.

Dellavigna, S., J. A. List, U. Malmendier, and G. Rao. 2016. "Voting to Tell Others". *The Review of Economic Studies*. 84(1): 143–81.

Downs, A. 1957. "An Economic Theory of Political Action in a Democracy". *Journal of Political Economy*. 65(2): 135–50.

Enos, R. D. and A. Fowler. 2014. "Pivotality and Turnout: Evidence from a Field Experiment in the Aftermath of a Tied Election". *Political Science Research and Methods*. 2(2): 309–19.

Enos, R. D. and A. Fowler. 2018. "Aggregate Effects of Large-Scale Campaigns on Voter Turnout". *Political Science Research and Methods*. 6(4): 733–51.

Esponda, I. and D. Pouzo. 2017. "Conditional Retrospective Voting in Large Elections". *American Economic Journal: Microeconomics*. 9(2): 54–75.

Fauvelle-Aymar, C. and A. Francois. 2006. "The Impact of Closeness on Turnout: An Empirical Relation Based on a Study of a Two-Round Ballot". *Public Choice*. 127(3/4): 469–91.

Feddersen, T. and A. Sandroni. 2006. "A Theory of Participation in Elections". *American Economic Review*. 96(4): 1271–82.

Finnegan, J. J. 2021. "Changing Prices in a Changing Climate: Electoral Competition and Fossil Fuel Taxation". *APSA Preprints*.

Fraga, B. L. 2016. "Redistricting and the Causal Impact of Race on Voter Turnout". *The Journal of Politics*. 78(1): 19–34.

Fraga, B. L., D. J. Moskowitz, and B. Schneer. 2021. "Partisan Alignment Increases Voter Turnout: Evidence from Redistricting". *Political Behavior*.

Fujiwara, T., K. Meng, and T. Vogl. 2016. "Habit Formation in Voting: Evidence from Rainy Elections". *American Economic Journal: Applied Economics*. 8(4): 160–88.

Gerber, A. S., D. P. Green, and R. Shachar. 2003. "Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment". *American Journal of Political Science*. 47(3): 540–50.

Gerber, A. S., G. A. Huber, and S. J. Hill. 2013. "Identifying the Effect of All-Mail Elections on Turnout: Staggered Reform in the Evergreen State". *Political Science Research and Methods*. 1(1): 91–116.

Gerber, A., M. Hoffman, J. Morgan, and C. Raymond. 2020. "One in a Million: Field Experiments on Perceived Closeness of the Election and Voter Turnout". *American Economic Journal: Applied Economics*. 12(3): 287–325.

Ghitza, Y., A. Gelman, and J. Auerbach. 2022. "The Great Society, Reagan's Revolution, and Generations of Presidential Voting". *American Journal of Political Science*.

Green, D. P., M. C. McGrath, and P. M. Aronow. 2013. "Field Experiments and the Study of Voter Turnout". *Journal of Elections, Public Opinion and Parties*. 23(1): 27–48.

Hamlin, A. and C. Jennings. 2011. "Expressive Political Behaviour: Foundations, Scope and Implications". *British Journal of Political Science*. 41(3): 645–70.

Harris v. McCrory. 2016. 159 F.Supp.3d 600. U.S. District Court for the Middle District of North Carolina.

Hayes, D. and S. C. McKee. 2012. "The Intersection of Redistricting, Race, and Participation". *American Journal of Political Science*. 56(1): 115–30.

Healy, A. and N. Malhotra. 2013. "Retrospective Voting Reconsidered". *Annual Review of Political Science*. 16(1): 285–306.

Henderson, J. A., J. S. Sekhon, and R. Titiunik. 2016. "Cause or Effect? Turnout in Hispanic Majority-Minority Districts". *Political Analysis*. 24(3): 404–12.

Highton, B. 2004. "Voter Registration and Turnout in the United States". *Perspectives on Politics*. 2(3): 507–15.

Hill, D. and S. C. McKee. 2005. "The Electoral College, Mobilization, and Turnout in the 2000 Presidential Election". *American Politics Research*. 33(5): 700–25.

Indridason, I. 2008. "Competition and Turnout: The Majority Run-off as a Natural Experiment". *Electoral Studies*. 27: 699–710.

Kanazawa, S. 1998. "A Possible Solution to the Paradox of Voter Turnout". *The Journal of Politics*. 60(4): 974–95.

Kaplan, E. and H. Yuan. 2020. "Early Voting Laws, Voter Turnout, and Partisan Vote Composition: Evidence from Ohio". *American Economic Journal: Applied Economics*. 12(1): 32–60.

Kawai, K., Y. Toyama, and Y. Watanabe. 2021. "Voter Turnout and Preference Aggregation". *American Economic Journal: Microeconomics*.

Lindgren, E. and P. Southwell. 2014. "The Effect of Electoral Competitiveness on Ideological Voting Patterns in the U.S. House, 2002–10". *Politics & Policy*. 42(6): 905–18.

McCarty, N., J. Rodden, B. Shor, C. Tausanovitch, and C. Warshaw. 2019. "Geography, Uncertainty, and Polarization". *Political Science Research and Methods*. 7(4): 775–94.

McDonald, M. P. and C. J. Tolbert. 2012. "Perceptions vs. Actual Exposure to Electoral Competition and Effects on Political Participation". *Public Opinion Quarterly*. 76(3): 538–54.

Meredith, M. 2009. "Persistence in Political Participation". *Quarterly Journal of Political Science*. 4(3): 187–209.

Morton, R. B., D. Muller, L. Page, and B. Torgler. 2015. "Exit Polls, Turnout, and Bandwagon Voting: Evidence from a Natural Experiment". *European Economic Review*. 77: 65–81.

Moskowitz, D. J. and B. Schneer. 2019. "Reevaluating Competition and Turnout in U.S. House Elections". *Quarterly Journal of Political Science*. 14(2): 191–223.

Paola, M. D. and V. Scoppa. 2014. "The Impact of Closeness on Electoral Participation Exploiting the Italian Double Ballot System". *Public Choice*. 160(3/4): 467–79.

Riker, W. H. and P. C. Ordeshook. 1968. "A Theory of the Calculus of Voting". *The American Political Science Review*. 62(1): 25–42.

Sekhon, J. S. and R. Titiunik. 2012. "When Natural Experiments are Neither Natural nor Experiments". *American Political Science Review*. 106(1): 35–57.

Shachar, R. and B. Nalebuff. 1999. "Follow the Leader: Theory and Evidence on Political Participation". *American Economic Review*. 89(3): 525–47.

Simonovits, G. 2012. "Competition and Turnout Revisited: The Importance of Measuring Expected Closeness Accurately". *Electoral Studies*. 31(2): 364–71.

Spenkuch, J. L. and D. Toniatti. 2018. "Political Advertising and Election Results". *The Quarterly Journal of Economics*. 133(4): 1981–2036.

Stephanopoulos, N. O. and C. Warshaw. 2020. "The Impact of Partisan Gerrymandering on Political Parties". *Legislative Studies Quarterly*. 45(4): 609–43.

Washington, E. 2012. "Do Majority-Black Districts Limit Blacks' Representation? The Case of the 1990 Redistricting". *The Journal of Law and Economics*. 55(2): 251–74.