

THE ECONOMICS OF PARTISAN GERRYMANDERING

ANTON KOLOTILIN

School of Economics, UNSW Business School

ALEXANDER WOLITZKY

Department of Economics, MIT

We study the problem of a partisan gerrymanderer who assigns voters to equipopulous districts to maximize his party's expected seat share. The designer faces both aggregate, district-level uncertainty (how many votes his party will receive) and idiosyncratic, voter-level uncertainty (which voters will vote for his party). *Segregate-pair districting*, where weaker districts contain one type of voter, while stronger districts contain two, is optimal for the gerrymanderer. The optimal form of segregate-pair districting depends on the designer's popularity and the relative amounts of aggregate and idiosyncratic uncertainty. When idiosyncratic uncertainty dominates, a designer with majority support pairs all voters, while a designer with minority support segregates opposing voters and pairs more favorable voters; these plans resemble uniform districting and “packing-and-cracking,” respectively. When aggregate uncertainty dominates, the designer segregates moderate voters and pairs extreme voters; this “matching slices” plan has received some attention in the literature. Estimating the model using precinct-level returns from recent U.S. House elections shows that, in practice, idiosyncratic uncertainty dominates. We discuss implications for redistricting reform, political polarization, and detecting gerrymandering. Methodologically, we exploit a formal connection between gerrymandering—partitioning voters into districts—and information design—partitioning states of the world into signals.

KEYWORDS: Gerrymandering, pack-and-crack, segregate-pair, idiosyncratic versus aggregate uncertainty, information design.

1. INTRODUCTION

LEGISLATIVE DISTRICT BOUNDARIES are drawn by political partisans under many electoral systems (Bickerstaff (2020)). In the United States, the significance of partisan districting has grown with the rise of computer-assisted districting (Newkirk (2017)), together with intense partisan efforts to gain and exploit control of the districting process. These trends culminated in “The Great Gerrymander of 2012” (McGhee (2020)), where the Republican party's Redistricting Majority Project (REDMAP), having previously targeted state-level elections that would give Republicans control of redistricting, aggressively redistricted several states, including Michigan, Ohio, Pennsylvania, and Wisconsin. The resulting districting plans are widely viewed as contributing to the outcome of the 2012 general election, where Republican congressional candidates won a 33-seat majority in the House of Representatives with 49.4% of the two-party vote (McGann, Anthony Smith, Latner, and Keena (2016)). In light of these developments—along with

Anton Kolotilin: akolotilin@gmail.com

Alexander Wolitzky: wolitzky@mit.edu

We thank Nikhil Agarwal, Garance Genicot, Ben Golub, Richard Holden, Gary King, Hongyi Li, Nolan McCarty, Stephen Morris, Ben Olken, and Ken Shotts, as well as many seminar and conference participants, for helpful comments. We thank Eitan Sapiro-Gheiler and Nancy Wang for excellent research assistance. Anton Kolotilin acknowledges support from the ARC DE160100964, DP240103257, FT240100731, and MIT Sloan's Program on Innovation in Markets and Organizations. Alexander Wolitzky acknowledges support from NSF CAREER Award 1555071 and Sloan Foundation Fellowship 2017-9633.

the Supreme Court ruling in *Rucho v. Common Cause* (2019) that partisan gerrymanders are not judiciable in federal court and the continued prominence of gerrymandering in the 2020 U.S. redistricting cycle (Rakich and Mejia (2022))—partisan gerrymandering is likely to remain an important feature of American politics for some time.

This paper studies the problem of a partisan gerrymanderer (the “designer”) who assigns voters to a large number of equipopulous districts so as to maximize his party’s expected seat share.¹ This problem approximates the one facing many partisan gerrymanderers in the United States, where the constraint that districts must be equipopulous is strictly enforced.² In practice, gerrymanderers also face additional constraints, including the federal requirements that districts are contiguous and do not discriminate on the basis of race, and various state-level restrictions including “compactness” requirements, requirements to respect political sub-divisions such as county lines, and requirements to represent racial or ethnic groups or other communities of interest. While these complex constraints can be important, we believe that often they are not as binding as they might seem, and also that they are more productively considered on a case-by-case basis rather than as part of a general theoretical analysis.³ We therefore follow much of the literature (discussed below) in focusing on the simpler problem with only the equipopulation constraint.

When the designer has perfect information, the solution to this problem is well known. If the designer’s party is supported by a minority of voters of size $m < 1/2$, he “packs” $1 - 2m$ opposing voters in districts where he receives zero votes and “cracks” the remaining $2m$ voters in districts which he wins with 50% of the vote. If the designer has majority support, he can win all districts by making them identical. Thus, under perfect information, *pack-and-crack* is optimal for a designer with minority support, while *uniform districting* is optimal for a designer with majority support. We instead consider the more general and realistic case where the designer must allocate a variety of types of voters (or, more realistically, groups of voters such as census blocks or precincts) under uncertainty. The goal of this paper is to characterize optimal partisan gerrymandering in this setting and to draw implications for broader legal and political economy issues surrounding gerrymandering.

In outline, our model and results are as follows. We assume that the designer faces both aggregate, district-level uncertainty (how many votes his party will receive) and idiosyncratic, voter-level uncertainty (which voters will vote for his party). Aggregate uncertainty is parameterized by a one-dimensional aggregate shock, while voters are parameterized by a one-dimensional type that determines a voter’s probability of voting for the designer’s party for each value of the aggregate shock. A *district* is identified with a distribution of voter types, and a district’s *strength* is defined as the cutoff aggregate shock below which the designer’s party wins the district, so the designer wins stronger districts with higher probability. We assume that the distributions of the aggregate and idiosyncratic shocks are centrally unimodal with log-concave densities drawn from the same location-scale

¹We hasten to add that studying this problem does not endorse gerrymandering, any more than studying monopolistic behavior endorses monopoly.

²In *Karcher v. Daggett* (1983), the Supreme Court rejected a districting plan in New Jersey with less than a 1% deviation from population equality, finding that “there are no *de minimus* population variations, which could practically be avoided, but which nonetheless meet the standard of Article I, Section 2 [of the U.S. Constitution] without justification.”

³An exception is the requirement to respect county lines, which we address in Section 6.2. See Friedman and Holden (2008) for discussion of the other constraints. For example, contiguity is not as severe a constraint as it might seem, because contiguous districts can have highly irregular shapes.

family. This assumption lets us cleanly compare the “amounts” of aggregate and idiosyncratic uncertainty.

Our analysis is based on a technical result (Lemma 1) that describes the designer’s value from assigning any type of voter to a district of any possible strength and states that an optimal districting plan must assign all voters to districts where they generate the highest value. Intuitively, assigning more voter types to stronger districts lets the designer create more such districts (by the equipopulation constraint), but also weakens these districts (by adding less favorable voters), and this tradeoff resolves differently for different voter types and district strengths.

Our first substantive result (Theorem 1) is that optimal districting takes a *segregate-pair* form.⁴ Under segregate-pair districting, the designer creates weaker districts that contain a single voter type (which are analogous to the packed districts under pack-and-crack) and stronger districts that contain two voter types (analogous to the cracked districts). The class of segregate-pair plans admits a tight characterization but is rich enough to cover a variety of districting plans, including refinements of all of the main plans proposed in the prior literature. The optimality of segregate-pair districting is thus a key organizing result.

We then turn to our main results, which characterize optimal districting as a function of the designer’s popularity and the relative amounts of idiosyncratic and aggregate uncertainty. First, we show that if the designer has strong support from all voter types, then a *negative assortative districting* (NAD) plan is optimal, where extreme left- and right-wing voters are paired together. Conversely, if the designer has weak support from all voter types and idiosyncratic uncertainty is larger than aggregate uncertainty, then a *segregation* plan is optimal, where each district contains only a single voter type (Theorem 2). Second, if aggregate uncertainty is very small, optimal districting for a designer with majority support approximates NAD, while optimal districting for a designer with minority support approximates a *segregate-opponents-and-pair* (SOP) plan, where unfavorable voters are segregated and more favorable voters are paired in a negatively assortative manner (Theorem 3). The former result is analogous to the optimality of uniform districting for a designer with majority support without uncertainty, because NAD plans are versions of uniform districting that pair voter types rather than pooling all types together; similarly, the latter result is analogous to the optimality of pack-and-crack districting for a designer with minority support without uncertainty, because SOP plans are versions of pack-and-crack districting that segregate unfavorable voters and pair more favorable voter types rather than pooling them. Indeed, while exactly optimal districting with small aggregate uncertainty approximates NAD or SOP, much simpler uniform districting or pack-and-crack plans are approximately optimal (with majority or minority designer support, respectively). Third, if idiosyncratic uncertainty is very small, optimal districting approximates NAD with a 50–50 voter type split in each district (Theorem 4).⁵ Fourth, in the intermediate region where both the designer’s support among voters and the ratio of aggregate and idiosyncratic uncertainty are balanced, mixed plans can be optimal (as well as a *segregate-moderates-and-pair* plan, where moderate voters are segregated and extreme left- and right-wing voters are paired), and we can numerically trace out the parameter regions where each plan is optimal.

As we discuss in Section 6, the form of optimal partisan districting has significant implications for several political and legal issues, including redistricting reform, intra- and

⁴We prove this result when idiosyncratic uncertainty is larger than aggregate uncertainty. We conjecture that it also holds more generally.

⁵This result refines the main result of Friedman and Holden (2008).

inter-district political polarization, and measuring gerrymandering. Since our results show that the ratio of idiosyncratic and aggregate uncertainty determines the form of optimal districting, it is therefore important to understand whether idiosyncratic or aggregate uncertainty is larger in practice. We answer this question using precinct-level returns from the 2016, 2018, and 2020 U.S. House elections. The data clearly show that idiosyncratic uncertainty is much larger than aggregate uncertainty. Intuitively, this holds because most precinct vote splits are much closer to 50–50 (the vote split under high idiosyncratic uncertainty) than 100–0 or 0–100 (the vote splits under high aggregate uncertainty).⁶ Therefore, for realistic parameters, exactly optimal districting approximates NAD (for a designer with majority voter support) or SOP (for a designer with minority support), while uniform districting or pack-and-crack is approximately optimal. This finding can help explain why actual gerrymandering usually resembles uniform districting or pack-and-crack.

Methodologically, we establish a formal connection between gerrymandering—partitioning voters into districts—and information design—partitioning states of the world into signals. With perfect information, pack-and-crack is equivalent to [Kamenica and Gentzkow’s 2011](#) solution to their “prosecutor-judge” example. Our general model is equivalent to a non-linear Bayesian persuasion problem with a one-dimensional state (corresponding to one-dimensional voter partisanship in our model), a one-dimensional action for the receiver (corresponding to the one-dimensional strength of a district), and state-independent sender preferences (as the designer only cares about how many districts he wins and not directly about the districts’ composition). Our key lemma (Lemma 1) is implied by Theorem 1 of [Kolotilin, Corrao, and Wolitzky \(2025\)](#) in the context of this persuasion problem.⁷ Conversely, our theorems are novel in the persuasion context, so this paper directly contributes to information design as well as gerrymandering. More importantly, we establish a tight connection between these topics.⁸

Related Literature. The closest prior papers on optimal partisan gerrymandering are [Owen and Grofman \(1988\)](#), [Friedman and Holden \(2008\)](#), and [Gul and Pesendorfer \(2010\)](#). [Owen and Grofman’s](#) model is equivalent to the special case of our model with two voter types. [Gul and Pesendorfer](#) study competition between two designers who each control districting in some area and aim to win a majority of seats.⁹ A simplified version of their model with a single designer is equivalent to the special case of our model with uniform idiosyncratic shocks. [Friedman and Holden](#) consider a model similar to ours (although with finitely many districts, rather than a continuum as in our model and [Gul and Pesendorfer](#)), but their main results concern the case where idiosyncratic uncertainty is much smaller than aggregate uncertainty. In contrast, we do not restrict the relative amounts of aggregate and idiosyncratic uncertainty, and we show empirically that the practically relevant case is that where idiosyncratic uncertainty dominates (i.e., the opposite of the case emphasized by [Friedman and Holden](#)).

The broader literature on gerrymandering and redistricting addresses a wide range of issues, including geographic constraints on gerrymandering ([Sherstyuk \(1998\)](#), [Shotts \(2001\)](#), [Puppe and Tasnádi \(2009\)](#)), gerrymandering with heterogeneous voter turnout

⁶This observation also implies that simple models with only two types of voters or precincts (e.g., [Owen and Grofman \(1988\)](#)) cannot closely approximate the problem facing actual gerrymanderers, who must assign many different types of precincts.

⁷In turn, [Kolotilin, Corrao, and Wolitzky \(2025\)](#) was inspired by [Friedman and Holden \(2008\)](#).

⁸Contemporaneous papers by [Lagarde and Tomala \(2021\)](#) and [Gomberg, Pancs, and Sharma \(2023\)](#) also emphasize connections between gerrymandering and information design, albeit in less general models: [Lagarde and Tomala](#) assume two voter types, while [Gomberg, Pancs, and Sharma](#) assume no aggregate uncertainty.

⁹[Friedman and Holden \(2020\)](#) study designer competition in the model of their 2008 paper.

(Bouton, Genicot, Castanheira, and Stashko (2024), Gomberg, Pancs, and Sharma (2024)), gerrymandering with party primaries (Moscarriello (2025)), socially optimal districting (Gilligan and Matsusaka (2006), Coate and Knight (2007), Bracco (2013)), measuring district compactness (Chambers and Miller (2010), Fryer and Holden (2011), Ely (2022)), the interaction of redistricting and policy choices (Shotts (2002), Besley and Preston (2007), Groll and O’Halloran (2024)), measuring gerrymandering (King and Browning (1987), McGhee (2014), Stephanopoulos and McGhee (2015), Deford, Duchin, and Solomon (2021), Gomberg, Pancs, and Sharma (2023)), and assessing the consequences of redistricting (Gelman and King (1994b), McCarty, Poole, and Rosenthal (2009), Hayes and McKee (2009), Jeong and Shenoy (2024), Sabet and Yuchtman (2025)). As the partisan gerrymandering problem interacts with many of these issues, our analysis may facilitate future research in these areas.

Outline. The paper is organized as follows: Section 2 presents the model. Section 3 establishes general properties of optimal districting plans that hold regardless of the designer’s popularity or the amount of aggregate and idiosyncratic uncertainty. Section 4 contains our main theoretical and numerical results, which characterize optimal districting as a function of these parameters. Section 5 contains our empirical results, which estimate which parameters are practically relevant. Section 6 discusses policy implications. Section 7 concludes. Proofs are deferred to Appendix A. Details of our estimation procedure are contained in Appendix B. Empirical robustness and goodness-of-fit considerations are discussed in Appendix C. All appendices are in the Supplemental Material (Kolotilin and Wolitzky (2026)).

2. MODEL

We consider a standard electoral model with one-dimensional voter types (parameterizing voter partisanship) and one-dimensional aggregate uncertainty in each district-level race (parameterizing the vote share for the designer’s party).

Voters and Vote Shares. There is a continuum of voters. A voter votes for the designer’s party (for short, “votes for the designer”) iff $s \geq r + t$, where

- $s \in [\underline{s}, \bar{s}]$, with $\underline{s} < \bar{s}$, is the voter’s type, which is observed by the designer and is the object of districting. The population distribution of s is denoted by F .
- $r \in \mathbb{R}$ is the aggregate shock in the voter’s district, which realizes after districting. The distribution of r in each district is denoted by G .¹⁰
- $t \in \mathbb{R}$ is an idiosyncratic, voter-specific “taste shock,” which also realizes after districting. The distribution of t is denoted by Q .

Thus, the share of type s voters who vote for the designer in a district where the aggregate shock takes value r equals $Q(s - r)$.

Note that the designer faces two kinds of uncertainty at the time of districting: aggregate, district-level uncertainty, r , and idiosyncratic, voter-level uncertainty, t . Many of our results involve comparing the “amount” of each kind of uncertainty. To facilitate this comparison, we assume that G and Q have the same shape, in that there exists $\eta > 0$ such that $G(r) = Q(\eta r)$.¹¹ We also define $\gamma = \eta^2 / (1 + \eta^2) \in (0, 1)$, so the ratio of the variances of

¹⁰The correlation among district-level aggregate shocks is irrelevant for our analysis. However, we do estimate it empirically.

¹¹Mathematically, this says that G and Q lie in the same location-scale family. While this assumption is convenient, most of our results do not require it, including Lemma 1, Lemma 2, Theorem 2.1, Theorem 2.3, Theorem 3, and Theorem 4.

r and t is $(1 - \gamma)/\gamma$. The parameter γ thus captures the *share of idiosyncratic uncertainty*. We say that aggregate uncertainty is *larger* than idiosyncratic uncertainty if $\gamma < 0.5$, while idiosyncratic uncertainty is larger if $\gamma > 0.5$.

The model is now fully parameterized by the distributions F and Q and the parameter $\gamma \in (0, 1)$. We assume that F and Q admit strictly positive densities f and q that are three-times differentiable. We also assume that Q is centrally unimodal (with median/mode normalized to 0), and q is strictly log-concave: $Q(0) = 1/2$, $d^2 \ln(q(t))/dt^2 < 0$ for all t , and $q'(0) = 0$. This implies that Q is strictly convex below 0 and strictly concave above 0.

Log-concavity of q is a key assumption. This standard property is satisfied by many distributions (Bagnoli and Bergstrom (2005)) and is similar to Friedman and Holden’s 2008 “Informative Signal Property.” Substantively, it captures the realistic feature that moderate voters are more sensitive to the aggregate shock than more extreme voters.¹²

Central unimodality of Q is also important for many of our results (and is also assumed by Friedman and Holden 2008). This is another standard property, which says that Q is unimodal and is not systematically skewed to either side, so its median and mode coincide.

Districing Plans. The designer assigns voters to a continuum of equipopulous districts based on their types s , and thus determines the distribution P of s in each district.¹³ A district is characterized by the distribution P of voter types s it contains. Thus, a *districing plan*—which specifies the measure of districts with each voter type distribution P —is a distribution \mathcal{H} over distributions P of s , such that the population distribution of s is given by F : that is, $\mathcal{H} \in \Delta\Delta[\underline{s}, \bar{s}]$ and $\int P(s) d\mathcal{H}(P) = F(s)$ for all s .¹⁴ For example, under *uniform districing*, where all districts are the same, \mathcal{H} assigns probability 1 to $P = F$. In the opposite extreme case of *segregation*, where each district consists entirely of one type of voter, every distribution P in the support of \mathcal{H} takes the form $P = \delta_s$ for some $s \in [\underline{s}, \bar{s}]$, where δ_s denotes the degenerate distribution on voter type s . Finally, the best-known districing plan is *pack-and-crack*, where there is a cutoff voter type $s^* \in (\underline{s}, \bar{s})$ such that $\text{supp}(\mathcal{H}) = \{P, P'\}$ and P and P' are the lower and upper truncations of F at s^* .

We say that a districing plan is *pure* if (almost) each voter type s is assigned to only one kind of district (so there is a unique $P \in \text{supp}(\mathcal{H})$ such that $s \in \text{supp}(P)$) and *mixed* otherwise. Since the distribution of voter types F is continuous, it is natural to expect pure districing to be optimal, but we will see that this is not always the case.

Designer’s Problem. The designer wins a district iff he receives a majority of the district vote. Thus, the designer wins a district with voter type distribution P (henceforth, “district P ”) iff the district’s aggregate shock r satisfies $\int Q(s - r) dP(s) \geq 1/2$. Since $Q(s - r)$ is decreasing in r and $Q(0) = 1/2$, this occurs iff r falls below the point $r^*(P)$ where

$$\int Q(s - r^*(P)) dP(s) = Q(0).$$

We refer to $r^*(P)$ as the *strength* of district P . Since the aggregate shock has the same distribution in all districts, the designer wins stronger districts with higher probability. Note

¹²As Rakich and Silver (2018) put it in describing the “elasticity scores” in FiveThirtyEight.com’s forecasting model, “Voters at the extreme end of the spectrum—those who have close to a 0 percent or a 100 percent chance of voting for one of the parties—don’t swing as much as those in the middle.”

¹³We follow Gul and Pesendorfer (2010) in assuming a continuum of districts. Since districing plans in the U.S. are drawn at the state level, this implicitly assumes that each state contains a large number of districts. Of course, this is a better approximation for state legislative districts and for congressional districts in large states than it is for congressional districts in small states. Introducing integer constraints on the number of districts, while interesting and realistic, would complicate the analysis and obscure our insights.

¹⁴Throughout, for any compact metric space X , ΔX denotes the set of probability measures on X , endowed with the weak* topology. For any $\mu \in \Delta X$, its support $\text{supp}(\mu)$ is the smallest compact set of measure 1.

that the strength of a segregated district $P = \delta_s$ is simply s , while in general a district's strength lies somewhere in the convex hull of $\text{supp}(P)$.

We assume that the designer maximizes his party's expected seat share.¹⁵ Thus, the designer's problem is

$$\begin{aligned} & \max_{\mathcal{H} \in \Delta\Delta[s, \bar{s}]} \int G(r^*(P)) d\mathcal{H}(P) \\ & \text{s.t. } \int P d\mathcal{H}(P) = F, \end{aligned} \tag{1}$$

where, recall, $r^*(P)$ is defined as the solution to $\int Q(s - r^*(P)) dP(s) = Q(0)$. This problem is similar to that of [Friedman and Holden \(2008\)](#), which in turn nests [Owen and Grofman \(1988\)](#) and a single-designer version of [Gul and Pesendorfer \(2010\)](#).¹⁶

REMARK 1—Equivalence to Bayesian Persuasion: The designer's problem (1) is equivalent to a Bayesian persuasion problem where the designer splits a prior distribution F into posterior distributions P and obtains utility $G(r^*(P))$ from inducing P . Formally, to map the designer's problem to a Bayesian persuasion problem, voter types s map to states of the world; the population distribution of voter types F maps to the prior over states; districts P map to posterior beliefs; a districting plan \mathcal{H} maps to a disclosure policy (i.e., a distribution over induced posteriors); district strengths $r^*(P)$ map to receiver actions; the distribution of idiosyncratic uncertainty $Q(s - r)$ maps to (1/2 plus) the receiver's marginal utility from increasing her action r in state s ; and the distribution of aggregate uncertainty $G(r)$ maps to the designer's utility from inducing action r .¹⁷ In particular, the designer's problem lies in the translation-invariant subcase of the state-independent sender preferences case of the persuasion problem studied by [Kolotilin, Corrao, and Wolitzky \(2025\)](#), who specialize the general Bayesian persuasion problem of [Kamenica and Gentzkow \(2011\)](#) by assuming that the state and the receiver's action are one-dimensional, the receiver's utility is supermodular and concave in her action, and the sender's utility is increasing in the receiver's action. This observation lets us use technical results from [Kolotilin, Corrao, and Wolitzky \(2025\)](#) to help analyze the designer's problem.

3. OPTIMAL PARTISAN GERRYMANDERING: GENERAL PROPERTIES

We first establish three properties of optimal districting plans that hold regardless of the designer's popularity or the amount of aggregate and idiosyncratic uncertainty. The first property is an optimality condition that describes the designer's value from assigning a voter to a district and says that optimal plans assign voters to the highest-value districts. The second property, *single-dippedness*, says that in a certain sense, more extreme vot-

¹⁵See Section 7 and [Kolotilin and Wolitzky \(2020\)](#) for discussion of more general designer objectives.

¹⁶[Friedman and Holden](#) assume a finite number of districts rather than a continuum and do not assume that G and Q have the same shape. [Owen and Grofman](#) assume binary voter types. [Gul and Pesendorfer](#) consider a majoritarian objective with both state-level and district-level aggregate shocks; however, after conditioning on the pivotal value of the state-level shock, their problem reduces to maximizing expected seat share with only district-level shocks.

¹⁷The constraint that a feasible districting plan \mathcal{H} satisfies $\int P(s) d\mathcal{H}(P) = F(s)$ for all s maps to the "Bayes plausibility" constraint in Bayesian persuasion.

ers are assigned to stronger districts. The third property, *segregate-pairedness*, says that segregated districts are weaker than paired ones.¹⁸

3.1. Optimality Condition

The following condition is our key tool for characterizing optimal districting plans.

LEMMA 1: *An optimal districting plan exists. In addition, there exists a function $\lambda : [\underline{s}, \bar{s}] \rightarrow \mathbb{R}$ such that, for any optimal districting plan \mathcal{H} , any district $P \in \text{supp}(\mathcal{H})$, any voter type $s \in \text{supp}(P)$, and any district strength $r \in [\underline{s}, \bar{s}]$, we have*

$$G(r^*(P)) + \lambda(r^*(P))(Q(s - r^*(P)) - Q(0)) \geq G(r) + \lambda(r)(Q(s - r) - Q(0)) \quad (2)$$

and

$$\lambda(r^*(P)) = \frac{g(r^*(P))}{\int q(s - r^*(P)) dP(s)}. \quad (3)$$

Lemma 1 follows from Theorem 1 of [Kolotilin, Corrao, and Wolitzky \(2025\)](#). Intuitively, the lemma just says that the designer assigns each voter type s to a district P where it generates the highest value. However, this value involves a few terms.

First, the value of assigning voter s to district P is the sum of two terms: the probability of winning district P —which equals $G(r^*(P))$ —and the impact of the assignment on the probability of winning district P . This impact is in turn the product of two terms: the impact of an extra vote on the probability of winning district P —which we denote by $\lambda(r^*(P))$ —and the net number of extra votes beyond $1/2 = Q(0)$ provided by voter s at the pivotal aggregate shock $r^*(P)$ —which equals $Q(s - r^*(P)) - Q(0)$.¹⁹ Next, the impact of an extra vote on the probability of winning district P equals the product of the impact of increasing $r^*(P)$ on the probability of winning district P —which equals $g(r^*(P))$ —and the impact of an extra vote on $r^*(P)$ —which equals the derivative of $r^*(P)$ with respect to ε in the equation $\int Q(s - r^*(P)) dP(s) = Q(0) - \varepsilon$. Finally, the latter derivative equals $1/\int q(s - r^*(P)) dP(s)$, by the implicit function theorem.

3.2. Single-Dippedness

We now show that optimal districting plans are *strictly single-dipped*, in that more extreme voters are assigned to stronger districts: formally, any district $P \in \text{supp}(\mathcal{H})$ containing any two voter types $s < s''$ is stronger than any district $P' \in \text{supp}(\mathcal{H})$ containing any intervening voter type $s' \in (s, s'')$, in that $r^*(P) > r^*(P')$.²⁰ Note that if districting is strictly single-dipped, then each district contains at most two distinct voter types. Thus, any district P in the support of a strictly single-dipped districting plan \mathcal{H} is either *segregated* (if $|\text{supp}(P)| = 1$) or *paired* (if $|\text{supp}(P)| = 2$). For example, segregation is strictly single-dipped, but uniform districting and pack-and-crack are not.

¹⁸We prove this last result for the case where idiosyncratic uncertainty is larger than aggregate uncertainty. See the discussion following Theorem 1.

¹⁹The net number of extra votes beyond $1/2$ is the relevant quantity for influencing $r^*(P)$, as adding a new voter to district P who votes for the designer with probability $1/2$ at aggregate shock $r^*(P)$ would leave $r^*(P)$ unchanged.

²⁰We say that a district P “contains” a voter type s if $s \in \text{supp}(P)$.

LEMMA 2: *Any optimal districting plan is strictly single-dipped.*

Lemma 2 recapitulates Lemma 1 of Friedman and Holden (2008) in our continuum-district model.²¹ To see the intuition, suppose a districting plan creates two equally strong districts, 1 and 2, where District 1 contains moderate voters and District 2 contains a mix of left-wing and right-wing extremists. Since q is log-concave, the vote share is more sensitive to the aggregate shock in District 1 than in District 2 (i.e., District 1 is more “swingy”), which implies that a marginal voter is more likely to be pivotal in District 2. The designer can then profitably exploit this asymmetry by re-assigning some unfavorable voters to District 1 and re-assigning some favorable voters to District 2, thus weakening the swingy, moderate District 1 and strengthening the less swingy, extreme District 2. Breaking all ties in favor of extreme districts in this manner leads to strictly single-dipped districting.

Lemma 2 follows from Lemma 1 and some additional arguments. First, if an optimal districting plan creates two districts P and P' such that district P contains two voter types $s < s''$ and district P' contains an intervening voter type $s' \in (s, s'')$, then Lemma 1 implies that the value of assigning type s or s'' voters to a district of strength $r^*(P)$ is greater than the value of assigning them to a district of strength $r^*(P')$, while the opposite is true for type s' voters. The proof of Lemma 3 in Appendix A.1 shows that, when q is strictly log-concave, these inequalities imply that $r^*(P) \geq r^*(P')$. Intuitively, this holds because extreme voters must be assigned to stronger and less swingy districts.

Second, Lemma 4 in Appendix A.1 shows that any district in an optimal districting plan contains at most two voter types. This follows from Lemma 1, because the values of assigning three distinct voter types to the same district cannot all be equal when q is strictly log-concave. Intuitively, this holds because such districts could be profitably split into a mixture of weaker, swingy districts containing moderate voters and stronger, less swingy districts containing more extreme voters. Together, Lemmas 3 and 4 imply Lemma 2.

Lemma 4 also shows why simple districting plans like uniform districting or pack-and-crack are sub-optimal when q is strictly log-concave: they neglect the gain from splitting equally strong districts containing many voter types into a mixture of weaker, swingy districts containing moderate voters and stronger, less swingy districts containing more extreme voters. However, we will see that these gains are small in the realistic case where idiosyncratic uncertainty is much larger than aggregate uncertainty: see Remark 2.

3.3. Segregate-Pairedness

A strictly single-dipped districting plan can contain a mix of segregated and paired districts of varying strengths. If such a plan \mathcal{H} has the further property that every segregated district is weaker than every paired district (i.e., for any $P, P' \in \text{supp}(\mathcal{H})$ such that $|\text{supp}(P)| = 1$ and $|\text{supp}(P')| = 2$, we have $r^*(P) < r^*(P')$), we say that \mathcal{H} is *segregate-pair*.

A segregate-pair plan \mathcal{H} can be described in a simple way. There is a *bifurcation point* $r^b \in (\underline{s}, \bar{s}]$ that divides the segregated and paired districts, so that $r^*(P) \leq r^b$ for all segregated districts $P \in \text{supp}(\mathcal{H})$, and $r^*(P) > r^b$ for all paired districts $P \in \text{supp}(\mathcal{H})$.²² The assignment of voters to paired districts is then described by a decreasing function s_1

²¹Kolotilin, Corrao, and Wolitzky (2025) give more general sufficient conditions for single-dippedness, allowing state-dependent designer preferences.

²²Formally, we define the bifurcation point as the infimum of $r^*(P)$ over all paired districts $P \in \text{supp}(\mathcal{H})$.

and an increasing function s_2 , where the two types in a paired district P are $s_1(r^*(P))$ and $s_2(r^*(P)) > s_1(r^*(P))$. Stronger paired districts thus contain more extreme voters, as single-dippedness requires. Lemma 5 in Appendix A.2 formalizes the description of a segregate-pair plan by a bifurcation point r^b and functions s_1 and s_2 .

Despite this tight characterization, a range of important districting plans are segregate-pair, including the following:

Segregation, where all voters are segregated: $P = \delta_{r^*(P)}$ for all $P \in \text{supp}(\mathcal{H})$, or equivalently $r^b = \bar{s}$.

Segregate-Moderates-and-Pair (SMP), where moderate voters are segregated and extreme voters are paired in a negatively assortative manner: \mathcal{H} is pure, $r^b \in (\underline{s}, \bar{s})$, and there exists $\hat{s} \in (\underline{s}, r^b)$ such that a district $P \in \text{supp}(\mathcal{H})$ is segregated iff $r^*(P) \in [\hat{s}, r^b]$.

Segregate-Opponents-and-Pair (SOP), where unfavorable voters are segregated and more favorable voters are paired in a negatively assortative manner: \mathcal{H} is pure, $r^b \in (\underline{s}, \bar{s})$, and there exists $\hat{s} \in (\underline{s}, r^b)$ such that a district $P \in \text{supp}(\mathcal{H})$ is segregated iff $r^*(P) \in [\underline{s}, \hat{s}] \cup \{r^b\}$.

Negative Assortative Districting (NAD), where all voters (except for $s = r^b$) are paired in a negatively assortative manner: $r^b = \inf_{P \in \text{supp}(\mathcal{H})} r^*(P)$.

These four plans feature prominently in our results and warrant some discussion. For a visualization of these plans, see Figure 1, where Panels (a)–(c) are SMP plans (with Panel (a) approaching the extreme case of NAD), and Panels (g)–(i) are SOP plans; or Figure 2, where Panels (a) and (d) are SMP plans (which are both similar to the extreme case of NAD), Panel (c) is a SOP plan, and Panel (f) is NAD.

Segregation and NAD are the extreme segregate-pair plans where all voter types are segregated and where only a single type is segregated. There is a unique segregation plan, but there is a continuum of NAD plans, depending on the weights on the different voter types in each paired district. (Similarly, there is also a continuum of SMP and SOP plans.) NAD plans can be viewed as “strictly single-dipped versions” of uniform districting: starting from uniform districting and splitting the pool of voters into pairs in a strictly single-dipped manner yields NAD.

Similarly, SOP plans are strictly single-dipped version of [Gul and Pesendorfer’s 2010](#) “ p -segregation” plan, where unfavorable voters are segregated and more favorable voters are pooled: starting from p -segregation and splitting the pool into pairs yields SOP. SOP can also be obtained from pack-and-crack districting by first splitting the weak districts into segregated ones (yielding p -segregation) and then splitting the strong districts into pairs.

Finally, SMP is the same as [Friedman and Holden’s 2008](#) “matching slices” plan, with the difference that [Friedman and Holden](#) assume a finite number of districts and do not mention the possibility of segregating a non-trivial interval of moderate voter types.

An instructive example of a plan that can be strictly single-dipped but not segregate-pair is “Segregate-Supporters-and-Pair,” where favorable voters are segregated and less favorable voters are paired. This plan can be obtained from pack-and-crack by splitting weak districts into pairs and splitting strong districts into segregated ones.

Our first main result is that segregate-pair districting is optimal if idiosyncratic uncertainty is larger than aggregate uncertainty.

THEOREM 1: *If idiosyncratic uncertainty is larger than aggregate uncertainty, there is a unique optimal districting plan, which is segregate-pair.*

Numerically, segregate-pair districting is also optimal when idiosyncratic uncertainty is smaller than aggregate uncertainty, but we were not able to prove this.²³ However, Theorem 1 covers the empirically relevant case, as we will estimate that γ is much greater than 0.5.

A rough intuition for Theorem 1 is as follows. First, log-concavity of g implies that G , the distribution of the aggregate shock, is first convex and then concave. Second, convexity of G favors segregation (as splitting a district with strength r^* into districts with strengths $r^* - \varepsilon$ and $r^* + \varepsilon$ increases the expected seat share if G is convex and Q is linear), while concavity of G favors pairing. Thus, when G is first convex and then concave, weak districts should be segregated and strong districts should be paired. This is precisely segregate-pair.

The proof of Theorem 1 has two steps. First, Lemma 8 in Appendix A.3 shows that if $\gamma \geq 0.5$, then it is sub-optimal to create a paired district P where $r^*(P)$ is below 0 (the inflection point of G).²⁴ In turn, Lemma 8 follows from Lemma 1 and strict log-concavity of q . Intuitively, when $\gamma \geq 0.5$ (and G and Q have the same shape, as we assume throughout), the curvature of G dominates that of Q , so the strength r of any paired district must lie on an interval where G is concave, that is, $r > 0$.

Next, note that a strictly single-dipped plan is *not* segregate-pair iff there exist $s < r < s' \leq s''$ such that voter types $s < s'$ are paired in a district P with $r^*(P) = r$ and voter type s'' is segregated. Let us suppose toward a contradiction that such a plan is optimal. Then, applying equations (2) and (3) for district P , voter type s , and district strength $s'' = r^*(\delta_{s''})$, we have

$$G(r) + \lambda(r)(Q(s - r) - Q(0)) \geq G(s'') + \frac{g(s'')}{q(0)}(Q(s - s'') - Q(0)).$$

(Intuitively, this inequality says that it is better to assign a type s voter to the paired district P with strength $r^*(P) = r$ than to the segregated district $\delta_{s''}$ with strength s'' .) We also have $r > 0$, by Lemma 8. But then we have

$$\begin{aligned} & G(s'') + \frac{g(s'')}{q(0)}(Q(s - s'') - Q(0)) - G(r) - \lambda(r)(Q(s - r) - Q(0)) \\ & > \frac{g(r)}{q(0)}(Q(0) - Q(s - r)) + G(s'') - G(r) - \frac{g(s'')}{q(0)}(Q(0) - Q(s - s'')) \\ & > \frac{g(s'')}{q(0)}[Q(0) - Q(s - r) + q(0)(s'' - r) - (Q(0) - Q(s - s''))] > 0, \end{aligned}$$

where the first inequality holds because $\lambda(r) > g(r)/q(0)$ by (3) (as q is uniquely maximized at 0), the second inequality holds by strict concavity of G on $[r, s'']$, and the third inequality holds by strict convexity of Q on $[s - s'', 0]$. (Intuitively, this inequality says that it is strictly better to move a few type s voters from district P to district $\delta_{s''}$.) This con-

²³See Figures 1–3, where all optimal plans are segregate-pair. In these figures, G and Q are normal. Using Lemma 6 in Appendix A.3, we have checked numerically that segregate-pair is also optimal when G and Q are logistic. The normal and logistic families are the only standard location-scale families we are aware of with symmetric and strictly log-concave densities on \mathbb{R} .

²⁴This is our first result that uses central unimodality.

tradition shows that any non-segregate-pair plan is sub-optimal, so the optimal plan is segregate-pair.²⁵

Theorem 1 is a fundamental result: it is optimal to make paired districts stronger than segregated ones. However, while segregate-pair plans admit a tight characterization, we have seen that a wide variety of plans are segregate-pair. The next section characterizes optimal plans as a function of F , Q , and γ .

As an aside, we note that Theorem 1 also contributes to the Bayesian persuasion literature. A disclosure policy that frequently arises in this literature is *upper censorship*, where states below a cutoff are disclosed and states above the cutoff are pooled. Upper censorship is often optimal in “linear” persuasion problems, where a posterior can be summarized by its mean (Kolotilin (2018), Kolotilin, Mylovanov, and Zapechelnyuk (2022)). However, in non-linear persuasion problems, a version of strict single-dippedness often holds, so disclosure policies that pool more than two states (like upper censorship) cannot be optimal (Kolotilin, Corrao, and Wolitzky (2025)). This raises the question of when a strictly single-dipped version of upper censorship—such as a segregate-pair policy—is optimal. Theorem 1 is the first result to give sufficient conditions for such policies to be optimal.

4. OPTIMAL PARTISAN GERRYMANDERING IN DIFFERENT PARAMETER REGIMES

We now present our results on optimal districting as a function of the designer’s popularity and the ratio of idiosyncratic and aggregate uncertainty. First, NAD is optimal if the designer has strong support from all voter types, and segregation is optimal if the designer has weak support from all voter types and idiosyncratic uncertainty is larger than aggregate uncertainty (Theorem 2). Second, optimal plans approximate NAD or SOP with equally strong paired districts if aggregate uncertainty is small (Theorem 3). Since we will estimate that aggregate uncertainty is small empirically, Theorem 3 is our most practically relevant result. However, while exactly optimal plans approximate NAD or SOP, uniform districting or pack-and-crack districting is also approximately optimal. Third, optimal plans approximate NAD with a 50–50 voter type split in each district if idiosyncratic uncertainty is small (Theorem 4). Fourth, in the intermediate region where both the designer’s support among voters and the ratio of idiosyncratic and aggregate uncertainty are balanced, mixed versions of SOP and SMP can be optimal, and we can numerically trace out the parameter regions where each plan emerges (Theorem 5 and the subsequent numerical results). Overall, we give a fairly complete picture of how optimal districting varies with the designer’s support and the ratio of idiosyncratic and aggregate uncertainty, which we illustrate in Figure 3 at the end of this section.

4.1. *Optimal Districting With Imbalanced Voter Support*

We first investigate optimal districting when voter support is highly imbalanced between the parties. We say that the designer has *uniformly strong support* if $\underline{s} \geq 0$. This means that, when the aggregate shock takes its modal value of 0, *all* voter types vote for the designer with probability at least 50%. Thus, a designer with uniformly strong support can only lose a district when the aggregate shock lands in the right (unfavorable) tail. Conversely, the designer has *uniformly weak support* if $\bar{s} \leq 0$, so he can only win a district when the aggregate shock lands in the left tail. Finally, the designer has *balanced support* if $r^*(F) = 0$, so the overall vote is 50–50 when the aggregate shock takes its modal value.

²⁵Uniqueness of the optimal plan is established by Lemma 6 in Appendix A.3.

THEOREM 2: *The following hold:*

1. *If the designer has uniformly strong support, there is a unique optimal districting plan, which is NAD.*
2. *If the designer has uniformly weak support and idiosyncratic uncertainty is larger than aggregate uncertainty, there is a unique optimal districting plan, which is segregation.*
3. *If the designer has balanced support, NAD and segregation are both sub-optimal.*

Since NAD plans are strictly single-dipped versions of uniform districting, the optimality of NAD in case 1 is akin to the optimality of uniform districting for a designer with majority support in the absence of aggregate uncertainty. To see why NAD is optimal, recall that any strictly single-dipped plan that never segregates two distinct voter types is NAD. So, since $\underline{s} \geq 0$, it suffices to show that it is sub-optimal to segregate any two voter types $s < s'$ that lie in an interval where G is concave. But this follows from Lemma 1 because

$$G(s') + \frac{g(s')}{q(0)}(Q(s - s') - Q(0)) > G(s') + g(s')(s - s') > G(s), \quad (4)$$

where the first inequality is by strict convexity of Q on $[s - s', 0]$, and the second inequality is by strict concavity of G on $[s, s']$. (Intuitively, this inequality says that it is strictly better to move a few type s voters from δ_s to district $\delta_{s'}$.) Thus, segregation can be optimal only on an interval where G is convex.

The optimality of segregation in case 2 follows from Lemma 8 in Appendix A.3, which shows that if $\gamma \geq 0.5$, then it is sub-optimal to create a paired district of strength $r \leq 0$. This implies that segregation is optimal if $\gamma \geq 0.5$ and $\bar{s} \leq 0$, as if $\bar{s} \leq 0$, then $r^*(P) \leq 0$ for any feasible district P .

In case 3, neither NAD nor segregation is optimal, so the optimal plan creates a mix of segregated and paired districts.²⁶ The intuition is that the designer prefers to pool any two positive voter types (as in case 1, since $G(s)$ is concave on $s > 0$), so segregation is sub-optimal; but at the same time, for any strictly single-dipped NAD plan, there exist nearby voter types s, s' that are paired in a district of strength $r < 0$, and the designer prefers to segregate these types because the convexity of G dominates the curvature of Q when s and s' are close to r . (More precisely, this follows because if $s < r < 0$ and $r - s$ is sufficiently small, then

$$\frac{G(r) - G(s)}{g(r)} < \frac{Q(0) - Q(s - r)}{q(0)},$$

because $g'(r) > 0 = q'(0)$ and g and q are strictly positive and three-times differentiable; and therefore

$$G(s) > G(r) + \frac{g(r)}{q(0)}(Q(s - r) - Q(0)) > G(r) + \lambda(r)(Q(s - r) - Q(0)), \quad (5)$$

which implies that the designer prefers to segregate type s voters rather than assigning them to a district of strength r , by Lemma 1.)

²⁶Proposition 1 of [Friedman and Holden \(2008\)](#) shows that SMP (“matching slices”) is optimal when idiosyncratic uncertainty is sufficiently small, but their discussion focuses on NAD. In contrast, Proposition 2 shows that NAD is never optimal with symmetric parties and a continuum of districts.

4.2. Optimal Districting With Small Aggregate or Idiosyncratic Uncertainty

We now consider optimal districting when either aggregate or idiosyncratic uncertainty is small. We will see that the small aggregate uncertainty case is the empirically relevant one. We include the small idiosyncratic uncertainty case for completeness and also to show how the main result of [Friedman and Holden \(2008\)](#) fits in our framework.²⁷

Aggregate uncertainty is small when F and Q are fixed and $\gamma \rightarrow 1$, so the aggregate shock r is close to 0 with high probability. When aggregate uncertainty is small and $r^*(F) > 0$ (so the designer has majority support at the modal aggregate shock), the designer's expected seat share is close to 1 under uniform districting. But since uniform districting is not strictly single-dipped, it cannot be exactly optimal for any $\gamma < 1$, by [Lemma 2](#). Instead, we show that optimal districting approximates NAD with equally strong paired districts. An interpretation is that, when aggregate uncertainty is small and $r^*(F) > 0$, optimal districting starts from uniform districting and then splits pooled districts into equally strong paired districts in a negatively assortative manner.

When aggregate uncertainty is small and $r^*(F) < 0$, the designer's optimal expected seat share is approximately $1 - F(s^*(0))$, where, for any $r \in (r^*(F), \bar{s})$, $s^*(r)$ is defined so that the designer's vote share among voter types $s \geq s^*(r)$ at aggregate shock r is 50%.²⁸ This expected seat share can be approximated by a pack-and-crack plan where, for a small $\varepsilon > 0$, voter types $s < s^*(\varepsilon)$ are assigned to identical weak districts that the designer loses with high probability, and voter types $s \geq s^*(\varepsilon)$ are assigned to identical strong districts that the designer wins with a vote share close to 50% with high probability.²⁹ However, since pack-and-crack districting is not strictly single-dipped, it cannot be exactly optimal. Instead, we show that optimal districting approximates SOP with equally strong paired districts. An interpretation is that, when aggregate uncertainty is small and $r^*(F) < 0$, optimal districting starts from pack-and-crack and then splits the packed districts into segregated districts and splits the cracked districts into equally strong paired districts in a negatively assortative manner.

To state this result, let \mathcal{H}^* be the unique districting plan that segregates types below $s^*(0)$ (if $s^*(0) > \underline{s}$, which holds when $r^*(F) < 0$) and pairs types above $s^*(0)$ in equally strong districts in a negatively assortative manner. Formally, letting $r_+^*(F) = \max\{0, r^*(F)\}$, \mathcal{H}^* is the unique plan \mathcal{H} such that, for any $P \in \text{supp}(\mathcal{H})$, either (a) $\text{supp}(P) = \{s(P)\}$ such that $s(P) \in [\underline{s}, s^*(0)] \cup \{r_+^*(F)\}$, or (b) $\text{supp}(P) = \{s_1(P), s_2(P)\}$ such that $r^*(P) = r_+^*(F)$, $s^*(0) \leq s_1(P) < r_+^*(F) < s_2(P) \leq \bar{s}$, and

$$\int_{[s^*(0), s_1(P)] \cup [s_2(P), \bar{s}]} (Q(s - r_+^*(F)) - Q(0)) dF(s) = 0.$$

THEOREM 3: *As aggregate uncertainty vanishes ($\gamma \rightarrow 1$ with F and Q fixed), the optimal expected seat share converges to $1 - F(s^*(0))$, and the optimal districting plan converges to \mathcal{H}^* .³⁰ Thus, when aggregate uncertainty is small, optimal districting approximates NAD with equally strong paired districts if $r^*(F) \geq 0$ and approximates SOP with equally strong paired districts if $r^*(F) < 0$.*

²⁷The results in this subsection, [Theorems 3](#) and [4](#), do not require the assumption that G and Q lie in the same location-scale family, although this assumption facilitates the exposition.

²⁸Formally, we define $s^*(r)$ as the smallest $\tilde{s} \in [\underline{s}, \bar{s}]$ such that $\int_{\tilde{s}}^{\bar{s}} (Q(s - r) - Q(0)) dF(s) \geq 0$. Note that $s^*(r) = \underline{s}$ if $r \leq r^*(F)$, $s^*(r) = \bar{s}$ if $r \geq \bar{s}$, and $s^*(r) \in (\underline{s}, \bar{s})$ if $r \in (r^*(F), \bar{s})$.

²⁹This is shown in [Lemma 11](#) in [Appendix A.6](#).

³⁰The latter convergence is in the weak* topology.

Intuitively, when aggregate uncertainty vanishes, the designer creates as many districts P where $r^*(P) \geq 0$ as possible and wins these districts with probability close to 1, while losing the remaining districts (where $r^*(P) < 0$) with probability close to 0. The designer segregates the districts P where $r^*(P) < 0$ to maximize the (small) probability of winning the strongest of these districts.³¹ Conversely, the designer equalizes $r^*(P)$ among districts where $r^*(P) \geq 0$ to minimize the (small) probability of losing the weakest of these districts.³² This follows from log-concavity, which implies that if $r^*(P) > r^*(P') > 0$ and $\gamma \rightarrow 1$, then the probability that the designer loses district P' is infinitely higher than the probability that he loses district P , so he would be better off weakening district P and strengthening district P' . The optimality of equalizing the strengths of paired districts is the basis for the gerrymandering test we propose in Section 6.4.

We now turn to the case where idiosyncratic uncertainty is small. Here, F and G are fixed and $\gamma \rightarrow 0$, so each idiosyncratic shock t is close to 0 with high probability. In this case, whether the designer wins a district P at aggregate shock r is essentially determined by the median voter type s^P in district P : for any $\varepsilon > 0$, when γ is close enough to 0, the designer loses districts where $s^P < r - \varepsilon$ and wins districts where $s^P > r + \varepsilon$. Therefore, any optimal districting plan must approximate the highest feasible distribution of district median voters, which is attained by pairing each voter type s above the population median $s^m = F^{-1}(1/2)$ with below-median types, with 50% weight on the above-median type. Under such a plan with an extra ε weight on the above-median type in each district, the designer's expected seat share is approximately $2 \int_{s^m}^{\bar{s}} G(r) dF(r)$ when γ is small. Moreover, for such a plan to be strictly single-dipped, all voter types must be paired in a negatively assortative manner. The resulting districting plan approximates NAD with a 50–50 voter type split in each district.

To state this result, let \mathcal{H}^{**} be NAD with a 50–50 split in each district. Formally, \mathcal{H}^{**} is the unique plan \mathcal{H} such that, for any $P \in \text{supp}(\mathcal{H})$, we have either (a) $\text{supp}(P) = \{s^m\}$, or (b) $\text{supp}(P) = \{s_1(P), s_2(P)\}$ such that $\underline{s} \leq s_1(P) < s^m < s_2(P) \leq \bar{s}$, and $F(s_1(P)) = 1 - F(s_2(P))$.

THEOREM 4: *As idiosyncratic uncertainty vanishes ($\gamma \rightarrow 0$ with F and G fixed), the optimal expected seat share converges to $2 \int_{s^m}^{\bar{s}} G(r) dF(r)$, and the optimal districting plan converges to \mathcal{H}^{**} . Thus, when idiosyncratic uncertainty is small, optimal districting approximates NAD with a 50–50 voter type split in each district.*

Theorem 4 is similar to [Friedman and Holden's 2008](#) main result. With finitely many districts, [Friedman and Holden](#) show that, when idiosyncratic uncertainty is sufficiently small, optimal districting is a discrete version of SMP, where some number of moderate districts are segregated.³³ Theorem 4 adds that, in the limit as idiosyncratic uncertainty vanishes, only one district is segregated—so SMP collapses to NAD—and the voter type split in all other districts is 50–50.

Comparing Theorems 3 and 4, we see that varying the ratio of idiosyncratic and aggregate uncertainty leads to completely different districting plans. When aggregate uncertainty is small, pack-and-crack is approximately optimal, and the exactly optimal plan is close to NAD or SOP with equally strong paired districts. When idiosyncratic uncertainty is small, pack-and-crack is far from optimal, and the optimal plan is close to NAD with a

³¹This is shown in Lemma 14 in Appendix A.6.

³²This is shown in Lemma 13 in Appendix A.6.

³³[Friedman and Holden's](#) proof relies on perturbation arguments, while ours relies on Lemma 1.

50–50 voter type split in each district. In particular, while a NAD plan can arise in either case, these plans are very different: NAD with equally strong paired districts is similar to uniform districting with small idiosyncratic uncertainty, while NAD with a 50–50 split (or, away from the limit, a $50 - \varepsilon$ – $50 + \varepsilon$ split in favor of the higher type) in each district is very different from uniform districting with small idiosyncratic uncertainty, as $r^*(P)$ is much higher in $50 - \varepsilon$ – $50 + \varepsilon$ districts with more extreme voter types. (For example, compare Figures 2(d) and 2(f).) The critical role of the ratio of idiosyncratic and aggregate uncertainty motivates estimating this parameter in Section 5.

The distinction between optimal districting under small aggregate uncertainty and small idiosyncratic uncertainty relates to results in the probabilistic voting literature. When aggregate uncertainty is small, the probability that the designer wins a district is approximately determined by the mean voter type in the district, as in probabilistic voting models with partisan taste shocks (e.g., Hinich (1977), Lindbeck and Weibull (1993)). Optimizing the distribution of district means against a unimodal aggregate shock then requires segregating opposing voters and pooling more favorable voters, as in p -segregation or SOP or NAD with equally strong paired districts. In contrast, when idiosyncratic uncertainty is small, the probability that the designer wins a district is approximately determined by the median voter type in the district, as in probabilistic voting models with an uncertain median bliss point (e.g., Wittman (1983), Calvert (1985)). The distribution of district medians is then optimized by pairing above-population-median and below-population-median voter types, as in NAD with a 50–50 voter type split in each district.³⁴

4.3. The Balanced Case and Regime Transitions

Finally, we analyze optimal districting in the intermediate case where neither the parties' supporters nor the amounts of aggregate and idiosyncratic uncertainty are highly imbalanced. Here, optimal districting will take the form of either SOP, SMP, or a mixed version of these districting plans that we call “Y-districting.” We say that a segregate-pair plan \mathcal{H} is *Y-districting* if it satisfies the following conditions:

1. The plan creates a range of districts with strengths just below and just above the bifurcation point: formally, there exists $\varepsilon > 0$ such that, for all $r \in [r^b - \varepsilon, r^b + \varepsilon]$ (where r^b is the bifurcation point), there exists $P \in \text{supp}(\mathcal{H})$ such that $r^*(P) = r$.
2. The voter types in the weakest paired districts converge to the voter type in the strongest segregated district: formally, the functions s_1 and s_2 describing the voter types in paired districts are twice differentiable and satisfy $\lim_{r \downarrow r^b} s_1(r) = \lim_{r \downarrow r^b} s_2(r)$.³⁵

For a visualization of Y-districting, see Panels (d)–(f) of Figure 1, and Panels (b) and (e) of Figure 2.

³⁴The distinction between mean- and median-dependence applies to several related strands of the literature. In gerrymandering, Owen and Grofman (1988) and Gul and Pesendorfer (2010) study the mean-dependent case, while Friedman and Holden (2008) study an approximately median-dependent case. In persuasion, Gentzkow and Kamenica (2016), Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017), Kolotilin (2018), Dworczak and Martini (2019), and Kleiner, Moldovanu, and Strack (2021) study the mean-dependent case, while Kolotilin, Corrao, and Wolitzky (2025) study a general case nesting both the mean- and quantile (e.g., median)-dependent case, and Yang and Zentefis (2024) and Kolotilin and Wolitzky (2024) study the quantile-dependent case.

³⁵Differentiability is used in the proof of Theorem 5. It may be possible to drop it.

Note that Y-districting encompasses a mixed version of SOP, where there exists $\hat{s} \in (\underline{s}, r^b)$ such that voter types in $[\underline{s}, \hat{s})$ are always segregated and types in (\hat{s}, r^b) are sometimes segregated and sometimes paired, as well as a mixed version of SMP, where there exists $\hat{s} \in (\underline{s}, r^b)$ such that types in $[\underline{s}, \hat{s})$ are always paired and types in (\hat{s}, r^b) are sometimes segregated and sometimes paired.³⁶ We will see that, with balanced voter support, SOP is optimal when idiosyncratic uncertainty is “substantially” larger than aggregate uncertainty, SMP is optimal when aggregate uncertainty is larger than idiosyncratic uncertainty, and Y-districting is optimal in the intermediate range.

To analyze these cases, let J be the distribution with variance 1 satisfying $Q(t) = J(t/\sqrt{\gamma})$ and $G(r) = J(r/\sqrt{1-\gamma})$, so that the variances of t and r are γ and $1-\gamma$. For example, if Q and G are normal, then J is the standard normal distribution. By varying γ while fixing J , we can simultaneously approximate the low-aggregate uncertainty and low-idiosyncratic uncertainty limits analyzed in Theorems 3 and 4, as r is almost constant as $\gamma \rightarrow 1$ and t is almost constant as $\gamma \rightarrow 0$.

Our analytic result in this section is modest: if Y-districting is optimal, then the ratio of idiosyncratic and aggregate uncertainty must fall in an intermediate range. However, numerically it appears that this result actually fully characterizes optimal districting when voter support is balanced: at least when J is normal and F is uniform, our necessary conditions for optimality of Y-districting are also approximately sufficient, and when the ratio of idiosyncratic uncertainty to aggregate uncertainty is below (resp., above) the range where Y-districting is optimal, then SMP (resp., SOP) is optimal.

THEOREM 5: *If Y-districting is optimal, then $r^b = 0$ and $\gamma \in (0.5, \sqrt{3} - 1 \approx 0.732]$.*

Intuitively, for it to be optimal to pair nearby voter types around r , G must be weakly concave at r ; and for it to be optimal to segregate voter types just below r , G must be weakly convex at r . Hence, the bifurcation point of an optimal Y-districting plan must equal 0, the inflection point of G . In addition, if we take parameters where Y-districting is optimal and increase aggregate uncertainty, it eventually becomes optimal to always segregate voter types just below 0 rather than pairing them with higher voter types, at which point optimal districting becomes SMP (with a bifurcation point below 0)—the proof of Theorem 5 shows that this happens when $\gamma \leq 0.5$. On the other hand, if we take parameters where Y-districting is optimal and decrease aggregate uncertainty, it eventually becomes optimal to always pair voter types just below 0 with higher voter types rather than segregating them, at which point optimal districting becomes SOP (with a bifurcation point above 0)—the proof of Theorem 5 shows that this happens when $\gamma > 0.732$.

This intuition suggests that, with balanced voter support, SMP is optimal when $\gamma \leq 0.5$, Y-districting is optimal when $\gamma \in (0.5, 0.732)$, and SOP is optimal when $\gamma \geq 0.732$. Figure 1 presents numerical solutions that verify this conjecture. In the figure, J is standard normal and F is uniform on $[-1, 1]$.³⁷ Voter types are on the x -axis, and the strengths of the districts to which each voter type is assigned are on the y -axis. (Thus, segregated districts lie on the 45° line, while paired districts straddle the 45° line.) For mixed districting

³⁶In contrast, under SOP there exists $\hat{s} \in (\underline{s}, r^b)$ such that types in $[\underline{s}, \hat{s})$ are always segregated and types in (\hat{s}, r^b) are always paired, while under SMP there exists $\hat{s} \in (\underline{s}, r^b)$ such that types in $[\underline{s}, \hat{s})$ are always paired and types in (\hat{s}, r^b) are always segregated.

³⁷To create the figure, we approximated the designer’s problem by a finite-dimensional linear program and solved it using Gurobi Optimizer. Our approximation specifies that s is uniformly distributed on $\{-1, -0.99, \dots, 0.99, 1\}$ and that the designer is constrained to create districts P satisfying $r^*(P) \in \{-1, -0.99, \dots, 0.99, 1\}$.

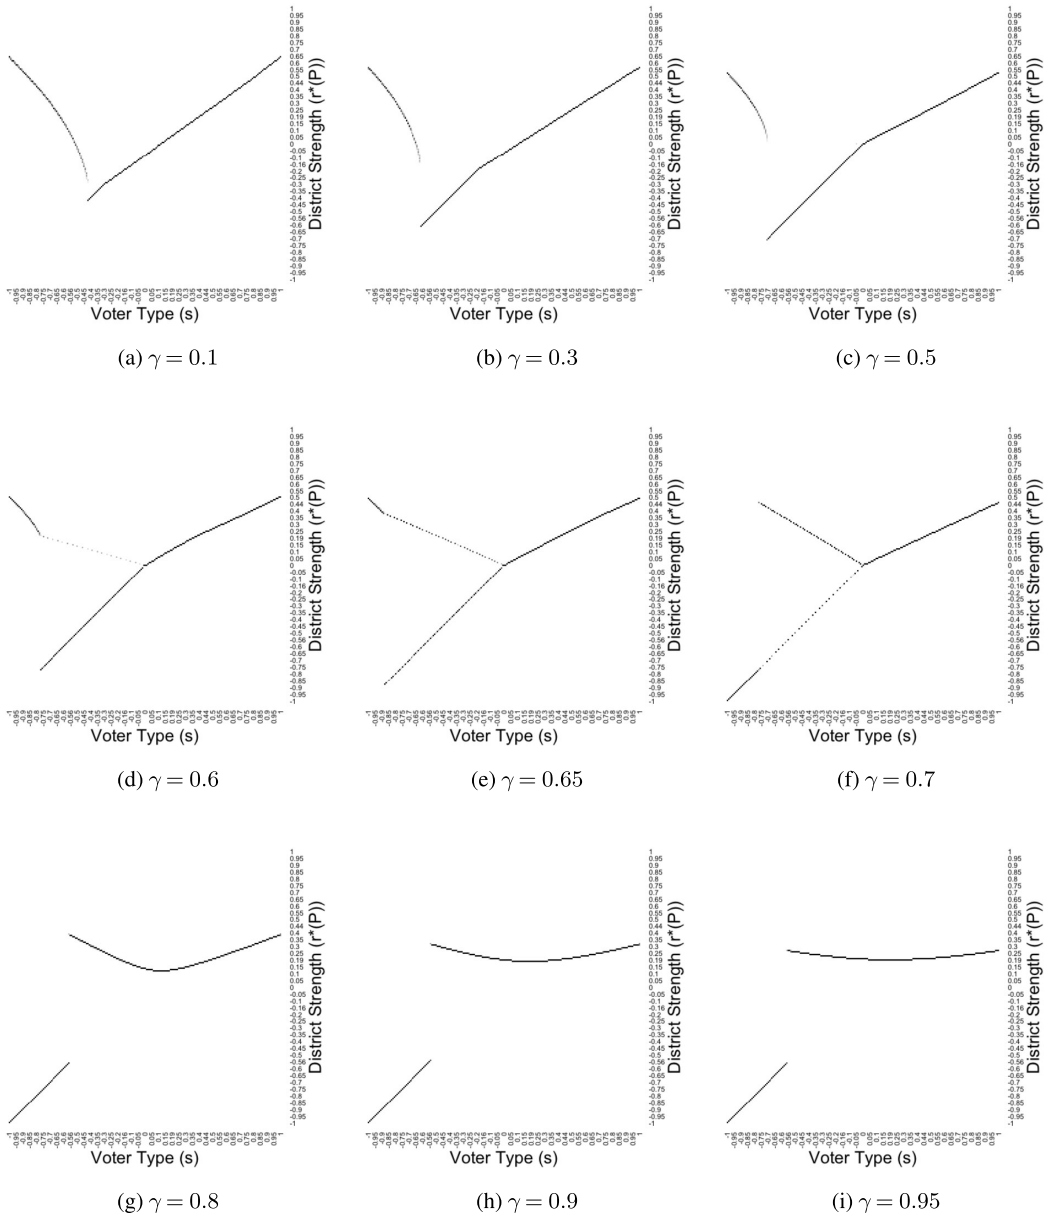


FIGURE 1.—Optimal districting with balanced voter support as share of idiosyncratic uncertainty varies. *Notes:* The optimal districting plan is SMP for $\gamma \in \{0.1, 0.3, 0.5\}$, Y-districting for $\gamma \in \{0.6, 0.65, 0.7\}$ (and, specifically, mixed SMP for $\gamma \in \{0.6, 0.65\}$ and mixed SOP for $\gamma = 0.7$), and SOP for $\gamma \in \{0.8, 0.9, 0.95\}$. Our empirical estimates of γ in Section 5 are above 0.96 for all U.S. states.

plans (i.e., Y-districting, the middle row of the figure), the shading intensity indicates the probability that a voter type is assigned to each district. We see that optimal districting takes the conjectured form: SMP is optimal for $\gamma \in \{0.1, 0.3, 0.5\}$, Y-districting is optimal for $\gamma \in \{0.6, 0.65, 0.7\}$, and SOP is optimal for $\gamma \in \{0.8, 0.9, 0.95\}$. The highest value of γ in the figure, $\gamma = 0.95$, is the value closest to our empirical estimates. When $\gamma = 0.95$, SOP

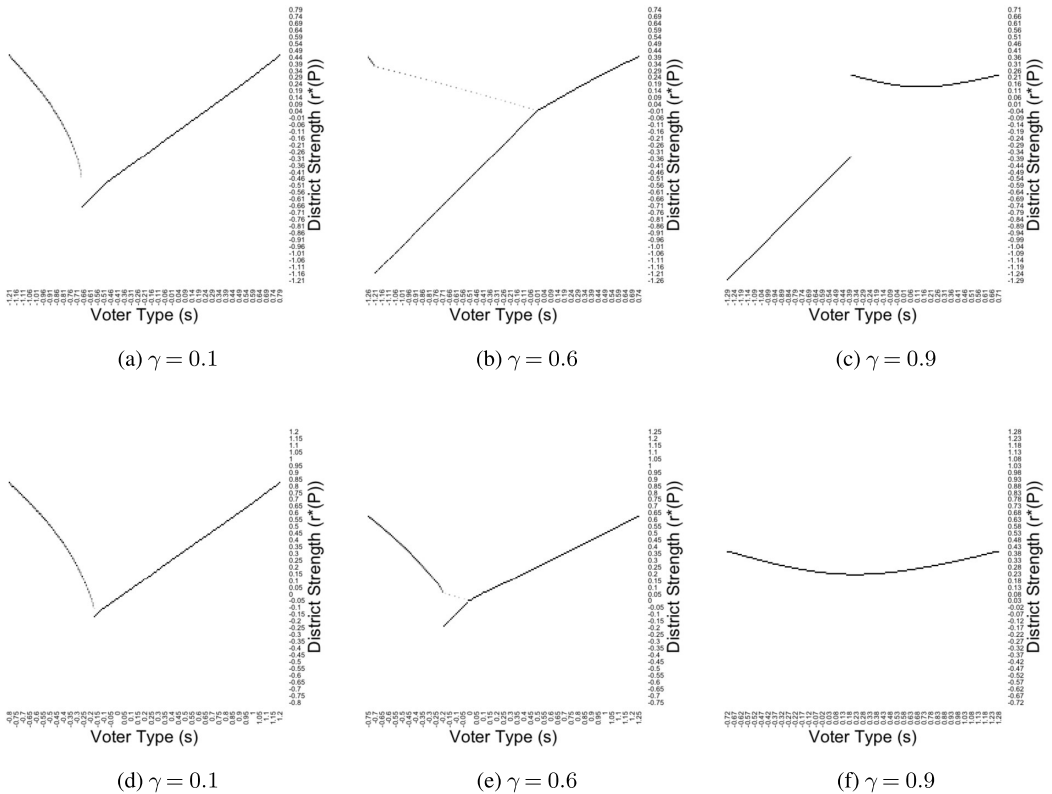


FIGURE 2.—Optimal districting with imbalanced voter support as share of idiosyncratic uncertainty varies. *Notes:* In the top and bottom panels, the designer’s expected vote share is 40% and 60%, respectively.

remains optimal but now closely resembles p -segregation (where unfavorable voters are segregated and favorable voters are pooled). Thus, for what we will see is the empirically relevant parameter range, p -segregation is approximately optimal.

Figure 2 illustrates optimal districting for the same parameters as Figure 1, except that now voter types are uniform on $[x - 1, x + 1]$, where x is scaled to give an expected vote share of 40% (top panels) or 60% (bottom panels). The figure shows that a less popular designer segregates more unfavorable voters, while a more popular designer pools more voters. The last panel shows that NAD (approximating uniform districting) is optimal for a designer with a 60% expected vote share and $\gamma = 0.9$.

Figure 3 illustrates the form of optimal districting as a function of the designer’s expected vote share and γ . The figure continues to assume that J is standard normal and F is uniform on $[x - 1, x + 1]$, where x is scaled so that the designer’s expected vote share ranges from 0 to 1. The figure shows that segregation is optimal for an unpopular designer (unless aggregate uncertainty dominates), NAD is optimal for a popular one, and optimal districting ranges from SMP to Y-districting to SOP as γ ranges from 0 to 1 with balanced voter support. These results match Theorems 2–5.³⁸

³⁸Due to numerical error, it is difficult to confidently classify optimal plans within one or two grid points of the boundaries between the regions where different plans are optimal in Figure 3. (By continuity, plans of




FIGURE 3.—Optimal districting as designer’s popularity and share of idiosyncratic uncertainty vary. *Notes:* Each U.S. state is located at its point estimate in Table I in Section 5.

Figure 3 also plots our point estimates of a Republican designer’s expected vote share and γ for every U.S. state. (The data and estimation procedure are described in the next section.) The most important observation is that γ is close to 1 in every state: the mean estimate of γ is 0.986, and the lowest estimate (for North Carolina) is 0.962. These estimates are all far above the cutoff of 0.732 above which SOP is optimal with balanced voter support. Thus, NAD (approximating uniform districting) is optimal for a Republican designer in Republican states like Oklahoma and Louisiana, while SOP (approximating pack-and-crack) is optimal for a Republican designer in swing states like Michigan and North Carolina, as well as in Democratic states like New York and Maryland (in the fanciful event that the Republicans found themselves controlling districting in such states).³⁹

different forms are both approximately optimal near the boundary.) The boundaries should thus be viewed as approximations.

³⁹A caveat is that Figure 3 is a two-dimensional plot and thus neglects heterogeneity in the variance of s across states, which we also estimate. It turns out that assuming that the variance of s is $1/\sqrt{3} \approx 0.577$ in all states—which is how we construct Figure 3—yields the correct classification of optimal districting for every state except Hawaii, where Republican-optimal districting is actually segregation (see Table I in Section 5).

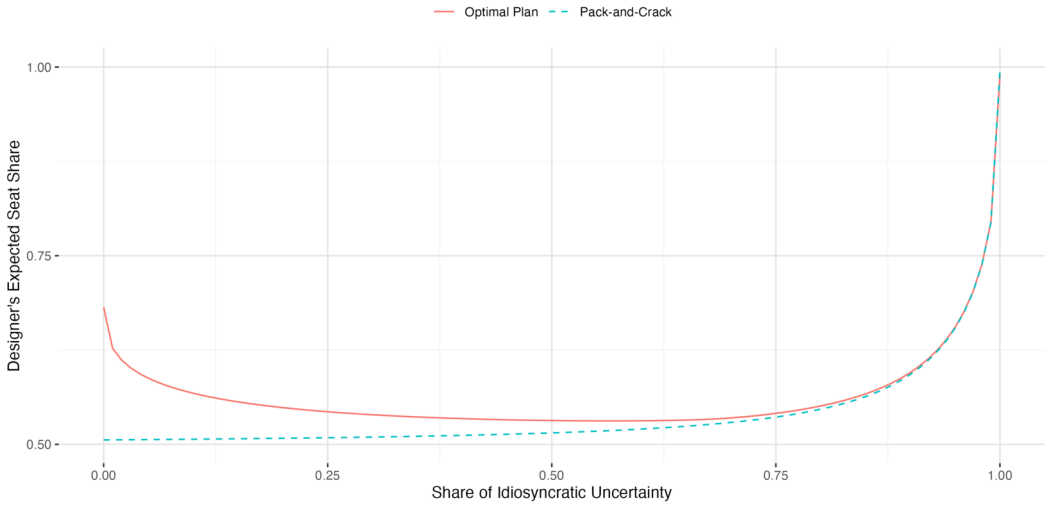


FIGURE 4.—Expected seat share under optimal districting and optimal pack-and-crack districting..

REMARK 2—Approximate Optimality of Pack-and-Crack: Lemma 11 in Appendix A.6 shows that uniform districting (for a designer with majority support) or pack-and-crack districting (for a designer with minority support) is approximately optimal with small aggregate uncertainty. The intuition is simple: with small aggregate uncertainty, a designer with minority support can obtain an expected seat share of approximately $1 - F(s^*(0))$ by creating slightly fewer than $1 - F(s^*(0))$ identical districts each with an expected vote share slightly greater than $1/2$, and $1 - F(s^*(0))$ is the optimal expected seat share in the limit. Indeed, pack-and-crack districting is approximately optimal for realistic parameters. For the same parameters as in Figure 1, Figure 4 plots the expected seat share under the optimal pack-and-crack plan and under the unconstrained optimal plan. The figure shows that the unconstrained expected seat share never exceeds the pack-and-crack expected seat share by more than 0.1% for any value of γ above 0.95. (Recall that our lowest estimate of γ for any U.S. state is 0.962.) We also estimate that the maximum loss from pack-and-crack relative to optimal districting in any U.S. state (accounting for unbalanced voter support) is 0.56% (see Table I in Section 5).⁴⁰

5. ESTIMATION

We have seen that the form of optimal districting depends on the designer's expected vote share and the parameter γ (the share of idiosyncratic uncertainty). We estimate these parameters using precinct-level returns from recent U.S. House elections.

5.1. Data and Empirical Model

Our data are the precinct-level returns from the U.S. House elections in 2016, 2018, and 2020, which were recently standardized by Baltz et al. (2022). For each precinct n and

⁴⁰The maximum loss is attained by Rhode Island, a Democratic state where Republicans are very unlikely to ever control districting. If we exclude the Democratic states of Rhode Island, Massachusetts, and Maine, the maximum estimated loss from pack-and-crack relative to optimal districting in any U.S. state is 0.09%.

election year $y \in \{2016, 2018, 2020\}$, we observe the total two-party vote k_{ny} and the share of the two-party vote for the Republican candidate v_{ny} .⁴¹ The data are a repeated cross-section rather than a panel, because there is no general way to match precincts across elections (Baltz et al. (2022), p. 6). We drop all districts with an uncontested House race in any of 2016, 2018, or 2020 (which drops 25% of all districts).⁴² This results in dropping South Dakota and Vermont, as in these states the single at-large district was uncontested in 2020 and 2016, respectively; we also drop Pennsylvania, as it was redistricted between 2016 and 2018. For each election, we also drop precincts with fewer than 50 total votes (which drops 0.14% of all votes) or where the Republican vote share is 0 or 1 (which drops an additional 0.015% of votes).

To take the model to these data, we assume that s indexes precincts, so that $Q(s - r)$ is the designer's vote share in a type s precinct at aggregate shock r . Formally, this is equivalent to assuming that all voters in a precinct have the same type. (As we clarify below, this does not mean that all voters in a precinct vote the same way.) We also assume that precincts are relatively large (in the data, the mean precinct vote count is 794 with standard deviation 1434, after dropping precincts with fewer than 50 total votes or a 0 or 1 vote share), and idiosyncratic voter taste shocks are normally distributed with variance γ .⁴³ By the law of large numbers, this implies that the designer's vote share in a precinct n with type s_n in district d and election y is given by

$$Q(s_n - r_{dy}) = \Phi\left(\frac{s_n - r_{dy}}{\sqrt{\gamma}}\right), \quad (6)$$

where Φ is the standard normal distribution. To see this, recall that each voter i in precinct n votes for the designer's party iff $s_n \geq r_{dy} + t_{iy}$, where t_{iy} is the voter's normally distributed idiosyncratic taste shock with variance γ , and hence votes for the designer's party with probability $\Phi((s_n - r_{dy})/\sqrt{\gamma})$.⁴⁴ Applying the law of large numbers at the precinct level gives (6).

We emphasize that this empirical model does not allow precinct-level aggregate shocks: the vote share $Q(s_n - r_{dy})$ in precinct n in district d and election y is given by (6), which is a deterministic function of the persistent precinct type s_n and the district-level aggregate shock in election y , r_{dy} .

⁴¹A "precinct" is the smallest election-reporting unit in a state, which typically corresponds to a geographic area where all voters vote at the same polling place. Maine and New Jersey report election returns only at the township level, so for these states n indexes townships rather than precincts. For some elections where a nominally third-party candidate runs in place of an official Democratic or Republican candidate, we manually re-label this candidate as a Democrat or Republican. For example, in New York, we re-assign Working Families Party candidates as Democrats and re-assign Conservative Party candidates as Republicans. Throughout, we focus on the two-party vote k_{ny} and the Republican share of the two-party vote v_{ny} , ignoring third parties.

⁴²Keeping these districts would bias our estimate of γ , because the relevant vote shares are for contested elections, and if these districts were contested their vote shares would be different from 0 or 1. Keeping a district with one or two uncontested elections only for the elections where it is contested would also bias our estimate of γ , by distorting the estimated swing across elections. Dropping uncontested districts does likely bias our estimate of the voter type distribution F , as uncontested districts are presumably more extreme; however, this bias is irrelevant for our main goal of estimating γ .

⁴³Our estimates are not sensitive to assuming normality: because we will find that γ is close to 1, the taste shock distribution is approximately uniform over the relevant range, so specifying any smooth taste shock distribution leaves our estimates almost unchanged. For example, our point estimate of γ for the U.S. as a whole is 0.986 with normal taste shocks, 0.987 with logistic taste shocks, 0.989 with Laplace taste shocks, and 0.981 with uniform taste shocks. See Appendix C.

⁴⁴In this section, as in Section 4.3, we assume that $Q(t) = \Phi(t/\sqrt{\gamma})$ and $G(r) = \Phi(r/\sqrt{1-\gamma})$.

To interpret the assumption that all voters in a precinct have the same type, note that a voter’s type and taste shock enter only through their difference $s_n - t_{iy}$. If we call this difference the voter’s “preference,” our assumption is that voter preferences in precinct s_n are normally distributed with mean s_n and variance γ . Also, while voter preferences must be independent across voters in each district to justify (6), the correlation of each voter’s preference across elections is arbitrary. Thus, voters in a precinct can differ in their persistent tastes for the parties as well as in their election-specific tastes.

REMARK 3—What if Precincts Can be Split?: Our estimation assumes that the smallest “districtable unit” is a precinct, which is the smallest election-reporting unit in a state. In practice, the smallest districtable unit is usually not a precinct but a census block, which is the smallest geographic unit for which the U.S. Census tabulates complete data. Census blocks are usually much smaller than precincts. However, [Bouton, Genicot, Castanheira, and Stashko \(2024\)](#) report that only 2% of precincts are split across proposed congressional districts in their sample. In addition, in Section 6.2, we redo our estimation under the assumption that designers can only assign counties rather than precincts and find that this increases our estimate of γ by only about 0.001. The difference in size between a county and a precinct is roughly similar to that between a precinct and a census block: there are around 50 times as many precincts as counties in the United States, and around 50 times as many census blocks as precincts. This suggests that our estimates are reasonably robust to letting designers split precincts.

However, precincts (and even census blocks) sometimes are split, and some designers strive to split them as finely as possible. For example, in *Dickson v. Rucho* (2014), plaintiffs alleged that a Republican-drawn map in North Carolina “divides 563 of the state’s 2692 precincts into more than 1400 sections” ([Newkirk \(2017\)](#)). If designers can split precincts extremely finely, so as to isolate individual voters or very small groups of voters, this could substantially affect our estimate of γ and our conclusions about the form of gerrymandering. We are not aware of evidence that fine precinct-splitting or census block-splitting is widespread, but we acknowledge that, to the extent that this is the case or may become so in the future, our analysis would have to be redone at the level of the smallest districtable unit (subject to data limitations, as our estimation already uses the finest currently available data).

5.2. Estimates

We estimate the key parameter γ , as well as the other parameters. Since districting plans in the U.S. are drawn at the state level, we estimate parameters separately for each U.S. state.⁴⁵ We assume that aggregate shocks are jointly normally distributed across districts and independent across elections, so that the variance of r_{dy} is $1 - \gamma$; the correlation between r_{dy} and $r_{d'y}$ is ρ for each $d \neq d'$ and y ; and the correlation between r_{dy} and $r_{d'y'}$ is 0 for each d, d' , and $y \neq y'$. Recall that the results in Section 4.3 show that, with balanced voter support, SMP is optimal if $\gamma \leq 0.5$, Y-districting is optimal if $\gamma \in (0.5, 0.732)$, and SOP is optimal if $\gamma \geq 0.732$. Thus, a key question of interest is which of these three regions contains our estimate of γ .

We estimate γ for each state by method of moments. Recall that v_{ny} is the Republican share of the two-party vote in precinct n and election y . Let $w_{ny} = \Phi^{-1}(v_{ny})$, the corresponding standard normal quantile. Let $T = 3$ denote the number of elections, D the

⁴⁵While our model assumes a large number of districts, we estimate parameters for all states (including ones with only one congressional district) to give as complete parameter estimates as possible.

number of districts in the state, and \mathcal{N}_{dy} the set of precincts in district d and election y . Next, define

$$w_{dy} = \sum_{n \in \mathcal{N}_{dy}} k_{ny} w_{ny} / \sum_{n \in \mathcal{N}_{dy}} k_{ny} \quad \text{and} \quad w_{d\bullet} = \frac{1}{T} \sum_y w_{dy}.$$

That is, w_{dy} is the average value of w_{ny} over precincts in district d , weighted by the number of votes in each precinct; and $w_{d\bullet}$ is the average value of w_{dy} over elections y . In Appendix B, we show that a consistent estimator of γ is given by

$$\hat{\gamma} = 1 / \left(1 + \frac{1}{D(T-1)} \sum_{d,y} (w_{dy} - w_{d\bullet})^2 \right),$$

and we also construct a confidence interval for γ , as well as a point estimator of the correlation among the district-level aggregate shocks, and point estimators of the mean and standard deviation of the distribution of precinct types. (We discuss the robustness and goodness-of-fit of our empirical model in Appendix C.)

Table I displays the resulting estimates for each U.S. state, as well as for the state average weighted by the number of districts included in the analysis (row WS) and the U.S. as a whole (row US). The states are ordered by column v , the designers expected vote share in the districts included in the estimation. Columns D_T and D_A are the total number of districts and the number of districts included in the analysis.

Columns γ and $\underline{\gamma}$ are our point estimate and the lower bound of a 95% one-sided confidence interval for γ . The confidence interval is wide because we only have data from three elections: $T = 3$. However, it is clear that γ is far above the critical value of 0.732. The lowest point estimate for γ for any state is 0.962 in North Carolina, and the weighted mean estimate for γ and the estimate for γ for the U.S. as a whole are both 0.986. Moreover, even with $T = 3$, the lower bound of a 95% one-sided confidence interval is above 0.732 for all available states except North Dakota, where the lower endpoint is 0.619. If we expand our data set to include the returns from the 2012 and 2014 elections (thus covering all five congressional elections held under the 2010 districting plans), the lower endpoints of the 95% confidence interval exceed 0.732 for all states, including North Dakota.⁴⁶ Together with the results in Section 4.3 (including Figure 3, which accounts for imbalances in voter support), this provides strong evidence that optimal gerrymandering is given by SOP (for a designer with minority support) or NAD (for a designer with majority support) for realistic parameters. Moreover, our estimates for γ are high enough that the optimal SOP plan approximates p -segregation and the optimal NAD plan approximates uniform districting (recall Figures 1–4), and that pack-and-crack (with minority support) or uniform districting (with majority support) is approximately optimal.

Columns v and σ_s are the designer's expected vote share and the standard deviation of s . The latter estimates are similar to those in Figure 1. However, our estimates of v and σ_s may be biased by dropping uncontested elections (unlike our estimates of γ , which remain

⁴⁶Precinct-level returns for 2012 and 2014 have been compiled by [Ansolabehere, Palmer, and Lee \(2014\)](#) but are less complete and less standardized than the [Baltz et al. \(2022\)](#) data we use, which only cover 2016, 2018, and 2020. We have checked that all of our empirical results are robust to including the 2012 and 2014 data.

TABLE I
ESTIMATES.

US	D_T	D_A	γ	$\underline{\gamma}$	v	σ_s	σ_s^c	V	\underline{V}	V^c	\mathcal{H}	\mathcal{H}^D
HI	2	2	0.972	0.839	0.250	0.181	0.076	0.001	0.001	0.000	Seg	NAD
NY	27	19	0.966	0.937	0.342	0.826	0.659	0.416	0.415	0.356	SOP	NAD
MD	8	8	0.990	0.978	0.346	0.728	0.624	0.456	0.456	0.410	SOP	NAD
RI	2	1	0.990	0.833	0.375	0.302	0.259	0.199	0.194	0.145	SOP	NAD
CT	5	5	0.995	0.987	0.379	0.377	0.327	0.328	0.328	0.263	SOP	NAD
ME	2	1	0.992	0.866	0.385	0.311	0.304	0.246	0.244	0.236	SOP	NAD
MA	9	1	0.998	0.956	0.385	0.233	0.211	0.142	0.139	0.098	SOP	NAD
DE	1	1	0.990	0.836	0.397	0.488	0.268	0.456	0.456	0.228	SOP	NAD
IL	18	13	0.984	0.962	0.399	0.737	0.545	0.560	0.560	0.463	SOP	NAD
NJ	12	12	0.981	0.962	0.402	0.590	0.445	0.492	0.492	0.411	SOP	NAD
CA	53	35	0.992	0.987	0.412	0.483	0.337	0.508	0.508	0.382	SOP	NAD
NM	3	3	0.979	0.925	0.436	0.543	0.412	0.548	0.548	0.476	SOP	NAD
NH	2	2	0.997	0.960	0.463	0.263	0.256	0.575	0.575	0.568	SOP	NAD
NV	4	4	0.998	0.992	0.467	0.449	0.310	0.741	0.741	0.662	SOP	NAD
MN	8	8	0.987	0.973	0.470	0.436	0.342	0.626	0.626	0.569	SOP	SOP
CO	7	7	0.989	0.953	0.470	0.527	0.429	0.691	0.691	0.649	SOP	SOP
OR	5	4	0.987	0.961	0.471	0.498	0.377	0.662	0.662	0.603	SOP	SOP
VA	11	8	0.985	0.934	0.491	0.548	0.428	0.726	0.726	0.684	SOP	SOP
WA	10	5	0.987	0.960	0.496	0.375	0.264	0.689	0.689	0.624	SOP	SOP
MI	14	13	0.990	0.980	0.501	0.596	0.503	0.803	0.802	0.772	SOP	SOP
GA	14	7	0.985	0.959	0.509	0.718	0.525	0.821	0.820	0.781	SOP	SOP
TX	36	23	0.989	0.978	0.514	0.645	0.472	0.841	0.841	0.811	SOP	SOP
IA	4	4	0.986	0.949	0.519	0.372	0.270	0.763	0.763	0.726	SOP	SOP
NC	13	11	0.962	0.933	0.526	0.560	0.373	0.740	0.739	0.699	SOP	SOP
AZ	9	6	0.990	0.974	0.537	0.402	0.286	0.868	0.868	0.852	SOP	SOP
FL	27	20	0.994	0.987	0.545	0.444	0.291	0.949	0.948	0.948	NAD	SOP
OH	16	16	0.984	0.967	0.552	0.635	0.469	0.908	0.908	0.896	NAD	SOP
AK	1	1	0.996	0.922	0.554	0.396	0.298	0.987	0.987	0.983	NAD	SOP
MT	1	1	0.993	0.884	0.556	0.490	0.325	0.973	0.973	0.973	NAD	SOP
SC	7	7	0.994	0.988	0.559	0.622	0.402	0.990	0.990	0.988	NAD	SOP
AR	4	1	0.985	0.773	0.566	0.629	0.391	0.950	0.949	0.945	NAD	SOP
NE	3	2	0.990	0.945	0.575	0.446	0.297	0.981	0.981	0.979	NAD	SOP
KY	6	4	0.991	0.968	0.584	0.548	0.408	0.994	0.994	0.992	NAD	SOP
MO	8	8	0.995	0.981	0.584	0.702	0.579	1.000	1.000	1.000	NAD	SOP
KS	4	4	0.978	0.905	0.598	0.463	0.355	0.968	0.968	0.971	NAD	SOP
IN	9	8	0.983	0.963	0.608	0.524	0.351	0.991	0.991	0.990	NAD	SOP
WI	8	5	0.989	0.970	0.617	0.301	0.228	0.998	0.998	0.999	NAD	SOP
AL	7	2	0.971	0.846	0.624	0.674	0.409	0.988	0.988	0.992	NAD	SOP
WV	3	3	0.971	0.883	0.646	0.340	0.252	0.989	0.989	0.989	NAD	SOP
UT	4	4	0.989	0.947	0.647	0.585	0.472	1.000	1.000	1.000	NAD	SOP
TN	9	8	0.991	0.976	0.650	0.691	0.526	1.000	1.000	1.000	NAD	SOP
MS	4	2	0.993	0.934	0.672	0.671	0.340	1.000	1.000	1.000	NAD	SOP
ID	2	2	0.987	0.930	0.673	0.462	0.365	1.000	1.000	1.000	NAD	SOP
OK	5	4	0.983	0.922	0.685	0.454	0.318	1.000	1.000	1.000	NAD	SOP
ND	1	1	0.969	0.619	0.696	0.426	0.335	0.999	0.999	0.999	NAD	SOP
WY	1	1	0.990	0.835	0.701	0.478	0.375	1.000	1.000	1.000	NAD	SOP
LA	6	4	0.974	0.898	0.725	0.595	0.288	1.000	1.000	1.000	NAD	SOP
WS	18	13	0.986	0.959	0.497	0.561	0.415	0.755	0.754	0.709	SOP	SOP
US	417	311	0.986	0.979	0.497	0.643	0.508	0.777	0.776	0.745	SOP	SOP

unbiased after dropping any set of districts).⁴⁷ Column σ_s^c is the standard deviation of s across counties rather than precincts. We discuss county-level estimates in Section 6.2.

Columns V and \underline{V} are the designer’s expected seat share under optimal unconstrained districting and optimal pack-and-crack districting, respectively. As illustrated in Figure 4, the shares are very similar. Column V^c is the expected seat share under optimal districting where the designer assigns counties rather than precincts; see Section 6.2.

Finally, Column \mathcal{H} is the form of the optimal districting plan for the estimated parameters γ , v , and σ_s , and Column \mathcal{H}^D is the form of the optimal districting plan with a Democratic districter (i.e., at the estimated parameters γ and σ_s but with $1 - v$ in place of v). The optimal plans are calculated assuming a continuum of districts in each state and assuming that F is uniform and G and Q are normal, so they reflect our state-specific estimates of γ , v , and σ_s but not the state-specific number of districts or state-specific features of the distribution of precinct types other than its mean and variance.⁴⁸ We estimate that if Republicans somehow found themselves in charge of districting Hawaii (or, more accurately, a hypothetical large state with the same values of γ , v , and σ_s as Hawaii), they would segregate the state. Otherwise, SOP is optimal (and pack-and-crack is approximately optimal) in states where the expected vote share for the designer’s party is less than 54%, and NAD is optimal (and uniform districting is approximately optimal) in states where the expected vote share for the designer’s party is greater than 54%.⁴⁹ This reflects the fact that, for our estimates of γ and F , the optimal pack-and-crack plan creates cracked districts where the designer’s expected vote share is around 54%. This feature matches well with the conventional wisdom that gerrymanderers typically target a vote share of around 55% in districts they intend to win (e.g., [Li and Leaverton \(2022\)](#), [Wasserman \(2023\)](#)).

6. DISCUSSION: WHY DOES THE FORM OF GERRYMANDERING MATTER?

We briefly discuss potential political and legal implications of our results. We consider three areas: implications for how regulations and restrictions on districting affect partisan representation; implications for how gerrymandering affects political competition and polarization; and implications for detecting and measuring gerrymandering.

6.1. *Effects of Districting Reform on Seat Shares I: Majority-Minority Districts*

The key U.S. federal laws regulating gerrymandering are the Equal Protection Clause of the Fourteenth Amendment and the Voting Rights Act of 1965. These laws have been interpreted as not only prohibiting adverse racial gerrymandering, but also as affirmatively requiring states to create electoral districts where racial or ethnic minority voters form either a majority (a so-called “majority-minority district”) or a large enough minority so as to have a strong opportunity to elect their candidate of choice (often called a “minority opportunity district”; e.g., [Canon \(2022\)](#)). The creation of such districts played a significant role in increasing Black representation in state legislatures and the U.S. Congress

⁴⁷We also estimate the correlation ρ among the district-level aggregate shocks to be 0.317 (at the country level). Since this estimate is not close to either 0 or 1, estimating a simpler empirical model where district-level shocks are either uncorrelated or perfectly correlated would yield biased estimates of γ .

⁴⁸The form of the estimated optimal plans is not sensitive to the assumption that F is uniform and G and Q are normal.

⁴⁹For a Democratic districter, the threshold expected vote share is 53%. The thresholds should be viewed as approximations, as it is difficult to confidently classify optimal plans near the threshold due to numerical error.

from the 1970s onward, especially in the South (Grofman and Handley (1991), Cox and Holden (2011)). However, the overall partisan impact of majority-minority and minority opportunity districts has long been contested, with some observers arguing that these districts effectively pack strong Democratic supporters and thus resemble a component of a Republican-optimal districting plan. This issue came to a head following the 1994 Republican takeover of the U.S. House, which many journalists and political scientists blamed in part on the creation of majority-minority districts in the 1990 redistricting cycle (but see Cox and Holden (2011), Washington (2012)).

Previous studies have observed that the impact of a requirement to create majority-minority or minority opportunity districts on overall partisan representation hinges on the form of optimal gerrymandering. The conventional view in the 1990s (what Cox and Holden (2011) call the “pack-and-crack consensus”) was that optimal gerrymandering packs opponents, and hence that a requirement to create majority-minority districts that pack strong Democratic supporters increases Republican representation.⁵⁰ Shotts (2001) adds an important caveat by noting that, since uniform districting is optimal for a designer with majority support (without aggregate uncertainty), majority-minority mandates hurt Republican designers in strongly Republican states. More dramatically, building on the results of Friedman and Holden (2008), Cox and Holden (2011) challenge the pack-and-crack consensus by arguing that optimal districting is given by NAD, and thus packs moderates rather than opponents. Since NAD does not create districts packed with strong Democratic supporters, Cox and Holden argue that a requirement to create such districts precludes NAD and therefore reduces Republican representation.

Our results contribute to this debate as follows. Cox and Holden’s argument that NAD is optimal in practice implicitly assumes that the low-idiosyncratic-uncertainty case studied by Friedman and Holden (2008) is representative. For example, Cox and Holden write, “In a world with diverse voter types, however, there is no plausible distribution of African American voters that would make it optimal for Republican redistricting authorities to create districts in which African Americans make up a supermajority of voters. Within the model, packing one’s opponents is never the optimal strategy” (p. 574). We instead show that, empirically, idiosyncratic uncertainty is much larger than aggregate uncertainty, and that this implies that packing opponents is optimal for a designer with minority voter support, while NAD is optimal for a designer with majority support. Thus, majority-minority mandates can increase Republican representation in closely divided states where SOP is optimal and pack-and-crack is approximately optimal (as in the pack-and-crack consensus), but are likely to decrease Republican representation in strongly Republican states where NAD is optimal and uniform districting is approximately optimal (as argued by Shotts (2001) in a model without aggregate uncertainty). Overall, our analysis gives a conclusion similar to that of Shotts (2001) and quite different from that of Cox and Holden (2011).

6.2. *Effects of Districting Reform on Seat Shares II: Respecting Political Subdivisions*

Among the practical restrictions on districting beyond equipopulation, one that is amendable to our analysis is a requirement not to split counties or other political subdivisions. Preserving counties or other subdivisions is one of the six traditional redistricting criteria according to the National Conference of State Legislators and is currently

⁵⁰Minority opportunity districts may or may not raise similar issues, depending on the share of strong Democratic supporters in these districts (Lublin, Handley, Brunell, and Grofman (2020)).

required in 29 of the 50 U.S. states.⁵¹ From the perspective of partisan gerrymandering, a requirement to preserve counties constrains the designer to choose among a coarser set of districting plans, where counties rather than census blocks or precincts become the object of districting.

We can assess the impact of a requirement to preserve counties by re-running our estimation of γ and F , taking the unit of districting as counties rather than precincts. Our estimates of γ are similar in both cases but are slightly higher with counties, because precinct vote shares swing slightly more from election to election than county vote shares: our mean precinct-based estimate of γ is 0.986, while our mean county-based estimate is 0.987. More importantly, our estimate of the standard deviation of F is considerably smaller with counties: the mean precinct-based estimate is 0.561, while the mean county-based estimate is 0.415. This gap is the key consequence of constraining the designer to assigning coarser units. Finally, this constraint significantly affects the designer's optimal expected seat share in closely divided states where SOP is optimal, as now fewer highly unfavorable units can be packed; however, it has only a small effect on the optimal seat share in states where the designer has strong support and NAD is optimal, as uniform districting (which is unaffected by a requirement to preserve counties) is approximately optimal in these states. In particular, our estimate of the reduction in a Republican designer's seat share from requiring him to preserve counties ranges from essentially 0 in strongly Republican states to 23% in Delaware, with a weighted average across states of 4.5%.⁵²

6.3. *Effects of Gerrymandering on Political Competition and Polarization*

An important debate concerns the impact of gerrymandering on the intensity of electoral competition (e.g., the fraction of “competitive” districts or the extent of incumbency advantage) and political polarization. Popular discourse often blames gerrymandering for reducing competition and increasing polarization (but see [Gelman and King \(1994a\)](#), [Abramowitz, Alexander, and Gunning \(2006\)](#), [McCarty, Poole, and Rosenthal \(2009\)](#), [Friedman and Holden \(2009\)](#)). Regardless of the size of the overall effects of gerrymandering on competition and polarization, the nature of these effects depends on the form of gerrymandering. In particular, under SOP, intra-district polarization is relatively low while inter-district polarization is relatively high; while under NAD or SMP, intra-district polarization is high and inter-district polarization is low. To see this, note that, with a right-wing designer, SOP or pack-and-crack creates a few strongly left-leaning districts and many slightly right-leaning districts, with a gap between the left-leaning and right-leaning districts: formally, there is a gap between the highest value of $r^*(P)$ among segregated districts and the lowest value of $r^*(P)$ among paired districts (see the last three panels in Figure 1). SOP also involves relatively low intra-district polarization within each district, since the lowest voter types in paired districts are “moderates” rather than extreme left-wingers. In contrast, NAD or SMP creates a continuum of districts ranging

⁵¹The other criteria are compactness, contiguity, preservation of communities of interest, preservation of the “cores” of previous districts, and avoiding incumbent pairing (<https://www.ncsl.org/elections-and-campaigns/2020-redistricting-criteria>).

⁵²Our weighted mean estimate of a Republican designer's optimal seat share is 70.9% under county-based districting and 75.4% under precinct-based districting. A limitation of this comparison is that our assumption that the designer assigns a continuum of units is more accurate when units are precincts rather than counties. The omitted integer constraint would bind more harshly for county-based districting, which biases our estimates of the designer's loss from being restricted to assigning counties downward.

from left-leaning to right-leaning—formally, the set $\{r : r = r^*(P) \text{ for some } P \in \text{supp}(\mathcal{H})\}$ is an interval (see the first three panels in Figure 1)—with less extreme left-leaning districts than under SOP. NAD or SMP also involves greater intra-district polarization than SOP, in that the maximum range of voter types that are pooled together under SMP is greater than under SOP.

Our model does not encompass endogenous political responses to districting, such as effects of districting on which politicians run for office and on what platforms. With this caveat, we can draw some tentative implications of the above features of optimal districting for political competition and polarization. First, since the distribution of district strengths $r^*(P)$ has a gap under SOP or pack-and-crack but not under NAD or SMP, SOP or pack-and-crack may lead to more polarized legislatures, where the packed districts elect left-wing representatives and the cracked districts elect right-leaning ones. Indeed, the possibility that packing opponents can increase polarization in this manner is a long-standing concern (e.g., [Cox and Holden \(2011\)](#), p. 595). In contrast, NAD or SMP may lead to less polarized legislatures. Second, SOP or pack-and-crack may produce more “uncompetitive,” far-left districts. Creating uncompetitive districts is usually viewed as a socially undesirable feature of a districting plan, but see [Buchler \(2005\)](#) and [Brunell \(2008\)](#) for opposing views. Finally, lower intra-district polarization under SOP or pack-and-crack may be socially desirable if voters benefit from being ideologically close to their representative, as in [Besley and Preston \(2007\)](#) and [Gomberg, Pance, and Sharma \(2023\)](#). These and other implications of optimal districting for political processes and outcomes can be studied more fully in models that endogenize additional aspects of political competition.

6.4. *Detecting and Measuring Gerrymandering*

A large literature proposes metrics that attempt to detect and measure gerrymandering. Most metrics compare a party’s seat share and its vote share, with a high seat share viewed as indicative of gerrymandering.⁵³ However, a limitation of this approach is that the acceptable range of seat shares for a given vote share is unclear. Indeed, the Supreme Court has objected that this class of metrics encodes a form of proportionality between seat and vote shares: as Justice Roberts wrote in *Rucho v. Common Cause*, “Partisan gerrymandering claims rest on an instinct that groups with a certain level of political support should enjoy a commensurate level of political power and influence. Such claims invariably sound in a desire for proportional representation, but the Constitution does not require proportional representation.”

Our results suggest an alternative test for gerrymandering that compares vote shares across districts, rather than comparing seat and vote shares. A novel and robust prediction of our analysis is that, in the realistic case with small aggregate uncertainty, optimal plans make favorable districts equally strong: a designer with majority support creates equally strong districts under NAD or uniform districting, while a designer with minority support creates packed districts that are lost with high probability and equally strong cracked districts that are won with high probability under SOP or pack-and-crack. In contrast, one would not expect favorable districts to be equally strong under non-gerrymandered districting. Thus, a proposed test for gerrymandering is whether a districting plan displays

⁵³Such measures include the partisan bias ([King and Browning \(1987\)](#)), efficiency gap ([Stephanopoulos and McGhee \(2015\)](#)), mean-median gap ([Wang \(2016\)](#)), and declination ([Warrington \(2018\)](#)). An alternative approach relies on statistical analysis of an ensemble of simulated maps ([Deford, Duchin, and Solomon \(2021\)](#)).

an unusually low variance in vote shares among districts won by the designer’s party. This test can be operationalized in future work.

7. CONCLUSION

This paper has developed a simple and general model of optimal partisan gerrymandering. Our main message has four parts. First, optimal districting is “segregate-pair”: weak districts are segregated; strong districts are paired. Second, the optimal form of segregate-pair districting depends on the gerrymanderer’s popularity and—more subtly—the relative amounts of aggregate and idiosyncratic uncertainty facing the gerrymanderer. Packing opposing voters is optimal when idiosyncratic uncertainty dominates, while packing moderate voters is optimal when aggregate uncertainty dominates. Third, empirically, idiosyncratic uncertainty dominates, implying that segregate-opponents-and-pair (SOP) districting is optimal for a designer with minority support, while negative assortative districting (NAD) is optimal for a designer with majority support. This finding also establishes that the relevant parameter range for future research on gerrymandering (and electoral competition more generally) is that where aggregate uncertainty is much smaller than idiosyncratic uncertainty. Fourth, estimated aggregate uncertainty is so small that a simple pack-and-crack plan is approximately optimal for a designer with minority support, while uniform districting is approximately optimal for a designer with majority support. This last observation helps rationalize the observed use of simple districting plans.

Methodologically, we develop and exploit a tight connection between gerrymandering and information design. We show that a general model of partisan gerrymandering is equivalent to a general Bayesian persuasion problem where the state of the world and the receiver’s action are both one-dimensional and the sender’s preferences are state-independent. This common framework nests the important prior contributions of [Owen and Grofman \(1988\)](#), [Friedman and Holden \(2008\)](#), and [Gul and Pesendorfer \(2010\)](#), and facilitates a more general and realistic analysis that allows diverse voter types and non-linear vote swings without restricting the relative amounts of aggregate and idiosyncratic uncertainty.

We hope our model can inform future research on various aspects of redistricting. We mention a few directions for future research.

First, we have assumed that the designer maximizes his party’s expected seat share. It may be more realistic to assume that the designer’s utility is non-linear in seat shares, for example due to a premium on winning a majority of seats. We examined this case in an earlier version of this paper ([Kolotilin and Wolitzky \(2020\)](#)). While non-linear designer utility introduces new complications, the extreme case where the designer simply maximizes the probability of winning a majority is straightforward: here, optimal districting maximizes seats conditional on the threshold aggregate shock at which the designer is barely able to attain a majority, and hence reduces to optimal districting without aggregate uncertainty.⁵⁴

Second, we have assumed that all voters always vote, or at least always vote at the same rate (as is equivalent). It would be interesting to incorporate heterogeneous turnout. Recently, [Bouton, Genicot, Castanheira, and Stashko \(2024\)](#) consider voters with a binary partisan type (as in [Owen and Grofman \(1988\)](#)), uniform aggregate shocks, and a continuous “turnout type,” which captures exogenous turnout heterogeneity across voters. It is

⁵⁴Among other results, [Kolotilin and Wolitzky \(2020\)](#) also show that, when G is linear, p -segregation is optimal for a designer with an S-shaped utility $W(x)$ for winning measure $x \in [0, 1]$ of seats.

promising to explore mutual generalizations of our models that allow more general forms of aggregate uncertainty as well as heterogeneous turnout. An alternative model, which captures endogenous turnout heterogeneity, would retain one-dimensional voter types but assume that voters abstain when they are close to indifferent between the parties. It is interesting to compare these models, as in practice turnout heterogeneity has both exogenous sources (e.g., education, race) and endogenous ones (e.g., almost-indifferent voters turn out less).

Finally, further questions include: What does the model imply for political competition and policy choices? What are the model's comparative statics—for example, what factors determine the proportion of packed and cracked districts?⁵⁵ And, what does the model imply about how gerrymandering should be measured and regulated? Understanding the form of optimal partisan gerrymandering can contribute to the study of such questions.

REFERENCES

- Abramowitz, Alan I., Brad Alexander, and Matthew Gunning (2006), “Incumbency, Redistricting, and the Decline of Competition in US House Elections.” *Journal of Politics*, 68, 75–88. [0098]
- Ansolabehere, Stephen, Maxwell Palmer, and Amanda Lee (2014), “Precinct-Level Election Data.” Harvard Dataverse. [0094]
- Bagnoli, Mark and Ted Bergstrom (2005), “Log-Concave Probability and Its Applications.” *Economic Theory*, 26 (2), 445–469. [0076]
- Baltz, Samuel et al. (2022), “American Election Results at the Precinct Level.” *Scientific Data*, 9 (651). [0091, 0092,0094]
- Besley, Timothy and Ian Preston (2007), “Electoral Bias and Policy Choice: Theory and Evidence.” *Quarterly Journal of Economics*, 122, 1473–1510. [0075,0099]
- Bickerstaff, Steve (2020), *Election Systems and Gerrymandering Worldwide*. Springer. [0071]
- Bouton, Laurent, Garance Genicot, Miguel Castanheira, and Allison Stashko (2024), “Pack-Crack-Pack: Gerrymandering With Differential Turnout.” Working Paper. [0075,0093,0100]
- Bracco, Emanuele (2013), “Optimal Districting With Endogenous Party Platforms.” *Journal of Public Economics*, 104, 1–13. [0075]
- Brunell, Thomas (2008), “Redistricting and Representation: Why Competitive Elections Are Bad for America.” Routledge. [0099]
- Buchler, Justin (2005), “Competition, Representation, and Redistricting: The Case Against Competitive Congressional Districts.” *Journal of Theoretical Politics*, 17, 431–463. [0099]
- Calvert, Randall L. (1985), “Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence.” *American Journal of Political Science*, 29 (1), 69–95. [0086]
- Canon, David (2022), “Race and Redistricting.” *Annual Review of Political Science*, 25, 509–528. [0096]
- Chambers, Christopher P. and Alan D. Miller (2010), “A Measure of Bizarreness.” *Quarterly Journal of Political Science*, 5, 27–44. [0075]
- Coate, Stephen and Brian Knight (2007), “Socially Optimal Districting: A Theoretical and Empirical Exploration.” *Quarterly Journal of Economics*, 122, 1409–1471. [0075]
- Cox, Adam B. and Richard Holden (2011), “Reconsidering Racial and Partisan Gerrymandering.” *University of Chicago Law Review*, 78, 553–604. [0097,0099]
- Deford, Daryl, Moon Duchin, and Justin Solomon (2021), “Recombination: A Family of Markov Chains for Redistricting.” *Harvard Data Science Review*, 3. [0075,0099]
- Dworzak, Piotr and Giorgio Martini (2019), “The Simple Economics of Optimal Persuasion.” *Journal of Political Economy*, 127 (5), 1993–2048. [0086]
- Ely, Jeffrey (2022), “A Cake-Cutting Solution to Gerrymandering.” Working Paper. [0075]
- Friedman, John N. and Richard Holden (2008), “Optimal Gerrymandering: Sometimes Pack, but Never Crack.” *American Economic Review*, 98 (1), 113–144. [0072-0074,0076,0077,0079,0080,0083-0086,0097, 0100]
- Friedman, John N. and Richard Holden (2009), “The Rising Incumbent Reelection Rate: What's Gerrymandering Got to Do With It?” *Journal of Politics*, 593–611. [0098]

⁵⁵Kolotilin and Wolitzky (2020) analyze comparative statics with binary voter types, showing for example that more districts are packed when the distribution of the aggregate shock is less favorable.

- Friedman, John N. and Richard Holden (2020), "Optimal Gerrymandering in a Competitive Environment." *Economic Theory Bulletin*, 1–21. [0074]
- Fryer, Roland G. and Richard Holden (2011), "Measuring the Compactness of Political Districting Plans." *Journal of Law and Economics*, 54, 493–535. [0075]
- Gelman, Andrew and Gary King (1994a), "Enhancing Democracy Through Legislative Redistricting." *American Political Science Review*, 88, 541–559. [0098]
- Gelman, Andrew and Gary King (1994b), "A Unified Method of Evaluating Electoral Systems and Redistricting Plans." *American Journal of Political Science*, 38, 514–554. [0075]
- Gentzkow, Matthew and Emir Kamenica (2016), "A Rothschild-Stiglitz Approach to Bayesian Persuasion." *American Economic Review, Papers & Proceedings*, 106, 597–601. [0086]
- Gilligan, Thomas W. and John G. Matsusaka (2006), "Public Choice Principles of Redistricting." *Public Choice*, 129, 381–398. [0075]
- Gomberg, Andrei, Romans Pancs, and Tridib Sharma (2023), "Electoral Maldistricting." *International Economic Review*, 64 (3), 1223–1264. [0074,0075,0099]
- Gomberg, Andrei, Romans Pancs, and Tridib Sharma (2024), "Padding and Pruning: Gerrymandering Under Turnout Heterogeneity." *Social Choice and Welfare*, 63 (2), 401–415. [0075]
- Grofman, Bernard and Lisa Handley (1991), "The Impact of the Voting Rights Act on Black Representation in Southern State Legislatures." *Legislative Studies Quarterly*, 16, 111–128. [0097]
- Groll, Thomas and Sharyn O'Halloran (2024), "Redistricting and Representation: The Paradox of Minority Power." Working Paper. [0075]
- Gul, Faruk and Wolfgang Pesendorfer (2010), "Strategic Redistricting." *American Economic Review*, 100 (4), 1616–1641. [0074,0076,0077,0080,0086,0100]
- Hayes, Danny and Seth C. McKee (2009), "The Participatory Effects of Redistricting." *American Journal of Political Science*, 53, 1006–1023. [0075]
- Hinich, Melvin J. (1977), "Equilibrium in Spatial Voting: The Median Voter Result Is an Artifact." *Journal of Economic Theory*, 16 (2), 208–219. [0086]
- Jeong, Dahyeon and Ajay Shenoy (2024), "The Targeting and Impact of Partisan Gerrymandering: Evidence From a Legislative Discontinuity." *Review of Economics and Statistics*, 106 (3), 814–828. [0075]
- Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian Persuasion." *American Economic Review*, 101, 2590–2615. [0074,0077]
- King, Gary and Robert X. Browning (1987), "Democratic Representation and Partisan Bias in Congressional Elections." *American Political Science Review*, 81 (4), 1251–1273. [0075,0099]
- Kleiner, Andreas, Benny Moldovanu, and Philipp Strack (2021), "Extreme Points and Majorization: Economic Applications." *Econometrica*, 89 (4), 1557–1593. [0086]
- Kolotilin, Anton (2018), "Optimal Information Disclosure: A Linear Programming Approach." *Theoretical Economics*, 13, 607–636. [0082,0086]
- Kolotilin, Anton and Alexander Wolitzky (2020), "The Economics of Partisan Gerrymandering." Working Paper. [0077,0100,0101]
- Kolotilin, Anton and Alexander Wolitzky (2024), "Distributions of Posterior Quantiles via Matching." *Theoretical Economics*, 19 (4), 1399–1413. [0086]
- Kolotilin, Anton and Alexander Wolitzky (2026), "Supplement to 'The Economics of Partisan Gerrymandering'." *Econometrica Supplemental Material*, 94, <https://doi.org/10.3982/ECTA23609>. [0075]
- Kolotilin, Anton, Roberto Corrao, and Alexander Wolitzky (2025), "Persuasion and Matching: Optimal Productive Transport." *Journal of Political Economy*, 133 (4), 1334–1381. [0074,0077-0079,0082,0086]
- Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnuk, and Ming Li (2017), "Persuasion of a Privately Informed Receiver." *Econometrica*, 85, 1949–1964. [0086]
- Kolotilin, Anton, Tymofiy Mylovanov, and Andriy Zapechelnuk (2022), "Censorship as Optimal Persuasion." *Theoretical Economics*, 17 (2), 561–585. [0082]
- Lagarde, Antoine and Tristan Tomala (2021), "Optimality and Fairness of Partisan Gerrymandering." *Mathematical Programming*, 203 (1), 9–45. [0074]
- Li, Michael and Chris Leaverton (2022), "Gerrymandering Competitive Districts to Near Extinction." Brennan Center for Justice. [0096]
- Lindbeck, Assar and Jorgen Weibull (1993), "A Model of Political Equilibrium in a Representative Democracy." *Journal of Public Economics*, 51, 195–209. [0086]
- Lublin, David, Lisa Handley, Thomas Brunell, and Bernard Grofman (2020), "Minority Success in Non-Majority Minority Districts: Finding the 'Sweet Spot'." *Journal of Race, Ethnicity, and Politics*, 5, 275–298. [0097]
- McCarty, Nolan, Keith T. Poole, and Howard Rosenthal (2009), "Does Gerrymandering Cause Polarization?" *American Journal of Political Science*, 53, 666–680. [0075,0098]

- McGann, Anthony J., Charles Anthony Smith, Michael Latner, and Alex Keena (2016), *Gerrymandering in America: The House of Representatives, the Supreme Court, and the Future of Popular Sovereignty*. Cambridge University Press. [0071]
- McGhee, Eric (2014), “Measuring Partisan Bias in Single-Member District Electoral Systems.” *Legislative Studies Quarterly*, 39 (1), 55–85. [0075]
- McGhee, Eric (2020), “Partisan Gerrymandering and Political Science.” *Annual Review of Political Science*, 23, 171–185. [0071]
- Moscariello, Paola (2025), “Gerrymandering With Endogenous Candidates.” Working Paper. [0075]
- Newkirk, Vann R. (2017), “How Redistricting Became a Technological Arms Race.” *The Atlantic*, 28 October. [0071,0093]
- Owen, Guillermo and Bernard Grofman (1988), “Optimal Partisan Gerrymandering.” *Political Geography Quarterly*, 7 (1), 5–22. [0074,0077,0086,0100]
- Puppe, Clemens and Attila Tasnádi (2009), “Optimal Redistricting Under Geographical Constraints: Why “Pack and Crack” Does not Work.” *Economics Letters*, 105, 93–96. [0074]
- Rakich, Nathaniel and Elena Mejia (2022), “Did Redistricting Cost Democrats the House?” <https://fivethirtyeight.com/features/redistricting-house-2022/>. [0072]
- Rakich, Nathaniel and Nate Silver (2018), “Election Update: The Most (and Least) Elastic States and Districts.” <https://fivethirtyeight.com/features/election-update-the-house-districts-that-swing-the-most-and-least-with-the-national-mood/>. [0076]
- Sabet, Navid and Noam Yuchtman (2025), “Identifying Partisan Gerrymandering and Its Consequences: Evidence From the 1990 US Census Redistricting.” *Journal of Political Economy: Microeconomics*, 3 (4), 649–690. [0075]
- Sherstyuk, Katerina (1998), “How to Gerrymander: A Formal Analysis.” *Public Choice*, 95, 27–49. [0074]
- Shotts, Kenneth W. (2001), “The Effect of Majority-Minority Mandates on Partisan Gerrymandering.” *American Journal of Political Science*, 120–135. [0074,0097]
- Shotts, Kenneth W. (2002), “Gerrymandering, Legislative Composition, and National Policy Outcomes.” *American Journal of Political Science*, 398–414. [0075]
- Stephanopoulos, Nicholas O. and Eric M. McGhee (2015), “Partisan Gerrymandering and the Efficiency Gap.” *University of Chicago Law Review*, 82, 831–900. [0075,0099]
- Wang, Samuel (2016), “Three Tests for Practical Evaluation of Partisan Gerrymandering.” *Stanford Law Review*, 69, 1263–1321. [0099]
- Warrington, Gregory (2018), “Quantifying Gerrymandering Using the Vote Distribution.” *Election Law Journal*, 17, 39–57. [0099]
- Washington, Ebonya (2012), “Do Majority-Black Districts Limit Blacks’ Representation? The Case of the 1990 Redistricting.” *Journal of Law and Economics*, 55, 251–274. [0097]
- Wasserman, David (2023), “Realignment, More Than Redistricting, Has Decimated Swing House Seats.” Cook Political Report. [0096]
- Wittman, Donald (1983), “Candidate Motivation: A Synthesis of Alternative Theories.” *American Political Science Review*, 77 (1), 142–157. [0086]
- Yang, Kai Hao, and Alexander Zentefis (2024), “Monotone Function Intervals: Theory and Applications.” *American Economic Review*, 114 (8), 2239–2270. [0086]

Co-editor Leeat Yariv handled this manuscript.

Manuscript received 22 October, 2024; final version accepted 14 September, 2025; available online 6 October, 2025.

All authors assume responsibility for all aspects of the paper.