American Political Science Review (2026) 1–20
doi:10.1017/S0003055426101610

Redistricting Reforms Reduce Gerrymandering by Constraining Partisan Actors

CORY McCARTAN *Pennsylvania State University, United States*
CHRISTOPHER T. KENNY *Princeton University, United States*
TYLER SIMKO *University of Michigan, United States*
EMMA EBOWE *College of William and Mary, United States*
MICHAEL Y. ZHAO *Harvard College, United States*
KOSUKE IMAI *Harvard University, United States*

Political actors often manipulate redistricting plans to gain electoral advantages, a process known as gerrymandering. Several states have implemented institutional reforms to address this problem, such as establishing map-drawing commissions. Estimating the impact of these reforms is challenging because each state structures its processes and rules differently. We model redistricting as a sequential game whose equilibrium solution summarizes multistep institutional interactions as a univariate score. We argue that this score measures the leeway political actors have over the partisan lean of the final plan. Our differences-in-differences analysis demonstrates that reforms reduce partisan bias and increase competitiveness by constraining partisan actors. We perform a counterfactual policy analysis to estimate the effects of enacting recent reforms nationwide. Though commissions are likely to reduce bias, reforms that restrict partisan actors in multiple ways, such as removing veto points (Michigan), are more effective than commissions where parties retain some control (Ohio).

INTRODUCTION

Democratic institutions play a crucial role in preventing political actors from pursuing policies that prioritize their own interests over the broader public good (Federalist Papers 53 and 54: Madison 1788a; 1788b). Further, increasing political polarization and the erosion of democratic norms in many countries have encouraged more recent attention to placing additional institutional constraints on policy makers (Levitsky and Ziblatt 2023; Little and Meng 2024; McCarty 2019). To safeguard institutions, advocates have suggested reforms that insulate democratic processes from partisan control.[1] For example, the Electoral Count Reform and Presidential Transition Improvement Act of 2022 represents an effort to make it more difficult for partisan actors to manipulate the presidential electoral certification process in the United States.

We study reform in American congressional redistricting, a political process often exploited by partisan actors to enact districting plans that favor their own party. This manipulation, known as *partisan gerrymandering*, has been widespread in the past two redistricting cycles (Kenny et al. 2023; 2024; Warshaw, McGhee, and Migurski 2022). Redistricting plans that disproportionately favor a certain party can limit how responsive a party's share of seats in the legislature is to changes in its vote share and can reduce the electoral power of racial minorities (e.g., Canon 2022; 1999; Grofman and Handley 1991; Polsby and Popper 1991).

Methodological Challenges and Proposed Approach

Reform efforts to limit gerrymandering are often designed to constrain partisan map drawers. They include the establishment of independent map-drawing commissions and the introduction of court oversight over proposed plans (Cain 2012). Estimating the causal

Corresponding author: Cory McCartan, Assistant Professor, Department of Statistics, Pennsylvania State University, United States, mccartan@psu.edu.
Christopher T. Kenny, Postdoctoral Research Associate, Data-Driven Social Science, Princeton University, United States, ctkenny@princeton.edu.
Tyler Simko, Assistant Professor, Department of Political Science, University of Michigan, United States, tsimko@umich.edu.
Emma Ebowe, Assistant Professor, Department of Government, College of William and Mary, United States, evebowe@wm.edu.
Michael Y. Zhao, Alumnus, Harvard College, United States; now at Google DeepMind, michaelzhao@alumni.harvard.edu.
Kosuke Imai, Professor, Department of Government and Statistics, Harvard University, United States, imai@harvard.edu.

Handling editor: Andrew Eggers.

Received: February 11, 2025; revised: August 10, 2025; accepted: February 24, 2026.

[1] For example, see "People Not Politicians Oregon" or "Ohio's Citizens Not Politicians."

impact of these institutional reforms, however, is challenging for three reasons. The first is the problem of *treatment complexity*. Redistricting reform efforts must intervene in a complex and multidimensional process. The specific rules governing each state's redistricting process differ in many ways, including who proposes initial maps and whether or not courts can intervene.

Second, states that adopt redistricting reforms may differ in both observable and unobservable ways from states that do not. These differences include statewide variations in institutional characteristics and political contexts. This *confounding bias* problem is common to any observational study and must be addressed to estimate credible causal effects.

Finally, the outcome of the institutional process—a redistricting plan—is also complex. For example, partisan features may be confounded by other factors, such as a state's geography and demographics (e.g., Cottrell 2019). If a large number of Democratic voters live in cities, any redistricting plan with compact districts may end up creating a small number of heavily Democratic-leaning districts rather than efficiently allocating Democratic votes across more districts. We must overcome this *outcome complexity* to accurately measure the partisan bias of each enacted plan.

In this article, we propose a new methodological approach that addresses the above challenges to study the causal impact of redistricting reform. Figure 1 summarizes our methodology. To deal with the *treatment complexity*, we first collect and standardize information about each state's relevant laws, starting from the initial map drawer (i.e., the redistricting commission or legislature) and following through varying stalemate processes and opportunities for court intervention. We then reduce the differing districting procedures across states to a single theoretically informed parameter by modeling redistricting reform as a zero-sum sequential game. We use our original dataset of institutional procedures to characterize the players and available moves in the game. Analyzing the game allows us to measure the ability of partisan players to maximize the partisan lean of a redistricting plan. The Nash equilibrium of the game is a measure of the "leeway" that a single party has over the final redistricting plan.

We use two variations of each state's game to produce two measures of leeway. The first is the *realized* leeway, which uses the observed parties of the players to compute the equilibrium. The second is the *maximum* leeway, which instead computes the equilibrium under one-party control. The latter allows us to measure the strength of institutions separately from party control, which varies across states and over time. We emphasize that both of these leeway variables are summary measures of our high-dimensional treatment variables of interest (i.e., institutional features) and are not functions of the outcome variable of interest. Therefore, our methodology places an observational study of institutional changes under the design-based approach of causal inference (Rubin 2008).

Specifically, we address the *confounding bias* problem by using our leeway measures as a continuous treatment variable in a differences-in-differences (DiD) design applied to the 2010 and 2020 redistricting cycles. For example, states with Democratic control tend to enact more liberal policies (Caughey, Xu, and Warshaw 2017). This approach addresses potential confounding by comparing changes in states that have enacted reforms to those in similar states that have not, under a parallel trends assumption (Callaway, Goodman-Bacon, and Sant'Anna 2024). We use this model to estimate how changes in the map drawer's leeway influence resulting

FIGURE 1. Schematic Summary of Our Methodology

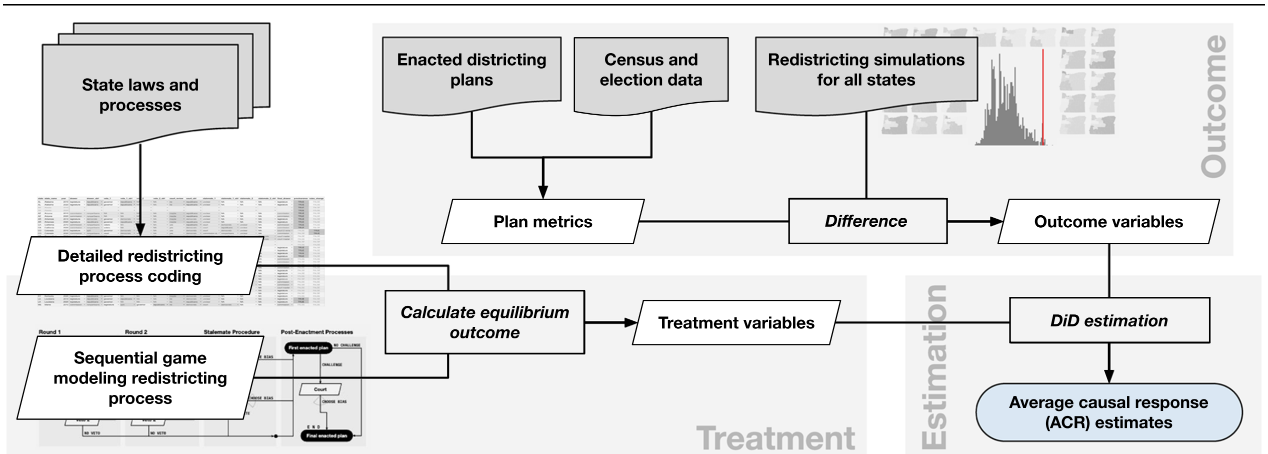

Note: Our approach is designed to address three key methodological challenges in the study of redistricting reform. First, we address *treatment complexity* by modeling the redistricting process as a zero-sum sequential game to estimate theoretically informed parameters that serve as our treatment. Second, we address *outcome complexity* by generating representative distributions of simulated redistricting plans for each state, which adjust for state-specific changes in political geography. Finally, we address *confounding bias* in causal effect estimation with a difference-in-differences design that uses simulated alternative to strengthen the credibility of the parallel trends assumption.

plans across a set of both partisan and nonpartisan outcomes, such as the responsiveness to swings in partisan preferences and the number of expected seats per party.

Our approach uses a formal model to summarize a multidimensional treatment variable, and then applies a causal inference methodology to adjust for bias due to unobserved confounding. However, we do not adopt a structural modeling approach, which would require modeling a state's decision to adopt particular reforms for their redistricting process. The reason is that, while we are able to measure all the institutional features of redistricting processes, it is impossible to observe all the factors which affect a state's decision to adopt redistricting reforms. Instead, we use a DiD design to address the issue of unobserved confounding.

Lastly, we deal with the issue of *outcome complexity* by separating the causal effect of state-specific institutional reforms from changes in political geography over time. To do this, we generate a sample of alternative redistricting plans via a simulation algorithm (McCartan and Imai 2023) following each state's rules and political geography, but without regard to any partisan information (Kenny et al. 2024; McCartan et al. 2022). We use these simulated plans as a nonpartisan baseline for both 2010 and 2020 redistricting cycles and compute the difference in outcome variables between the enacted and simulated plans. This adjusts for state-specific changes in political geography, making the parallel trends assumption more plausible.

Summary of Findings

We find that more restrictive redistricting processes reduce partisan bias by constraining map drawers. For example, changing from a single-party legislature to an independent commission leads to a reduction of about 0.5 excess seats. These effects are substantively large, as congressional gerrymandering generally results in gains of less than one or two seats per state (Kenny et al. 2023). Similarly, we find smaller but positive effects of constraining reforms on electoral responsiveness. We estimate that a similar change from legislature to commission would increase the share of competitive seats within a state from 25% to 38% on average.

A key advantage of our methodological approach is the ability to perform a counterfactual analysis of institutional reforms. We investigate how enacting recent procedural reforms nationwide could reduce widespread partisan gerrymandering. We quantify how the partisan bias and responsiveness of adopted plans could counterfactually change if all states nationwide adopted three kinds of commission structures currently enacted in several states: (1) an Ohio-style approach that requires supermajorities and uses a bipartisan backup commission, (2) a New York-style commission with a nonpartisan map drawer but several partisan veto points later in the process, and (3) a Michigan-style commission with a nonpartisan commission drawer, no partisan veto points, and the potential for court review.

We find that commissions are likely to reduce existing pro-Republican bias, but the details of commission structure matter. In particular, unlike reforms adopted in Ohio and New York, a Michigan-style nonpartisan commission has no partisan veto points. We find that implementing Michigan-style reforms nationwide leads to an additional 7.8 Democratic seats, on average. All three styles of reforms increase electoral responsiveness. We emphasize that a counterfactual policy analysis including ours involves some degree of extrapolation and should be interpreted with caution. In particular, our findings do not necessarily imply that nonpartisan commissions will always advantage Democrats. The reason is that our findings depend on the current and past political geography, electoral environments, and modeling assumptions.

Contributions to the Literature

We make two primary substantive contributions to the literature on redistricting reform. First, we help address a debate in the literature about whether, and how, reforms impact redistricting plans. Existing work on the effectiveness of redistricting reforms presents conflicting arguments and findings, with Nelson (2023) concluding that "the efficacy of redistricting reforms is contested in political science" (207). For example, some studies argue that commissions produce fairer plans, largely by removing self-serving electoral incentives from legislative map drawers (Carson and Crespin 2004; Carson, Crespin, and Williamson 2014; Edwards et al. 2017; Keena et al. 2021; Lindgren and Southwell 2013; Litton 2012; McDonald 2004; Nelson 2023). In contrast, others reach much less optimistic conclusions about the efficacy of redistricting commissions and find limited or no impact of reforms on outcomes like competitiveness and partisan bias (Cottrill 2012; Henderson, Hamel, and Goldzimer 2018; Kousser, Phillips, and Shor 2018; Miller and Grofman 2013; Seabrook 2017). We contribute to this debate by formalizing and presenting an approach that focuses on how reforms impact *leeway*—the relative control political actors have over the resulting redistricting plan.

Some work has focused on the long-term decrease of partisan bias in redistricting. Caughey and Warshaw (2022) highlight how the size of the partisan bias has largely decreased since the 1940s, especially after the reapportionment revolution, though the trend has flattened in recent cycles. We focus on this most recent period. While we do not explain the long-term decrease, we offer policy evaluations on how restrictive reforms could help restart the decrease.

Second, we contribute to the literature on redistricting reform by presenting the most comprehensive empirical evidence yet available about how the full process of redistricting shapes political outcomes. Most existing studies examine reform structures by comparing outcomes from single aspects (e.g., who draws initial maps or whether courts can intervene) of much more complex redistricting process (Carson and Crespin 2004; Carson, Crespin, and Williamson 2014; Edwards et al. 2017; Nelson 2023). We argue that redistricting is best examined by analyzing its entire process. Crucially, our approach allows us to demonstrate that various reforms implemented together can be more effective than a single reform alone.

This process-based approach requires several methodological innovations for studying redistricting reform. We believe that the proposed new methodological approach can also be applied to the study of other institutional systems. Reform efforts in redistricting have produced diverse institutional changes across states. While most scholars have classified reforms into different categories (Cain 2012; Edwards et al. 2017; Nelson 2023; Warshaw, McGhee, and Migurski 2022), such an approach may miss important nuances and potential interactions between features of these institutional changes.

In contrast, we use formal modeling to place these complex institutional characteristics on a continuous univariate scale and summarize how they constrain partisan actors. This theoretically driven approach, which we demonstrate accurately predicts empirical patterns, makes it possible to apply a DiD strategy to our complex setting for credible causal inference.

We also advance the literature on redistricting reform by addressing the aforementioned three methodological challenges in our unified approach. First, most studies of redistricting electoral reforms have ignored *treatment complexity* by focusing on only one or two aspects of the reform process at a time. For example, some researchers use an indicator for the existence of a redistricting commission (e.g., Carson and Crespin 2004; Carson, Crespin, and Williamson 2014); while others account for different types of commissions (e.g., Edwards et al. 2017; Nelson 2023). Similarly, much of the existing causal research on redistricting has focused on a single aspect of the process, such as being directly impacted by the *Shelby* decision (Komisarchik and White 2025) or being placed in a packed district (Fraga, Moskowitz, and Schneer 2022). In contrast, we model the entire process of redistricting using a unified formal theoretic framework. This approach enables us to examine how various institutional features affect redistricting outcomes rather than studying each feature in isolation.

Second, much of prior work has been descriptive in nature and primarily relied upon cross-sectional comparisons (e.g., Best et al. 2021; Carson and Crespin 2004; Carson, Crespin, and Williamson 2014; Edwards et al. 2017; Keena et al. 2021; Nelson 2023; Warshaw, McGhee, and Migurski 2022). This can lead to *confounding bias* if states that adopted redistricting reforms politically differ from those that did not. Notably, many medium-to-large Democratic states have adopted these reforms. We address confounding bias by examining the over-time changes in institutional rules and employing a DiD design.

Finally, we address *outcome complexity* through the use of redistricting simulations that account for factors such as political geography. Specifically, these simulated alternative districts allow us to differentiate the impact of redistricting reforms from that of changes in political geography over time. The period between redistricting cycles in the United States is a full decade, allowing political geography to undergo meaningful changes. Most studies, with some exceptions (e.g., Best et al. 2021; Warshaw, McGhee, and Migurski 2022), do not account for changes in underlying political geography when studying redistricting reform. Even studies that address the outcome complexity do not address the above treatment complexity issue.

Beyond redistricting, our methodological approach can be seen as a general strategy for the causal analysis of complex institutional reforms. By combining a game-theoretic treatment model with causal inference methods, we are able to leverage the strengths of these two approaches. A game-theoretic approach is a substantively effective way to model formal institutions with specific rules, making it possible to map multidimensional policies to a univariate summary in a theoretically-informed way. Once this summarization of various institutional reforms is done, we can apply standard causal inference methods to estimate the effects of counterfactual policies by mapping them directly to this univariate treatment variable.

Our work relates to a broader methodological literature in the social sciences. In economics, for example, Chetty (2009) advocates the combined use of structural and reduced-form approaches via a sufficient statistic in a manner similar to our approach. In sociology, Lundberg, Johnson, and Stewart (2021) emphasize the importance of theoretically informed quantities of interest in causal analysis. Finally, in political science, Canen and Ramsay (2024) call for the integration of rigorous theoretical and empirical approaches in causal research. We demonstrate how such an analysis can be done in the estimation of causal effects of redistricting reforms.

OVERVIEW OF REDISTRICTING PROCESSES AND REFORMS

Before introducing our methodological framework, we provide a brief overview of redistricting processes and reforms. Every decade following the U.S. Census, states and localities use many different procedures to redraw their legislative district boundaries. For example, states differ on whether legislatures or independent redistricting commissions propose initial plans for new Congressional district maps. If these actors fail to produce a plan, state laws further vary on how these stalemates are handled. Some states pass responsibilities to a court (e.g., Virginia), a backup commission (e.g., Ohio), or a group of state party leaders (e.g., Iowa). Even plans that pass the proposal step can face a veto from other actors, such as the governor and state legislature.

As explained below, we collect information about the redistricting process used in each state for the 2010 and 2020 cycles. Figure 2 summarizes this dataset and illustrates the diversity of 2020 redistricting procedures across states. For example, 29 states drew initial plans in their legislature, while 15 states used an independent redistricting commission instead. There is a significant variation in the procedures following this initial draw, with seven distinct veto mechanisms across all the

FIGURE 2. Summarized Redistricting Procedures in 2020

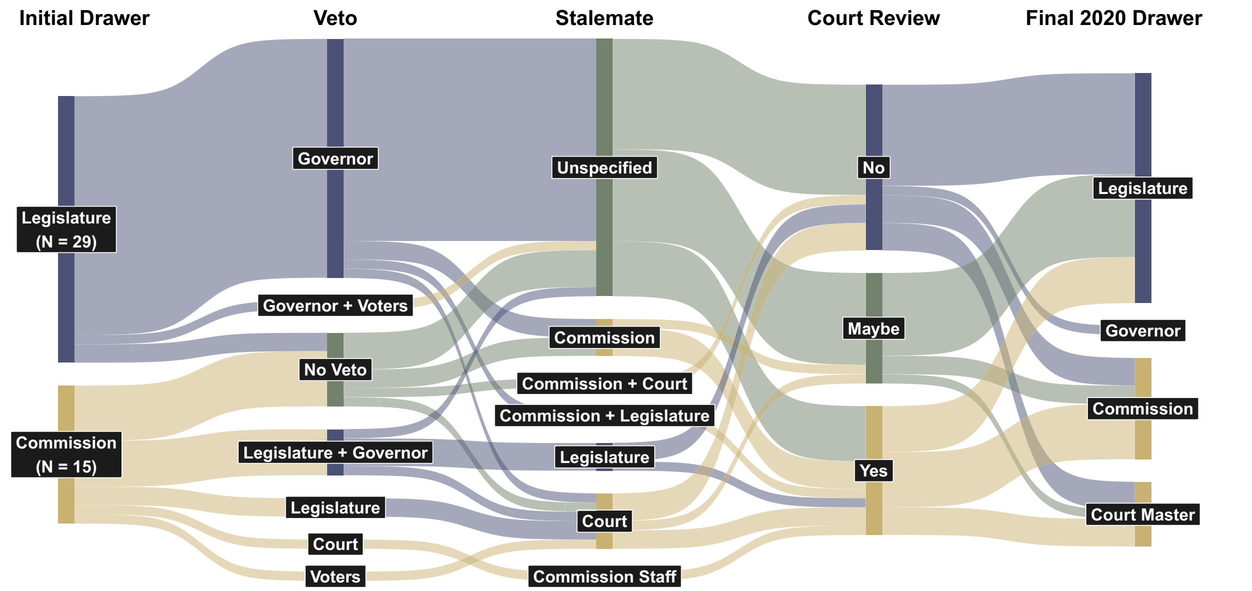

Note: Redistricting procedures for all 44 states with more than one district in 2020. Each vertical column indicates a separate step in the redistricting process, and nodes indicate different procedures that each state can adopt at that step. The width of each area connecting the nodes is proportional to the number of states with that specific combination of procedure at both ends. Yellow nodes indicate actors or institutions that are not explicitly partisan, while blue nodes indicate explicitly partisan actors or choices. Green nodes indicate cases where the procedure is not known or does not exist. Here, we collapse multiple potential stalemate and veto procedures into one step for visual clarity (e.g., Governor + Voters indicates the possibility of a first veto by a governor, and a second by the voters).

states. Much of this variation originates from state-specific policy campaigns that have adopted widely different goals driven by local actors (Keena et al. 2021). For example, California created a redistricting commission that has the power of drawing congressional district boundaries through a series of ballot propositions in the early 2000s. In recent years, several states have also implemented redistricting reforms in response to organized efforts to limit partisan influence. They include Michigan whose voters approved a constitutional amendment to establish an independent redistricting commission in 2018.

These procedural differences make it difficult to attribute causal effects to particular institutional designs or reforms. Most existing work has turned to classification schemes that simplify this variation by assigning the control of map redrawing to a single drawer, typically the creator of the initial or final plan (Carson, Crespin, and Williamson 2014; Edwards et al. 2017). A more complex alternative would be to create indices that count the number of times a particular action appears in the state procedure (e.g., the number of veto points). For example, as shown in Figure 2, even states that have a commission propose their initial plan vary drastically in how the process continues afterwards.

While these simplifications make the study of complex procedures tractable, they risk oversimplifying complex procedures in two ways. First, attributing institutional outcomes to a single actor ignores the fact that the final redistricting plan is produced through a series of steps. Take, for example, the New York redistricting process in 2020. An independent bipartisan commission had the power to propose an initial plan, but failed to agree on a single plan. State law required the stalemated process to move to the legislature, which adopted a plan that Republican and civil rights groups criticized as a Democratic gerrymander. These groups challenged the adopted plan in a series of lawsuits, and in 2022, the New York Court of Appeals struck down the plan and tasked a court-appointed special master with drawing a remedial plan. Though the final plan scores well on quantitative fairness metrics (Kenny et al. 2023), simply classifying a "court" as the sole map drawer for New York in 2020 ignores the partisan interests involved, failing to capture the complexity of the procedures that led to the final plan.

Second, classification schemes can overlook strategic interactions, where the behavior of certain actors depends on the presence or characteristics of others. For example, a commission might draw a different map if it knew that the map could later be reviewed by a court. Or, the potential of a governor's veto may limit the likelihood of a partisan gerrymander by the state legislature, but not in cases where the governor and legislative majority share a partisan affiliation. More detailed coding of procedural schemes can account for some of these interaction effects and increase realism, but will necessarily decrease statistical power, making it more difficult to estimate causal effects.

FIGURE 3. Prototypical Game Tree Used to Model Redistricting

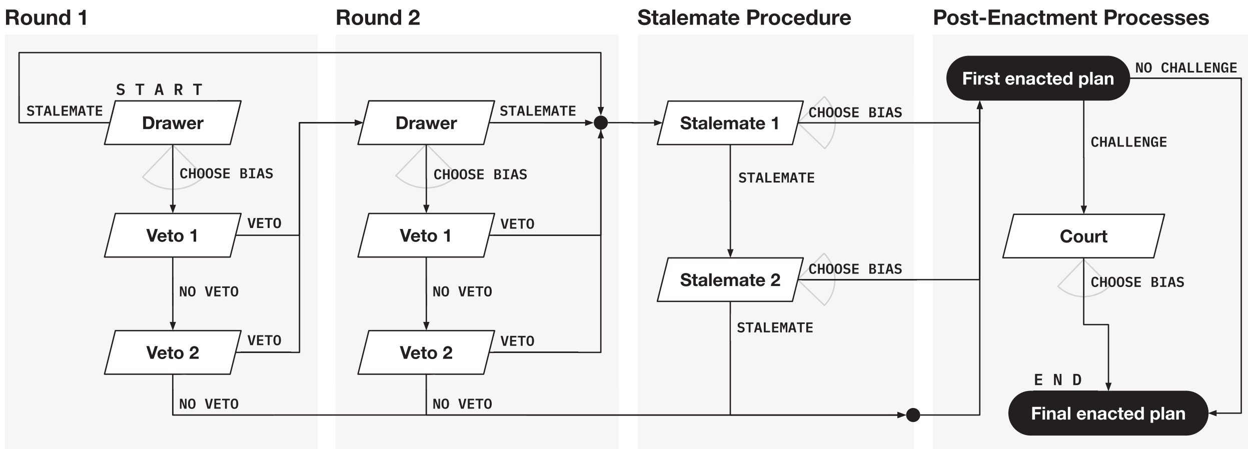

Note: States differ in which party, if any, controls each node, and which nodes are present in the state's process.

Thus, we face a methodological dilemma: while some simplification of institutional features is necessary, common approaches to doing so obscure critical characteristics of redistricting processes. In the next section, we propose a theoretically grounded approach that models redistricting as a zero-sum sequential game and uses its Nash equilibrium as a treatment variable to summarize these complex institutions.

A THEORETICAL MODEL OF INSTITUTIONAL LEEWAY

To study the impact of redistricting reforms and processes, we must first address the problem of *treatment complexity*. We develop a standardized set of 14 institutional features that capture the most important actors in congressional redistricting processes across states. We collect data on these features so that it is possible to compare redistricting processes consistently across states, despite the differences in the laws and bodies governing redistricting in each state. This step corresponds to the left top corner of Figure 1.

Using these data, we develop a sequential redistricting game to summarize these features in a theoretically informed manner (see the "treatment" box of Figure 1). Specifically, we use the Nash equilibrium of this game as a one-dimensional measure of the institutional leeway political actors have over the partisan lean of the resulting redistricting plan. Finally, we empirically validate the proposed measure by demonstrating that it is not particularly sensitive to model specification and predicts redistricting outcomes well.

We emphasize that the proposed measure is a theoretically motivated summary of our high-dimensional treatment variable of interest. Since this measure does not use outcome variables, we are able to apply standard causal inference methodology to adjust for observed and unobserved confounding factors.

Data on the Relevant Institutional Features

Our standardized coding of relevant institutional features is based on the prototypical redistricting process shown in Figure 3. First, an initial map drawer proposes a plan that may be vetoed by other actors. If the plan is not vetoed, it can be challenged in court. If it is vetoed, there is another round of map drawing. If a plan is vetoed twice, or if the initial map drawer cannot agree on a plan, then a different institution (often a court) must resolve the stalemate and adopt a plan.

Every state's process can be described as a subset of this prototype. For example, in Michigan, a commission draws congressional districts, resolves any stalemates, and there is an explicit mechanism for state court review. Thus, Michigan's process would be described by the "Drawer" step in Round 1 and "Stalemates" steps only in Figure 3.

For each state, we record which institutional body, if any, acts at each step and which party, if any, controls that institution. We also collected additional information relevant to modeling redistricting processes and court review, such as whether the state's redistricting plans were subject to DOJ preclearance before 2013,[2] and which institutional body ended up drawing the plan that was used in the first postcensal elections. Most procedural details are straightforward, and the information for coding is readily available in public data from each state (see, e.g., https://redistricting.lls.edu/). Furthermore, many states have similar redistricting procedures that are easily classified under the categories we defined earlier.

Section S1 of the Supplementary Material explains in detail how each of these variables was coded and describes special cases. Figure 2 of the previous

[2] The *Shelby* decision was handed down in 2013, ending DOJ preclearance.

section graphically summarizes this dataset for 2020 while Table S1.1 in the Supplementary Material presents all the variables across states in both 2010 and 2020.

The Redistricting Game

Coding the details of each state's redistricting process preserves important procedural information, compared to categorizing each state into a small number of groups, such as "legislature-controlled" and "independent commission." However, the detailed coding presents a challenge for causal inference, since the treatment—a state's redistricting process—is now high dimensional. Out of the 87 state-decade processes we code,[3] there are 58 distinct combinations of procedural variables, 41 of which are completely unique.

The combination of high-dimensional treatment and a limited sample size means that there may not be enough information to estimate the causal effect of changing from one specific configuration of procedural variables to another, without further assumptions.

We address this methodological challenge by leveraging two basic substantive assumptions about the redistricting process. First, each party aims to draw a map that favors it as much as possible, and second, the parties are constrained by statutory and constitutional rules in doing so. Specifically, we treat the redistricting process depicted in Figure 3 as a sequential zero-sum game[4] with two players—the Democratic and Republican parties—each trying to maximize the degree to which the drawn plan favors their party.

In each state, the nodes in the game tree can be controlled by different parties, or by neither party (e.g., when a supermajority is required to adopt a plan and neither party has supermajority control). Any split-control nodes, as well as a node for state court action, are considered moves by nature. Moves at one node do not affect the set of actions available at other nodes. Some nodes involve discrete choices, such as whether to veto a plan or not, others are labeled "choose bias," meaning that the player at that node draws a plan with a chosen amount of partisan bias favoring either party. This bias is exactly the utility for the party of the chosen plan: a plan with bias x is worth x to the Republicans and $-x$ to the Democrats. We need not quantify exactly what the utility measure is as a function of a specific plan chosen; it suffices to let the parties try to maximize an abstract univariate measure of partisan bias. We let the bias score range from -4, indicating a maximum Democratic advantage, to $+4$, indicating a maximum Republican advantage.

For concreteness, consider the case of redistricting in Oregon. In 2020, the first move belonged to the Democratic controlled legislature. It has to pick the amount of partisan bias $-4 < x < 4$ in the plan that it adopts. If this plan is ultimately the final enacted plan, Republicans receive utility x and Democrats receive utility $-x$. If the legislature fails to adopt a plan, the first stalemate move in Oregon belongs to a nonpartisan commission.

If a plan is adopted in the first round in Oregon, it proceeds to face a potential veto by the Democratic governor or a possible court review. The court may decide to accept a legal challenge, decide in favor of the plaintiffs, and redraw the map; this choice is considered a move by nature. If the court review results in a redrawn map with partisan bias x', then the Republicans receive utility x' and the Democrats receive utility $-x'$. In Oregon, court review is explicitly allowed on partisan grounds and challenges under the federal Voting Rights Act are possible, so there is moderate probability that the commission-adopted plan will be overturned. If no plan is enacted by the commission in the first round, the commission is again tasked with drawing a plan. If the commission fails to enact a plan again, courts must step in and redraw district lines to ensure compliance with federal constitutional "one person, one vote" apportionment requirements. This is also considered a move by nature.

To complete the description of the game, we must specify the rules for determining the expected outcomes for moves by nature. There are two kinds of moves by nature: map drawers controlled by neither party exclusively, and the results of court challenges. The full specification of these moves can be found in Section S2 of the Supplementary Material, but we briefly summarize them here. We make three assumptions: (1) nonpartisan map-drawers whose choices are subject to veto will generate maps which favor the party controlling the veto after a first veto has been made, (2) stalemate map-drawing will produce a map that is moderately balanced but tends to be influenced by any biases present in the most recent redistricting proposal, and (3) split-control map-drawers will stalemate with some probability and produce similar results to nonpartisan map-drawers the rest of the time. Similar assumptions about partisan map-drawers are not needed, since partisan drawers are assumed to act strategically within the game.

Finally, we decompose the court challenge process into five components: the probability that a legal challenge is possible, the probability that a challenge is made when possible, the probability that a court sides with plaintiffs, the expected remedy a court orders in those cases, and the probability and expected effect of a challenge based on the federal VRA. The latter applies only to states previously subject to DOJ preclearance. Each of these five components has a parametric specification, detailed in Section S2 of the Supplementary Material. The partisan control of state courts is accounted for in the specification of these various components, thus allowing judicial polarization to enter the picture without assuming absolute strategic coordination between the state party and its allies in the judiciary. Overall, the court challenge process is not modeled strategically, since non-party actors are more often than not the ones that initiate litigation and have different incentives than the party actors.

[3] Only 43 states had two or more congressional districts in the 2010 cycle; for 2020, Montana did as well, bringing the total to 44.

[4] The zero-sum assumption rules out cooperation between the parties to, for example, protect incumbents. Given the role of the game here as a summary of institutional features rather than a structural model that predicts the outcome, we believe this is an acceptable limitation.

All in all, the game specification depends on 19 parameters which govern the moves by nature, with most of these parameters relating to the court challenge process. Rather than fix these parameters to arbitrary constants, we place a prior distribution on each parameter over a range of probable values. We simulate 200 different draws from this joint prior distribution; each draw generates a slightly different game specification. We then average our results across the random draws.[5]

Equilibrium Solution as a Treatment Variable

To effectively summarize a high-dimensional treatment, we use the subgame perfect Nash equilibrium of the redistricting game. The utility of the equilibrium solution captures the expected partisan bias of plans that arise out of a state's redistricting process, under the current party control of the state's institutions. All of the multistep institutional interactions and negotiations that might happen as part of the redistricting process are thus reduced to a univariate score. The upshot is that the 14 procedural variables can be reduced into a single variable that measures the leeway political actors have over the partisan lean of the final redistricting plan.

To calculate the equilibrium itself, we numerically solve the game via backward induction. This requires up to four levels of nested optimization in some states. A later section describes a full example of this process for Michigan, and Section S2.5 of the Supplementary Material walks through a more complex, multistep example for Alabama. This solution process is automated across all of the states and is carried out on each of the 200 different draws from the prior on the game's parameters. This process requires the use of several tuning parameters described at length when we present the full details of the game in Section S2 of the Supplementary Material. Further, Section S8 of the Supplementary Material shows that the values of our calculated equilibria are not sensitive to the choice of initial values for these tuning parameters.

For our causal analysis, we use three treatment measures. The first treatment measure is exactly the game's Nash equilibrium, averaged across the prior. This measure depends on which parties control each node in the game tree. In our analysis of nonpartisan outcomes, we also use the absolute value of the equilibrium to capture the magnitude but not the sign of the expected bias. Finally, we also generate a third treatment measure that does not depend on the current party control. This is calculated by assigning a single party (here, the Democrats[6]) to control each node in the game tree that belongs to a partisan actor—legislature, governor, or partisan commission—and then recalculating the average Nash equilibrium. We refer to this equilibrium as the *maximum leeway* of a state's redistricting process, since it captures the expected bias under a worst-case partisan outcome where all of the levers of state government are controlled by one party. The first and second treatment measures we refer to as the *realized leeway* and *absolute realized leeway*, since they depend on the realized values of the party control for each state institution.

Realized and Maximum Leeway Scores

Figure 4 visualizes these treatment measures for all the states we study; an arrow indicates a change in the treatment value from 2010 to 2020. The realized leeway scores (game equilibria) cover the entire range of possible partisan biases, from West Virginia in 2010 with complete Democratic control of state government, to Wisconsin, Iowa, Kansas, Utah, and Nebraska, where Republican trifectas were unconstrained by Democrats or by the VRA in both cycles. The maximum leeway scores likewise span the range of possible bias, but take on fewer values, since the various actual combinations of partisan control of state institutions are no longer considered.

Intuitively, states introducing a commission (e.g., New York, Virginia, or Colorado) drastically reduced their leeway to match states with similar sets of rules in 2020. States that had intervening litigation to clarify the interpretation of state redistricting rules (e.g., New Mexico, North Carolina, or Pennsylvania) see appropriate changes in leeway, even absent a commission. Further, in our party-signed treatment, states that had a flip in party control of the whole system see large changes in leeway in the correct direction (e.g., West Virginia). Finally, states with minor changes to the total control of state institutions see minor changes in leeway (e.g., North Carolina).

Full Process Example: Michigan

Here, we build further intuition for our approach by summarizing each step described above for Michigan, a state that experienced a large change in leeway from 2010 to 2020 by adopting an independent commission. Section S2.5 of the Supplementary Material walks through a more complex example for this process in Alabama.

[5] We average, rather than propagating each draw through the estimation step. This has computational benefits and reduces the complexity of an outcome model (the outcome model would need to account for a different scale of treatment values across different draws). More importantly, however, the variation in the estimated equilibria, as described below, is quite small compared to other sources of uncertainty.

[6] The game is completely symmetric between parties, with one exception: due to racial polarization in the United States, challenges to redistricting plans under the VRA are much more likely for Republican-leaning plans, and when these challenges prevail the remedy almost always results in a plan more favorable to Democrats. In keeping with the goal of calculating the *maximum* leeway in each state, we therefore assign states to Democratic control under the counterfactual scenario here. Assigning to Republican control instead would slightly reduce the calculated leeway for those states previously subject to DOJ preclearance.

FIGURE 4. Summary of Treatment Values for All States

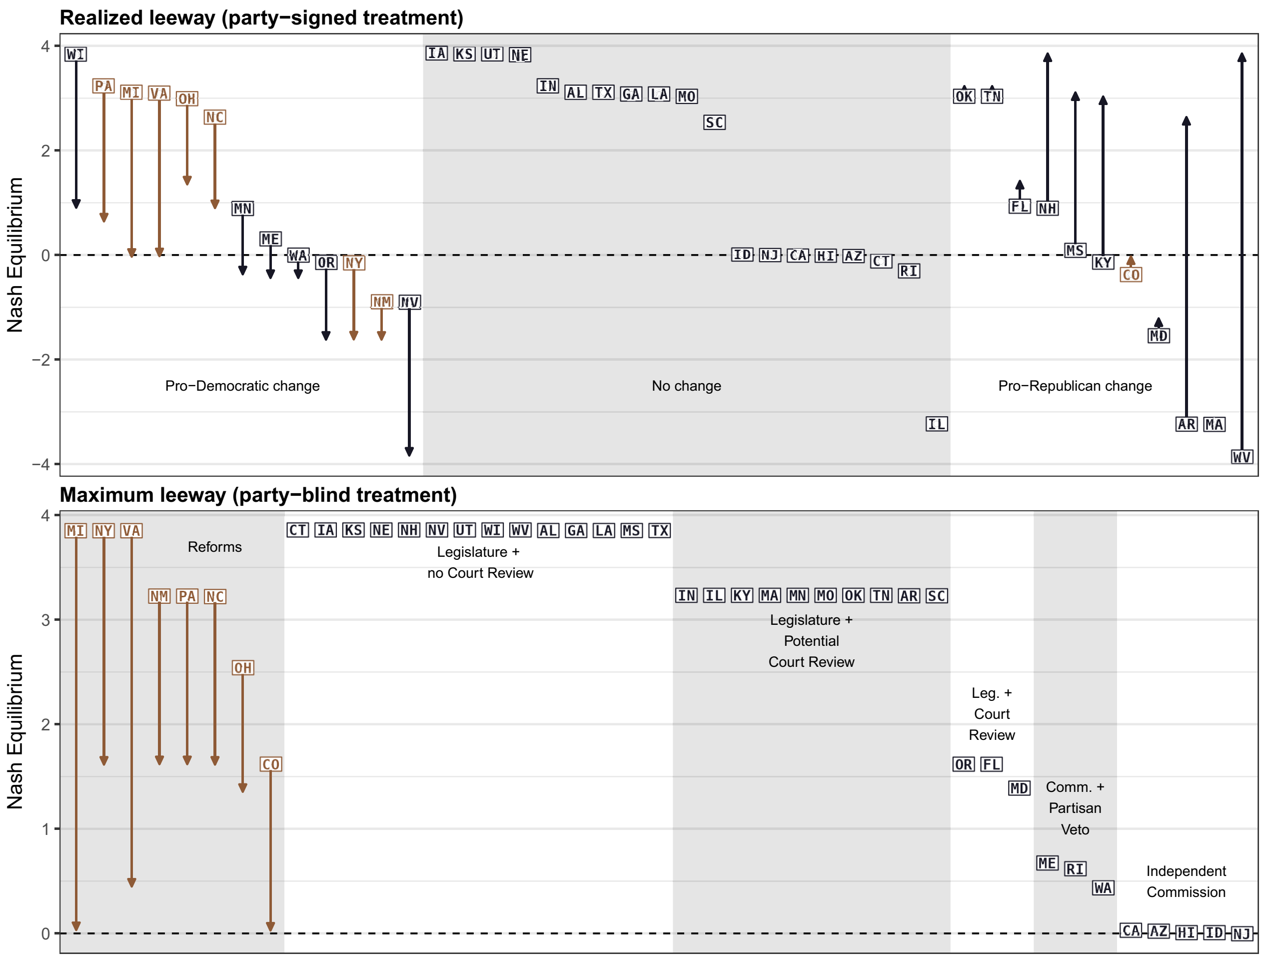

Note: Treatment values for each state in 2010, with values for 2020 indicated by arrows, where different. States in orange are those which experienced a reform to their redistricting procedures, either by legislation, constitutional amendment, or a court ruling that allowed for state court review of alleged partisan gerrymanders.

First, we outline the institutional structure of the redistricting process to establish the steps in the game. In Michigan, an independent commission drew the initial map in 2020. See Section S1 of the Supplementary Material for details on all steps for each state. Michigan's commission makes the first move, which is to pick the amount of partisan bias $-4 < x < 4$ in the plan that it adopts. As above, the bias in this plan translates into partisan utility—if this proposal is adopted, Republicans receive utility x and Democrats receive utility $-x$.

The following steps also depend on the state's specific institutional structure. In Michigan, if this initial proposal is adopted, the plan faces potential court review. Alternatively, if the commission's first proposal fails, the stalemate move also belongs to the commission (see Section S1 of the Supplementary Material). While Michigan has no partisan veto points, plans in other states could also face potential vetoes at this point (or, in states like Iowa, could even face vetoes from two sources).

With this game structure established, we solve the equilibrium as above using backward induction. Michigan's redistricting process requires one optimization step due to its lack of partisan veto points, though see Section S2.5 of the Supplementary Material for a more complex example with multiple subgames. This means that the last partisan move in the game is the initial (Round 1) commission proposal described above.

Our overall equilibrium solution for Michigan's 2020 structure is a bias of –0.03, suggesting very little bias in favor of either party. This near-zero bias reflects the lack of partisan control in this particular structure. Because our solution approach reflects institutional designs at each time period, our equilibria estimates also capture intuitive changes in leeway over time. For example, in 2010, Michigan redistricting was entirely controlled by the legislature and governor, who were both Republican. This game structure results in the highest possible value of maximum leeway (party-

FIGURE 5. Treatment Values versus Measures of Partisan Advantage

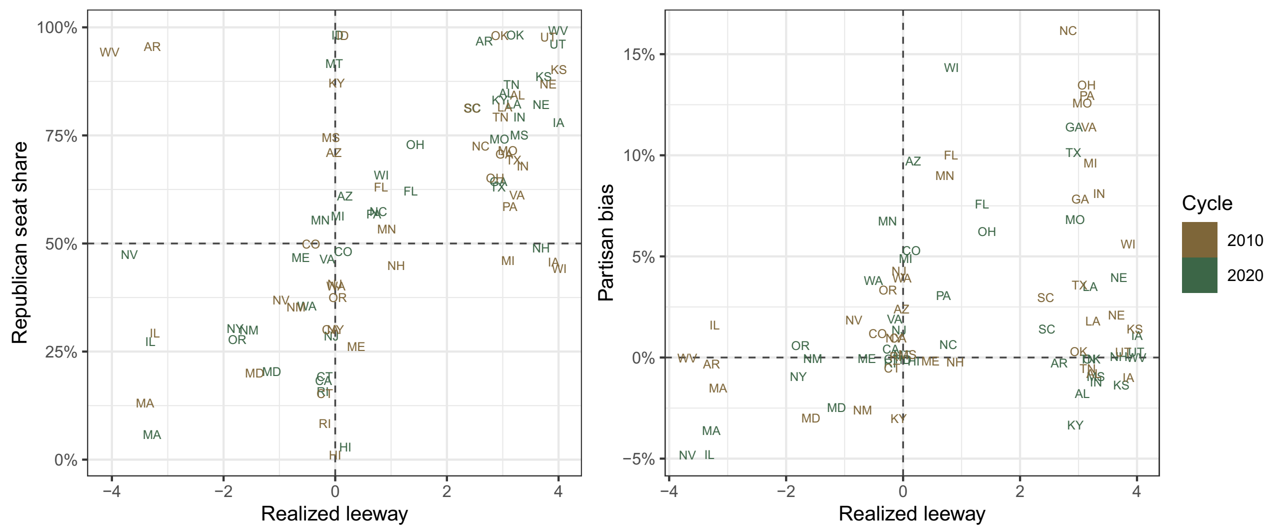

Note: Measures of partisan advantage versus treatment values for states' enacted plans for the 2010 and 2020 redistricting cycles. Points are slightly jittered to avoid overplotting.

blind) near 4, as a single party controls all important game steps. However, our near-zero 2020 equilibria reflects the adoption of a strong independent commission. In the new process, the commission is bound by strict criteria, the process allows for court review, and even the commissioners themselves are selected in part by a lottery system. With the reform, Michigan's maximum leeway changed from the highest observed value to zero, meaning that partisan actors have no control over redistricting outcomes, on average.

Empirical Validation of the Proposed Treatment Variables

Building treatment variables through a game-theoretic model provides significant advantages in dimension reduction and interpretability, but it also comes with the risk that the model is misspecified. Although it is difficult to completely verify the validity of each component of the model, we take several steps toward empirically validating the model as a whole, showing that the resulting treatment variable predicts the observed outcomes well. We also later conduct a robustness analysis for model misspecification in the context of causal effect estimation (see Section S6.2 of the Supplementary Material).

The model was designed to be flexible enough to capture the actual redistricting processes in each state. Any misspecification is therefore due to how the moves by nature are specified, or in the underlying setup of two opposing partisan actors competing on a single zero-sum dimension. First, we find that the treatment values are not sensitive to specific parameter values, up to monotonic transformations. Across the 200 random draws from the prior, the average pairwise Spearman correlation between the Nash equilibria for each state is above 0.99.[7] This also gives us confidence that the specific choice of the prior is not influencing the results.

Second, as shown in Figure 5, there is a substantial correlation in the expected direction between the treatment values and the measures of partisan advantage. The left panel plots the expected share of seats won by Republicans versus the realized leeway measure; recall that the leeway measure is positive for plans that favor Republicans. The right panel shows the partisan bias (King and Browning 1987), measured at the state's baseline vote share, versus the realized leeway. The greater values of partisan bias correspond to plans that systematically favor one party, with positive values favoring Republicans.

For both outcome measures, the purely *a priori* model predictions of the expected partisan bias do in fact correlate with the actual partisan bias of the plans that come out of each state's redistricting process. To what extent these correlations can tell us about the causal effects of reforms is of course another matter, and the primary question addressed in the next section.

We also compare the equilibrium path in the model with the actual institution that drew the final map in each state. Because the moves by nature in the game are random, the actual equilibrium path may be a probabilistic mixture over multiple possible paths. We average these probabilities over the 200 draws of the game to arrive at an overall probability that a legislature, commission, or court draws the final map in each state, given its procedure and party control.

Table 1 compares the most likely final map-drawer, according to these probabilities, to the actual institution

[7] That is, for each pair of random draws, we calculate the Spearman correlation between the two vectors of calculated state equilibria, then average these correlations across all pairings.

TABLE 1. Accuracy of Model Predictions of Final Map-Drawers

	Most likely in equilibrium			
Final drawer	Legislature	Commission	Court	Total
Legislature	26.9	0.0	20.1	47
Commission	1.9	18.1	3.0	23
Court	5.8	0.9	10.3	17
Total	34.6	19.0	33.4	87

Note: Correspondence between the institution that drew the final redistricting plan for each state and the most likely outcome based on the equilibrium path of the redistricting game.

that drew the final redistricting plan. The correspondence is generally excellent, considering that the predictions are made based on theory alone. There is a tendency, however, for the model to predict court-drawn maps when they are in fact drawn by legislatures. The overall likelihood of a court intervention is itself a model parameter that is varied across random draws, and as noted above the ranking of the states by leeway is basically unchanged by different parameter values. Therefore, we expect this tendency to overpredict court intervention to at most slightly understate the equilibrium value for states with court review, since a higher likelihood of intervention would tend to pull the equilibrium toward zero.

Finally, it is natural to wonder whether a more accurate model could be obtained by estimating the model parameters from the observed data rather than specifying their prior distributions. That is, one could find the parameter values that maximize the correlation between the treatment and the outcome. We do not take this approach to cleanly separate the estimation of causal effects from the construction of a treatment variable (Rubin 2008). This enables us to employ a causal identification strategy that does not assume the correct specification of the game-theoretic model. Instead, we use the model to summarize a high-dimensional treatment variable in a theoretically informed way without looking at the outcome variable.[8]

ESTIMATING THE CAUSAL EFFECTS OF INSTITUTIONAL LEEWAY

We now discuss how our approach addresses the remaining challenges regarding outcome complexity and confounding bias in causal effect estimation (see the "estimation" and "outcome" boxes of Figure 1). We use a DiD design with the continuous treatment variable of institutional leeway to adjust for both observed and unmeasured state-specific confounding factors. We also use a simulation approach to adjust for changes in political geography that have happened during this time period. Specifically, we generate a representative sample of nonpartisan redistricting plans for both 2010 and 2020 redistricting cycles to quantify state-specific changes in political geography. By accounting for such changes, we are able to better isolate the impact of redistricting reforms from that of political geography.

Difference-in-Differences Design

To estimate the causal effects of changes in leeway on redistricting outcomes, we use a DiD design with a continuous treatment variable. This strategy addresses potential confounding by comparing changes in states that have enacted reforms to changes in similar states that have not modified their redistricting process across the 2010 and 2020 redistricting cycles. We assume that, in the absence of reforms, the states with institutional changes would have experienced the same trend in outcome variables as those states without reforms.

Formally, let $\mathbf{Z}_{it}$ be the 14-dimensional vector of institutional features for state i at time t discussed in the data section above. We use $t = 0$ and $t = 1$ to denote 2010 and 2020 redistricting cycles, respectively. Let $u^*(\mathbf{z})$ represent the equilibrium outcome (game utility) for the average Nash equilibrium of the redistricting game described by the process $\mathbf{z}$. Then, we can define a univariate treatment variable $D_{it} = u^*(\mathbf{Z}_{it})$ for each state-decade, which represents the "dose" of institutional leeway given to political actors.

The key assumption made by our use of the Nash equilibria as the treatment variable is that the institutional features affect the outcome only through this treatment variable, that is,

$$Y_{it}(\mathbf{z}) = Y_{it}(\mathbf{z}') \text{ for any } \mathbf{z}, \mathbf{z}' \text{ with } u^*(\mathbf{z}) = u^*(\mathbf{z}'),$$

where $Y_{it}(\mathbf{z})$ denotes a generic potential outcome variable for state i at time t with institutional features $\mathbf{z}$.[9] Under this *sufficiency* assumption, we can simply write the potential outcomes as $Y_{it}(d) = Y_{it}(u^*(\mathbf{z}))$ for the treatment dose $d = u^*(\mathbf{z})$. While the assumption is not directly testable with the data we have, we evaluate robustness to its violations in Section S6.2 of the Supplementary Material by additionally controlling for $\mathbf{z}$ in effect estimation; the resulting estimates are not statistically distinguishable from the main specification's estimates.[10]

[8] We did attempt to fit the game parameters to the observed data for the 2010 cycle alone via maximum likelihood, but we found that the observed data were only minimally informative about most of the parameter values. This is good news, insofar as it implies that a wide range of parameter values are all compatible with the observed data.

[9] For identification, it in fact suffices to assume only that $\mathbb{E}[Y_{it}(\mathbf{z})|\mathbf{X} = \mathbf{x}] = \mathbb{E}[Y_{it}(\mathbf{z}')|\mathbf{X} = \mathbf{x}]$ for any $\mathbf{z}, \mathbf{z}'$ with $u^*(\mathbf{z}) = u^*(\mathbf{z}')$; however, the stronger individual-level assumption is invariant under transformations and ensures that no additional assumptions are required for estimation.

[10] It is also possible that we may have not measured all relevant features of redistricting institutions. The existence of such unmeasured features may alter the interpretation of causal effects of interest (VanderWeele and Hernán 2013). For example, the causal effect of a change in redistricting institutions would be interpreted as the effect of such a change averaging over other unmeasured institutional changes that are associated with it.

Our target estimand is the conditional average treatment effect (CATE) for a change from the treatment level d to d' given a set of covariates $\mathbf{X}_i$ that are not affected by the treatment. The CATE is defined as

$$\text{CATE}_{\mathbf{x}}(d', d) = \mathbb{E}[Y_{i1}(d') - Y_{i1}(d) | \mathbf{X}_i = \mathbf{x}].$$

Estimating $\text{CATE}_{\mathbf{x}}(d', d)$ for any pair of values d', d and covariate values $\mathbf{x}$ requires the identification of a full dose–response curve. Our identification approach relies on a strong conditional parallel trends assumption introduced by Callaway, Goodman-Bacon, and Sant'Anna (2024). Specifically, we assume that the average change in outcomes for states experiencing a change in treatment dosage from d to d' depends only on the observed covariates, and not the observed dosage value.

Our covariates include a range of variables that existing work in the redistricting literature identifies as predicting redistricting outcomes and are not affected by redistricting reforms. We include (pretreatment) 2008 Democratic presidential vote share, an indicator for being in the South,[11] the logarithm of the number of districts in 2020, the change in the number of districts between 2010 and 2020, the logarithm of the average number of state corruption convictions by year between 2000 and 2010, and an indicator for whether states allow ballot initiatives. For the maximum leeway treatment, which uses no information about partisan control of state institutions, we also control for changes in partisan control of the map-drawing body and of the state supreme court.

Formally, we require the following conditional parallel trends assumption for all d, d':

$$\begin{aligned}&\mathbb{E}[Y_{i1}(d') - Y_{i0}(d) | \mathbf{X}_i = \mathbf{x}] \\ &\quad = \mathbb{E}[Y_{i1}(d') - Y_{i0}(d) | \mathbf{X}_i = \mathbf{x}, D_{i0} = d, D_{i1} = d'].\end{aligned}$$

This assumption is analogous to the parallel trends assumption in the traditional binary DiD design, but differs in that it refers to changes in dosage across a continuous measure, rather than a single level change from zero. Compared to a traditional selection-on-observables assumption, the conditional parallel trends does not rule out confounding factors that are constant over time. With this assumption, we can identify our estimand as

$$\begin{aligned}\text{CATE}_x(d, d') = &\ \mathbb{E}[Y_{i1} - Y_{i0} | \mathbf{X}_i = \mathbf{x}, D_{i0} = d, D_{i1} = d'] \\ &- \mathbb{E}[Y_{i1} - Y_{i0} | \mathbf{X}_i = \mathbf{x}, D_{i0} = d, D_{i1} = d],\end{aligned}$$

where $Y_{it} = Y_{it}(D_{it})$ is the observed outcome. Averaging over the marginal distribution of $\mathbf{X}_i$, we can also estimate the average treatment effect (ATE).

[11] Southern states are defined as Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, South Carolina, Tennessee, Texas, Virginia, and West Virginia. This set corresponds to Confederate states, plus Kentucky, and West Virginia, which also had a history of Democratic dominance in the twentieth century and notable differences between party preferences at the state and national levels.

Evaluating the plausibility of the parallel trends assumption here is complicated by a lack of public data on precinct-level election returns for the 2000 redistricting cycle or earlier cycles. Without these data, we cannot generate simulated districts for decades before 2010, which prevents us from conducting placebo checks using the pre-trends, a common practice for DiD studies. We are able, however, to conduct a placebo check using an outcome that should be unaffected by redistricting reform, as we discuss below.

Thus, at least partially, we are left to justify the assumption based on substantive grounds. To violate the conditional parallel trends assumption, the change in outcomes (such as Democratic seats) that would have been observed in Michigan, had it not adopted reforms, would have to systematically differ from the corresponding change in other non-reform states like Wisconsin, after controlling for covariates. The covariates we include are some of the strongest predictors of treatment adoption and redistricting outcomes, which substantially increases the plausibility of this assumption.

Second, aspects of the treatment and outcome themselves weigh against the possibility of a violation. Changes in redistricting procedure are usually advanced outside of legislatures, often via citizen referenda funded by nonpartisan good-governance groups who are not responsible for map-drawing. The political parties also coordinate redistricting strategy at a national level. Further, state legislatures and other map-drawing bodies experience turnover between decades.

Together, these factors suggest that the change in Michigan's outcomes would have looked much like that of a non-reform state's, had it not adopted reforms. In the specific example of Wisconsin, political developments since 2023 further justify this conclusion. A narrow state supreme court victory by a justice who campaigned, in part, in opposition to gerrymandering, led to a successful court challenge to state legislative maps that resulted in a significant decrease in the partisan bias of those maps (on common bias metrics). Challenges to the congressional maps are currently pending.

Adjusting for Political Geography

To further increase the credibility of the conditional parallel trends assumption above, we adjust for the changes in political geography between the two redistricting cycles. Partisan preferences and the geographic distribution of voters can change in different ways within each state over time. For example, Michigan, which experienced reform, became more Republican during our sample period, while Georgia, which did not, became more Democratic. Without this adjustment, an estimated effect of the Michigan reform would be biased by the differential change in the states' political geographies.

To make the adjustment, we use representative sets of nonpartisan redistricting plans within each state that respect each state's specific redistricting rules. These simulation samples were separately generated for 2020 (McCartan et al. 2022) and 2010 (Kenny et al. 2024) using the algorithm of McCartan and Imai (2023). A detailed discussion of these simulations and their limitations may be found in Section S3 of the Supplementary Material. We subtract the mean outcome in each state's simulated sample, denoted by $\tilde{Y}_{it}$, from the observed outcome Y_{it} before estimating the causal effects under the DiD design. We argue that this subtraction of the outcomes based on simulated baseline plans from the observed outcomes accounts for the change in political geography in each state.

One limitation of basing simulations off of each state's rules is that any changes in outcome variables from 2010 to 2020 that are due to changes in rules would be subtracted away. Since redistricting rules are sometimes changed as part of larger reform efforts, part of the overall effect of redistricting reform would not be included in the measured effect, attenuating the overall estimated effects. However, these rules changes occurred in only a few states, and their effect on partisan outcomes was measured by Kenny et al. (2023) and found to be minimal.

Thus, we assume that once we adjust for the state-specific change in political geography, states with different changes in institutional features have parallel trends in the potential outcomes. Formally, if we denote the difference between the potential outcome and the simulated outcome as $\Delta Y_{it}(d) = Y_{it}(d) - \tilde{Y}_{it}$, our conditional parallel trends assumption becomes

$$\begin{aligned}&\mathbb{E}[\Delta Y_{i1}(d') - \Delta Y_{i0}(d) | \mathbf{X}_i = \mathbf{x}] \\ &\quad = \mathbb{E}[\Delta Y_{i1}(d') - \Delta Y_{i0}(d) | \mathbf{X}_i = \mathbf{x}, D_{i0} = d, D_{i1} = d']\end{aligned}$$

for all d, d'. Then, as above, the CATE is identified as

$$\begin{aligned}\mathrm{CATE}_{\mathbf{x}}(d, d') = {} & \mathbb{E}[\Delta Y_{i1} - \Delta Y_{i0} | \mathbf{X}_i = \mathbf{x}, D_{i0} = d, D_{i1} = d'] \\ & - \mathbb{E}[\Delta Y_{i1} - \Delta Y_{i0} | \mathbf{X}_i = \mathbf{x}, D_{i0} = d, D_{i1} = d],\end{aligned}$$

where $\Delta Y_{it} = Y_{it} - \tilde{Y}_{it}$.

Estimation of Causal Effects

We estimate this causal estimand with a Bayesian linear regression model, where the response is the change in each simulation-adjusted outcome between 2010 and 2020, $\Delta Y_{i1} - \Delta Y_{i0}$ (Section S7 of the Supplementary Material for the descriptive analysis of raw changes). The predictors are the change in treatment level $D_{i1} - D_{i0}$, the baseline treatment level D_{i0}, the covariates $\mathbf{X}_i$, and the interaction of the treatment change with baseline treatment and with each of the covariates. This is a regression-adjusted DiD estimator (Heckman, Ichimura, and Todd 1997). In contrast to the common two-way fixed effects estimator, this approach allows for heterogeneous treatment effects. The difference from most regression-adjusted estimators is the continuous treatment, which necessitates a modeling choice on the form of the dose–response curve (Callaway, Goodman-Bacon, and Sant'Anna 2024).

Given the small sample size ($n = 87$), the moderate number of covariates ($p = 15$ with interactions for the modal specification), and the high noise level, we believe that the linear specification and Bayesian estimation are appropriate. The coefficient priors help avoid overfitting to a few samples, and uncertainty in the ATE is automatically quantified; details on these priors are included in Section S5 of the Supplementary Material. A more flexible regression model beyond linear would be unlikely to increase predictive power, given the small sample size, and would further risk overfitting. However, as a robustness check, we also fit a nonparametric Bayesian Additive Regression Trees (BART) model (Chipman, George, and McCulloch 2010) and include those results in Section S6.1 of the Supplementary Material. The results are qualitatively the same, but the estimated effect magnitudes are attenuated.

ESTIMATED CAUSAL EFFECTS OF REDISTRICTING REFORMS

The outcome of the redistricting process is a complete congressional districting plan. A plan can be evaluated in a number of ways—how many seats each party is expected to win, how many of the seats are competitive, how well the partisan composition of the delegation matches the voters' preferences, and so on. We begin this section by introducing quantitative measures of several partisan and nonpartisan aspects of districting plans. These constitute the outcome variables for the causal estimates we present in the second half of the section. Section S7 of the Supplementary Material presents the results of descriptive analyses that are largely consistent with the main results shown below.

Outcome Measures

Our measures can be divided into two buckets: nonpartisan outcomes that quantify how much gerrymandering is present or how responsive a districting plan is to shifts in public opinion, and partisan outcomes that capture which party is advantaged by a districting plan. For the nonpartisan outcomes, we also use the absolute value of the realized leeway as an additional treatment variable so that neither treatment nor outcome considers the partisan direction.

All of these measures are calculated from district-level election results. Since there is a significant exogenous variation in election results due to swings in the national political environment, rather than using actual 2012 and 2022 House election results to evaluate districting plans, we use a statistical model to estimate the distribution of election results across future hypothetical elections. A statistical election model also allows us to apply these same measures to the simulated districting plans, under which no elections have taken place.

We adopt the model from Kenny et al. (2023), which assumes district election outcomes can be decomposed into a baseline district-level vote share plus district-specific and national swings.[12] The model is closely

[12] The full details of this approach can be found in the Methods section and SI Appendix, Section B of Kenny et al. (2023).

related to the stochastic uniform partisan swing model of Gelman and King (1994) and the congressional model of Ebanks, Katz, and King (2023). As the baseline district-level vote, we use the 2008 presidential results for the 2010 redistricting cycle and an average of the 2016 and 2020 presidential results for the 2020 cycle, each logit-shifted so that the national partisan vote is exactly 50/50 (Voting and Election Science Team 2018; 2020). As in Kenny et al. (2023), we use historical congressional election data since 1976 to estimate the variance of district and national swings. We calculate exact expectations of all of our outcome measures against the predictive distribution of the model via numerical integration.

Our primary partisan measure of districting plans is the expected number of seats won by the Republican party. As discussed in the estimation section, we subtract the average outcomes in the simulated baseline redistricting plans (Kenny et al. 2024; McCartan et al. 2022) from the observed outcome in the enacted plan to adjust for the state-specific change in political geography, which is a potential confounding factor (Cottrell 2019). After this adjustment, the seat outcome ranges from −1.57 in Illinois in 2020 (favoring Democrats) to 1.78 in Texas in 2020 (favoring Republicans). These simulation-differenced seats can be directly interpreted as a measurement of bias due to partisan gerrymandering, with positive values indicating a Republican bias beyond what would be expected based on the state's political geography alone.

The expected number of Republican seats is an interpretable measure, but may not be completely comparable across states. Depending on political geography and especially the total number of districts in each state, the natural variation in the number of seats outcome may vary significantly between states. To address this, we also include as an outcome measure the simulation z-score of Republican seats, which is calculated by taking the simulation-differenced seats outcome and dividing it by the standard deviation of Republican seats in the simulation set (Kenny et al. 2023). This puts all the states' outcomes on a common scale, increasing the plausibility of the parallel trends assumption and the homoskedasticity assumption of the estimation model.

For a nonpartisan outcome measure, we take the absolute value of these two partisan measures. Both the absolute difference and the absolute z-score of Republican seats measure how far the enacted plan deviate from the nonpartisan simulation baseline in terms of partisan composition.[13]

Finally, we also measure the responsiveness of districting plans to changes in the national electoral environment. Responsiveness is measured as the rate of change in the share of Republican seats given an infinitesimal change in Republican vote share nationwide. Responsiveness is closely linked to the presence of competitive seats; the more competitive seats there are, the larger the change in seat share will be for a given shift in vote share. Indeed, electoral responsiveness is the primary motivation for the creation of competitive congressional districts.

The responsiveness measure ranges from 0.044 in Wyoming in 2020 to 7.91 in New Hampshire in 2020, with most plans' values lying between 0.5 and 3. The interpretation of these values is as follows; in New Hampshire, a 1 pp increase in a party's vote share leads to a 7.91 pp change in the party's expected seat share. This relatively large increase makes sense in context—both of the state's congressional districts have Republican vote share within a few points of 50%. As with the other measures, we subtract the mean responsiveness of the simulated plans from the enacted plan's responsiveness when estimating causal effects.

In Section S4 of the Supplementary Material, we extend the analysis to a series of alternative measures of partisan bias, such as the efficiency gap (Stephanopoulos and McGhee 2015), and find qualitatively similar results.

Empirical Findings

We first study the effects of changes in maximum leeway on Republican seats. Recall that maximum leeway represents the worst expected bias under a single-party control. The left panel of Figure 6 shows the estimated coefficients of the estimation model fit to the simulation-adjusted seats variable. All estimated coefficients for this model and all other models in the article are contained in Section S10 of the Supplementary Material. The coefficients in Figure 6 are shown on the scale of the outcome, where a positive coefficient indicates a relationship with a positive change in Republican seats from 2010 to 2020. The model R^2 is around 0.4, highlighting the potential for the control variables to confound any observed correlation between reform and the outcome measure.

Given the fitted outcome model, we can estimate the CATE for any given starting and final value of maximum leeway and any combination of covariates. The right panel of Figure 6 plots these CATEs for each state that experienced a change in maximum leeway between 2010 and 2020, using the specific covariates of that state. There is a clear dose–response relationship between the decrease in maximum leeway and the decrease in the expected number of Republican seats.[14] A state with Michigan's dose (nearly 4) has an estimated CATE of around −1.25 , with 95% credible interval (−2.4, 0.5), and so would be expected to gain just over one Democratic seat, on average. Similarly, a state with Ohio's dose (just over 1) and covariates would be expected to gain about half of a Democratic seat (−0.7, −0.2). Motivated by this pattern, we also include in the right figure a dose–response curve

[13] For the transformed outcomes based on z-score and absolute value, we do not interpret the estimated causal effects in terms of the original outcome variable (i.e., the expected number of Republican seats). Instead, the estimates should be interpreted as the effects on the transformed outcomes. This also means that the parallel trends assumption is assumed to hold on the transformed outcome variable.

[14] This pattern was also observed in robustness checks that used a flexible BART model for estimation.

FIGURE 6. Sample Model Fit and Estimated Conditional Average Treatment Effects

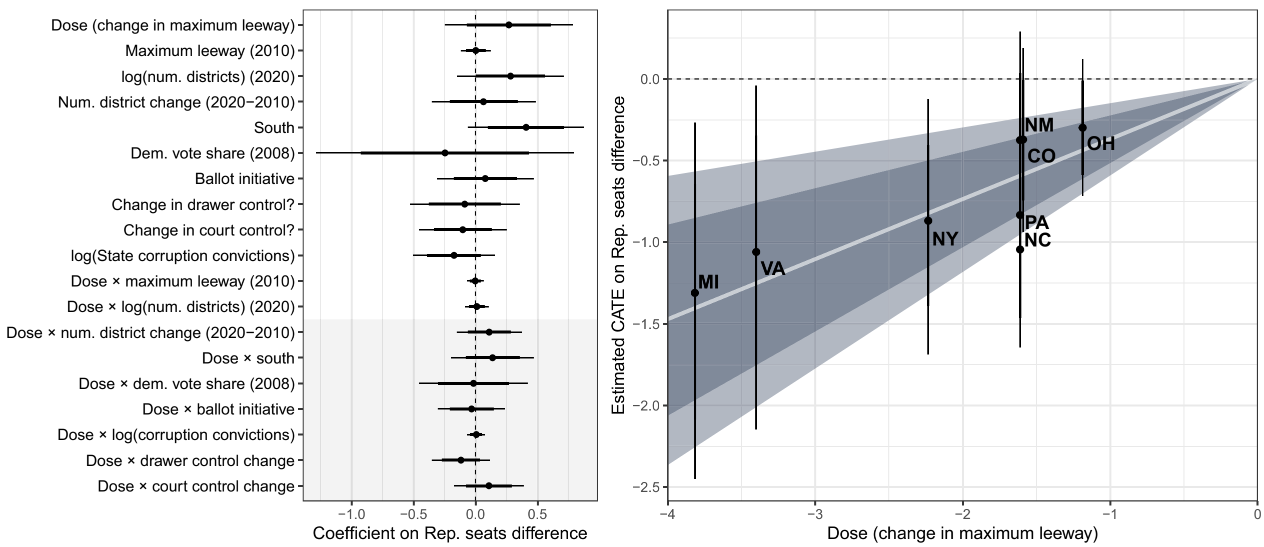

Note: Fitted model coefficient estimates for the Republican seat outcome measure (left) using the maximum leeway treatment, and estimated conditional average treatment effects for each reformed state's covariate combination plotted against the state's dose (right). The model-based dose–response curve is underlaid in blue. 80% and 95% credible intervals are shown throughout.

calculated by calculating a CATE for dose level and averaging over the observed covariate distribution.[15]

Since the dose–response pattern is well approximated by a linear relationship, we can summarize it by its slope, a quantity also known as the average causal response (ACR). The product of the ACR and a given dose corresponds to the estimated causal effect of that dose, averaged across states.

Figure 7 presents the ACR for Republican seats and the other outcome variables, using both the maximum and realized leeway treatment values. As mentioned above, for the nonpartisan outcomes, we also use the absolute value of the realized leeway as a treatment variable. The right side of the figure plots the ACR estimates in terms of the standard deviation of the respective outcome variable per unit change in treatment. This means that a point estimate of 0.5 is interpreted as an increase of 0.5 standard deviations of the outcome variable for a 1-unit change, or an increase of two standard deviations for a 4-unit change.

Across the three nonpartisan outcomes measuring the amount of gerrymandering, we find a consistent effect of leeway. The estimated ACRs are all positive, meaning that an increase in leeway leads to a greater gap between the enacted plan and the simulated baseline. That is, states that reduce leeway through reform efforts are expected to rein in partisan gerrymandering. The magnitude of the effect is substantial: a reform effort like Michigan's, with a decrease in the maximum leeway of four units, would be expected to reduce the difference in Republican seats from the baseline by about 0.6 seats (0.06, 1.20), or 2.6 simulation standard deviations (–1, 6)—enough to move an extreme partisan gerrymander to a fair plan. The ACRs for the absolute realized leeway are very similar to those for the maximum leeway; the ACRs for the fully signed realized leeway are positive but smaller in magnitude and have enough uncertainty that the 80% credible intervals cross zero. Thus, there is some weak evidence that giving Democrats more control of the redistricting process by increasing their leeway or reducing Republican leeway would also reduce gerrymandering.

We also find that reforms increase electoral responsiveness. The point estimates of the ACRs for both the maximum leeway and the realized leeway agree; the posterior probability that the effect is negative is more than 99% for both measures. The estimated effect for the absolute realized leeway is near zero. We estimate that a reduction in the maximum or realized leeway of four units would increase responsiveness by 1.5 (0.28, 2.67), meaning that a 1 pp increase in a party's vote share would translate to an additional 0.6 pp in seat share on top of the average 2 pp increase in seat share across all states.

Another way to interpret the effect on responsiveness is in terms of the share of competitive seats, where a seat that is counted as competitive in proportion to how close its baseline partisanship is to 50%. This is equivalent to a linear rescaling of the responsiveness outcome.[16] In terms of competitive seats, we estimate

[15] Because of the interaction between the dose variable and the covariates (which do not have mean zero), the overall slope estimate and its variance differ from the estimate and variance of the coefficient on Dose in the linear model.

[16] Specifically, we measure the competitiveness of a single seat as the derivative of the election model's win probability for that seat at its

FIGURE 7. Average Causal Response (ACR) of Leeway on Redistricting Outcomes

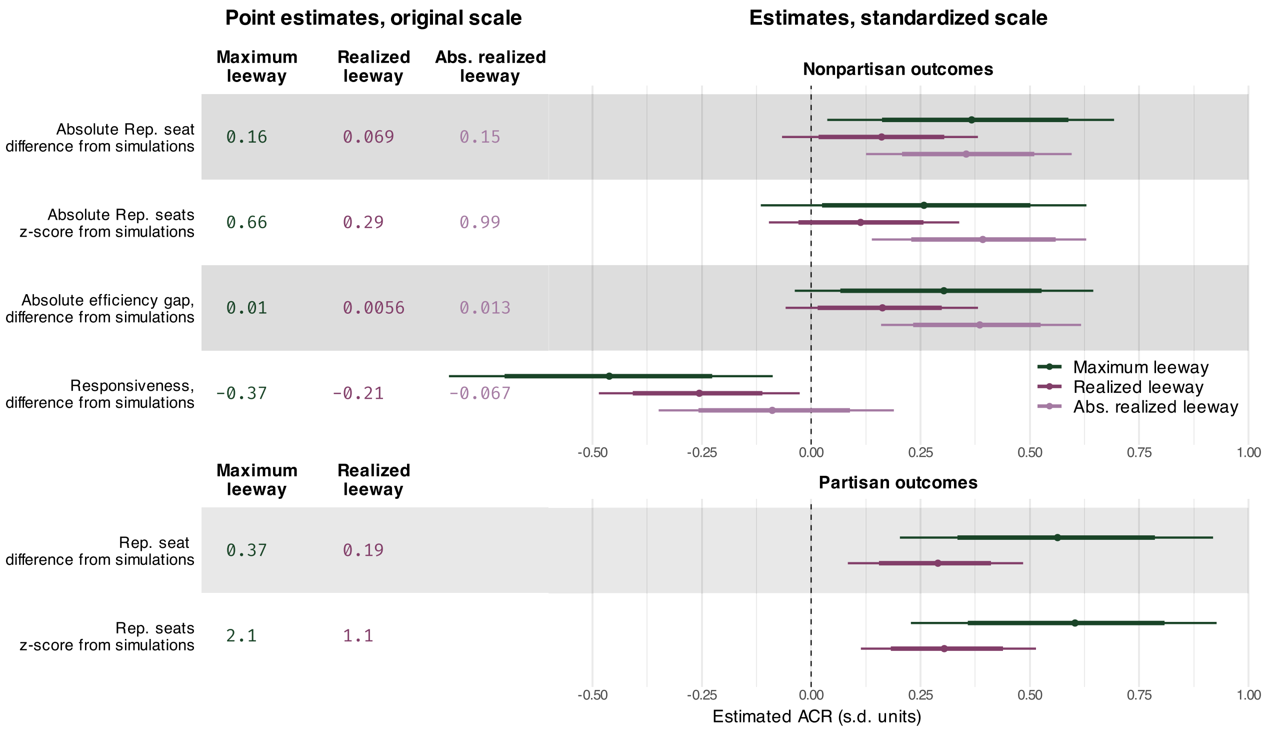

Note: The points correspond to the mean estimated ACR, while the lines represent 80% and 95% credible intervals. Intervals are colored by the treatment variable used. The numbers in the columns display the mean ACR on each outcome's response scale. The estimates and intervals on the right are displayed in units of outcome standard deviations, to allow for comparability between outcomes. For partisan outcomes, a positive number indicates a pro-Republican effect and a negative number indicates a pro-Democratic effect for a positive dose.

that a reduction in leeway of four units would increase the share of competitive seats in an average state from 25% to 43% (28%, 58%).

In terms of partisan outcomes, the estimated ACRs on Republican seats are all positive, for both realized and maximum leeway. This is in line with an interpretation that *increasing* leeway benefits the party in control, as one might expect. Recall that the doses from Figure 4, for the realized leeway estimates, a dose of 4 represents approximately going from a Democratic legislature to an independent commission or an independent commission to a Republican legislature. For such a dose of 4, we would expect Republicans to gain around 0.75 seats (0.22, 1.27), or 8.6 simulation standard deviations (3.3, 13.5)—again, a substantial effect.

Placebo Check

As a further validation of our overall approach to estimating causal effects, we conduct a placebo check by applying the same estimation procedure to a political outcome variable that should not be affected by redistricting reform but is closely related to the outcome variable of interest. For this placebo outcome variable, we use the Democratic two-party vote share in the first presidential election following each census (2012 and 2024). As national contests, presidential elections are insulated from state-specific reform efforts. Moreover, issues like the economy, immigration, and health care dominate presidential campaigns, not state-specific governance issues such as redistricting reform. Thus, we expect there to be no effect of redistricting reforms on presidential vote share.

Figure 8 displays the estimated ACRs for the placebo outcome using each of the three treatment variables. All three estimates are nearly zero, with credible intervals that easily cover zero. If the estimates were significantly different from zero, we would suspect the validity of the sufficiency assumption, the parallel trends assumption, or the estimation method. That all placebo estimates are close to zero is therefore strong evidence of the plausibility of these assumptions and our estimation strategy overall.

REDISTRICTING COMMISSIONS: COUNTERFACTUAL POLICY ANALYSIS

What if redistricting reforms were adopted nationwide? Redistricting commissions are the most commonly

baseline vote, divided by the maximum derivative, which is obtained at a 50/50 baseline vote. Based on our election model, a district with 50/50 baseline vote is counted as 1.0 competitive seats, and a district with 60/40 baseline vote is counted as just 0.14 competitive seats. This approach avoids arbitrary cutoffs in what counts as a competitive district.

FIGURE 8. Average Causal Response (ACR) of Leeway on Placebo Outcome

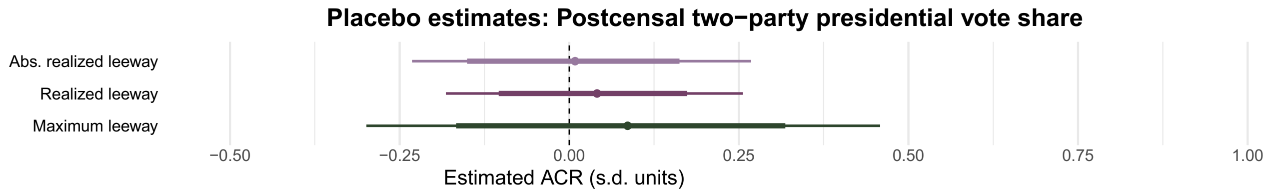

Note: The points correspond to the mean estimated ACR, while the lines represent 80% and 95% credible intervals. Positive numbers correspond to an increase in Democratic vote share.

adopted map-drawing reform. A total of 15 states, including Michigan, New Jersey, and Arizona, currently use some form of commission. Advocates often argue that commissions can lead to fairer redistricting plans by further removing legislators from the process (e.g., American Academy of Arts and Sciences 2020).[17]

In this section, we conduct a counterfactual policy analysis of the nationwide adoption of redistricting commissions. Though hypothetical, counterfactual policy analyses are useful tools for evaluating institutional reform proposals (Cervas and Grofman 2019). Our approach is well suited for this analysis, in part because commissions can vary drastically in the extent to which they limit the influence of partisan actors. We caution, however, that the results of our counterfactual policy analysis may not be applicable, for example, if political geography and electoral environments change in the future.

Analysis Procedure

To conduct the analysis, we apply fitted causal regression models based on the party-aware treatment to predict electoral outcomes under a series of hypothetical scenarios, in which every state adopts a redistricting commission with particular institutional powers, but partisan control of state institutions remains unchanged. Although each of the 15 states uses a redistricting commission with different structures and rules, our model is able to characterize these different commissions in terms of their partisan leeway and use these treatment values to predict electoral outcomes under a given commission structure.

We caution that any counterfactual policy analysis of this type involves substantial extrapolation from the observed data, and its results should therefore be interpreted as exploratory rather than confirmatory. While the regression models do allow for effect heterogeneity, the critical sufficiency assumption rules out unmodeled heterogeneity that may in practice lead to different reform effects.

We study three types of commission structures currently enacted in several states: (1) a New York-style commission with a nonpartisan map drawer and several partisan veto points; (2) an Ohio-style reform with legislature-drawn map and several partisan and bipartisan veto points and stalemate procedures, including a bipartisan commission stalemate procedure; (3) and a Michigan-style reform, with a nonpartisan commission, no partisan veto points, and the potential for court review.[18] Though each of these states has a commission of some kind, they differ in important ways—for example, Michigan's elimination of partisan veto points removes the ability of partisan actors to veto unfavorable plans.

We visualize the impact of each reform structure using a seats–votes curve. Seats–votes curves are often used to study a measure of partisan bias, operationalized as the difference in the share of seats and votes at a given point on the curve (Katz, King, and Rosenblatt 2020; King and Browning 1987; Tufte 1973). A seats–votes curve for which a 50% vote share translates exactly to a 50% seat share represents the baseline for partisan bias. We approximate the seats–votes curve around 50% vote share as a line and estimate the effect of reforms on its slope and intercept separately. Up to a linear transformation, the slope corresponds to responsiveness and the intercept to seat share. These two outcome variables are differenced using the simulation sets, and a regression model is fit to these differenced outcomes, as discussed in the estimation section.

Findings

We find that nationwide adoption of commissions would generally reduce partisan bias, and that reforms that place additional restrictions on partisan actors (as in Michigan) are generally more effective. Figure 9 shows a seats–votes curve for each counterfactual commission scenario (the orange lines), which display the predicted number of total Democratic seats in the U.S. House (y-axis) for a given national Democratic vote share (x-axis). Figure S9.19 in the Supplementary Material presents the estimated state-level effects, which show that the reform effects in states with opposite partisan control (e.g.,

[17] See also "People Not Politicians Oregon" or "Ohio's Citizens Not Politicians."

[18] Although states could change their redistricting processes in the future, we use shorthands like "Michigan-style" to refer to the current (as of 2024) institutional design of each state's commission structure.

FIGURE 9. Predicted Seats–Votes Curves under Hypothetical Reforms

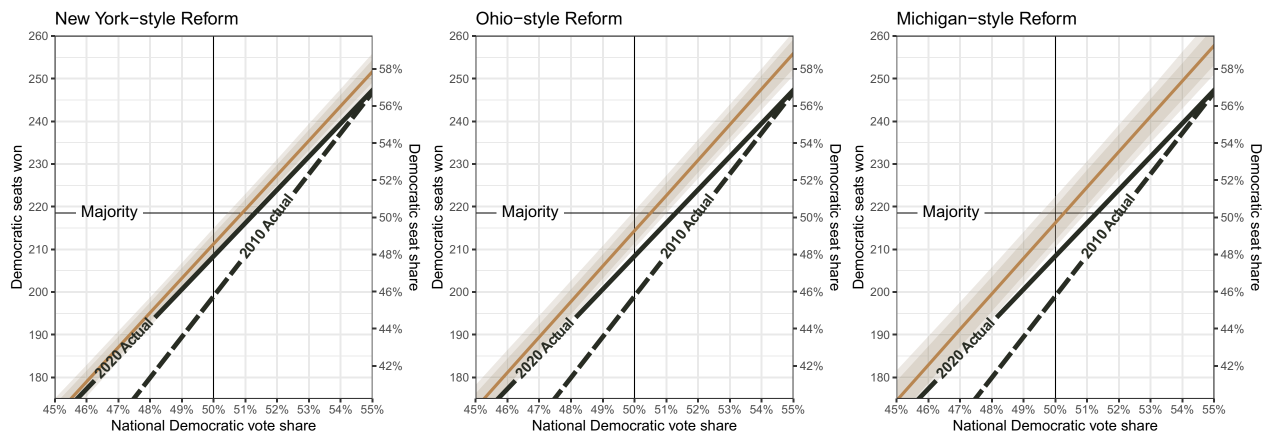

Note: The figure shows three predicted seats–votes curves if all U.S. states adopted new redistricting institutions with: (1) a New York-style commission with a nonpartisan map drawer and several partisan veto points; (2) an Ohio-style legislature-drawn map and several partisan and bipartisan veto points; and (3) a Michigan-style reform, with a nonpartisan commission, no partisan veto points, and the potential for court review. Hypothetical commission structures are plotted as orange lines (with 80% and 95% credible intervals), with reference lines for actual plans for both 2020 and 2010 in black.

Texas vs. Illinois) may, in part, cancel out. The model predicts that all commission structures would reduce the existing Republican advantage in the House driven by both the inefficient geographic distribution of Democratic voters (Chen and Cottrell 2016; Chen and Rodden 2015) and net partisan manipulation favoring Republicans (Kenny et al. 2023).

We also compare these estimated relationships to the actual seats–votes curves for 2010 and 2020 (the black lines). We note that the lines for 2010 and 2020 intersect the 50% vote line below the 50% seat line, since a combination of the underlying geographic distribution of Democratic votes and gerrymandering efforts nationwide disadvantages the Democratic party. These results, however, do not necessarily imply that commissions would always advantage Democrats. For example, if political geography and electoral environments change in the future, the establishment of commissions may have different partisan impacts.

We find that New York- and Ohio-style reforms produce relatively moderate improvements in partisan bias and responsiveness. The point estimates for the number of Democratic seats gained under these reforms is 3.0 for New York-style reforms, with 95% credible interval (−0.17, 6.60), and 6.0 for Ohio-style reforms (1.8, 10.7). Responsiveness for New York- and Ohio-style reforms is also similar as for every 1 pp increase in vote share, Democrats gain an estimated 8.1 seats with New York-style reforms, and 8.3 seats with Ohio-style reforms. These responsiveness values are only slightly higher than the actual responsiveness following 2020 redistricting. Ohio-style reforms appear more effective overall, though the substantive differences are small; we are more than 99% confident that Ohio-style reforms produce a larger effect on partisan bias than New York-style reforms, and 89% confident that Ohio-style reforms increase responsiveness more than New York-style reforms do.

Further, we find that Michigan-style reforms have the greatest effects in both responsiveness and Democratic seats, since partisan actors are the most constrained by the presence of a nonpartisan commission, no partisan veto points, and the potential for court review. We estimate that the number of Democratic seats gained under Michigan-style reforms is 7.8 with 95% credible interval (2.1, 13.9), and that for every 1 pp increase in vote share, Democrats gain an estimated 8.3 seats (7.7, 8.9), an increase of 0.54 over the baseline. We are 91% confident that Michigan-style reforms have a larger effect than Ohio reforms, and more than 99% confident in a larger effect than New York-style reforms. We are also 94% confident that Michigan-style reforms increase responsiveness.

All three proposed reforms reduce the deviation from partisan symmetry. In this case, Democrats gain net seats which counteract the existing Republican advantage. However, greater effects are produced when, as in Michigan-style reforms, multiple nodes of the redistricting game constrain leeway by ensuring partisan nodes are nonpartisan, and do not precede partisan vetoes. Meanwhile, examples of New York-style and Ohio-style reforms demonstrate how constraining partisan actors at different nodes, through different reforms, may produce substantively similar effects.

CONCLUSION

In this article, we propose a methodology for estimating causal effects of complex institutional reforms and apply it to study the impacts of redistricting reforms in the United States. Although redistricting reforms differ across states in important procedural details, we show how to obtain a theoretically informed univariate summary measure of such reforms through a game-theoretic model. We then combine the results of this formal model with a standard causal inference research design strategy to obtain credible causal effect estimates even with a limited sample size. We find that redistricting reforms are likely to reduce the partisan bias of enacted plans by constraining the leeway of partisan actors.

Our methodological approach enables us to go beyond the estimation of causal effects. Specifically, we conduct a counterfactual policy analysis to infer the consequences of adopting a series of different redistricting reforms nationwide. We find that adopting redistricting commissions generally reduces the current Republican advantage, and this reduction is substantially larger when the reforms place greater restrictions on partisan actors. For example, a Michigan-style reform (which combines a nonpartisan commission with no partisan veto points) would yield a greater pro-Democratic effect than commissions adopted in Ohio and New York (which maintain the potential for some partisan control). While we apply this approach to three commission reforms, the same model can be applied to other reforms of interest. Future research should also apply our methodology to other redistricting problems, such as the one for state legislatures.

Our analysis consistently shows that reducing the leeway of partisan actors leads to fairer redistricting plans across a variety of measures. For measures with partisan signs, the reduction in bias favors Democrats. This is in part because the redistricting plans passed in 2010 and 2020 each favored Republicans, due to both geography and gerrymandering efforts (Kenny et al. 2023). Thus, the findings of this article should not be interpreted as evidence that reducing leeway inherently advantages Democrats. Rather, we show that reducing leeway constrains the ability of the party in power to gerrymander in its favor.

Like most studies of state-level institutions in the United States, our findings are based on a dataset with a small sample size. To address this fundamental limitation, we have developed a methodological approach that summarizes high-dimensional institutional features using a univariate statistic based on a game-theoretic model. This measurement strategy, along with our identification strategy and a variety of empirical validation and placebo checks, enable us to estimate credible causal effects despite our limited sample size. Nevertheless, collecting and analyzing more data, either going further back to 2000 and 1990 or updating our analysis to include future redistricting reform efforts, could further test the external validity of our theoretical arguments.

The key idea behind our methodology is also applicable to studies of other complex institutions. For example, most scholars analyze democratic institutions by using the Polity score, which is a 21-point scale index based on six "component variables" concerning executive recruitment, executive constraints, and political participation. Instead of summing these scores to obtain the overall Polity score, our proposed methodology suggests that researchers consider developing a more theoretically informed measure of democracy, based directly on their substantive questions of interest. We believe that such an approach preserves critical aspects of institutional complexity while improving the credibility of causal analysis.

SUPPLEMENTARY MATERIAL

The supplementary material for this article can be found at https://doi.org/10.1017/S0003055426101610.

DATA AVAILABILITY STATEMENT

Research documentation and data that support the findings of this study are openly available at the American Political Science Review Dataverse: https://doi.org/10.7910/DVN/6HGLTN.

ACKNOWLEDGEMENTS

We acknowledge helpful comments from Christian Fong, Shiro Kuriwaki, Jacob Montgomery, Daniel Thompson, Seth McKee, Peter Buisseret, and an anonymous reviewer from the Alexander and Diviya Magaro Peer Pre-Review Program at Harvard's Institute for Quantitative Social Science.

CONFLICT OF INTEREST

The authors declare no ethical issues or conflicts of interest in this research.

ETHICAL STANDARDS

The authors affirm this research did not involve human participants.

REFERENCES

American Academy of Arts and Sciences. 2020. "Our Common Purpose: Reinventing American Democracy for the 21st Century." Final Report and Recommendations from the Commission on the Practice of Democratic Citizenship.

Best, Robin E., Steve B. Lem, Daniel B. Magleby, and Michael D. McDonald. 2021. "Do Redistricting Commissions Avoid Partisan Gerrymanders?" *American Politics Research* 50 (3): 379–95.

Cain, Bruce E. 2012. "Redistricting Commissions: A Better Political Buffer." *Yale Law Journal* 121: 1808–44.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2024. "Difference-in-Differences with a Continuous Treatment." NBER Working Paper. https://doi.org/10.3386/w32117.

Canen, Nathan, and Kristopher Ramsay. 2024. "Quantifying Theory in Politics: Identification, Interpretation, and the Role of Structural Methods." *Journal of Theoretical Politics* 36 (4): 301–27.

Canon, David T. 1999. *Race, Redistricting, and Representation: The Unintended Consequences of Black Majority Districts*. Chicago, IL: University of Chicago Press.

Canon, David T. 2022. "Race and Redistricting." *Annual Review of Political Science* 25: 509–28.

Carson, Jamie L., and Michael H. Crespin. 2004. "The Effect of State Redistricting Methods on Electoral Competition in United States House of Representatives Races." *State Politics & Policy Quarterly* 4 (4): 455–69.

Carson, Jamie L., Michael H. Crespin, and Ryan D. Williamson. 2014. "Reevaluating the Effects of Redistricting on Electoral Competition, 1972–2012." *State Politics & Policy Quarterly* 14 (2): 165–77.

Caughey, Devin, and Christopher Warshaw. 2022. *Dynamic Democracy: Public Opinion, Elections, and Policymaking in the American States*. Chicago, IL: University of Chicago Press.

Caughey, Devin, Yiqing Xu, and Christopher Warshaw. 2017. "Incremental Democracy: The Policy Effects of Partisan Control of State Government." *The Journal of Politics* 79 (4): 1342–58.

Cervas, Jonathan R., and Bernard Grofman. 2019. "Are Presidential Inversions Inevitable? Comparing Eight Counterfactual Rules for Electing the US President." *Social Science Quarterly* 100 (4): 1322–42.

Chen, Jowei, and David Cottrell. 2016. "Evaluating Partisan Gains from Congressional Gerrymandering: Using Computer Simulations to Estimate the Effect of Gerrymandering in the U.S. House." *Electoral Studies* 44: 329–40.

Chen, Jowei, and Jonathan Rodden. 2015. "Cutting through the Thicket: Redistricting Simulations and the Detection of Partisan Gerrymanders." *Election Law Journal: Rules, Politics, and Policy* 14 (4): 331–45.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods." *Annual Review of Economics* 1: 451–88.

Chipman, Hugh A., Edward I. George, and Robert E. McCulloch. 2010. "BART: Bayesian Additive Regression Trees." *The Annals of Applied Statistics* 4 (1): 266–98.

Cottrell, David. 2019. "Using Computer Simulations to Measure the Effect of Gerrymandering on Electoral Competition in the US Congress." *Legislative Studies Quarterly* 44 (3): 487–514.

Cottrill, James B. 2012. "The Effects of Non-Legislative Approaches to Redistricting on Competition in Congressional Elections." *Polity* 44 (1): 32–50.

Ebanks, Danny, Jonathan N. Katz, and Gary King. 2023. "How American Politics Ensures Electoral Accountability in Congress." Preprint. https://gking.harvard.edu/files/gking/files/acc.pdf.

Edwards, Barry, Michael Crespin, Ryan D. Williamson, and Maxwell Palmer. 2017. "Institutional Control of Redistricting and the Geography of Representation." *The Journal of Politics* 79 (2): 722–6.

Fraga, Bernard L., Daniel J. Moskowitz, and Benjamin Schneer. 2022. "Partisan Alignment Increases Voter Turnout: Evidence from Redistricting." *Political Behavior* 44 (4): 1883–910.

Gelman, Andrew, and Gary King. 1994. "A Unified Method of Evaluating Electoral Systems and Redistricting Plans." *American Journal of Political Science* 38 (2): 514–54.

Grofman, Bernard, and Lisa Handley. 1991. "Identifying and Remedying Racial Gerrymandering." *Journal of Law & Politics* 8(2): 345–404.

Heckman, James J., Hidehiko Ichimura, and Petra Todd. 1997. "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme." *The Review of Economic Studies* 64 (4): 605–54.

Henderson, John A., Brian T. Hamel, and Aaron M. Goldzimer. 2018. "Gerrymandering Incumbency: Does Nonpartisan Redistricting Increase Electoral Competition?" *The Journal of Politics* 80 (3): 1011–6.

Katz, Jonathan N., Gary King, and Elizabeth Rosenblatt. 2020. "Theoretical Foundations and Empirical Evaluations of Partisan Fairness in District-Based Democracies." *American Political Science Review* 114 (1): 164–78.

Keena, Alex, Michael Latner, Anthony J. McGann, and Charles Anthony Smith. 2021. *Gerrymandering the States: Partisanship, Race, and the Transformation of American Federalism*. Cambridge: Cambridge University Press.

Kenny, Christopher T., Cory McCartan, Tyler Simko, Shiro Kuriwaki, and Kosuke Imai. 2023. "Widespread Partisan Gerrymandering Mostly Cancels Nationally, but Reduces Electoral Competition." *Proceedings of the National Academy of Sciences* 120 (25): e2217322120.

Kenny, Christopher T., Tyler Simko, Cory McCartan, Emma Ebowe, Philip W. O'Sullivan, Aneetej Arora, Taran Samarth, et al. 2024. "Simulated Redistricting Plans for the 2010 Redistricting Cycle." https://doi.org/10.7910/DVN/SLCD3E.

King, Gary, and Robert X. Browning. 1987. "Democratic Representation and Partisan Bias in Congressional Elections." *American Political Science Review* 81 (4): 1251–73.

Komisarchik, Mayya, and Ariel White. 2025. "Throwing Away the Umbrella: Minority Voting after the Supreme Court's Shelby Decision." *Quarterly Journal of Political Science* 20 (2): 269–305.

Kousser, Thad, Justin Phillips, and Boris Shor. 2018. "Reform and Representation: A New Method Applied to Recent Electoral Changes." *Political Science Research and Methods* 6 (4): 809–27.

Levitsky, Steven, and Daniel Ziblatt. 2023. *Tyranny of the Minority: Why American Democracy Reached the Breaking Point*. New York: Crown.

Lindgren, Eric, and Priscilla Southwell. 2013. "The Effect of Redistricting Commissions on Electoral Competitiveness in US House Elections, 2002-2010." *Journal of Politics and Law* 6(2): 13–8.

Little, Andrew T., and Anne Meng. 2024. "Measuring Democratic Backsliding." *PS: Political Science & Politics* 57(2): 149–61.

Litton, Noah. 2012. "The Road to Better Redistricting: Empirical Analysis and State-Based Reforms to Counter Partisan Gerrymandering." *Ohio State Law Journal* 73(4): 839–81.

Lundberg, Ian, Rebecca Johnson, and Brandon M. Stewart. 2021. "What Is Your Estimand? Defining the Target Quantity Connects Statistical Evidence to Theory." *American Sociological Review* 86 (3): 532–65.

Madison, James. 1788a, "The Apportionment of Members among States." The New York Packet.

Madison, James. 1788b. "The Same Subject Continued: The House of Representatives." The New York Packet.

McCartan, Cory, and Kosuke Imai. 2023. "Sequential Monte Carlo for Sampling Balanced and Compact Redistricting Plans." *Annals of Applied Statistics* 17 (4): 3300–23.

McCartan, Cory, Christopher T. Kenny, Tyler Simko, George Garcia III, Kevin Wang, Melissa Wu, Shiro Kuriwaki, et al. 2022. "Simulated Redistricting Plans for the Analysis and Evaluation of Redistricting in the United States." *Scientific Data* 9 (1). https://doi.org/10.1038/s41597-022-01808-2.

McCartan, Cory, Christopher T Kenny, Tyler Simko, Emma Ebowe, Michael Y Zhao and Kosuke Imai. 2026. "Replication Data for: Redistricting Reforms Reduce Gerrymandering by Constraining Partisan Actors." Harvard Dataverse. Dataset. https://doi.org/10.7910/DVN/6HGLTN.

McCarty, Nolan. 2019. *Polarization: What Everyone Needs to Know*. Oxford: Oxford University Press.

McDonald, Michael P. 2004. "A Comparative Analysis of Redistricting Institutions in the United States, 2001-02." *State Politics & Policy Quarterly* 4 (4): 371–95.

Miller, Peter, and Bernard Grofman. 2013. "Redistricting Commissions in the Western United States." *UC Irvine Law Review* 3(3): 637–68.

Nelson, Matthew. 2023. "Independent Redistricting Commissions Are Associated with more Competitive Elections." *PS: Political Science & Politics* 56 (2): 207–12.

Polsby, Daniel D., and Robert D. Popper. 1991. "The Third Criterion: Compactness as a Procedural Safeguard against Partisan Gerrymandering." *Yale Law & Policy Review* 9: 301–53.

Rubin, Donald B. 2008. "For Objective Causal Inference, Design Trumps Analysis." *The Annals of Applied Statistics* 2 (3): 808–40.

Seabrook, Nicholas R. 2017. *Drawing the Lines: Constraints on Partisan Gerrymandering in US Politics*. Ithaca, NY: Cornell University Press.

Stephanopoulos, Nicholas O., and Eric M. McGhee. 2015. "Partisan Gerrymandering and the Efficiency Gap." *The University of Chicago Law Review* 82 (2): 831–900.

Tufte, Edward R. 1973. "The Relationship between Seats and Votes in Two-Party Systems." *American Political Science Review* 67 (2): 540–54.

VanderWeele, Tyler J., and Miguel A. Hernán. 2013. "Causal Inference under Multiple Versions of Treatment." *Journal of Causal Inference* 1 (1): 1–20.

Voting and Election Science Team. 2018. "2016 Precinct-Level Election Results." Harvard Dataverse. https://doi.org/10.7910/DVN/NH5S2I.

Voting and Election Science Team. 2020. "2020 Precinct-Level Election Results." Harvard Dataverse. https://doi.org/10.7910/DVN/K7760H.

Warshaw, Christopher, Eric McGhee, and Michal Migurski. 2022. "Districts for a New Decade—Partisan Outcomes and Racial Representation in the 2021–22 Redistricting Cycle". *Publius: The Journal of Federalism* 52 (3): 428–51.