

# Votemandering: Strategies and Fairness in Political Redistricting

SANYUKTA DESHPANDE, University of Illinois Urbana-Champaign, USA

IAN G. LUDDEN, Rose-Hulman Institute of Technology, USA

SHELDON H. JACOBSON, University of Illinois Urbana-Champaign, USA

Gerrymandering, the deliberate manipulation of electoral district boundaries for political advantage, is a persistent issue in U.S. redistricting cycles. In this work, we introduce and analyze *Votemandering*, a strategic blend of gerrymandering and targeted political campaigning devised to gain more seats by circumventing fairness measures. *Votemandering* leverages accurate demographic and socio-political data, bolstered by advancements in technology and data analytics, to influence voter decisions in pursuit of subtle gerrymandering strategies. We formulate *votemandering* as a Mixed Integer Program (MIP) that performs fairness-constrained gerrymandering over multiple election rounds. To combat *votemandering*, we present a computationally efficient heuristic for creating and testing district maps that more robustly preserve voter preferences. We analyze the influence of various redistricting constraints and parameters on *votemandering* efficacy. We explore the interconnectedness of gerrymandering, substantial campaign budgets, and strategic campaigning, illustrating their collective potential to generate biased electoral maps. A case study of Wisconsin State Senate redistricting substantiates our findings on real data, demonstrating how major parties can secure additional seats through *votemandering*. Our findings underscore the practical implications of these manipulations, stressing the need for informed policy and regulation to safeguard democratic processes.

## 1 INTRODUCTION

*Partisan gerrymandering* is the manipulation of voting district lines for political gain. There is numerous evidence of gerrymandering in the US electoral history, giving unfair political advantage to various parties in power (Bickerstaff et al. 2020). In an effort to detect and quantify gerrymandering, political scientists have devised various *fairness measures*, some of which incorporate historical voting data. If a proposed district plan has a fairness measure outside of a typical range, or, more robustly, is an outlier with respect to a fairness measure over a sample of feasible plans, this anomaly offers evidence of partisan gerrymandering. However, federal courts have refrained from endorsing proposed fairness measures as gerrymandering litmus tests, indicating a need for further research on the robustness and trade-offs of such tests (Rucho v. Common Cause 2019).

In addition to gerrymandering, political parties seek to enhance their political representation through huge campaign budgets (Evers-Hillstrom 2021, Horncastle 2020). Although campaigning alone cannot change the party inclinations of voters, it supports the Get Out The Vote (GOTV) cause, increasing voter turnout (Imai and Strauss 2011, Karp et al. 2008). Recent GOTV campaigns carefully target specific audiences for maximum impact, leveraging advanced machine learning algorithms that use voter data (collected through geographical surveys and the available telemetric data) to deliver information about the political inclination of the audience (Zarouali et al. 2020). Once the targets are clear, personalized campaigns are delivered through direct messages or via social media advertisements. Such campaign efforts have been used in both the 2016 and 2020 U.S. presidential elections, where clear evidence of the effectiveness of the advertisements as well as research scrutinizing the implications of presenting the social choice surfaced (Brodnax and Sapiezynski 2022, Liberini et al. 2020). The implications of such precise campaign efforts become critical, as historical election data (influenced by the campaigns) are often used to judge the fairness of proposed maps. A question of interest is then studying how smart campaign strategies can simultaneously affect immediate elections and future redistricting, and help in securing even higher political representation.

This paper aims to investigate the robustness of fairness measures to strategic campaigning and traditional gerrymandering, which we term *votemandering*. *Votemandering* is based on the idea that a party can strategically campaign in an election to alter the voting data and then draw a new district plan that appears fair for a fixed fairness measure, but gives them an unfair advantage in the next election. The focus is on identifying patterns of selective and disproportionate amendments to the representation of social choice through voting, to circumvent fairness measures for redistricting. This manipulation can be critical, particularly when slight deviations in election data can lead to significantly different fairness measure evaluations. On this background, the paper seeks to address the following research questions:

- How vulnerable are popular fairness measures to *votemandering*?
- How might voter turnout levels and political geography exacerbate or thwart *votemandering*?
- Can careful combinations of fairness measures and legal constraints promote district plans that are more robust to *votemandering*?

These questions delve into both social choice theory and practical public policy considerations. A significant concern arises when technological advancements enable greater access to detailed data on voters' preferences and increase the capacity to influence decisions, thus allowing strategic actors to target specific communities in ways that undermine fairness and equity.

We next present a brief description of the problem: Consider two election rounds with a redistricting cycle falling in between. The majority party in the state legislature, referred to as the "majority party," campaigns in the first election, winning the maximum number of seats while simultaneously

ensuring they can draw a desired district plan for the second election, which appears fair. Fairness is measured by a metric that uses past election data, such as the efficiency gap (EG), which is influenced by campaigning. Assuming complete information about the opponent party's Get Out The Vote (GOTV) campaign, the goal of the majority party is to maximize the number of seats won in both rounds. We refer to this as votemandering and formulate an optimization framework that identifies the best campaign strategies with the combined objective of securing maximum wins in both rounds and drawing a desired map that remains valid for many years. Motivated by practical and often legal constraints on redistricting in the US, we also analyze the special case of imposing proximity constraints for the proposed maps, i.e., making the least changes to the original plan while proposing a new plan, calling it *local votemandering*. Through this research, we aim to shed light on the unreliability in the process of redistricting (and detecting gerrymandering), and further point at measures that ensure more robust maps in general.

Key takeaways from this work include:

- (1) We demonstrate that fairness measures can be susceptible to data manipulation, leading to an indirect form of gerrymandering called votemandering. Therefore, the quality of a fairness measure can also be defined by its robustness against strategic amendments to the vote-share data. We formally model this phenomenon and discuss the case of the efficiency gap.
- (2) We show the fragility of district maps concerning votemandering and establish sufficient conditions for a party to benefit from it. We show how campaign budgets and access to opponents' campaign information facilitate; high voter turnout and stricter compactness bounds curtail; and voter clustering patterns have little effect on votemandering.
- (3) We lay the groundwork for creating and evaluating district plans that strongly preserve social choice, providing computationally efficient votemandering solutions. Our work is applicable to real-world data, as demonstrated by the case studies.

The remainder of the paper is structured as follows. Section 2 summarizes literature from various disciplines that connect methodologically or philosophically. Section 3 formally defines votemandering, expounding the model and methods. Section 4 proves the efficacy and computability of votemandering specific to the efficiency gap and further explores its sensitivity to state-specific factors such as voter distribution and nonpartisan redistricting constraints. Section 5 defines and analyzes local votemandering, a variation with the constraint that the new district plan is close to the original. Section 6 applies votemandering to the case of state senate redistricting in Wisconsin, demonstrating votemandering strategies for both major parties. Finally, Section 7 concludes and outlines directions for future work.

## 2 RELATED LITERATURE

This paper connects to a rich body of work from the perspectives of social choice, game theory, optimization, and statistics.

*Social Choice Theory.* Social choice theory studies and evaluates the translation of individual preferences or votes to collective societal decisions (Sen 1986). In our work, we examine the impact of strategic campaigning on political redistricting, which may be easily translated to a form of strategic voting aimed at manipulating social choice. The pure form of strategic voting has been studied for decades, although the focus has been more on various voting mechanisms and their evaluation using strategy-proofness, Pareto efficiency, independence of irrelevant alternatives, etc (Lackner and Skowron 2018, Myatt 2007). As famously shown by Gibbard (1973) and Satterthwaite (1975), no voting system for more than two players is strategy-proof. Bartholdi et al. (1989) came up with a voting rule where it is NP-complete for manipulative voters to perform strategic voting,

and also noted that many voting rules including the plurality rule can be manipulated with only polynomial computational effort. As we see within our framework, finding optimal strategies for votemandering is hard, but good solutions can be achieved with little computational effort.

*Manipulations within Plurality Voting.* Within the domain of plurality voting, such problems have also been studied from a computational theory point of view, while making a few abstractions on the redistricting part. Cohen-Zemach et al. (2018) study the problem of gerrymandering over graphs and show that the problem of dividing a social network into connected components is NP-complete, and Ito et al. (2021) build over their settings. Eiben et al. (2020), Lewenberg et al. (2017) have studied another variant involving geographic manipulation of borders and location of districts. In another interesting work by Stewart et al. (2019), information gerrymandering has been studied where the structure of the influence network manipulates the voting outcomes, along with newly placed zealots. Lev and Lewenberg (2019) study reverse gerrymandering in multi-group decision-making systems, where agents move across units to maximize their influence. The game of allocating optimal resources for campaigning has been modeled as the classic Colonel Blotto game, and its complexity, as well as equilibria, are studied (Behnezhad et al. 2018, 2017, Macdonell and Mastronardi 2015), although without examining the subsequent consequences on redistricting.

*Quantifying District Plan Fairness.* Lately, with a lot of research being done on finding ways to fairly draw the district boundaries and on knowing if a particular map is gerrymandered (Benadè et al. 2021, Chikina et al. 2017, Landau et al. 2009, Swamy et al. 2023), there has been a growing interest in defining measures to judge the fairness of a map. With multiple redistricting processes reaching the Supreme Court (Royden and Li 2017), and the latter relying on ongoing research for the mathematical analysis (Pennsylvania Case 2022), we ask if there are any strategies for fooling the measures while drawing the politically motivated map boundaries. We study a different form of strategic voting, where the strategies are implemented by the political parties, although carried out through a section of voters. In our work, we introduce a new criterion for the evaluation of voting mechanisms as well as the fairness of the district maps, stressing on the fact that the representation of social choice through voting is inherently connected to redistricting.

*Fooling Fairness Measures.* The idea of fooling the measures that are actually designed for achieving fairness is not new. Starting with Adsul et al. (2010), there has been a lot of work in the field of fair division in algorithmic game theory (Babaioff et al. 2021, Brânzei et al. 2017). By manipulating the preference data of buyers, the fairness criteria of allocation results in higher utility for the strategic players. In the field of redistricting, our work is philosophically the closest to Brubach et al. (2020), where the authors study the effects of fairness measurements on voting strategies. Using the outlier detection method, the work heuristically studies the game of strategic voting where loyal voters alter their votes as directed by their political party. Building on this work and also addressing some open questions raised, we demonstrate our results using indirect manipulation of voter turnout through selective campaigning, and we use a popular fairness measure called the efficiency gap.

*The Efficiency Gap and its Shortcomings.* Stephanopoulos and McGhee (2015) introduce the *efficiency gap* (EG) fairness measure to quantify partisan gerrymandering. EG is a fairly straightforward measure that computes the difference between the wasted votes of two major parties and labels a map as unfair if a party disproportionately wastes more votes than the other. It has been widely used because of its simplicity, intuition, and the use of actual voter preference data from the elections (Gill v Whitford 2018, Missouri Constitution 2022).

With the widespread use of EG, there has also been growing literature on the shortcomings of EG, typically focusing on its implications and the nature of it being a single-dimensional number

trying to capture the complex forms of gerrymandering (Bernstein and Duchin 2017, Kean 2018). In particular, Chambers et al. (2017) have majorly criticized the philosophical implications like the possible increase in polarization, the problematic ranking of maps and technical implications like discouragement in proportional representation. Moreover, the EG has also been criticized for favoring uncompetitive elections and voter suppression (Plener Cover 2018) and volatility in competitive elections leading to a high number of asymmetric wasted votes as well as for terming 3 to 1 victories as neutral (Bernstein and Duchin 2017). Tam Cho (2017) have observed the problematic variations of EG implications across states with comparable vote shares; and further have talked about the limited number of values EG can take for any fixed vote shares. Alexeev and Mixon (2018) show that sometimes only bizarrely shaped districts satisfy both population balance and EG constraints. Other philosophical shortcomings include the definition and weights of wasted votes (winner’s surplus wasted votes, loser’s all votes are weighted the same), incorrect reporting of the social choice, and bias to the winning party (Barton 2018, Nagle 2019). Numerous updates are proposed to the current computation of the EG (Barton 2018, Tam Cho 2017), while also criticizing the implication of wasted votes being improperly biased towards districts with higher voting turnout (Wallin 2017). We note that the main criticism offered by our work is fundamentally independent of the previous work done on evaluating the EG, and our main focus is on the sensitivity of EG and its susceptibility to getting fooled in the broader context of votemandering.

### 3 METHODOLOGY

In this section, we formally discuss the votemandering model and our methodology. Section 3.1 sets the premise with a high-level description of the problem and Section 3.2 formally expounds the model. Section 3.3 presents votemandering as an optimization problem, applicable to a general fairness metric using past-election data. Finally, Section 3.4 outlines a two-stage heuristic approach to solving the votemandering optimization problem and describes the specific case of EG.

#### 3.1 High-level Votemandering Model

We begin by defining a function,  $\mathcal{E} : \mathcal{D} \times \mathcal{V} \rightarrow \mathbb{N}$ , to determine state-wide election winners. This function maps a district plan,  $D \in \mathcal{D}$ , and a set of voter ballots,  $V \in \mathcal{V}$ , to the number of districts won by party  $A$  in the election. The election function,  $\mathcal{E}$ , represents a specific electoral system, such as single-member districts with first-past-the-post voting. Although the voting data,  $V$ , may be influenced by stochastic processes like migration and political dialogue,  $\mathcal{E}$  is deterministic.

In this framework, partisan gerrymandering involves replacing  $D$  with  $\tilde{D}$  to win more districts, i.e.,  $\mathcal{E}(\tilde{D}, V) > \mathcal{E}(D, V)$ . Similarly, election campaigning alters  $V$  to  $\tilde{V}$  to secure more districts:  $\mathcal{E}(D, \tilde{V}) > \mathcal{E}(D, V)$ . Note that election campaigning is generally considered fair within the confines of the Federal Election Campaign Act.

Existing approaches to limit partisan gerrymandering involve calculating a fairness measure,  $f : \mathcal{D} \times \mathcal{V} \rightarrow \mathbb{R}$ , and rejecting a district plan  $D$  if and only if  $f(D, V_0) > \delta$ . Here,  $V_0$  represents historical voting data, and  $\delta$  is a predetermined threshold. This fairness constraint aims to reduce the strategic impact of gerrymandering on election outcomes. However, as noted by Brubach et al. (2020), partisan agents may manipulate voting data in one election to make a future gerrymandered district plan appear fair. Let  $D_0$  represent the current district plan. The manipulative partisan agent, party  $A$ , attempts to solve the optimization problem:

$$\begin{aligned} & \underset{\tilde{D} \in \mathcal{D}, \tilde{V} \in \mathcal{V}}{\text{maximize}} && \mathcal{E}(D_0, \tilde{V}) + \mathcal{E}(\tilde{D}, V_0) \\ & \text{subject to} && f(\tilde{D}, \tilde{V}) \leq \delta. \end{aligned} \tag{1}$$

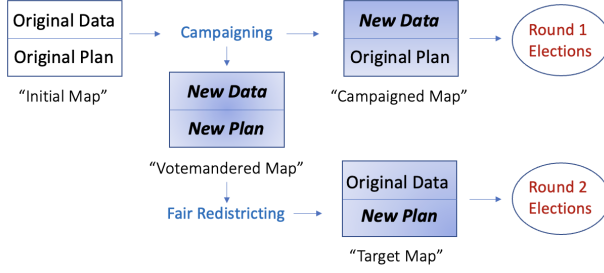


Fig. 1. The model highlighting various stages of votemandering

Previous research on elections and redistricting has focused on the effects of either  $\tilde{D}$ ,  $\tilde{V}$ , or  $f$ . In contrast, this paper investigates the efficacy of votemandering, which combines gerrymandering and past or present campaigning, primarily in opposition to a specific partisan bias measure, such as the efficiency gap (EG). The votemandering framework assumes translation of campaign budgets to improved voter turnout and an access to other party’s budget allocation information, although it is fairly robust to overcome small uncertainties within the data, as discussed in Section 4.3.

### 3.2 Model Details and Terminology

Consider two political parties: the (state legislative) majority party,  $A$ , and the minority party,  $B$ . Party  $A$  is assumed to be in-charge of the redistricting process, in line with requirements of majority of the states in the US (Center 2022). Suppose parties  $A$  and  $B$  compete in two rounds of elections with a redistricting cycle in between. By examining this narrow time window, our model studies only short-term implications of campaigning, affecting the round-1 election and the subsequent map-drawing process.

Recall the high-level votemandering optimization problem (1). Set the electoral system,  $\mathcal{E}$ , as single-member districts with first-past-the-post voting. Function  $f$  represents the fairness measure, such as the EG. We distinguish between *plan* and *map*, with the former indicating unit-to-district assignments and the latter encompassing both a district plan and unit-level voter data.

Figure 1 illustrates the stages of votemandering. In round-1, elections use the existing district plan,  $D_0$ , with voter ballots  $\tilde{V}$  resulting from GOTV campaign efforts. We refer to  $V_0$  as the *original data* and  $\tilde{V}$  as the *new data*. We label  $(D_0, V_0)$  the *initial map* and  $(D_0, \tilde{V})$  the *campaigned map*. Following round-1, party  $A$  creates a new district plan,  $\tilde{D}$ , satisfying fairness constraints using voter data from the round-1 elections, i.e., the new data. Round-2 elections employ the new plan,  $\tilde{D}$ , but with the original data,  $V_0$ . We designate  $(\tilde{D}, \tilde{V})$  as the *votemandered map* and  $(\tilde{D}, V_0)$  as the *target map*. The reversion to  $V_0$  in round-2 implicitly assumes party  $A$  can precisely match party  $B$ ’s GOTV budget allocation, negating any increases in voter turnout. We do not model campaign budget strategies in round-2 to avoid added complexity and, more importantly, to concentrate on showcasing the ability to manipulate vote shares for generating a desired map while still appearing to uphold fairness.

Party  $A$ ’s strategic GOTV campaign in round-1 influences their seat-share in both election rounds: directly through wins in the campaigned map,  $\mathcal{E}(D_0, \tilde{V})$ , and indirectly through wins in the target map,  $\mathcal{E}(\tilde{D}, V_0)$ . To examine the strategies of the majority party, we fix party  $B$ ’s budget allocation across all units and consider party  $A$ ’s optimization problem (1) of maximizing their total number of seats. Wins in round-2 are critical because the target map will remain in effect until the next redistricting phase. To cover a complete redistricting cycle (such as a 10-year period in

U.S. elections), the model can be extended by adjusting the weight of round-2 wins accordingly. Furthermore, the model accommodates the inclusion of aggregated historical data from multiple elections by appropriately adjusting the weight attributed to campaign influence.

It is essential to emphasize that votemandering is fundamentally different from both strategic campaigning and gerrymandering due to its interactions between stages. As Section 6 demonstrates, even modest budget allocations can lead to significant votemandering outcomes, setting it apart from traditional campaigning by incorporating additional gerrymandering tactics. An example in Appendix B.1 illustrates the votemandering process.

### 3.3 Optimization Framework for Votemandering

The votemandering model motivates an optimization framework for exploring potential campaign and redistricting strategies for party  $A$ . Table 4 in Appendix A lists the notation built.

**3.3.1 State Characteristics.** Let  $K$  denote the set of units in a state with  $n$  districts. Each district designates one unit as its *center*. The district assignment of each unit  $j \in K$  in each round  $r \in \{1, 2\}$  is represented by the indicator variables

$$z_{ij}^r = \begin{cases} 1, & \text{if unit } j \text{ is assigned to the district centered at unit } i \text{ in round } r \\ 0, & \text{otherwise.} \end{cases}$$

Moreover,  $z_{ii}^r = 1$  if  $i$  is a district center in round  $r$ . The original district plan,  $D_0$ , determines the values of  $z_{ij}^1$ , and all  $z_{ij}^2$  are decision variables. The following constraints enforce the proper formation of districts in round-2.

$$\sum_{k \in K} z_{ki}^2 = 1 \quad \forall i \in K \quad (2)$$

$$\sum_{k \in K} z_{kk}^2 = n \quad (3)$$

Constraint (2) ensures every unit is assigned to some district, and constraint (3) ensures exactly  $n$  units are chosen as district centers.

**3.3.2 Budget and Campaigning.** Assume complete information about unit populations and the corresponding party affiliations, i.e., the maximum number of voters for each party in each unit. The maximum vote counts for party  $A$  ( $B$ ) are given by  $v_{init,k}^A$  ( $v_{init,k}^B$ ) in unit  $k \in K$ , with total unit population  $p_k = v_{init,k}^A + v_{init,k}^B$ . Let  $\alpha \in [0, 1]$  denote the fractional baseline voter turnout, assumed constant across all units. The number of party  $A$  votes is the sum of  $\alpha v_{init,k}$  and the votes through GOTV campaigning in unit  $k$ . The vote shares  $v_{init,k}$  and  $\alpha$  are fixed for all rounds, and the actual voting turnout varies depending on campaigning.

Let  $\mathcal{B}^A$ ,  $\mathcal{B}^B$  denote the parties' GOTV campaign budgets in terms of the total number of their supporters they can convince to show up to the polls. Budget allocations in unit  $k$  by party  $P \in \{A, B\}$  may push their actual number of votes above the baseline turnout  $\alpha v_{init,k}^P$ , but their total number of votes cannot exceed  $v_{init,k}^P$ . (Hence if  $\alpha = 1$ , then GOTV budget allocations have no effect.) This constraint is implemented by defining  $v_k^A$  as the actual voter turnout for party  $A$  in unit  $k$ , with

budget  $b_k^A$  spent satisfying

$$v_k^A = \alpha v_{init,k}^A + b_k^A \quad \forall k \in K \quad (4)$$

$$v_k^B = \alpha v_{init,k}^B + b_k^B \quad \forall k \in K \quad (5)$$

$$b_k^A \leq (1 - \alpha) v_{init,k}^A \quad \forall k \in K \quad (6)$$

$$\sum_{k \in K} b_k \leq \mathcal{B}^A \quad (7)$$

By assumption, party  $B$ 's GOTV campaign allocation, and therefore the values  $v_k^B$ , are known to party  $A$ .

**3.3.3 Winning Districts.** A party must win more than half of the votes in a district to secure a win. We use indicator variables  $\hat{s}_i^1$  and  $\hat{s}_i^2$  and the big- $M$  method for incorporating the wins in campaigned and target maps respectively

$$1 - M(1 - \hat{s}_i^1) \leq \sum_{k \in K} z_{ik}^1 (v_k^A - v_k^B) \leq M\hat{s}_i^1 \quad \forall i \in K \quad (8)$$

$$1 - M(1 - \hat{s}_i^2) \leq \sum_{k \in K} z_{ik}^2 (v_{init,k}^A - v_{init,k}^B) \leq M\hat{s}_i^2 \quad \forall i \in K \quad (9)$$

Note that constraints (8) and (9) are both linear:  $z_{ik}^1$  show the unit to district assignments in the initial map and are given, although variables  $v_k^A$  depend on the budget spent. For (9), we know the values of  $v_{init,k}^A, v_{init,k}^B$ , but variables  $z_{ik}^2$  depend on the plan that we make for round-2.

**3.3.4 The Votemandering MIP.** The objective function of (1) is now represented by the sum of individual district wins in both rounds, i.e.,  $\hat{s}_i^1$  and  $\hat{s}_i^2$  for every unit  $i$ . As described in (1), a fairness measure constraint  $f(\tilde{D}, \tilde{V}) \leq \delta$  is implemented, here precisely represented as a function of the first round variables (updated vote shares  $v_k^A, v_k^B$ ) as well as the second-round assignment variables ( $z_{ij}^2$ ). Using our notation, this constraint refers to the fairness constraint on the votemandered map. Furthermore, the round-2 plan, i.e.,  $z_{ij}^2$  also needs to satisfy the contiguity, population, and/or compactness constraints for making districts. We omit these nonpartisan constraints for brevity and refer the reader to (Swamy et al. 2023) for implementation details.

Finally, given  $v_{init,k}^A, v_{init,k}^B, z_{ij}^1, v_k^B, \mathcal{B}^A$  and  $\delta$ , a mixed-integer program (MIP) formulation of party  $A$ 's optimization problem is

$$\begin{aligned} & \text{maximize} && \sum_{i \in K} \hat{s}_i^1 + \sum_{i \in K} \hat{s}_i^2 \\ & \{b_k^A\}_k, \{z_{ik}^2\}_{i,k} && \\ & \text{subject to} && \text{constraints (2) - (9),} \\ & && z_{ik}^2, \hat{s}_i^1, \hat{s}_i^2 \in \{0, 1\} \quad \forall i, k \in K, \\ & && b_i^A \geq 0 \quad \forall i \in K, \\ & && f(z_{ik}^2, v_k^A, v_k^B) \leq \delta, \\ & && \{z_{ik}^2\} \text{ satisfy nonpartisan constraints.} \end{aligned} \quad (10)$$

This concludes the description of the optimization problem (10). It is evident that the problem is computationally challenging due to the complex map-making constraints. For most fairness constraints, an exact approach to solving this optimization problem is only feasible for very small-sized grids (on the order of  $3 \times 4$ ).



### 3.4 A Sampling-Based Votemandering Heuristic

The complexity of the optimization problem (10) arises from the interplay between the four votemandering stages. Campaigning decisions depend on the first and third stages (the initial and votemandered maps), whereas the objective depends on the second and fourth stages (the campaigned and target maps). Additionally, (10) accounts for the budget, voter turnout, and feasible map-making constraints, making it difficult to determine the best direction to improve the objective and find an optimal solution.

To address this complexity, the problem is split into two parts leading to an efficient heuristic approach: find a promising target map, then increase round-1 wins while maintaining the apparent fairness of the votemandered map. For a fixed target map defined by  $z_{ik}^2$  variables, solving (10) reduces to finding an optimal budget allocation  $b_k^A$  while maintaining feasibility (if possible).

A brute-force method of checking all possible new plans  $\tilde{D} \in \mathcal{D}$  is computationally infeasible due to the size of  $\mathcal{D}$ , i.e., the combinatorial explosion of possible redistricting plans. Instead, sampling is used to reduce the new plan search space from the set  $\mathcal{D}$  of all district plans satisfying nonpartisan redistricting constraints to a smaller pool,  $\mathcal{P} \subset \mathcal{D}$ . To quickly sample a small but diverse pool  $\mathcal{P}$ , we implement the popular recombination Markov chain (DeFord et al. 2021).

The proposed algorithm considers each candidate plan in  $\mathcal{P}$  according to a priority order, stopping when a pool-optimal plan,  $D^*$ , is found. Note that an optimal solution within the pool may not be unique, and experiments suggest a large number of pool-optimal plans exist. The number of wins for party  $A$  in  $D^*$  with the original data,  $\mathcal{E}(\tilde{D}, V_0)$ , is a valid lower bound on the global optimum across all of  $\mathcal{D}$ . Although recombination sampling may miss optimal new plans, this two-stage heuristic is tractable for standard-sized instances and provides practical solutions that effectively utilize votemandering strategies, showing improvements in the number of seats won.

Let  $\mathcal{P} \subset \mathcal{D}$  be a pool of  $N$  candidate new plans, i.e.,  $\mathcal{P} \equiv \{D_1, D_2, \dots, D_N\}$ . The choice of new plan  $D_i$  combined with the original voter data  $V_0$  determines the number of wins in the target map,  $\mathcal{E}(D_i, V_0)$ . Hence the best new plan for  $A$  is determined by finding, for each plan  $D_i \in \mathcal{P}$ , the maximum number of round-1 wins for  $A$  (via spending budget  $\mathcal{B}^A$ ) such that the votemandered map with plan  $D_i$  fools the fairness constraint. By decoupling the round-1 and round-2 contributions to the objective function of (10), this heuristic efficiently returns the optimal new plan from the pool. Algorithm 1 provides a high-level description of the heuristic.

**PROPOSITION 3.1.** *Algorithm 1 returns a district plan in  $\mathcal{P}$  that, when used for the votemandered and target maps, maximizes the total number of wins for party  $A$  across the two election rounds.*

**PROOF.** See Appendix B.2. □

The main computational effort in Algorithm 1 occurs in Line 10. With the target map fully determined, the objective of (10) simplifies to maximize the campaigned map (round-1) wins over all possible budget allocations  $\{b_k^A\}_k \in K$  while maintaining the fairness of the votemandered map. We call Line 10 the *fairness step*, because the goal is to maximize wins conditioned on plan  $D_i$  appearing fair as the votemandered map. Henceforward, we use a specific fairness measure, the efficiency gap (EG), which we formally define in Section 3.4.2. We next expound on the simplified version of (10) solved with the fairness step specific to EG.

**3.4.1 Additional Notation.** Let  $\mathcal{I} = \{I_1, \dots, I_n\}$  be the set of districts  $I_i$  in the original plan (round-1) such that each  $I_i$  is a set of units from  $K$ . Sets  $I_i$  satisfy  $I_i \cap I_j = \emptyset$  as no unit can belong to two districts in any round. Let  $(V_{init,I}^A, V_{init,I}^B)$  and  $(V_I^A, V_I^B)$  denote pre-campaigning and post-campaigning votes, respectively, in district  $I$ . Similarly,  $\mathcal{J} = \{J_1, \dots, J_n\}$  is the set of districts in the new plan (round-2), with  $(V_{init,J}^A, V_{init,J}^B)$  and  $(V_J^A, V_J^B)$  denoting pre and post-campaigning votes in district  $J \in \mathcal{J}$ . Let

**ALGORITHM 1:** Votemandering Heuristic: Select Optimal Plan from a Pool

---

```

1: Input: Pool  $\mathcal{P} = \{D_1, D_2, \dots, D_N\}$  of candidate new plans
2: Sort  $\mathcal{P}$  in decreasing order of  $\mathcal{E}(D_i, V_0)$ , relabeling from  $D_1$  to  $D_N$ 
3:  $s_{\max}^1 \leftarrow$  maximum number of round-1 wins for party  $A$  by spending campaign budget  $\mathcal{B}^A$ 
4: best_plan  $\leftarrow$  NULL
5: best_obj  $\leftarrow -\infty$ 
6: for all  $D_i \in \mathcal{P}$  do
7:   if  $s_{\max}^1 + \mathcal{E}(D_i, V_0) < \text{best\_obj}$  then
8:     break
9:   end if
10:  obj  $\leftarrow$  solve (10) with  $z_{ik}^2$  variables fixed to encode  $D_i$ , returning  $-\infty$  if infeasible
11:  if obj  $>$  best_obj then
12:    best_plan  $\leftarrow D_i$ 
13:    best_obj  $\leftarrow$  obj
14:  end if
15: end for
16: Output: best_plan

```

---

$\hat{x}_I$ , indexed using sets  $I_i \in \mathcal{I}$  and  $\hat{y}_J$ , indexed using sets  $J_j \in \mathcal{J}$  be the indicator variables denoting the wins in the campaigned map (round-1) and the votemandered map (round-2), respectively. Note that we resort to the set notation ( $\hat{x}_I$  using  $I, J$ ) unlike that of the original optimization problem (i.e.,  $\hat{s}_i^1$ ), as we now have assignments of both initial and target maps, allowing lesser notation.

3.4.2 *Incorporating EG into the Fairness Step.* Using the definition for EG, the difference between wasted votes for each district  $I \in \mathcal{I}$  (denoted henceforth by  $\mathcal{W}(B - A)$ ) is given by

$$\mathcal{W}(B - A)(I) = \begin{cases} \frac{3V_I^B - V_I^A}{2}, & \text{if } V_I^A > V_I^B \text{ (A wins)} \\ \frac{V_I^B - 3V_I^A}{2}, & \text{if } V_I^A < V_I^B \text{ (B wins)} \end{cases} \quad (11)$$

Using this definition, we further write the constraint of EG less than a particular constant, say 8% (Stephanopoulos and McGhee 2015).

$$\text{efficiency gap of the state} = \left| \sum_{I \in \mathcal{I}} \mathcal{W}(B - A)(I) / \left( \sum_{I \in \mathcal{I}} V_I^A + V_I^B \right) \right| \leq 0.08 \quad (12)$$

Next, we describe the MIP we use to ensure the fairness of the proposed map in round-2, i.e., the votemandered map. Letting  $\tau_J, \forall J \in \mathcal{J}$  denote the difference between wasted votes in  $J$ 's district, i.e.,  $\mathcal{W}(B - A)(I)$ , we can write:

$$\begin{aligned} & \max \sum_{I \in \mathcal{I}} \hat{x}_I \\ & \text{s.t. Constraints (6), (7)} \\ & V_I^A = \sum_{k \in I} \alpha v_{init,k}^A + b_k^A \quad \forall I \in \mathcal{I} \\ & V_J^A = \sum_{k \in J} \alpha v_{init,k}^A + b_k^A \quad \forall J \in \mathcal{J} \\ & 1 - M(1 - \hat{x}_I) \leq (V_I^A - V_I^B) \leq M\hat{x}_I \quad \forall I \in \mathcal{I} \quad (13) \\ & 1 - M(1 - \hat{y}_J) \leq (V_J^A - V_J^B) \leq M\hat{y}_J \quad \forall J \in \mathcal{J} \quad (14) \end{aligned}$$

$$0 \leq -\tau_J + \left( \frac{3V_J^B - V_J^A}{2} \right) \leq M(1 - \hat{y}_J) \quad \forall J \in \mathcal{J} \quad (15)$$

$$0 \leq \tau_J - \left( \frac{V_J^B - 3V_J^A}{2} \right) \leq M\hat{y}_J \quad \forall J \in \mathcal{J} \quad (16)$$

$$-0.08 \leq \sum_J \tau_J / \left( \sum_{J \in \mathcal{J}} V_J^A + V_J^B \right) \leq 0.08 \quad (17)$$

$$\hat{x}_I, \hat{y}_J \in \{0, 1\}, \quad b_k^A, \tau_J \geq 0 \quad \forall k \in K, J \in \mathcal{J} \quad (18)$$

Algorithm 1 finds an optimal solution within a pool of target maps, but one may question about the probability that a globally optimal solution exists within a pool generated by running a recombination chain for  $N$  steps. However, given the hardness of finding a globally optimal solution, it is unlikely that a bound on this probability can be determined. In practice, the algorithm is computationally efficient as shown in Section 4 (Theorem 4.3), and the returns diminish as the size of the pool increases. It is important to note that the primary goal of this paper is to establish the mechanism of votemandering and study its dependence on various crucial factors that affect redistricting, as opposed to finding the optimal votemandering strategies.

As the algorithm works given any inputs of the initial map and campaign budget, it establishes a framework that can be used to test the robustness of any district plan or pool of maps against votemandering. This framework is used in later sections to compare the effects of various state characteristics and external redistricting conditions on votemandering. An ideal map would have a lower objective when tested against a standard pool of target maps. The higher the budget required to votemandeer, the better the robustness.

## 4 RESULTS AND ANALYSIS

This section presents the efficacy and efficiency of votemandering under various conditions. Using EG as our fairness measure, we begin by examining the impact of campaigning on votemandering objective and fairness in Section 4.1. We show that under certain general conditions, votemandering can always occur. In Section 4.2, we establish the polynomial-time complexity of Algorithm 1. Finally, in Section 4.3, we experimentally analyze the dependence of various factors on votemandering, such as the budget of Party A and Party B, compactness, voter turnout, and the concentration index Moran's I.

### 4.1 Sufficient Conditions for Votemandering

As a build-up to this question, we analyze the strategy space of party A: it can add new votes via campaigning in round-1; effectively gerrymander to shift votes from a winning (W) district to a losing (L) district or vice versa. We discuss their key implications on fairness in Lemma 4.1.

**LEMMA 4.1.** *The actions of campaigning and vote shifts have an impact on the difference between wasted votes, i.e.,  $\mathcal{W}(B - A)$ , as given in Table 1.*

**PROOF.** See Appendix C.1. □

Next, we discuss the total impact on the change in wasted votes  $\mathcal{W}(B - A)$ , and thereby, the efficiency gap, as a new district plan gets drawn over the same vote data. As the total number of votes does not change in this case, this change can be tracked just through a reshuffle of units into winning and losing districts. For district assignment  $I$  in round-1, the difference between the

**Table 1.** Change in the difference between wasted votes (EG) as new votes are added/shifted

|   | Action  | Impact on $\mathcal{W}(B - A)$ |
|---|---|--------------------------------|
| 1 | Wasting an additional vote on a losing district     | $-3/2$                         |
| 2 | Wasting an additional vote on a winning district    | $-1/2$                         |
| 3 | Winning a district $I$ through campaigning          | $(3V_I^A + V_I^B)/2$           |
| 4 | Shift $x$ votes from a winning to a losing district | $-x$                           |
| 5 | Shift $x$ votes from a losing to a winning district | $x$                            |

wasted votes  $\mathcal{W}(B - A)_I$  is expressed as:

$$\begin{aligned} \mathcal{W}(B - A)_I &= \alpha \left( \sum_{I \in \mathcal{I}(W)} \frac{3V_{init,I}^B - V_{init,I}^A}{2} \right) + \alpha \left( \sum_{i \in \mathcal{I}(L)} \frac{V_{init,I}^B - 3V_{init,I}^A}{2} \right) \\ &= \alpha \left( \sum_{I \in \mathcal{I}} \frac{V_{init,I}^B - V_{init,I}^A}{2} \right) + \alpha \left( \sum_{i \in \mathcal{I}(W)} V_{init,I}^B - \sum_{I \in \mathcal{I}(L)} V_{init,I}^A \right) \end{aligned} \quad (19)$$

where  $\mathcal{I}(W)$  and  $\mathcal{I}(L)$  are the sets of winning and losing districts, respectively. Then, after reshuffling to district assignment  $\mathcal{J}$  in round-2, the change in  $\mathcal{W}$  (defined by  $\Delta \mathcal{W}(B - A)_{I \rightarrow \mathcal{J}}$ ) and the final  $\mathcal{W}$  is given as:

$$\begin{aligned} \Delta \mathcal{W}(B - A)_{I \rightarrow \mathcal{J}} &= \alpha \left( \sum_{j \in \mathcal{J}(W)} V_{init,j}^B - \sum_{j \in \mathcal{J}(L)} V_{init,j}^A \right) - \alpha \left( \sum_{i \in \mathcal{I}(W)} V_{init,i}^B - \sum_{i \in \mathcal{I}(L)} V_{init,i}^A \right) \\ \mathcal{W}(B - A)_{\mathcal{J}} &= \mathcal{W}(B - A)_I + \Delta \mathcal{W}(B - A)_{I \rightarrow \mathcal{J}} + [\text{Any wasted votes through campaign}] \end{aligned} \quad (20)$$

To conclude, Table 1 and Eq. (20) show that a campaign budget can be allotted (thereby updating  $V_{init,I}^A, V_{init,I}^B$  to  $V_I^A, V_I^B$ ) to achieve fairness of a target map, given that the allocation also satisfies the budget and voter-turnout constraints. Thus, votemandering can potentially include at least two (interdependent) ways: (1) Fixing plan  $\mathcal{J}$  and allotting appropriate budget to satisfy the fairness bound, whilst benefiting from campaigning in round-1, and (2) Designing a target map with  $\mathcal{J}$  that leads to a higher number of wins in round-2, maintaining fairness. To measure the efficacy of votemandering, we define *votemandering bonus* i.e.,  $\Delta$  which measures the gain in the number of wins after enabling votemandering. For a target plan  $\tilde{D}$  and campaigning resulting with  $\tilde{V}$ ,

$$\text{Votemandering bonus } \Delta = \mathcal{E}(D_0, \tilde{D}) + \mathcal{E}(\tilde{D}, V_0) - 2\mathcal{E}(D_0, V_0) \quad (21)$$

Using this definition, a positive votemandering bonus would indicate that we have successfully votemandered. We now characterize the sufficient conditions for successful votemandering using the second way of improving the objective. Intuitively, we need a target map better than the initial map, and a baseline voter turnout to allow GOTV efforts to take place.

**THEOREM 4.2.** *For any vote-share distribution and a corresponding fair initial map with assignment  $\mathcal{I}$ , the existence of strategies leading to a positive votemandering bonus is guaranteed if*

- (1) *A feasible, contiguous map with assignment  $\mathcal{J}$  exists with a higher number of wins than the initial map.*
- (2) *The voter turnout  $\alpha$  satisfies:*

$$\alpha \leq 1 - \left( \frac{2\Delta \mathcal{W}(B - A)_{I \rightarrow \mathcal{J}}}{\sum_{J \in \mathcal{J}(W)} \sum_{j \in J} v_j^A} \right)$$

where  $\Delta W(B - A)_{\mathcal{I} \rightarrow \mathcal{J}}$  is the change in the difference between wasted votes from assignment  $\mathcal{I}$  to  $\mathcal{J}$ , and  $v_j^A$  are party  $A$ 's votes in unit  $j$ .

PROOF. We prove this by achieving fairness of the map with assignment  $\mathcal{J}$  using a strategic allocation of budget, thereby establishing the existence of strategies. We primarily satisfy Eq. (20). We continue with the proof details in Appendix C.2.  $\square$

In practice, it is generally much easier to votemander (as we demonstrate in Section 4.3 and through case studies in Section 6), except under highly specific conditions such as near 100% voter turnout and nearly all voters favoring a single party. The first way of votemandering, as discussed above, also allows for a positive bonus to be achieved through an increase in wins in the first round, as long as the voter turnout allows for such campaigning to occur in a fair way. Its campaigning effects on fairness, as translated from the additional number of wins, can be dissolved in the votemandered map through reorganization of the campaigned map. While this first way is easier to see in practice, its dependence on the specificity of  $\mathcal{J}$  makes it difficult to establish sufficient conditions for votemandering as it demands map making, given an initial assignment  $\mathcal{I}$ . We explore specific votemandering strategies using both ways in detail in Section 5.

## 4.2 Efficiency of Votemandering Heuristic

We now establish the polynomial time complexity of the votemandering heuristic. Recall that it takes a pool of maps  $\mathcal{P}$  as an input, and outputs the target map in  $\mathcal{P}$  maximizing the votemandering objective with respect to the given initial map, i.e., it finds the target map with the maximum votemandering bonus. Proposition 3.1 confirms the correctness of Algorithm 1 in its convergence to the optimal target map. Theorem 4.3 now shows that this may be achieved efficiently.

**THEOREM 4.3.** *Let  $\mathcal{P}$  be a pool of  $N$  candidate target district plans such that  $\mathcal{P}$  is a subset of feasible, but not necessarily fair,  $n$ -districts plans. A plan in  $\mathcal{P}$  that maximizes the votemandering bonus may be found in  $\text{poly}(N, n)$  time.*

PROOF. Following Algorithm 1, see that the only complicated part is the fairness step (10) in the MIP in Section 3.4.2. We show that each target map can be checked in polynomial time, enabling us to move through the pool quickly until convergence. We sketch the proof here and provide details in Appendix C.3.

- (1) For each district in round-1, we decompose its space into *pieces* that each belongs to a district in round-2. For each such piece, we define its capacity =  $\min(\text{its voter-turnout capacity, budget needed to win the round-2 district it is part of (only if part of a losing district)})$ .
- (2) Given party  $B$ 's investment and the original vote shares of  $A$  and  $B$ , we next find the win/lose ( $W/L$ ) status of the districts in the votemandered map and compose a linear program to find the maximum wins in the campaigned map while constraining on the status of the districts, implemented by the variables for investment in the pieces.
- (3) If the ( $W, L$ ) constraint for district  $k$  (in the votemandered map) is tight in the optimal solution, do: i) mark  $k$ 's status as a win and update the fairness of the votemandered map, ii) add a constraint for allocating the budget needed to win  $k$ . We then solve the updated linear program, and if the objective increases, we repeat all the steps with updated  $W/L$  status and constraints until the objective stops increasing. This converges in polynomial time since there is a predefined number of districts with an  $L$  status, bounded above by  $n$ .

These three steps suffice to prove Theorem 4.3.  $\square$




Fig. 2. Party A's votemandering bonus, with increase in both budgets

### 4.3 Analysis of factors impacting votemandering

Although the conditions for votemandering may be easily satisfied in practice, the required budget to ensure a positive bonus can vary. The efficacy of votemandering depends on several factors, including the initial vote share distribution across the state, the initial district assignment plan, and the available campaign budget for parties  $A$  and  $B$ . Furthermore, it is influenced by various externally imposed constraints on the redistricting process, such as EG, compactness, the number of majority-minority districts, proportionality, etc. As a result, we opt for a randomized approach, i.e., Algorithm 1, for a pool of maps to examine the dependence on these factors and, in turn, demonstrate the efficacy of votemandering under various conditions. The pool of plans is randomly generated using recombination and contains plans that all satisfy the externally imposed constraints. We fix a randomly generated vote share distribution across a grid with  $20 \times 20$  units such that each unit  $i$  has a population  $p_i$  uniformly chosen between 350 – 400 and vote shares  $(v_{init,i}^A/p_i, v_{init,i}^B/p_i)$  between 20 – 80% for each party. Each feasible map from the pool provides a unit-to-district assignment, mapping the 400 units to 10 districts, with district populations allowed to deviate 1% from the average district population.

**4.3.1 Impact of increasing budget.** We present our first key result through Figure 2, which tracks how increasing the campaign budget strengthens the ability to votemand. Recall that the budget equals the number of votes that can be influenced above the baseline voter turnout, with an upper bound given by total party affiliation shares. Figure 2 plots Party  $A$ 's votemandering bonus as the parties increase their budget uniformly. In experiments for Figure 2a, Party  $B$ 's budget  $\mathcal{B}^B$  is fixed at 400, and Party  $A$ 's budget  $\mathcal{B}^A$  is varied, whereas for Figure 2b,  $\mathcal{B}^B$  is varied, and  $\mathcal{B}^A$  is fixed at 400. In both experiments, we plot the majority party  $A$ 's votemandering bonus coming from its strategic investment. Recall that we do not assume any campaigning strategies from Party  $B$ . Given any budget allocation of  $B$ , if  $A$  has access to the allocation information, then the algorithm finds the best strategies for  $A$ . Here, we let  $B$  invest most straightforwardly, making its budget investment proportional to each unit's population, allowing fractional investments.

Figure 2a shows a steady increase in the bonus through the means and medians shifting upwards with the increase in  $\mathcal{B}^A$ . The bonus for  $\mathcal{B}^A = 100$  indicates the objective that can be attained by accessing Party  $B$ 's budget investment information while  $A$  puts in a little campaigning effort itself. Note that the increase in bonus in Figure 2a is not linear. Improving the allowed budget has diminishing returns in the form of objectives. This is expected since the objective, and therefore the bonus, is capped by the total number of seats available in both rounds.

Most interestingly, the bonus has a counterintuitive relation with increasing  $\mathcal{B}^B$  as shown in Figure 2b. As opposed to a clear steady increase in 2a, increasing  $\mathcal{B}^B$  does not ensure a steady decrease in the bonus. Although increasing  $\mathcal{B}^B$  may impact Party  $A$ 's chances of winning in the first




Fig. 3. Party A's votemandering objective with lower and higher compactness bounds

round,  $A$  may use this information to create better districts in the second round. This is achieved by letting Party  $B$  win some districts in the votemandered map, only to lose those in the target map as the campaign effects diminish. Because of this trade-off, the decreasing trend is not obvious: Party  $A$ 's votemandering bonus remains largely unaffected until  $\mathcal{B}^B$  reaches 500.

**4.3.2 Impact of increasing compactness.** The metric of compactness is generally not perceived as a fairness notion and is usually imposed to be in line with the (older) belief of creating districts that minimize the physical distance between units within a district. The salamander-shaped district in the first gerrymander suggests rejecting freehand-shaped districts and asking for compactness as a proxy for partisan neutrality (Polsby and Popper 1991). However, compactness is often deemed orthogonal to fairness measures (Gurnee and Shmoys 2021). Contrary to this belief, we demonstrate that imposing tighter compactness bounds limits the ability of votemandering, leading to better (robust) maps in general. We achieve this by comparing votemandering objectives on two separate pools of maps, generated through recombination: one with looser and one with tighter compactness constraints. For ease of handling, compactness is expressed through the number of cut edges, as done in the foundational work on recombination (DeFord et al. 2021). The number of cut edges is defined as the number of edges in a state's unit adjacency graph with endpoints belonging to different districts. For instance, a  $20 \times 20$  grid graph with each district composed of two adjacent columns—making 10 districts overall—will have  $9 \times 20 = 180$  cut edges. For showing the effects of compactness, the first pool has the maximum number of cut edges equal to  $2 \times 180 = 360$ , and the second pool has a bound of  $0.75 \times 180 = 135$  cut edges.

The results are given in Figure 3, which show that more compact plans lead to a lower number of seats achieved through votemandering. Recall Lemma 4.1, which shows that investing in a losing district is 3 times more beneficial in achieving fairness, while winning through campaigning marks the investment as a winning-district specific. Then, campaigning in targeted units is usually followed by their re-assignments to losing districts, as a votemandering strategy to achieve fairness benefits. Compactness limits this scope of targeting units for the campaign and subsequent reassigning, by disallowing arbitrary shapes. We elaborate more on the intuition behind this phenomenon, as we discuss the local votemandering strategies in Section 5.3. Note that Figure 3 compares votemandering objectives, as opposed to bonuses shown in Figure 2, as here two different pools are used, which also significantly affects the distribution of wins in the initial maps, and thereby the bonuses.

**4.3.3 Impact of voter turnout.** Intuitively, the ability to votemandeer is a function of how efficiently and thus also, how disproportionately we can allocate budget across the units. The parameter  $\alpha$  captures the natural voter turnout and the selective campaigning by a party strategically adds




Fig. 4. Increase in A's votemandering bonus with voter turnout




Fig. 5. Votemandering bonus with increasing Moran's I

more party votes, bounded above by the natural vote share of that party in every unit. Hence, it is straightforward to see that an increase in  $\alpha$  would provide less flexibility to votemander, and will more accurately represent the true vote share, putting less weight on the campaigned votes. We rigorously show this in effect in Figure 4 where we plot the objective with respect to increasing  $\alpha$ . With both the means and medians shifting downwards with increasing  $\alpha$ , this demonstrates that a higher voter turnout supports a better representation of social choice through not only higher volume and election credibility, but also through disallowing political parties to votemander.

**4.3.4 Impact of spatial autocorrelation of voters (Moran's I).** One may question if the current experimental setting of a geographically uniformly spread voter population is reasonable, as we often see clusters of societies divided across political and geographic lines. In this experiment, we show that clustering of this data does not have a very significant effect on the votemandering bonus. To demonstrate this, we use a popular spacial auto-correlation metric called *Moran's I* (Duchin and Walch 2021). This is a measure of the overall clustering of the spatial data. For  $(v_1, \dots, v_{|K|})$  as the vector of vote shares,  $\bar{v}$  as the average vote share,  $y_{ij}$  as a binary variable indicating adjacency of units  $i, j$  and  $Y = \sum_{i,j}^{|K|} y_{ij}$  as the number of total adjacencies, Moran's I is defined as

$$I = \frac{|K|}{Y} \frac{\sum_{i=1}^{|K|} \sum_{j=1}^{|K|} y_{ij} (v_i - \bar{v})(v_j - \bar{v})}{\sum_{i=1}^K (v_i - \bar{v})^2} \quad (22)$$

Moran's I is usually used to measure the segregation of geospatial data and it varies between  $[-1, 1]$  with -1 indicating anti-segregation, 0 with no segregation, and 1 with extreme segregation. The randomly generated voter patterns used in Section 4.3.1, 4.3.2, and 4.3.3 produce Moran's I



values in the range  $(-0.01, 0.01)$ . For the current set of experiments, we generate voter distributions with varying Moran’s  $I$  values, group those into nine bins, and plot the votemandering bonuses in Figure 5. Note that the overall party vote-shares are kept around the same value while increasing the clustering of data. The insensitivity of the votemandering bonus to variations in Moran’s  $I$  supports the experimental design decision of keeping Moran’s  $I$  near zero when varying the other redistricting factors in Sections 4.3.1, 4.3.2, and 4.3.3.

In summary, campaign budgets affect the votemandering bonus with diminishing marginal improvements, high voter turnout and stricter compactness bounds curtail votemandering, and voter clustering patterns have little effect on votemandering. Moreover, these factors together help establish the robustness of votemandering objectives to smaller uncertainties in voter-data. For the strategist party, assuming base voter inclinations at the lower end of their confidence intervals is sufficient to devise votemandering strategies. While access to  $B$ ’s budget allocation information is essential, its slow rate of effect helps if the information is stochastic. In further continuation with the discussion in Section 4.3.1, it is worth emphasizing that the fusion of gerrymandering and strategic campaigning separates the effects of pure campaigning strategies from votemandering strategies, enabling a more robust dependence of budgets and voter inclinations on votemandering.

## 5 LOCAL VOTEMANDERING

The votemandering methods in Section 3 allow district lines in the target map to deviate significantly from those in the initial map. However, in practice, considerations are made for maintaining the original community boundaries (e.g., retaining majority/minority districts) as well as structural boundaries (e.g., disallowing county splits) while drawing new district plans on the ground. Some states also demand that redistricting plans remain close to the existing plan. For example, Nebraska requires the new plan to “preserve the cores of prior districts” (Nebraska 2021), and the Wisconsin Supreme Court issued a similar “least-change” order for the 2020 cycle (Wisconsin 2021).

Inspired by these requirements, Section 5.1 introduces a *local votemandering heuristic* that conducts a local search within smaller map sections. This approach generates new plans that satisfy local proximity requirements while maintaining global fairness and budget constraints. The heuristic aims to maximize the votemandering objective, offering insights into crucial votemandering strategies as detailed in Section 5.2. Although local votemandering provides a lower bound on global votemandering performance, it is a faster method for achieving a positive votemandering bonus. Additionally, it is parallelizable and thus scales better with increasing state sizes. Section 5.3 further explores the relationship between local and global votemandering and emphasizes the significance of key strategies in understanding this connection.

### 5.1 Local Votemandering Methods

The local votemandering heuristic generates new target plans by applying small changes, or *local boundary perturbations*, to existing plans between pairs of neighboring districts. These perturbations involve exchanging units between two districts while satisfying external proximity and redistricting requirements. Unlike the top-down approach in Section 3 (hereafter referred to as the global votemandering heuristic), this bottom-up approach strategically employs local perturbations to increment the votemandering bonus by one. A district adjacency graph is formed, with edges that have a potential positive bonus—containing at least one district with initial status  $L$ —assigned weights representing (budget, fairness) costs associated with the perturbations.

The ultimate objective is to maximize the overall bonus by finding a maximum-sized matching that adheres to budget and fairness constraints. The resulting target plan is created by applying perturbations corresponding to the maximum matching. Matchings are utilized because they enable mutually exclusive perturbations between district pairs.

(a) Between neighboring districts with  $W, L$  status(b) Between neighboring districts with  $L, L$  status

Fig. 6. The three key strategies to locally improve the votemandering bonus. The shaded districts show an improvement in the number of wins of the strategist party.

The new target plan remains similar to the initial plan in terms of unit-to-district assignments, with alterations only involving districts that expect positive bonuses. Since matchings are used, the difference is quantified as, at most,  $\frac{n}{2}$  independent recombination steps away from the initial map, where  $n$  is the number of districts. The heuristic considers perturbations only between district pairs to avoid the complexity of perturbations within an arbitrary number of districts, to create plans closer to the initial plan, and to maintain the highest improvement ratio over the number of perturbed districts. Although it is technically feasible to generalize the heuristic to consider perturbing three or more districts simultaneously, the current approach focuses on pairs.

Generating bonuses within submaps—regions restricted to two adjacent districts—exposes the functionality of votemandering’s key strategies, as illustrated in Figure 6. Recall the two general votemandering ways discussed in Section 4.1. Stemming parallel to those, these strategies encompass securing a first-round win in the campaigned map and either of the second-round wins in the target map: with or without winning in the votemandered map. Depending on the initial map, any of the key strategies could be more cost-effective in terms of budget and/or fairness (or be infeasible). Given cost information about edges, we first formulate an optimization framework that produces the maximum-matching solution and subsequently discuss the formation and cost computation of the key strategies in Section 5.2.

**5.1.1 An optimization framework for local votemandering.** As illustrated above, the objective is to select the maximum number of mutually exclusive edges while adhering to the fairness and budget constraints. An edge in the optimal solution may correspond to any of the three strategies depicted in Figure 6. Without loss of generality, we assume that the budget is never sufficient to win in all districts in round-1; otherwise, finding efficient votemandering strategies becomes trivial.

Let  $n$  represent the set of nodes in the district adjacency graph, and let  $E$  denote the set of edges. Let  $E_{n_i}$  represent a set of edges (including those corresponding to all 3 strategies) incident on node  $n_i \in n$ . Let  $b_e, f_e$  denote budget and fairness costs associated with an edge  $e$ , and let  $b$  represent the budget spent to satisfy fairness constraints. Since we do not have a budget to win all districts,  $b$  can always be spent on losing districts, contributing  $3b/2$  to the fairness cost, as Lemma 4.1 demonstrates. The MIP formulation for the local votemandering heuristic can then be expressed as:

$$\begin{aligned} \max \quad & \sum_{e \in E} x_e \\ \text{s.t.} \quad & \sum_{e \in E} b_e x_e + b \leq \mathcal{B}^A \\ & \sum_{e \in E} f_e x_e - \frac{3}{2}b \leq \text{Fairness cost} \end{aligned}$$

$$\begin{aligned}
\sum_{e \in E_{n_i}} x_e &\leq 1 && \forall n_i \in n \\
x_e &\in \{0, 1\} && \forall e \in E
\end{aligned} \tag{23}$$

Note that (23) is equivalent to finding a maximum cardinality matching with two knapsack constraints. Although (23) is computationally challenging, it significantly simplifies the variable space by discarding unit-specific variables in (10). In practice, the adjacency graph instance has  $|n|$  nodes and only those edges with a positive bonus, making the instance sparse and the problem tractable. We provide an example of this heuristic in Appendix D.1.

## 5.2 Taxonomy of Local Votemandering Strategies

As shown in (1), securing a win in round-1 is only possible through budget investment, while winning in round-2 can only be achieved through target plan design, which considers the original vote shares. The fairness constraint is relevant because the new plan must satisfy it on the votemandered map, factoring in the vote shares after investments. A strategist party employs the strategies in Figure 6 to generate bonuses using its budget and redistricting abilities:

- (1) Strategy-1: Secure a votemandering bonus with an extra win only in the campaigned map, using  $A$ 's budget investment. The edge weight vector is of the form: [significant budget, insignificant fairness cost]. The implementation bottleneck is the budget needed for winning (i.e., the vote margin for  $L$  district) and finding perturbations that maintain the same  $W/L$  status in the votemandered map.
- (2) Strategy-2: Secure an extra win only in the target map using  $B$ 's budget investment. The edge weight vector is: [no budget, insignificant fairness cost]. The bottleneck is identifying perturbations that enable  $B$  to win in the votemandered map with a margin smaller than its investment, allowing  $A$  to win in the target map with original vote shares.
- (3) Strategy-3: Use boundary perturbations to secure an extra win in the target map while also achieving an additional win in the votemandered map. The edge weight vector is: [insignificant budget, significant fairness cost]. The bottleneck is the fairness cost resulting from the extra win in the votemandered map.

Other strategies, where  $A$  wins in a different set of maps than shown in Figure 6, may be considered. However, these are expensive and dominated by the key strategies when feasible. For example, winning in both campaigned and votemandered maps incurs significantly higher fairness costs than strategy-1 while providing the same bonus. Winning in both campaigned and target maps combines strategies 1 and 2, making it less likely than either. Therefore, for practical purposes, we only illustrate the three key strategies here and note the generalization of considering additional edge types in (23).

Consequently, up to three edges may exist between any two districts, each associated with a specific weight vector. Strategy-2 is particularly important because it only uses information about party  $B$ 's budget and virtually imposes no cost on party  $A$ . More details and exact calculations of edge weights can be found in Appendix D.2.

To determine the optimal edge weights as inputs for (23), i.e., costs  $b_e, f_e$  for all three strategies, optimal boundary perturbations must be specified while adhering to redistricting constraints such as compactness, contiguity, proximity, and population balance. To achieve this, we employ the randomized recombination technique once again. As with the global heuristic, we separate the problems of finding local perturbations and optimizing costs by creating a pool of plausible submaps for each edge.

For each strategy, we can now examine this pool and select the best edge—i.e., the perturbations leading to a new pair of districts with minimal costs. Finally, given a district adjacency graph, we

add up to three edges for each pair of neighboring districts, with costs equal to those of the best pairs from the corresponding pool.

### 5.3 Connection to Global Votemandering

As discussed in Section 3, the votemandering bonus  $\Delta$  describes how effectively we can votemandeer. Eq. (24) further divides  $\Delta$  into three sub-parts.

$$\begin{aligned}\Delta &= \mathcal{E}(D_0, \tilde{V}) + \mathcal{E}(\tilde{D}, V_0) - 2\mathcal{E}(D_0, V_0) \\ &= [\mathcal{E}(D_0, \tilde{V}) - \mathcal{E}(D_0, V_0)] + [\mathcal{E}(\tilde{D}, V_0) - \mathcal{E}(\tilde{D}, \tilde{V})] + [\mathcal{E}(\tilde{D}, \tilde{V}) - \mathcal{E}(D_0, V_0)]\end{aligned}\quad (24)$$

The first part of Eq. (24), i.e., the bonus resulting from the difference between the number of wins in campaigned and initial maps, depends on the optimal budget allocation of  $A$ . The second part, i.e., the difference between the target and votemandered maps, depends on  $B$ 's budget allocation. The third part, i.e., the difference between votemandered and initial maps, is essentially bounded as only a few discrete EG values are acceptable for the maps, meaning that the votemandered and initial maps can only have a small difference in their number of wins (Tam Cho 2017).

Comparing  $\Delta$  for global maxima to when the space is restricted for local votemandering, the bonus breakdown highlights the exact three areas where the local heuristic operates (approximating the optimal) through its key strategies. This also suggests that the efficiency of votemandering can be explained using the generalized versions of key strategies, i.e., without necessarily restricting to two districts. Considering the efficiency of local votemandering, note that its produced optimal plan is also achievable by the global heuristic if the plan is present in its pool of maps. In fact, relaxing its more-than-sufficient matching constraints also indicates a possibility of improvement. Despite this, the deeper and targeted local search between every pair of neighboring districts better explores the existence of key strategies and may produce outlier plans more effectively. Moreover, it works precisely by exploiting the information of the initial map. As seen in Section 6, the local votemandering heuristic works efficiently, even outperforming the global heuristic in one case.

Eq. (24) also explains why increasing compactness leads to a lower objective and, consequently, a decrease in the ability to votemandeer. Using the notion of cut edges, higher compactness means fewer cut edges, i.e., a smaller shared boundary between two districts. The efficiency of votemandering depends on the ease of unit exchange across borders via strategies 1 and 2, and more generally, via using Lemma 4.1 implications to invest to win in a campaigned map and re-structure to shift this investment to a losing district in the votemandered map. This implies a relationship between the number of cut edges and the ease of votemandering via its key strategies: With a mandated fewer number of cut edges between any two districts, it is more challenging to find units that can be exchanged, resulting in a positive votemandering bonus.

In conclusion, while the global heuristic bypasses the matching constraints, the local heuristic provides an efficient search at the district level, generating plans that closely resemble the original ones. The local search in submaps is entirely parallelizable, potentially yielding much faster results compared to the global heuristic as instance sizes grow. In practice, this runs very quickly. After performing the local search for all edges, the corresponding adjacency graph of the initial map incorporates budget and fairness bounds as variable inputs in the optimization program (23). This allows for a more targeted and efficient approach to votemandering, taking advantage of the districts in the initial map.




Fig. 7. Distribution of Votemandering Objectives

## 6 CASE STUDY: WISCONSIN STATE SENATE REDISTRICTING AFTER 2020

This section demonstrates the existence of practical votemandering strategies using Wisconsin state senate redistricting after the 2020 census. The redistricting cycle was delayed due to lawsuits, prompting the Wisconsin state legislature and the governor’s People’s Maps Commission to propose state and congressional district plans. The Wisconsin Supreme Court eventually approved the state senate and house maps drawn by the legislature (Ballotpedia 2022).

Wisconsin’s balanced partisan composition provides a suitable environment for exploring votemandering’s practical potential. The state senate has 33 seats, with approximately half up for election every two years. The Republican party (R) controls the state senate with a 21-11 majority (excluding one vacancy) and had a 51.1% statewide senate election vote share in previous elections. Conversely, the Democratic party (D) won the 2020 presidential election in Wisconsin by a 0.63% margin, and the state’s governor is a Democrat. The Democratic governor’s veto power over redistricting proposals from the Republican state legislature ensures that both parties influence the final state senate map. Additionally, the Wisconsin state constitution mandates that districts be compact, contiguous, and "bounded by county, precinct, town, or ward lines where possible" (Wisconsin 2022).

Throughout this section, we consider the 2021 governor’s office final state senate plan (GOV2021) as the initial map. Deemed fair by the Princeton Gerrymandering Project with a score of A (2022), GOV2021 serves as a reasonable starting point for our votemandering case study, despite not being enacted.

The GOV2021 plan incorporates state senate election data from 2018 and 2020 (for 17 and 16 seats, respectively). According to the 2020 census, the ideal population for each senate district is 178,598. We assume a statewide voter turnout of 65%, the average from the 2018 and 2020 elections. Both parties receive a budget to influence 13,734 voters, constituting a 1% total investment of the expected votes cast. The votemandering party strategizes its campaign investment, while the other party is assumed to allocate its investment proportionally to unit populations.

Section 6.1 investigates global votemandering effects for each major party, revealing that although both parties can benefit, the Democratic party achieves a more substantial votemandering bonus given GOV2021 as the initial map. Section 6.2 illustrates local votemandering strategies that better align with the Wisconsin Supreme Court’s goal of minimizing changes to the previous district plan. Both parties continue to benefit within this restricted strategy space, and the votemandering bonus even increases for one party.

### 6.1 Wisconsin Global Votemandering

The global votemandering heuristic expounded in Section 3 is applied to Wisconsin state senate redistricting, using a pool of 80,000 candidate target maps generated via recombination. Figure 7




Fig. 8. Distribution of Efficiency Gap measure over the pool of maps

depicts the pool’s distribution of votemandering bonus from each party’s perspective. Each histogram shows the number of maps yielding a specific votemandering bonus for the votemandering party when fixed as the target map. The EG of GOV2021, 0.1409, indicates more wasted votes for the Democratic party, giving the Democratic party more room to gain seats while maintaining a safe EG. Hence the Democratic party tends to achieve a greater votemandering bonus than does the Republican party, i.e., the distribution in Figure 7b is shifted to the right of the distribution in Figure 7a. Note the relatively high EG of GOV2021 does not necessarily imply the map is a partisan gerrymander. Figure 8 shows the distribution of EG for the pool of maps. Due to the spatial distribution of voters in Wisconsin, recombination tends to generate maps with positive (i.e., Republican-leaning) EG values. Based on the EG value of GOV2021 and the pool’s EG distribution, we impose a fairness bound of  $0.0 \leq EG \leq 0.15$  throughout the case study.


Table 2. Wisconsin Global Votemandering Characteristics: Republican and Democratic Parties

| Majority party   | Republican Votemandering |                | Democratic Votemandering |                |
|------------------|--------------------------|----------------|--------------------------|----------------|
|                  | Number of Wins           | Efficiency Gap | Number of Wins           | Efficiency Gap |
| Initial Map      | 21                       | 0.1409         | 12                       | 0.1409         |
| Campaigned Map   | 24                       | 0.2314         | 15                       | 0.0468         |
| Votemandered Map | 21                       | 0.1285         | 16                       | 0.0096         |
| Target Map       | 23                       | 0.1915         | 17                       | -0.0195        |


To show an illustration of how these votemandering objectives are attained, we next describe optimal target maps with both parties as strategists. Table 2 shows the characteristics of votemandering stages specific to our examples. Figures 9-10 visually show the initial and target maps, strategic investment of the Republican party, and the various stages of votemandering on the real state data of Wisconsin. Figures 11-12 show the same for the Democratic party.

For the Republican votemandering, the objective is 47 seats across two rounds, against 42 with no strategic investment. The Democratic votemandering bonus is 8, showing a larger improvement for the Democratic party as explained above. The proposed Republican map would show an EG of 0.1285 with the previous election data (apparently fairer than the initial map), with the actual value being 0.1915. This is reflected in the difference of 2 seats between the votemandered and target map. Most interestingly, although the investment leads to 3 new seats in the campaigned map, the new plan is such that it completely negates this effect, making the votemandered map win 21 seats.


In these cases, both parties choose outlier maps as the target maps and make the votemandered maps fair. The strategic investment for both cases is done across the map by smartly choosing units: the bottom-left part in Figure 11c is largely uniform (to make a difference only in round-1), the top-half part in both Figure 9c and 11c is non-uniform (to only affect units that remain part of




**Fig. 9.** Republican Global Votemandering: The initial map, the chosen target map, and the strategic investment of budget (with intensity indicated by the darker color)



**Fig. 10.** The Four Stages of Republican Global Votemandering, (with red and blue indicating the districts won by the Republican and Democratic parties, respectively)



**Fig. 11.** Democratic Global Votemandering: The initial map, the chosen target map, and the strategic investment of budget (with intensity indicated by the darker color)



**Fig. 12.** The Four Stages of Democratic Global Votemandering, (with red and blue indicating the districts won by the Republican and Democratic parties, respectively)

losing districts in the target map) and finally, the bottom-half part in Figure 9c is sporadic, with the right corner serving as an investment made to satisfy the fairness bound.


## 6.2 Wisconsin Local Votemandering

When forced to decide the final state senate district plan, the Wisconsin Supreme Court announced that it would seek to make as few changes as possible to the existing map. In accordance with this goal, this section applies the bottom-up local votemandering heuristic to find strategies for both parties which produce target maps close to GOV2021.

**Table 3.** Local Votemandering Characteristics: Republican and Democratic Parties

| Majority party   | Republican Votemandering |                | Democratic Votemandering |                |
|------------------|--------------------------|----------------|--------------------------|----------------|
|                  | Number of Wins           | Efficiency Gap | Number of Wins           | Efficiency Gap |
| Initial Map      | 21                       | 0.1409         | 12                       | 0.1409         |
| Campaigned Map   | 21                       | 0.1478         | 14                       | 0.1072         |
| Votemandered Map | 21                       | 0.1442         | 12                       | 0.1014         |
| Target Map       | 25                       | 0.2513         | 20                       | -0.0888        |

Both parties can gain advantages from local votemandering, although each party employs a significantly different approach. Figure 13 illustrates the GOV2021 district adjacency graph (13a), the best Republican strategy discovered (13b), and the best Democratic strategy discovered (13c). Node colors represent the party with a higher vote-share in each district. Edges with nonzero weights indicate strategy-1 edges, while 0 weights correspond to strategy-2, and ‘FC’ denotes fairness costs associated with strategy-3.



**Fig. 13.** Local Votemandering: The district adjacency graph and the maximum matching solutions

The maximum matching solution for the Republican party (Figure 13b) yields a votemandering bonus of four seats through strategy-2 improvements (i.e., without incurring monetary or fairness costs). Consequently, the Republicans’ local votemandering bonus, which results in 46 seats across the two elections, is one seat fewer than their global votemandering bonus of 47 seats. The maximum matching solution for the Democratic party (Figure 13c) attains a votemandering bonus of ten seats, surpassing their global votemandering bonus of eight, via two strategy-1, six strategy-2, and two strategy-3 edges. Strategy-1 edges consume a budget of 11,347 to boost Democratic voter turnout. Table 3 details the map characteristics for each stage of local votemandering by each party, while Figures 19-22 in Appendix E display the corresponding maps and campaign strategies.

Local votemandering strategies for both parties are not only computationally more efficient than their global votemandering counterparts but also generate target maps closer to GOV2021. As described in Section 5, the closeness parameter is adjustable, allowing for the optimization of the



bonus while producing maps as close to the initial map as desired. Notably, local votemandering relies extensively on strategy-2 improvements, which exploit the other party's campaign investments. Overall, the case study demonstrates the vulnerability of Wisconsin state senate redistricting to global and local votemandering.

## 7 CONCLUSIONS AND FUTURE DIRECTIONS

In this study, we introduce the concept of *votemandering*, a combination of strategic campaigning and gerrymandering employed to deceive fairness measures and increase seat-share across multiple elections. Focusing on the efficiency gap (EG) as a fairness metric, we establish sufficient conditions for a positive votemandering bonus (Theorem 4.2) and present an efficient heuristic for identifying votemandering strategies (Algorithm 1, Proposition 3.1, Theorem 4.3).

Through computational experiments, we investigate the impact of campaign budget, compactness, voter turnout, and spatial autocorrelation of voters on votemandering efficacy. Our findings indicate that enhancing voter turnout and compactness, parameters seemingly unrelated to partisan fairness, can potentially mitigate the influence of votemandering. A case study of Wisconsin state senate redistricting illustrates practical votemandering strategies for both parties, emphasizing its applicability beyond hypothetical scenarios and into real-world instances.

To further demonstrate the practicality of votemandering, we introduce *local votemandering*, which allows the party controlling redistricting to make minor adjustments to a limited number of district boundaries. Our heuristic efficiently discovers profitable district maps with minimal district pair recombination. In the Wisconsin case study, local votemandering yields a higher votemandering bonus for one party compared to the global, pool-based heuristic.

Future research may explore votemandering in the context of alternative fairness measures by utilizing the general framework outlined here. Determining the most effective fairness measure, or a combination thereof, to prevent votemandering remains an open question. Our votemandering model currently allows for strategic campaigning only in the first election; expanding it to include the second election would increase realism but also complexity. Further extensions might examine votemandering across more than two election rounds, which would necessitate accounting for migration patterns and shifts in voter sentiment. From a technical perspective, the computation of optimal votemandering strategies, perhaps within restricted settings, is an interesting problem. Moreover, allowing both parties to strategically allocate their campaign budgets would introduce a generalized version of the Colonel Blotto game in the redistricting context, presenting challenges such as breaking ties and addressing increased computational complexity.

Finally, our work carries significant policy and practical implications. We demonstrate that strategic campaigning can substantially impact redistricting outcomes despite the presence of fairness constraints and that the sole use of EG as a fairness measure may be insufficient. Our results advocate for additional measures that account for strategic behavior and policies that curtail manipulative campaigning practices. Ultimately, this study underscores the importance of considering strategic behavior when designing and evaluating redistricting processes, and the need for ongoing research on the consequences of political manipulation in democratic systems.

## REFERENCES

- Bharat Adsul, Ch Babu, Jugal Garg, Ruta Mehta, Milind Sohoni, et al. 2010. Nash equilibria in Fisher market. In *International Symposium on Algorithmic Game Theory*. Springer, Athens, Greece, 30–41.
- Boris Alexeev and Dustin G Mixon. 2018. An impossibility theorem for gerrymandering. *The American Mathematical Monthly* 125, 10 (2018), 878–884.
- Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. 2021. Competitive equilibrium with indivisible goods and generic budgets. *Mathematics of Operations Research* 46, 1 (2021), 382–403.
- Ballotpedia. 2022. Redistricting in Wisconsin after the 2020 census. [https://ballotpedia.org/Redistricting\\_in\\_Wisconsin\\_after\\_the\\_2020\\_census](https://ballotpedia.org/Redistricting_in_Wisconsin_after_the_2020_census) Jan 24, 2023.
- John J Bartholdi, Craig A Tovey, and Michael A Trick. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6, 3 (1989), 227–241.
- Jeffrey T Barton. 2018. Improving the efficiency gap. *Math Horizons* 26, 1 (2018), 18–21.
- Soheil Behnezhad, Avrim Blum, Mahsa Derakhshan, MohammadTaghi HajiAghayi, Mohammad Mahdian, Christos H Papadimitriou, Ronald L Rivest, Saeed Seddighin, and Philip B Stark. 2018. From battlefields to elections: Winning strategies of blotto and auditing games. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, Philadelphia, PA, 2291–2310.
- Soheil Behnezhad, Sina Dehghani, Mahsa Derakhshan, MohammadTaghi HajiAghayi, and Saeed Seddighin. 2017. Faster and simpler algorithm for optimal strategies of Blotto game. In *Thirty-first AAAI conference on Artificial Intelligence*. AAAI, San Francisco, CA, 369–375.
- Gerdus Benadè, Ariel D Procaccia, and Jamie Tucker-Foltz. 2021. You can have your cake and redistrict it too. *ACM Transactions on Economics and Computation* –, – (2021), 1–26.
- Mira Bernstein and Moon Duchin. 2017. A formula goes to court: Partisan gerrymandering and the efficiency gap. *Notices of the AMS* 64, 9 (2017), 1020–1024.
- Steve Bickerstaff et al. 2020. *Election Systems and Gerrymandering Worldwide*. Springer, Cham, Switzerland.
- Simina Brânzei, Vasilis Gkatzelis, and Ruta Mehta. 2017. Nash social welfare approximation for strategic agents. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, New York, NY, 611–628.
- NaLette M Brodnax and Piotr Sapiezynski. 2022. From Home Base to Swing States: The Evolution of Digital Advertising Strategies during the 2020 US Presidential Primary. *Political Research Quarterly* 75, 2 (2022), 460–478.
- Brian Brubach, Aravind Srinivasan, and Shawn Zhao. 2020. Meddling metrics: the effects of measuring and constraining partisan gerrymandering on voter incentives. In *Proceedings of the 21st ACM Conference on Economics and Computation*. ACM, New York, NY, 815–833.
- Brennan Center. 2022. Informational Brief of Redistricting. <https://www.brennancenter.org/our-work/research-reports/who-controlled-redistricting-every-state> May 9, 2023.
- Christopher P Chambers, Alan D Miller, and Joel Sobel. 2017. Flaws in the efficiency gap. *JL & Pol.* 33 (2017), 1.
- Maria Chikina, Alan Frieze, and Wesley Pegden. 2017. Assessing significance in a Markov chain without mixing. *Proceedings of the National Academy of Sciences* 114, 11 (2017), 2860–2864.
- Amittai Cohen-Zemach, Yoad Lewenberg, and Jeffrey S Rosenschein. 2018. Gerrymandering over graphs. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 274–282.
- Pennsylvania Supreme Court. 2022. Pennsylvania case. <https://www.pacourts.us/assets/opinions/Supreme/out/J-20-2022mo.pdf?cb=1> Jan 24, 2023.
- United States Supreme Court. 2018. Gill v. Whitford. [https://www.supremecourt.gov/opinions/17pdf/16-1161\\_dc8f.pdf](https://www.supremecourt.gov/opinions/17pdf/16-1161_dc8f.pdf) Jan 24, 2023.
- United States Supreme Court. 2019. Rucho v Common Cause. [https://www.supremecourt.gov/opinions/18pdf/18-422\\_9o11.pdf](https://www.supremecourt.gov/opinions/18pdf/18-422_9o11.pdf) Jan 24, 2023.

- D DeFord, M Duchin, and J Solomon. 2021. Recombination: A family of Markov chains for redistricting. Harvard Data Science Review.
- Moon Duchin and Olivia Walch. 2021. Political Geometry.
- Eduard Eiben, Fedor Fomin, Fahad Panolan, and Kirill Simonov. 2020. Manipulating districts to win elections: fine-grained complexity. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. AAAI Press, Palo Alto, CA, 1902–1909.
- Karl Evers-Hillstrom. 2021. *Most expensive ever: 2020 election cost 14.4 billion*. OpenSecrets News. <https://www.opensecrets.org/news/2021/02/2020-cycle-cost-14p4-billion-doubling-16/> Jan 27, 2023.
- Allan Gibbard. 1973. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society* 1, 1 (1973), 587–601.
- Wes Gurnee and David B Shmoys. 2021. Fairmandering: A column generation heuristic for fairness-optimized political districting. In *SIAM Conference on Applied and Computational Discrete Algorithms (ACDA21)*. SIAM, Philadelphia, PA, 88–99.
- William CR Horncastle. 2020. *The scale of US election spending explained in five graphs*. The Conversation. <https://theconversation.com/the-scale-of-us-election-spending-explained-in-five-graphs-130651> Jan 27, 2023.
- Kosuke Imai and Aaron Strauss. 2011. Estimation of heterogeneous treatment effects from randomized experiments, with application to the optimal planning of the get-out-the-vote campaign. *Political Analysis* 19, 1 (2011), 1–19.
- Takehiro Ito, Naoyuki Kamiyama, Yusuke Kobayashi, and Yoshio Okamoto. 2021. Algorithms for gerrymandering over graphs. *Theoretical Computer Science* 868 (2021), 30–45.
- Jeffrey A Karp, Susan A Banducci, and Shaun Bowler. 2008. Getting out the vote: Party mobilization in a comparative perspective. *British Journal of Political Science* 38, 1 (2008), 91–112.
- Sam Kean. 2018. The Flaw in America’s Holy Grail Against Gerrymandering. The Atlantic.
- Martin Lackner and Piotr Skowron. 2018. Approval-Based Multi-Winner Rules and Strategic Voting.. In *IJCAI*. IJCAI, Stockholm, Sweden, 340–346.
- Zeph Landau, Oneil Reid, and Ilona Yershov. 2009. A fair division solution to the problem of redistricting. *Social Choice and Welfare* 32, 3 (2009), 479–492.
- Omer Lev and Yoad Lewenberg. 2019. “Reverse Gerrymandering”: Manipulation in Multi-Group Decision Making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. AAAI Press, Palo Alto, CA, 2069–2076.
- Yoad Lewenberg, Omer Lev, and Jeffrey S Rosenschein. 2017. Divide and conquer: Using geographic manipulation to win district-based elections. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 624–632.
- Federica Liberini, Michela Redoano, Antonio Russo, Angel Cuevas, and Ruben Cuevas. 2020. Politics in the Facebook era-evidence from the 2016 US presidential elections.
- Scott T Macdonell and Nick Mastronardi. 2015. Waging simple wars: a complete characterization of two-battlefield Blotto equilibria. *Economic Theory* 58, 1 (2015), 183–216.
- Constitution Missouri. 2022. State of Missouri Constitution. <https://www.sos.mo.gov/CMSImages/Publications/CurrentMissouriConstitution.pdf?v=202212> Jan 24, 2023.
- David P Myatt. 2007. On the theory of strategic voting. *The Review of Economic Studies* 74, 1 (2007), 255–281.
- John F Nagle. 2019. What criteria should be used for redistricting reform? *Election Law Journal: Rules, Politics, and Policy* 18, 1 (2019), 63–77.
- Nebraska. 2021. Informational Brief of Redistricting. <http://news.legislature.ne.gov/lrd/redistricting/informational-brief-of-redistricting/> Jan 31, 2023.
- Benjamin Plener Cover. 2018. Quantifying partisan gerrymandering: An evaluation of the efficiency gap proposal. *Stan. L. Rev.* 70 (2018), 1131.

- Daniel D Polsby and Robert D Popper. 1991. The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale L. & Pol'y Rev.* 9 (1991), 301.
- Princeton Gerrymandering Project. 2022. Princeton Gerrymandering Project. <https://gerrymander.princeton.edu/redistricting-report-card/?planId=recALbYWseXmEg1eG> Jan 24, 2023.
- Laura Royden and Michael Li. 2017. *Extreme maps*. Brennan Center for Justice at New York University School of Law, New York, NY.
- Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory* 10, 2 (1975), 187–217.
- Amartya Sen. 1986. Social choice theory. *Handbook of mathematical economics* 3 (1986), 1073–1181.
- Nicholas O Stephanopoulos and Eric M McGhee. 2015. Partisan gerrymandering and the efficiency gap. *U. Chi. L. Rev.* 82 (2015), 831.
- Alexander J Stewart, Mohsen Mosleh, Marina Diakonova, Antonio A Arechar, David G Rand, and Joshua B Plotkin. 2019. Information gerrymandering and undemocratic decisions. *Nature* 573, 7772 (2019), 117–121.
- Rahul Swamy, Douglas M King, and Sheldon H Jacobson. 2023. Multiobjective optimization for politically fair districting: A scalable multilevel approach. *Operations Research* 71, 2 (2023), 536–562.
- Wendy K Tam Cho. 2017. Measuring Partisan Fairness: How Well Does the Efficiency Gap Guard Against Sophisticated as well as Simple-Minded Modes of Partisan Discrimination? *University of Pennsylvania Law Review Online* 166, 1 (2017), 2.
- Ray Wallin. 2017. *The Equal Districts Efficiency Gap: Fundamental Gerrymandering Theory and an Analysis and Modification of the Efficiency Gap*. SSRN. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3027578](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3027578) Jan 20, 2023.
- Wisconsin. 2022. Wisconsin constitution. [https://docs.legis.wisconsin.gov/constitution/wi\\_unannotated](https://docs.legis.wisconsin.gov/constitution/wi_unannotated) Jan 24, 2023.
- Supreme Court Wisconsin. 2021. Wisconsin Supreme Court. <https://www.wicourts.gov/sc/opinion/DisplayDocument.pdf?content=pdf&seqNo=459269> Feb 8, 2023.
- Brahim Zarouali, Tom Dobber, Guy De Pauw, and Claes de Vreese. 2020. Using a personality-profiling algorithm to investigate political microtargeting: assessing the persuasion effects of personality-tailored ads on social media. *Communication Research* 49, 8 (2020), 1066–1091.

## FUNDING AND ACKNOWLEDGMENTS

The authors are grateful to Kiera Dobbs and Raunak Sengupta for helpful discussions. The second author was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE–1746047, while being a student at University of Illinois Urbana-Champaign. The third author is supported in part by the Air Force Office of Scientific Research [Grant FA9550-19-1-0106]. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the Air Force Office of Scientific Research.

## A NOTATION

| Notation   | Definition  |
|--|---|
| $\mathcal{D}$  | Set of district plans $\{D\}$   |
| $\mathcal{V}$  | Set of voter ballots $\{V\}$  |
| $D_0, V_0$   | Original plan, original vote data   |
| $\bar{D}, \bar{V}$   | New plan, new vote data   |
| $\mathcal{E}$  | Election function $\mathcal{D} \times \mathcal{V} \rightarrow \mathbb{N}$ |
| $f$  | Fairness function $\mathcal{D} \times \mathcal{V} \rightarrow \mathbb{R}$ |
| $\delta$   | Fairness threshold; the map is fair if $f(D, V) \leq \delta$              |
| $K$  | Set of units  |
| $n$  | Number of districts   |
| $z_{ij}^r$   | Indicates 1 if unit $j$ is assigned to the district with center $i$       |
| $v_{init,k}^A, v_{init,k}^B$   | Vote shares of party $A$ and $B$ in unit $k$                              |
| $\alpha$   | Fractional baseline voter turnout   |
| $\mathcal{B}^A, \mathcal{B}^B$   | Party $A$ and $B$ 's GOTV budgets   |
| $b_k^A, b_k^B$   | Budget allocation by $A$ and $B$ in unit $k$                              |
| $\hat{s}_i^A, \hat{s}_i^B$   | Indicate 1 for $A$ 's wins in campaigned, target maps                     |
| $\mathcal{P}$  | Pool of district plans  |
| $\mathcal{I} = \{I_1, \dots, I_n\}, \mathcal{J} = \{J_1, \dots, J_n\}$ | District assignments in the original, new plans                           |
| $\hat{x}_{I_i}, \hat{y}_{J_j}$   | Indicate 1 for $A$ 's wins in the campaigned, votemandered maps           |
| $(V_{init,I}^A, V_{init,I}^B)$ and $(V_I^A, V_I^B)$                    | Pre-campaigning and post-campaigning votes in district $I$                |
| $\mathcal{W}(B - A)_I$   | Difference between wasted votes in district assignment $I$                |
| $\Delta$   | Votemandering bonus   |

Table 4. Table of Notation

## B SECTION 3 DETAILS

### B.1 Votemandering Illustration

As an illustration of the votemandering phenomenon, consider a hypothetical state comprised of a  $10 \times 10$  grid of equipopulous counties to be partitioned into  $k = 5$  districts. We start with the initial map with given unit vote shares as marked in Figure 14a. Both parties have close state-wide total vote proportions (51%, 49%), spread uniformly throughout the map. Out of the total vote shares in each unit, 50% votes are cast without any campaign influence, i.e., it is the baseline voter turnout, allowing the remaining to be added through GOTV campaigning. In the initial map, party  $A$  wins three districts out of 5, as marked by the units with the symbol 'A'. The resulting efficiency gap (EG) of the initial map is 0.072. In this case of 5 districts, we call a map 'fair' if its EG is less than or equal to 0.20 (Stephanopoulos and McGhee 2015). Thus, any map proposed by  $A$  for the second round must have its EG less than 0.20.

In this example, party  $B$  adds zero extra votes through campaigning. This is done to illustrate only the effect of  $A$ 's campaigning on the election results and the proposed plan for round-2. Party  $B$ 's campaigning has a different type of effect on votemandering as discussed in Section 4. If party  $A$  does not do any strategic campaigning or propose a different plan in round-2, it would win 3 seats in each round, making a total of 6 wins in two rounds.  $A$  may choose to propose the plan in Figure 16b, leading to 4 wins in round-2. However, the EG of this plan using the original vote shares, i.e., that of the target map is 0.28, making it an unfair proposal. We show that  $A$  can successfully votemand in this case: via strategic campaigning, it manages to claim the fairness of this plan




Fig. 14. Initial map and the proposed map for round-2

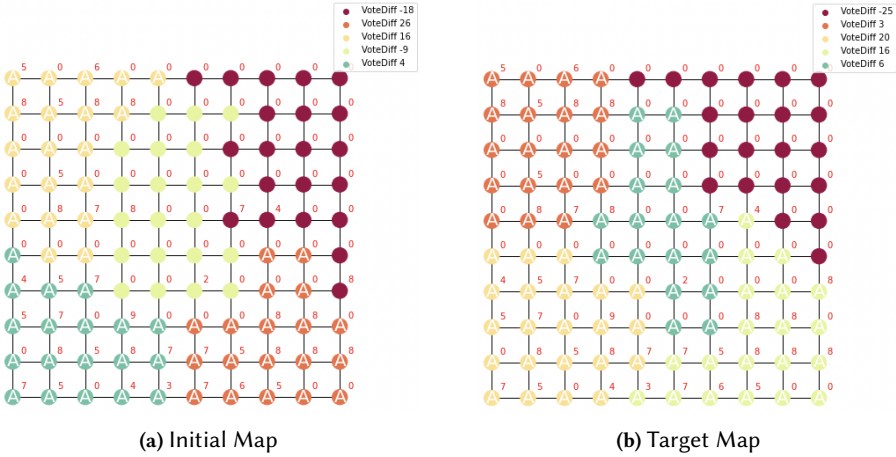


Fig. 15. Investments as seen on both maps

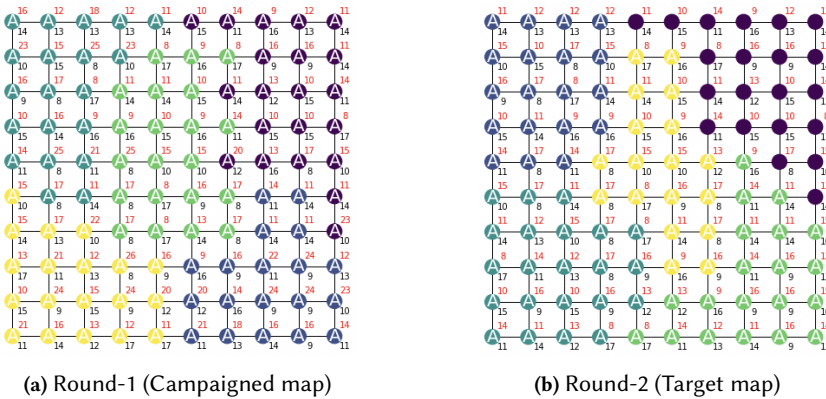


Fig. 16. Final results

(winning 4 seats in round-2), while simultaneously winning all 5 seats in round-1 elections. Hence, as opposed to 6 without strategic campaigning,  $A$  can win a total of 9 seats in two rounds.

We further illustrate the strategy implemented by  $A$  through Figure 15. The chosen plan for round-2 has multiple unit-to-district assignments same as that of the original plan, i.e., the two plans are not drastically different from each other. We see that  $A$  has majorly invested in the districts it was already winning in the initial map. It also invests a sufficient budget in the districts it initially loses on, thereby winning all 5 in round-1 elections. This investment is clever: there is no investment in the units that are part of the losing district in the target map. The budget is allocated in such a way that we win that district (top-right positioned) in round-1 elections and we again lose it in the target map as well as the votemandered map (after redrawing its boundaries). All investment then becomes a part of the winning districts in round-2, according to the new plan.

To maintain the EG bound in the votemandered map,  $A$  needs to lose in at least one district. In the target map (using original vote shares),  $A$  loses in one district and the map attains an EG score of 0.28, but the proposed votemandered map (using updated vote shares) gets an acceptable score of 0.195. The final results can be seen in Figure 16. Using a budget of 250 for a state with a total voter population 5000, i.e., by influencing just 5% of the total voter population,  $A$  can successfully votemandeer by winning 9/10 seats in two rounds. Moreover, if we extend this to 10 election rounds per redistricting cycle (as in the U.S.), the differences become starker, with 41 wins ( $5 + 4 \times 9$ ) through votemandering as opposed to 30 wins ( $3 \times 10$ ) without strategizing (out of 50).

## B.2 Proof of Proposition 3.1

PROOF. Algorithm 1 is a straightforward loop over all candidate plans in  $\mathcal{P}$ . For each candidate plan  $D_i$ , Line 10 computes the optimal objective value of (10) when  $D_i$  is used. The `best_plan` variable maintains the candidate plan with maximum objective value, `best_obj`, among plans considered so far.

It remains to prove the correctness of the termination condition in Line 7. Observe the objective of 10 decomposes into round-1 wins in the campaigned map,  $s^1 \equiv \sum_{i \in K} \hat{s}_i^1 \equiv \mathcal{E}(D_0, \tilde{V})$ , and round-2 wins in the target map,  $s^2 \equiv \sum_{i \in K} \hat{s}_i^2 \equiv \mathcal{E}(\tilde{D}, V_0)$ . Fixing  $\tilde{D} = D_i$  determines  $s^2$  via the  $z_{ik}^A$  variables, and  $s^1$  becomes a function of the budget allocation variables  $b_k^A$ . The value of  $s_{\max}^1$  computed in Line 3 is found by removing the fairness constraints and second objective term from (10), so  $s_{\max}^1$  is an upper bound on  $s_i^1$  for any given plan  $D_i$ .

Suppose the loop in Algorithm 1 terminates in iteration  $j$  with `best_plan` =  $D_{i^*}$  for some  $i^* < j$ . For any  $\ell \in \{j, j+1, \dots, N\}$ , the maximum objective value using new plan  $D_\ell$  is

$$\begin{aligned} s^2 + \max_{\{b_k^A\}_k: \text{feasible in (10) with } D_\ell} (s^1) &\leq s_{\max}^1 + s^2 \\ &= s_{\max}^1 + \mathcal{E}(D_\ell, V_0) \\ &\leq s_{\max}^1 + \mathcal{E}(D_j, V_0) && \text{(by sorting of } \mathcal{P} \text{)} \\ &< \text{best\_obj} && \text{(by Line 7)} \\ &= s_{i^*}^1 + \mathcal{E}(D_{i^*}, V_0). \end{aligned}$$

Hence none of the candidate plans omitted from consideration by loop termination could achieve a larger objective value for (10) than `best_plan`. This completes the proof of Proposition 3.1.  $\square$

## C SECTION 4 DETAILS

### C.1 Proof of Lemma 4.1

LEMMA 4.1. *The actions of campaigning and vote shifts have an impact on the difference between wasted votes, i.e.,  $\mathcal{W}(B - A)$ , as given in Table 1.*

**Table 5.** Change in the difference between wasted votes (EG) as new votes are added/shifted

|   | <i>Action</i>   | <i>Impact on <math>\mathcal{W}(B - A)</math></i> |
|---|---|--|
| 1 | <i>Wasting an additional vote on a losing district</i>                | $-3/2$   |
| 2 | <i>Wasting an additional vote on a winning district</i>               | $-1/2$   |
| 3 | <i>Winning a district <math>I</math> through campaigning</i>          | $(3V_I^A + V_I^B)/2$                             |
| 4 | <i>Shift <math>x</math> votes from a winning to a losing district</i> | $-x$   |
| 5 | <i>Shift <math>x</math> votes from a losing to a winning district</i> | $x$  |

PROOF. We discuss each action from Table 1 and its impact on the difference between wasted votes below.

- (1) Wasting an additional vote on a winning district  $I$ : Before adding the extra vote, with votes  $V_I^A$  and  $V_I^B$  for both parties, the wasted votes are as follows

$$w_I^A = V_I^A - \frac{V_I^A + V_I^B}{2}, \quad w_I^B = V_I^B \quad (25)$$

Then, the initial difference between wasted votes is

$$\mathcal{W}(B - A)_{init}(I) = V_I^B - V_I^A + \frac{V_I^A + V_I^B}{2} \quad (26)$$

After adding a vote to party  $A$ , the wasted votes are updated as

$$w_I^A = V_I^A + 1 - \frac{V_I^A + 1 + V_I^B}{2}, \quad w_I^B = V_I^B \quad (27)$$

Then,

$$\mathcal{W}(B - A)_{final}(I) = V_I^B - V_I^A + \frac{V_I^A + V_I^B}{2} - \frac{1}{2} = \mathcal{W}(B - A)_{init}(I) - \frac{1}{2} \quad (28)$$

Resulting in  $\Delta\mathcal{W}(B - A) = -\frac{1}{2}$ . In words, if party  $A$  adds a vote to a winning district, the difference between wasted votes changes by half a vote.

- (2) Wasting an additional vote on a losing district: Before adding the extra vote, with votes  $P_A$  and  $P_B$  for both parties, the wasted votes are as follows:

$$w_I^A = V_I^A, \quad w_I^B = V_I^B - \frac{V_I^A + V_I^B}{2} \quad (29)$$

Then,

$$\mathcal{W}(B - A)_{init}(I) = V_I^B - \frac{V_I^A + V_I^B}{2} - V_I^A \quad (30)$$

After adding a vote to party  $A$ , the wasted votes are updated:

$$w_I^A = V_I^A + 1, \quad w_I^B = V_I^B - \frac{V_I^A + V_I^B + 1}{2} \quad (31)$$

Then,

$$\mathcal{W}(B - A)_{final}(I) = V_I^B - \frac{V_I^A + V_I^B}{2} - V_I^A - \frac{3}{2} = \mathcal{W}(B - A)_{init}(I) - \frac{3}{2} \quad (32)$$

Resulting in  $\Delta\mathcal{W}(B - A) = -\frac{3}{2}$ . For each vote added by  $A$  in a losing district, the difference between wasted votes changes by  $-3/2$  votes. Clearly, if  $A$  wants to decrease  $W$  to satisfy the EG bound, it is more beneficial to waste votes in a losing district.



- (3) Winning a district through campaigning: To win district  $I$ ,  $A$  just needs to add  $V_I^B - V_I^A$  votes (assuming ties break in favor of  $A$ ). Initially, the wasted votes difference is:

$$w_I^A = V_I^A, \quad w_I^B = V_I^B - \frac{V_I^A + V_I^B}{2} \quad (33)$$

$$\mathcal{W}(B - A)_{init}(I) = V_I^B - \frac{V_I^A + V_I^B}{2} - V_I^A \quad (34)$$

After party  $A$  adds  $V_I^B - V_I^A$  votes to win the district, the difference between wasted votes is updated as:

$$w_I^A = 0 \quad w_I^B = V_I^B \quad (35)$$

$$\mathcal{W}(B - A)_{final}(I) = V_I^B \quad (36)$$

Then, the change in  $W$  is computed as:

$$\Delta \mathcal{W}(B - A) = V_I^B - \left( V_I^B - \frac{V_I^A + V_I^B}{2} - V_I^A \right) = \frac{3V_I^A + V_I^B}{2} \quad (37)$$

Thus, given the initial vote count for district  $I$ , we can compute the change in the difference between wasted votes as  $\frac{3V_I^A + V_I^B}{2}$ .

- (4) Shift  $x$  votes from a winning to a losing district: Using the analysis for  $\Delta \mathcal{W}(B - A)$  where  $A$  wastes votes on a winning district, we compute  $\mathcal{W}(B - A)$  when its reverse operation is performed. That is, when  $A$  removes  $x$  extra votes from district  $i$ ,  $\Delta \mathcal{W}(B - A)$  is:

$$\Delta \mathcal{W}(B - A) = \frac{x}{2} \quad (38)$$

Further, when  $A$  distributes these votes in  $B$ 's winning district  $J$ ,  $\Delta \mathcal{W}(B - A)$  is updated as:

$$\Delta \mathcal{W}(B - A) = \frac{x}{2} + \frac{-3}{2} \times x = -x \quad (39)$$

If  $A$  chooses to redistribute these votes to another  $A$ 's winning district  $l$ , we get:

$$\Delta \mathcal{W}(B - A) = \frac{x}{2} + \frac{-1}{2} \times x = 0 \quad (40)$$

- (5) Shift  $x$  votes from a losing to a winning district: Similar to the previous case, shift from a losing to a winning district results in a difference of  $\Delta \mathcal{W}(B - A) = x$ .

This concludes the discussion of the impact of  $A$ 's strategies on fairness.  $\square$

## C.2 Proof of Theorem 4.2

**THEOREM 4.2.** *For any vote-share distribution and a corresponding fair initial map with assignment  $I$ , the existence of strategies leading to a positive votemandering bonus is guaranteed if*

- (1) *A feasible, contiguous map with assignment  $\mathcal{J}$  exists with a higher number of wins than the initial map.*
- (2) *The voter turnout  $\alpha$  satisfies:*

$$\alpha \leq 1 - \left( \frac{2\Delta \mathcal{W}(B - A)_{I \rightarrow \mathcal{J}}}{\sum_{J \in \mathcal{J}(W)} \sum_{j \in J} v_j^A} \right)$$

where  $\Delta \mathcal{W}(B - A)_{I \rightarrow \mathcal{J}}$  is the change in the difference between wasted votes from assignment  $I$  to  $\mathcal{J}$ , and  $v_j^A$  are party  $A$ 's votes in unit  $j$ .

PROOF. Since assignment  $\mathcal{J}$  has a higher number of wins for the majority party compared to the initial map, the corresponding votemandering bonus is always positive. That is, even if we maintain the same number of wins in the first round after campaigning, the second round has an improvement. For a map with assignment  $\mathcal{J}$ , we then need to establish fairness by showing

$$\mathcal{W}(B - A)_{\mathcal{J}} = \mathcal{W}(B - A)_{\mathcal{I}} + \Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}} + [\text{campaign wasted votes}] \leq \text{EG bound} \quad (41)$$

We know that the initial map is fair, i.e.,  $\mathcal{W}(B - A)_{\mathcal{I}} \leq \text{EG bound}$ . Since  $\mathcal{J}$  has higher number of wins than  $\mathcal{I}$ , wlog,  $\Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}}$  is considered positive. If it's negative and makes  $\mathcal{W}(B - A)_{\mathcal{J}} \leq -\text{EG bound}$ , we naturally establish that the map gives more advantage to the minority party as compared to the initial map (and is acceptable), making a trivial case. If  $\mathcal{W}(B - A)_{\mathcal{J}} \geq \text{EG bound}$ , then we use Table 1 to spend budget to satisfy the bound. However, this budget allocation is not trivial: the allocation of budget should not change the  $W/L$  status of the districts in  $\mathcal{J}$  and it needs to satisfy the individual unit voter-turnout constraints. Intuitively, the constraints may get violated under special cases like very high voter turnout or the election mandate hugely tilting towards a party. We next deduce sufficient conditions that allow such budget allocation to occur.

For districts in  $\mathcal{J}(W)$ ,  $\mathcal{J}(L)$ , i.e., the winning and losing districts respectively, the capacities of budget allocation can be written as

$$\text{Total capacity of winning districts} = \sum_{J \in \mathcal{J}(W)} \sum_{j \in J} (1 - \alpha)v_i^A = c_1 \quad (42)$$

$$\text{Total capacity of losing districts} = \sum_{J \in \mathcal{J}(L)} \min\left\{ \sum_{j \in J} (1 - \alpha)v_i^A, V_j^B - V_j^A \right\} = c_2 \quad (43)$$

Both  $c_1$  and  $c_2$  can be computed using the assignment  $\mathcal{J}$ . Using Table 1, this translates to a bound on the effect on wasted votes:

$$\text{Max. achievable difference in } \mathcal{W}(B - A) \text{ through budget allocation} = -\frac{1}{2}c_1 - \frac{3}{2}c_2$$

Then, the sufficient condition for votemandering becomes

$$\mathcal{W}(B - A)_{\mathcal{I}} + \Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}} - \frac{1}{2}c_1 - \frac{3}{2}c_2 \leq \text{EG bound} \quad (44)$$

We can simplify this condition further to deduce a (comparatively stringent) sufficient condition on  $\alpha$  by asking if allocating only on  $\mathcal{J}(W)$  can satisfy the bound:

$$\begin{aligned} 0 &\geq \Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}} - \frac{1}{2}c_1 \\ c_1 &= \sum_{J \in \mathcal{J}(W)} \sum_{j \in J} (1 - \alpha)v_i^A \geq 2\Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}} \\ (1 - \alpha) &\geq \left( \frac{2\Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}}}{\sum_{J \in \mathcal{J}(W)} \sum_{j \in J} v_i^A} \right) \\ \therefore \alpha &\leq 1 - \left( \frac{2\Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}}}{\sum_{J \in \mathcal{J}(W)} \sum_{j \in J} v_i^A} \right) \end{aligned} \quad (45)$$

Thus, we can see that the lesser the difference between the wasted votes of  $\mathcal{I}$  and  $\mathcal{J}$ , i.e.,  $\Delta\mathcal{W}(B - A)_{\mathcal{I} \rightarrow \mathcal{J}}$ , the higher the voter turnout votemandering strategies can handle.  $\square$

### C.3 Proof of Theorem 4.3

THEOREM 4.3. *Let  $\mathcal{P}$  be a pool of  $N$  candidate target district plans such that  $\mathcal{P}$  is a subset of feasible, but not necessarily fair,  $n$ -districts plans. A plan in  $\mathcal{P}$  that maximizes the votemandering bonus may be found in  $\text{poly}(N, n)$  time.*

PROOF. We elaborate on each of the steps outlined in the main body. Careful observation of the fairness MIP reveals that the first set of constraints (4)-(7) track the effects of budget on updating vote shares, the second set of constraints (13)-(14) translates these effects into the  $W/L$  status of campaigned and votemandered map, the third set (15)-(16) further translates this into the fairness evaluation of the districts and finally the state in (17). The objective only uses the variables  $\hat{x}_I$  that indicates the  $W/L$  status in the campaigned map.

We use  $b_k$  variables (over the space of  $K$  units) for capturing the budget investment in each unit. Since all remaining constraints and objective use the win/lose variables ( $\hat{x}_I, \hat{y}_J$ ) using the set notation, we can simplify the  $b_k$  variables space to just capture the budget invested in *pieces* formed by overlapping each district in  $\mathcal{I}$  with each district in  $\mathcal{J}$ . We let variables  $z_{IJ}$  denote the budget invested in the set of units that belong to district  $I$  and  $J$  in round-1 and 2, respectively. As outlined in the idea sketch, the spending capacity of each piece is defined as  $c_{IJ} = \min(\text{division voter-turnout capacity, budget needed to win the round-2 district it is part of (if it's a losing district)})$ .

Further, when we fix the  $W/L$  status in the votemandered map, i.e., variables  $\hat{y}_J$ , we simplify the constraints by eliminating the  $\tau$  variables. Finally, we see that the MIP is reduced to the following linear formulation:

$$\begin{aligned}
& \max \sum_{I \in \mathcal{I}} \hat{x}_I \\
& \text{s.t. } 1 - M(1 - \hat{x}_I) \leq (V_I^A - V_I^B) \leq M\hat{x}_I && \forall I \in \mathcal{I} \\
& \sum_{I \in \mathcal{I}} z_{IJ} \leq V_J^B - V_J^A && \forall J \in \mathcal{J}(L) \\
& \sum_{I, J \in \mathcal{I}, \mathcal{J}} z_{IJ} \leq B^A && \forall I \in \mathcal{I} \\
& \frac{1}{2} \sum_{J \in \mathcal{J}(W)} Z_{IJ} + \frac{3}{2} \sum_{J \in \mathcal{J}(L)} Z_{IJ} \geq \mathcal{W}(B - A)_{\mathcal{J}} - \text{EG bound} \\
& z_{IJ} \leq c_{IJ} && \forall I, J \in \mathcal{I}, \mathcal{J} \\
& \hat{x}_I \in (0, 1), \quad z_{IJ} \geq 0 && \forall I, J \in \mathcal{I}, \mathcal{J} \quad (46)
\end{aligned}$$

After solving this linear program, we check if the optimal solution involves a tight  $\sum_{I \in \mathcal{I}} z_{IJ'} \leq V_{J'}^B - V_{J'}^A$  for a  $J' \in \mathcal{J}(L)$ . If not, this implies that the optimal investment leading to maximum seats in round-1 does not need to change the  $W/L$  status of any district in the votemandered map. Otherwise, a tight constraint implies a change in  $W/L$  status of a district, i.e.,  $\hat{y}_J$  towards optimality which changes the EG significantly and non-continuously. We thus make the required change in constraints and solve the updated linear program again. Corresponding to district  $J'$ , the updates include replacing  $\sum_{I \in \mathcal{I}} z_{IJ'} \leq V_{J'}^B - V_{J'}^A$  with

$$\sum_{I \in \mathcal{I}} z_{IJ'} \geq V_{J'}^B - V_{J'}^A$$

and replacing the fairness constraint with

$$\frac{1}{2} \sum_{J \in \mathcal{J}(W) \cup J'} Z_{IJ} - \frac{1}{2}(V_{J'}^B - V_{J'}^A) + \frac{3}{2} \sum_{J \in \mathcal{J}(W)/J'} Z_{IJ} \geq \mathcal{W}(B - A)_{\mathcal{J}} - \text{EG bound} + \frac{3V_{J'}^B + V_{J'}^A}{2}$$

or

$$\frac{1}{2} \sum_{J \in \mathcal{J}(W) \cup J'} Z_{IJ} + \frac{3}{2} \sum_{J \in \mathcal{J}(W)/J'} Z_{IJ} \geq \mathcal{W}(B - A)_{\mathcal{J}} - \text{EG bound} + 2V_{J'}^B$$



budget in a currently losing district  $D_1$  (change  $D_1$  status from  $L \rightarrow W$  in the first round elections i.e. in the campaigned submap) and propose perturbations with its neighbor  $D_2$  so that it again loses in the second round (mark  $D_1$  status as  $L$  in the second round i.e. in the votemandered/target submap). We thus increase the votemandering bonus by 1, via a campaigned map win at  $D_1$ . The budget required here is simply the margin with which  $D_1$  loses in the campaigned map ( $margin(D_1)$ ) with no investment. Naturally, the perturbations for the new submap involve exchanging units across districts  $D_1, D_2$  that nullify the effect of budget investment and regain the  $L$  status of  $D_1$  in the votemandered submap. Let this imply a movement of  $V_A^T$  party A votes and  $V_B^T$  party B votes from  $D_2$  to  $D_1$ . If  $(D_1, D_2)$  have status  $(L, W)$  in the initial submap, Lemma 4.1 implies:

$$\Delta\mathcal{W}(B - A)_i \text{ for } (L, W) = -V_A^T - V_B^T - \frac{a_1}{2} - \frac{3a_2}{2} + \frac{b_2}{2} + \frac{3b_1}{2} \quad (47)$$

where,  $a_i, b_i$  is the budget allocation in district  $i$  (in the initial submap) by party A and B respectively. Similarly, if the current status is  $(L, L)$ , we have:

$$\Delta\mathcal{W}(B - A)_i \text{ for } (L, L) = -V_A^T - V_B^T - \frac{a_1}{2} - \frac{a_2}{2} + \frac{3b_2}{2} + \frac{3b_1}{2} \quad (48)$$

Strategy-1 edge weight is then given as  $(margin(D_1), \Delta\mathcal{W}(B - A)_i)$ . Note the fairness cost from the initial map to the votemandered map is insignificant since the votemandered submap retains the  $W/L$  status of both districts.

#### Strategy-2:

In this strategy, we let party B win  $D_1$  in campaigned and the votemandered submap and propose perturbations between  $D_1, D_2$  such that  $D_1$  wins in the second round (mark from  $L \rightarrow W$  in the target submap). The target submap is chosen such that the votemandered submap leads to no change in the wins, but the target submap (using original vote shares) secures a votemandering bonus of 1. This is possible  $margin(D_1)$  is smaller than B's budget allocation in  $D_1$ . This also suggests that with an increase in B's budget allocation, A can specifically use strategy-2 to votemandeer more efficiently. The change in  $\Delta\mathcal{W}$  for this setting is exactly the same as for strategy-1, giving strategy-2 edge weight as  $(0, \Delta\mathcal{W}(B - A)_i)$ . Moreover, we don't expect  $\Delta\mathcal{W}$  to change a lot here as well, since there is no change in the number of wins where the fairness constraint is concerned, i.e. in the votemandered submap.

#### Strategy-3:

In this strategy, we let party B win a district in the first round but propose perturbations such that A wins in both the votemandered as well as the target submaps (mark from  $L \rightarrow W$  in the second round elections). The difference with strategy-2 is that here we don't depend on B's budget allocation to secure a win, but achieve only via the changes in the district boundaries i.e. the local perturbations. The target submap is chosen such that both the votemandered submap and the target submap have  $D_1$  as a winning district. Naturally, here we expect  $\Delta\mathcal{W}$  to change significantly that of the initial map, as A's number of wins increases in the votemandered submap.

If  $(D_1, D_2)$  have status  $(L, W)$  and  $V_A^1, V_B^1$  are the original vote shares of district  $D_1$ , Lemma 4.1 implies:


$$\Delta\mathcal{W}(B - A)_i \text{ for } (L, W) = -V_A^1 - V_B^1 - \frac{a_1}{2} - \frac{3a_2}{2} + \frac{b_2}{2} + \frac{3b_1}{2} \quad (49)$$

where,  $a_i, b_i$  are the budget allocations in district  $i$  by party A and B respectively. Similarly, if the status is  $(L, L)$ , we have:


$$\Delta\mathcal{W}(B - A)_i \text{ for } (L, L) = -V_A^1 - V_B^1 - \frac{a_1}{2} - \frac{a_2}{2} + \frac{3b_2}{2} + \frac{3b_1}{2} \quad (50)$$

**E SECTION 5 DETAILS**


We include the visualization of local votemandering strategies for both Republican and Democratic parties in Figures 19- 22.




**Fig. 19.** Republican Local Votemandering: The initial map, the created target map, and the strategic investment (here, zero investment followed by all strategy-2 improvements)



**Fig. 20.** The Four Stages of Republican Local Votemandering, (with red and blue indicating the districts won by the Republican and Democratic parties, respectively)



**Fig. 21.** Democratic Local Votemandering: The initial map, the created target map, and the strategic investment



**Fig. 22.** The Four Stages of Democratic Local Votemandering