

Not in My Backyard! Temporal Voting Over Public Chores

Edith Elkind¹, Tzeh Yuan Neoh^{2,3}, Nicholas Teh⁴

¹Northwestern University, USA

²Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore

³Centre for Frontier AI Research, Agency for Science, Technology and Research, Singapore

⁴University of Oxford, UK

edith.elkind@northwestern.edu, neohty@cfar.a-star.edu.sg, nicholas.teh@cs.ox.ac.uk

Abstract

We study a temporal voting model where voters have dynamic preferences over a set of public chores—projects that benefit society, but impose individual costs on those affected by their implementation. We investigate the computational complexity of optimizing utilitarian and egalitarian welfare. Our results show that while optimizing the former is computationally straightforward, minimizing the latter is computationally intractable, even in very restricted cases. Nevertheless, we identify several settings where this problem can be solved efficiently, either exactly or by an approximation algorithm. We also examine the effects of enforcing temporal fairness and its impact on social welfare, and analyze the competitive ratio of online algorithms. We then explore the strategic behavior of agents, providing insights into potential malfeasance in such decision-making environments. Finally, we discuss a range of fairness measures and their suitability for our setting.

1 Introduction

The local government is launching a multi-year initiative to enhance community programs and boost tourism, aiming to create a vibrant and sustainable future for the district. At the beginning of every year, the government will unveil a calendar of planned initiatives for the year ahead, featuring events for each month. These initiatives include popular recurring activities such as food and music festivals, farmers’ markets, fireworks displays, and sports tournaments. As part of this effort, residents are invited to participate in a voting process to voice their preferences on which initiatives should take place in the following year. However, residents living near proposed event locations often raise concerns about potential negative impacts on their quality of life. These concerns include crowding, noise pollution, traffic congestion, or disruptions to the community’s character. Consequently, some members of the community may oppose specific events, even those deemed beneficial for the district as a whole.

This phenomenon is often referred to as “*Not in My Backyard*” (*NIMBY*), and usually comes in the form of an organized effort by local residents or community groups to op-

pose certain developments or activities in their neighborhood, often due to perceived adverse effects on their local environment or daily lives. *NIMBY* groups often present a significant challenge to public policy design: while they may be driven by legitimate concerns, they can delay or obstruct important community goals. This tension between local concerns and broader societal needs makes such movements a complex and often polarizing issue to address.

Thus, the question that we focus on in this work is: how can we effectively handle such preferences to derive an outcome that is *good* for everyone and/or treats everyone *fairly*?

We approach this problem by viewing it through the lens of *temporal voting* [Chandak *et al.*, 2024; Elkind *et al.*, 2024c; Lackner, 2020]. However, instead of voters having approval preferences (indicating which candidates they would like to support) over a set of candidates, voters now express disapproval preferences (indicating which candidates they object to) over a set of candidates.

Disapproval preferences are also relevant in settings where there are too many candidates to consider, and agents only have strong opinions about candidates they do not want chosen (which could be due to proximity concerns as mentioned earlier), but are indifferent among the rest (because this choice does not affect them). However, results for existing models on temporal voting with “positively-valued” candidates (sometimes also known as *public goods* or issues [Conitzer *et al.*, 2017; Fain *et al.*, 2018; Skowron and Górecki, 2022]) do not automatically transfer to this setting of “negatively-valued” candidates (which we call *public chores*¹); we discuss this briefly at the end of the next subsection. This is also the case in the standard *fair allocation of indivisible items* model, whereby the fairness concepts or corresponding results for goods and chores could be vastly different. In these models, items (goods or chores) are considered to be *privately* allocated to agents, and are not shared among agents. Considering a model with agents that have disutilities over candidates is also relatively new to temporal voting (apart from *costs* associated with implementing projects in participatory budgeting, as opposed to disutilities here).

¹We feel that the term *public chores* reflect the property of these candidates that while they lead to positive outcomes for the community at large, they often require sacrifices that may not be immediately appreciated by some.

1.1 Our Contributions

We consider a novel variant of the temporal voting model, where agents express disapproval over candidates.

In Section 3, we investigate the computational complexity of two welfare objectives: minimizing the sum of agents’ disutilities (MIN-SUM) and minimizing the maximum disutility (MIN-MAX). We show that, while finding a MIN-SUM outcome is easy, the decision problem associated with MIN-MAX is NP-complete even in simple settings. Nevertheless, we provide several positive parameterized complexity results for the MIN-MAX objective. We also provide an approximation ratio for the corresponding optimization problem.

In Section 4, we analyze the price of enforcing temporal constraints with respect to both objectives; for both cases, we derive tight bounds.

In Section 5, we shift our focus to the study of strategic manipulation with respect to both welfare objectives. We show that while MIN-SUM is compatible with *strategyproofness* (i.e., no agent can strictly benefit from misreporting their preferences), this is no longer true for MIN-MAX. Additionally, we show that while MIN-SUM is incompatible with a stronger *group-strategyproofness* property, finding a ‘suitable manipulation’ is computationally hard. We present similar results for MIN-MAX, but with respect to strategyproofness.

In Section 6, we consider the online setting where information on future projects are unknown, and analyze the competitive ratio of online algorithms.

Finally, in Section 7, we discuss other popular fairness notions such as proportionality and equitability, and show that they are either not well-defined in the public chores setting, or comes at a very high cost to welfare.

We wish to emphasize that many of our results involve multiplicative upper and lower bounds, addressing challenges such as computational intractability (Section 3), the need for temporal fairness and its impact on welfare (Section 4), and the lack of information relating to future projects (Section 6). Importantly, the nature and analysis of these multiplicative bounds differs substantially between the ‘public goods’ setting (analogous to temporal voting) and the ‘public chores’ setting (our focus), and they are not directly transferable.

To illustrate, consider an example with 100 timesteps where the optimal outcome (with respect to the MIN-MAX objective) satisfies all agents in 95 timesteps. In the public goods’ setting, an outcome that satisfies all agents in 50 timesteps would be considered a $\frac{1}{2}$ -approximation. However, in the public chores’ setting, this same outcome only provides a 10-approximation. This distinction in bounds has significant implications for how agents perceive and engage in collective decision-making. Revisiting our motivating example of NIMBY movements, where agents derive limited personal benefit from approved projects, but experience considerable disutility from disapproved ones, it is more accurate to interpret the outcome above as agents being 10 times more unhappy, rather than $\frac{1}{2}$ as happy.

1.2 Related Work

Sequential Decision-Making Works in this area mostly consider positively-valued candidates. Elkind *et al.* [2024b] is the most relevant one, which looks into the computational

questions associated with welfare maximization, and its compatibility with strategyproofness and proportionality (a popular fairness measure) in the setting with positively-valued candidates. They also look extensively at the welfare cost of mandating proportionality. Our work also looks at analogous welfare objectives, and we also consider their compatibility with strategyproofness. However, on top of the distinctions raised earlier, several other key differences are (i) we consider a more general version of the decision problem for MIN-MAX, which imposes constraints involving the sequential nature of timesteps; (ii) for manipulability, we consider computational problems in manipulation and also a stronger version of strategyproofness; (iii) we analyze the online setting; (iv) we discuss why proportionality is not well-defined in our setting and consider equitability as an additional fairness measure.

Another relevant line of work in this area is that of *perpetual voting* [Lackner, 2020; Lackner and Maly, 2023], which focuses on temporal extensions of traditional multiwinner voting rules. Bulteau *et al.* [2021], Chandak *et al.* [2024], Elkind *et al.* [2025b], and Phillips *et al.* [2025] build on the temporal voting framework and consider temporal extensions of popular proportional representation (or group fairness) axioms. Kahana and Hazon [2023] study several popular allocation rules (round-robin, maximum Nash welfare, and leximin) with respect to approximately proportional fairness axioms. Kozachinskiy *et al.* [2025] study a similar model, but look into the conditions under which there is a sublinear growth of dissatisfaction. However, we note that their definition of ‘dissatisfaction’ is inherently different from our ‘disutility’ and studying disapprovals more generally: they consider a model where agents still express approval preferences, and the dissatisfaction of an agent is number of rounds where the project chosen was not approved by the agent.

Bredereck *et al.* [2020; 2022] and Zech *et al.* [2024] look at sequential committee elections, whereby an entire committee (set of candidates) is elected in each round, and impose constraints on the extent a committee can change, whilst ensuring that the candidates retain sufficient support from the electorate.

We refer the reader to the survey by Elkind *et al.* [2024c] for an analysis of other works in this area.

Public Decision-Making Conitzer *et al.* [2017] study several relaxations of proportionality and its axiomatic guarantees towards individual voters. Fain *et al.* [2018] considered the notion of an approximate core, whereas Skowron and Górecki [2022] proposed another variant of proportionality that guarantees fairness to groups of voters. Lackner *et al.* [2023] studied strategic considerations in the same setting. Alouf-Heffetz *et al.* [2022] consider a model of issue-by-issue voting, which can be viewed as a special case of the temporal voting framework. Again, all works in this setting look at positively-valued issues/candidates and the focus generally revolves around proportionality, which is ill-defined in the setting with negatively-valued candidates (as we demonstrate towards the end of our paper)

Temporal Fair Division Another related model is that of temporal fair division [Cookson *et al.*, 2025; Elkind *et al.*,

2025a; Igarashi *et al.*, 2024; Neoh *et al.*, 2025]. The key difference is that in this model, a single item (good or chore) is *privately* allocated to an agent at each round (i.e., rivalrous in ownership). This is in contrast to our model, which can be thought of as allocating *public* items (i.e., every agent gets a copy of the item chosen at each round).

Fair Allocation of Indivisible Chores Chores have also been studied in the field of fair division, but works in the area typically consider private chores that are rivalrous in ownership [Aziz *et al.*, 2022a; Bogomolnaia *et al.*, 2017; Dehghani *et al.*, 2018; Elkind *et al.*, 2024a] (see also recent survey by Aziz *et al.* [2022b]). The study of private and public goods is known to be vastly different, and this distinction can also be observed in the context of chores. In fact, even in the case of binary chores (i.e., agents having disutility in $\{0, 1\}$ for each chore), the difference is stark. Under binary valuations (which is a very popular model in several fair allocation settings [Halpern *et al.*, 2020; Suksompong and Teh, 2022]), the problem of allocating private chores is known to be easy: popular (approximate) notions of fairness such as EF1 can be trivially obtained. However, when considering public chores, the problem becomes non-trivial even under disapproval preferences.

2 Preliminaries

For each positive integer z , let $[z] := \{1, \dots, z\}$. Let $N = [n]$ be the set of n agents, let $P = \{p_1, \dots, p_m\}$ be the set of m projects (or candidates), and let $T = [\ell]$ be the set of ℓ timesteps. Symmetrically to the literature on multiwinner/temporal voting with approval preferences [Lackner and Skowron, 2023; Elkind *et al.*, 2024c], we will assume that voters have *disapproval preferences*. For each $i \in N$ and $k \in T$, let $D_{ik} \subseteq P$ denote the *disapproval set* of agent i at timestep k , and let the *disapproval vector* of an agent i be $\mathbf{D}_i = (D_{i1}, \dots, D_{i\ell})$. An instance of our problem is a tuple $(N, P, T, (\mathbf{D}_i)_{i \in N})$.

An *outcome* is a vector $\mathbf{o} = (o_1, \dots, o_\ell)$, where $o_k \in P$ for each $k \in T$. Let $\Pi(\mathcal{I})$ denote the space of all possible outcomes for an instance \mathcal{I} . For every $k \in T$ and $\mathbf{o} \in \Pi(\mathcal{I})$, the k -truncation of \mathbf{o} is the vector $\mathbf{o}^{(k)} = (o_1, \dots, o_k)$.

The *disutility* of an agent $i \in N$ from an outcome \mathbf{o} is given by $d_i(\mathbf{o}) = |\{k \in T : o_k \in D_{ik}\}|$. We extend this definition to truncated outcomes by writing $d_i(\mathbf{o}^{(k)}) = |\{t \in [k] : o_t \in D_{it}\}|$. A *mechanism* maps an instance $\mathcal{I} = (N, P, T, (\mathbf{D}_i)_{i \in N})$ to an outcome in $\Pi(\mathcal{I})$.

We assume that the reader is familiar with basic notions of classic complexity theory [Papadimitriou, 2007] and parameterized complexity [Flum and Grohe, 2006; Niedermeier, 2006]. All omitted proofs can be found in the appendix.

3 Social Welfare Optimization

Two commonly studied welfare objectives in collective decision-making are maximizing the sum of agents' utilities (i.e., *utilitarian* welfare) or maximizing the utility of the least happy agent (i.e., *egalitarian* welfare). In our context, these objectives translate to, respectively, minimizing the sum of agents' disutilities $\sum_{i \in N} d_i(\mathbf{o})$ (we will refer to

this as the MIN-SUM objective) or minimizing the maximum disutility $\max_{i \in N} d_i(\mathbf{o})$ (we will refer to this as the MIN-MAX objective)². Accordingly, given an instance \mathcal{I} , we refer to outcomes $\mathbf{o} \in \Pi(\mathcal{I})$ that minimize $\sum_{i \in N} d_i(\mathbf{o})$ (resp., $\max_{i \in N} d_i(\mathbf{o})$) as MIN-SUM (resp., MIN-MAX) outcomes.

We will now discuss the complexity of finding MIN-SUM and MIN-MAX outcomes, starting with the former.

3.1 Minimizing the Sum of Agents' Disutilities

To find a MIN-SUM outcome, we can greedily select, at each timestep, a project with the lowest number of disapprovals at that timestep. It is easy to observe that this greedy algorithm runs in polynomial time. However, just like in the case of positively-valued projects, the outcomes of the greedy algorithm may be unfair. Indeed, consider an instance with $2\kappa + 1$ agents $N = \{1, \dots, 2\kappa + 1\}$, two projects $P = \{p_1, p_2\}$, and ℓ timesteps. Let $\kappa + 1$ agents disapprove of p_2 at each timestep, and the remaining κ agents disapprove of p_1 at each timestep. Then, the MIN-SUM algorithm will select p_1 at every timestep, favoring the $\kappa + 1$ agents and disadvantaging the other κ agents. Arguably, this is not a fair outcome. This motivates us to explore another welfare objective—MIN-MAX—that specifically focuses on fairness.

3.2 Minimizing the Maximum Agents' Disutility

When considering the egalitarian welfare, we aim to take into account the temporal nature of our problem, by imposing constraints on agents' disutilities not just at timestep ℓ , but also at earlier checkpoints.

Formally, a *set of constraints* for an instance $\mathcal{I} = (N, P, T, (\mathbf{D}_i)_{i \in N})$ is a set $\mathbf{A} = \{(t_1, \lambda_1), \dots, (t_\tau, \lambda_\tau)\}$, where $t_j \in T$, $\lambda_j \in \{0, \dots, \ell\}$ for each $j \in [\tau]$, and $t_1 \leq \dots \leq t_\tau$, $\lambda_1 \leq \dots \leq \lambda_\tau$.

A pair $(\mathcal{I}, \mathbf{A})$ defines a decision problem as follows.

MIN-MAX-DEC

Input: A problem instance $(N, P, T, (\mathbf{D}_i)_{i \in N})$ and a set of τ constraints \mathbf{A} .

Question: Is there an outcome \mathbf{o} such that for each $(t, \lambda) \in \mathbf{A}$ it holds that $\max_{i \in N} d_i(\mathbf{o}^{(t)}) \leq \lambda$?

In words, each pair $(t_j, \lambda_j) \in \mathbf{A}$ mandates that at timestep t_j , the cumulative disutility of every agent should be at most λ_j . Having a single constraint (ℓ, λ) can be seen as a decision version of MIN-MAX. We note that for utilitarian welfare constraints of this form have no impact on the choice of outcome: if there is a solution that satisfies all constraints, then so does the greedy solution described in Section 3.1.

We will now investigate the complexity of MIN-MAX-DEC, both in the worst case, and from a parameterized perspective. In addition to the natural parameters of our problem, i.e., n , m and ℓ , we will consider a parameter $\gamma = \max_{k \in T, p \in P} |\{i : p \in D_{ik}\}|$, i.e., the maximum number of disapprovals that each project has at any timestep.

²Another popular welfare objective is Nash welfare, which maximizes the geometric mean of agents' utilities. However, it is well-known that Nash welfare is ill-defined for negative utilities.

As a warm-up, we first show that the special case $\gamma = 1$ admits a polynomial-time algorithm. The constraint $\gamma = 1$ is satisfied in settings where each project is “very local” and only affects a single agent (but an agent can still disapprove of multiple projects).

Theorem 3.1. *If $\gamma = 1$, there is a polynomial-time algorithm for MIN-MAX-DEC.*

Proof sketch. Consider a timestep $k \in T$. If there exists a project that receives no disapprovals at timestep k , we should select some such project. Thus, without loss of generality, since $\gamma = 1$, we can assume that at each timestep each project is disapproved by exactly one agent.

Construct a bipartite graph $G = (N \times [\lambda_\tau], T, E)$ with parts $N \times [\lambda_\tau]$ and T ; for all $(i, \lambda) \in N \times [\lambda_\tau]$, $k \in T$ the graph G contains the edge $\{(i, \lambda), k\}$ if and only if agent i disapproves a project at timestep k , and for all $\lambda' < \lambda$ and $k' \geq k$, we have $(k', \lambda') \notin A$. We claim that the maximum matching in G has cardinality $|T|$ if and only if there is an outcome \mathbf{o} satisfying MIN-MAX-DEC, and, given a size- $|T|$ matching in G , we can transform it into an outcome that satisfies all constraints in A in polynomial time; the proof is deferred to the appendix. \square

Notably, increasing γ from 1 to 2 leads to a hardness result for MIN-MAX-DEC, even if there are at most two projects per timestep and $\tau = |A| = 1$. Importantly, since our hardness proof works for $\tau = 1$, it follows that even finding MIN-MAX outcomes is computationally hard.

Theorem 3.2. *MIN-MAX-DEC is NP-complete, even with $m = 2$, $\gamma = 2$, and $\tau = 1$.*

On the positive side, if both the number of agents and the number of constraints are small, MIN-MAX-DEC becomes tractable. This provides a practical solution for small-group voting (e.g., when the voting process can be broken down into smaller districts, so that only the people that can be directly affected by the projects get to vote).

Theorem 3.3. *MIN-MAX-DEC is FPT with respect to $n + \tau$.*

If only the number of agents is small, we obtain a weaker tractability result.

Theorem 3.4. *MIN-MAX-DEC is XP with respect to n .*

We also prove a similar tractability result (XP) with respect to the number of timesteps ℓ . This result is applicable in settings where the planning timeline is limited (e.g., short-term policy cycles), even if the number of agents or projects can be substantial. Moreover, we show that our XP result cannot be strengthened to FPT unless $\text{FPT} = \text{W}[2]$.

Theorem 3.5. *MIN-MAX-DEC is XP and W[2]-hard with respect to ℓ .*

Having explored our problem from the parameterized complexity perspective, we would like to understand whether it admits an approximation algorithm. However, MIN-MAX-DEC is inherently a feasibility problem, so it does not have a natural optimization version. Nevertheless, we can provide some insights for the case $\tau = 1$, which corresponds to the standard notion of egalitarian welfare.

We first present an inapproximability result, derived as a corollary of the proof of Theorem 3.2: our reduction from

3-OCCUR-3SAT shows that we cannot hope to obtain an approximation ratio better than $3/2$ in polynomial time.

Corollary 3.6. *For any $\varepsilon > 0$, if there exists a polynomial-time algorithm for MIN-MAX-DEC with $\tau = 1$ and approximation ratio $\frac{3}{2} - \varepsilon$, then $\text{P} = \text{NP}$.*

Now, let ℓ^+ be the number of timesteps where every project is disapproved by some agent. Then, we give an algorithm that achieves a $\min(m, 1 + \frac{n^2}{\ell^+})$ approximation ratio, thereby providing an upper bound.

Notably, there are many settings (including the one described in our introductory example) where the number of projects is small (e.g., the projects have been shortlisted by the local government and put up for voting). Moreover, if ℓ^+ is large relative to n , the quantity $1 + \frac{n^2}{\ell^+}$ is close to 1.

Theorem 3.7. *There exists a $\min(m, 1 + \frac{n^2}{\ell^+})$ -approximation algorithm for the MIN-MAX objective.*

Proof. First, for the timesteps where some project receives no disapprovals, we can always choose some such project. Hence, it suffices to consider the ℓ^+ timesteps in which each project is disapproved by some agent. From now on, we will assume that $\ell = \ell^+$.

We construct a polynomial-size integer program for finding MIN-MAX outcomes as follows. For each $p \in P$ and $k \in T$, we define a variable $c_{(p,k)} \in \{0, 1\}$: $c_{(p,k)} = 1$ if and only if project p is selected at timestep k . Our constraints require that (1) for each $k \in T$, at least one project has to be chosen in timestep k : $\sum_{p \in P} c_{(p,k)} \geq 1$, and (2) the disutility of each agent $i \in N$ is at most η : $\sum_{k \in T} \sum_{p \in D_{ik}} c_{(p,k)} \leq \eta$. By relaxing the 0–1 variables $c_{(p,k)}$ to take values in \mathbb{R}_+ , we obtain the following LP relaxation:

$$\begin{aligned} & \text{minimize} && \eta && \text{(P1)} \\ & \text{subject to} && \sum_{p \in P} c_{(p,k)} \geq 1, && \text{for all } k \in T, \\ & && \sum_{k \in T} \sum_{p \in D_{ik}} c_{(p,k)} \leq \eta, && \text{for all } i \in N, \\ & && c_{(p,k)} \geq 0, && \text{for all } p \in P \text{ and } k \in T. \end{aligned}$$

Let $((c_{(p,k)}^*)_{p \in P, k \in T}, \eta)$ be an optimal solution to P1 that lies at a vertex of the respective polytope. Construct an outcome \mathbf{o} by selecting, at each timestep $k \in T$, a project $p^* \in \arg \max_{p \in P} c_{(p,k)}^*$. We first show that $\max_{i \in N} d_i(\mathbf{o}) \leq m \cdot \eta$ and then argue that $\max_{i \in N} d_i(\mathbf{o}) \leq \left(1 + \frac{n^2}{\ell}\right) \cdot \eta$. As the value of the optimal integer solution is at least η , this proves the theorem.

For each $i \in N$, let $T_i = \{k : o_k \in D_{ik}\}$ be the set of timesteps where i disapproves the outcome selected by \mathbf{o} . Consider a timestep $k \in T$. Since $\sum_{p \in P} c_{(p,k)}^* \geq 1$, we have $c_{(o_k, k)}^* \geq \frac{1}{m}$. Thus, for all $i \in N$ we have $d_i(\mathbf{o}) = |T_i| \leq \sum_{k \in T_i} m \cdot c_{(o_k, k)}^* \leq \sum_{k \in T} m \cdot c_{(o_k, k)}^* \leq m \cdot \eta$. Hence, $\max_{i \in N} d_i(\mathbf{o}) \leq m \cdot \eta$.

Next, we will show that $d_i(\mathbf{o}) \leq \left(1 + \frac{n^2}{\ell}\right) \cdot \eta$ for all $i \in N$, and thus $\max_{i \in N} d_i(\mathbf{o}) \leq \left(1 + \frac{n^2}{\ell}\right) \cdot \eta$. We first note that our linear program P1 has $m\ell + 1$ variables and $n + \ell + m\ell$ constraints. As we consider a vertex solution, there are at least $m\ell + 1$ constraints that are tight. Thus, at most $n + \ell$ variables of the form $c_{(p,k)}$ are non-zero. Let T' be the set of timesteps where at least two projects are assigned a positive weight by our solution to the LP. For each $k \in T$ we have $\sum_{p \in P} c_{(p,k)}^* \geq 1$, so each $k \in T$ contributes at least one non-zero variable. Thus, there are at most $(n + \ell) - |T| = n$ additional non-zero variables, i.e., $|T'| \leq n$. Moreover, $c_{(o_k,k)}^* = 1$ for all $k \in T \setminus T'$, so for each $i \in N$ we have

$$\begin{aligned} |\{k \in T \setminus T' : o_k \in D_{ik}\}| &= \sum_{k \in T \setminus T' : o_k \in D_{ik}} c_{(o_k,k)}^* \\ &\leq \sum_{k \in T \setminus T'} \sum_{p \in D_{ik}} c_{(p,k)}^* \leq \sum_{k \in T} \sum_{p \in D_{ik}} c_{(p,k)}^* \leq \eta, \end{aligned}$$

where the last transition follows since $((c_{(p,k)}^*)_{p \in P, k \in T}, \eta)$ is feasible for P1. Therefore, for each $i \in N$ we have $d_i(\mathbf{o}) = |\{k \in T : o_k \in D_{ik}\}| = |\{k \in T' : o_k \in D_{ik}\}| + |\{k \in T \setminus T' : o_k \in D_{ik}\}| \leq |T'| + \eta \leq n + \eta$.

Furthermore, as each project is disapproved by at least one agent at each timestep, we have

$$\begin{aligned} n \cdot \eta &\geq \sum_{i \in N} \sum_{k \in T} \sum_{p \in D_{ik}} c_{(p,k)}^* = \sum_{k \in T} \sum_{i \in N} \sum_{p \in D_{ik}} c_{(p,k)}^* \\ &\geq \sum_{k \in T} \sum_{p \in P} c_{(p,k)}^* \geq \ell, \end{aligned}$$

and hence $n^2/\ell \geq n/\eta$. Thus, for all agents $i \in N$ we have $d_i(\mathbf{o}) \leq \frac{n^2}{\ell} \cdot \eta = \left(1 + \frac{n^2}{\ell}\right) \cdot \eta$.

It is easy to see that the described algorithm runs in polynomial time with respect to n, m , and ℓ . \square

4 Price of Temporal Fairness

Next, we consider the impact of imposing temporal constraints \mathbf{A} on the agents' welfare. Our analysis belongs to the line of work on the *price of fairness*, initiated by Bei *et al.* [2021]. Similar questions have been considered by a number of authors in the multiwinner voting [Brill and Peters, 2024; Elkind *et al.*, 2022a; Lackner and Skowron, 2020] and temporal voting literature [Elkind *et al.*, 2024b]. Unsurprisingly, our setting, too, exhibits a fundamental tension between fairness and efficiency.

Our analysis applies to both MIN-SUM and MIN-MAX, but we only consider constraints of the form $\max_{i \in N} d_i(\mathbf{o}^{(t)}) \leq \lambda$. We note that constraints of the form $\sum_{i \in N} d_i(\mathbf{o}^{(t)}) \leq \lambda$ have no impact on the utilitarian welfare, but may reduce egalitarian welfare. While it may be interesting to investigate the impact of utilitarian constraints on egalitarian welfare, this question does not quite fit the price of fairness framework, so we leave it to future work.

Given a problem instance \mathcal{I} and a set of constraints $\mathbf{A} = \{(t_1, \lambda_1), \dots, (t_\tau, \lambda_\tau)\}$, let $\Pi_{\mathbf{A}}(\mathcal{I}) \subseteq \Pi(\mathcal{I})$ denote the set of

outcomes for \mathcal{I} that satisfy constraints in \mathbf{A} . We say that \mathbf{A} is *feasible* if $\Pi_{\mathbf{A}}(\mathcal{I}) \neq \emptyset$, i.e., if there is an outcome that satisfies all constraints in \mathbf{A} .

For the objectives MIN-SUM and MIN-MAX, let \sum and \max be the corresponding *welfare operation* (with respect to agents' disutilities). Then, for a welfare objective $W \in \{\text{MIN-SUM}, \text{MIN-MAX}\}$, we denote by $W_{\text{op}}(\mathbf{o})$ the W -value of an outcome \mathbf{o} , i.e., the result of applying the welfare operation corresponding to W to \mathbf{o} .

We now define the *price of temporal fairness*, which measures the cost of imposing (feasible) temporal constraints \mathbf{A} .

Definition 4.1 (Price of Temporal Fairness). For an objective $W \in \{\text{MIN-SUM}, \text{MIN-MAX}\}$, the *price of temporal fairness with respect to W* (PoTF_W) is the supremum over all instances \mathcal{I} and feasible constraints \mathbf{A} of the ratio between the minimum W -value of an outcome for \mathcal{I} that satisfies \mathbf{A} and the minimum W -value of an outcome for \mathcal{I} :

$$\text{PoTF}_W = \sup_{\mathcal{I}, \mathbf{A} : \mathbf{A} \text{ feasible}} \frac{\min_{\mathbf{o} \in \Pi_{\mathbf{A}}(\mathcal{I})} W_{\text{op}}(\mathbf{o})}{\min_{\mathbf{o} \in \Pi(\mathcal{I})} W_{\text{op}}(\mathbf{o})}.$$

We derive tight bounds for both welfare objectives. For MIN-SUM, the price of temporal fairness scales with the number of agents.

Theorem 4.2. $\text{PoTF}_{\text{MIN-SUM}} = \Theta(n)$.

Proof. We first prove the upper bound of $\mathcal{O}(n)$. Let $T^+ = \{k \in T : \cup_{i \in N} D_{ik} = P\}$; for each timestep k in T' each project is disapproved by at least one agent, so no matter how we select o_k , this will contribute to the disutility of some agent. Hence, for each outcome \mathbf{o} we have $\sum_{i \in N} d_i(\mathbf{o}) \geq |T^+|$. On the other hand, since \mathbf{A} is feasible, there is an outcome \mathbf{o} that satisfies all constraints in \mathbf{A} . We modify this outcome as follows: for each $k \in T \setminus T^+$ the set $P \setminus \cup_{i \in N} D_{ik}$ is not empty, so let o_k be some project in this set. For the modified outcome, the disutility of every agent is at most $|T^+|$, so the total disutility is at most $n \cdot |T^+|$. This completes the proof of the upper bound.

For the lower bound, consider an instance with two projects $P = \{p_1, p_2\}$ and two timesteps. Let $D_{11} = D_{12} = \{p_1\}$ and $D_{i1} = D_{i2} = \{p_2\}$ for all other agents $i \in N \setminus \{1\}$. Let \mathbf{A} contain a single constraint $(2, 1)$, which requires that at the end of timestep 2 the disutility of every agent is at most 1. The only outcomes that satisfy this constraint are (p_1, p_2) and (p_2, p_1) ; under either outcome the total disutility is n . On the other hand, the total disutility of (p_1, p_1) is 2. Thus, $\text{PoTF}_{\text{MIN-SUM}} \geq \frac{n}{2} = \Omega(n)$. \square

For MIN-MAX, the price of fairness is small if there are few constraints or if the number of agents is small.

Theorem 4.3. $\text{PoTF}_{\text{MIN-MAX}} = \Theta(\min(\tau, n))$.

Theorem 4.3 shows that one needs to proceed with caution when imposing fairness constraints, especially if the number of agents is large.

5 Strategic Manipulation

In our motivating example, ensuring that agents cannot engage in strategic manipulation is vital for maintaining trust

and participation in the voting process. To address these concerns, we will now focus on agents' strategic considerations with respect to both welfare objectives.

One popular concept in the social choice literature is that of *strategyproofness*, which states that no agent should be able to strictly benefit (in this case, decrease their disutility) by misreporting their preferences. It is formally defined as follows. Note that agent i 's disutility function d_i is computed with respect to her (truthful) disapproval vector \mathbf{D}_i .

Definition 5.1 (Strategyproofness). For each $i \in N$, let \mathcal{D}_{-i} denote the list of all disapproval vectors except that of agent i : $\mathcal{D}_{-i} = (\mathbf{D}_1, \dots, \mathbf{D}_{i-1}, \mathbf{D}_{i+1}, \dots, \mathbf{D}_n)$. A mechanism \mathcal{M} is *strategyproof* (SP) if for each instance $(N, P, T, (\mathbf{D}_i)_{i \in N})$, each agent $i \in N$ and each disapproval vector \mathbf{D}'_i it holds that $d_i(\mathcal{M}(\mathcal{D}_{-i}, \mathbf{D}_i)) \leq d_i(\mathcal{M}(\mathcal{D}_{-i}, \mathbf{D}'_i))$.

5.1 Mechanisms for MIN-SUM

We first show that the greedy algorithm for obtaining a MIN-SUM outcome, which chooses a project with the lowest number of disapprovals at each timestep (with lexicographical tie-breaking) is strategyproof. We will refer to this algorithm as GREEDY MIN-SUM.

Theorem 5.2. GREEDY MIN-SUM is strategyproof.

A natural follow-up question is whether the MIN-SUM objective is compatible with a stronger version of strategyproofness. A well-known generalization of strategyproofness is *group strategyproofness* (GSP). Intuitively, GSP states that no group of agents should be able to misreport their preferences so as to benefit every member of the group. Unfortunately, we show that the MIN-SUM objective is incompatible with GSP.

We start by presenting the formal definition of group strategyproofness. Recall that agent i 's disutility function d_i is computed with respect to her (true) disapproval vector \mathbf{D}_i .

Definition 5.3 (Group-strategyproofness). For each $S \subseteq N$, let \mathcal{D}_{-S} denote the list of all disapproval vectors except those of agents in S . A mechanism \mathcal{M} is *group-strategyproof* (GSP) if for each instance $(N, P, T, (\mathbf{D}_i)_{i \in N})$, each subset of agents $S \subseteq N$, each agent $i \in S$ and each list of disapproval vectors $(\mathbf{D}'_j)_{j \in S}$, it holds that $d_i(\mathcal{M}(\mathcal{D}_{-S}, (\mathbf{D}_j)_{j \in S})) \leq d_i(\mathcal{M}(\mathcal{D}_{-S}, (\mathbf{D}'_j)_{j \in S}))$.

Note that GSP reduces to SP if we only consider singleton groups. Then, our negative result is as follows.

Proposition 5.4. Let \mathcal{M} be a mechanism that always returns a MIN-SUM outcome. Then \mathcal{M} is not group-strategyproof.

Notably, while we only defined GSP (and SP) for deterministic mechanisms, the above negative result also applies to an analogous definition of GSP for *randomized* mechanisms. This is because the instance constructed in the counterexample has a unique MIN-SUM outcome, and any randomized mechanism behaves exactly like a deterministic one.

The above results indicate that the MIN-SUM objective is compatible with disincentivizing strategic manipulation by individuals (Theorem 5.2), but not by groups (Proposition 5.4). We further show that, while groups may have opportunities for manipulation, identifying a 'suitable manipulation' that strictly benefits every agent within the group can

be computationally intractable. Following the conventions of the literature on voting manipulation [Conitzer and Walsh, 2016], we assume the group has knowledge of the reported disapproval vectors of agents outside the group.

Theorem 5.5. Let \mathcal{M} be a mechanism that always returns a MIN-SUM outcome. Given an instance $(N, P, T, (\mathbf{D}_i)_{i \in N})$ and a subset of agents $S \subseteq N$, determining whether there exists disapproval vectors $(\mathbf{D}'_i)_{i \in S}$ such that $d_i(\mathcal{M}(\mathcal{D}_{-S}, (\mathbf{D}_i)_{i \in S})) > d_i(\mathcal{M}(\mathcal{D}_{-S}, (\mathbf{D}'_i)_{i \in S}))$ for all $i \in S$ is NP-complete.

5.2 Mechanisms for MIN-MAX

Next, we turn to the same questions for the MIN-MAX objective. However, in contrast, we show that MIN-MAX is fundamentally incompatible with strategyproofness. Intuitively, agents are incentivized to appear 'worse off' by misreporting disapproval for projects they do not actually disapprove of.

Proposition 5.6. Let \mathcal{M} be a mechanism that always returns a MIN-MAX outcome. Then \mathcal{M} is not strategyproof.

Again, while we only defined SP for deterministic mechanisms, the above negative result also applies to an analogous definition of SP for randomized mechanisms.

Despite the stronger negative result above, we show that for any mechanism returning a MIN-MAX outcome with ties broken lexicographically, it may be computationally intractable for any agent to find a manipulation that decreases their disutility. Our proof places this problem at the second level of the polynomial hierarchy; this is because computing a MIN-MAX outcome is already a hard problem (see Theorem 3.2).

Theorem 5.7. Let \mathcal{M}_{lex} be a mechanism that returns a MIN-MAX outcome with lexicographical tiebreaking. Given some instance $(N, P, T, (\mathbf{D}_i)_{i \in N})$ and an agent $i \in N$, determining whether there exists a disapproval vector \mathbf{D}'_i such that $d_i(\mathcal{M}_{\text{lex}}(\mathcal{D}_{-i}, \mathbf{D}_i)) > d_i(\mathcal{M}_{\text{lex}}(\mathcal{D}_{-i}, \mathbf{D}'_i))$ is Σ_2^P -complete.

6 Online Setting

In practice, multi-year plans often face discontinuities due to political shifts or other unforeseen changes, leading to potential failure in execution mid-way. Consequently, ensuring fairness for voters only at the conclusion of X years may be inadequate; voters might instead expect satisfaction at every consecutive timestep. Further, project plans (and hence availability/feasibility) themselves may evolve over time, making it difficult, if not impossible, to predict the project availability in advance. In both of these scenarios, it is natural to consider an *online* setting where agents' preferences for future projects are unknown, and decisions are made without access to future information. We utilize *competitive analysis* for online algorithms to measure the impact of this lack of future information on our ability to achieve optimal welfare.

Let \mathcal{B} be an *online algorithm* for our setting, i.e., an algorithm that for each $k \in T$ selects o_k based on $(D_{it})_{i \in N, t \in [k]}$. Let $\mathcal{B}(\mathcal{I})$ denote the output of \mathcal{B} on instance \mathcal{I} .

Definition 6.1 (Competitive Ratio). For an online algorithm \mathcal{B} and objective $W \in \{\text{MIN-SUM}, \text{MIN-MAX}\}$, the *competitive ratio* (CR) of \mathcal{B} with respect to W is the supremum over

all instances \mathcal{I} of the ratio between the W -value of $\mathcal{B}(\mathcal{I})$ and the minimum W -value of an outcome for \mathcal{I} :

$$\text{CR}_W(\mathcal{B}) = \sup_{\mathcal{I}} \frac{W_{\text{op}}(\mathcal{B}(\mathcal{I}))}{\min_{\mathbf{o} \in \Pi(\mathcal{I})} W_{\text{op}}(\mathbf{o})}.$$

Note that GREEDY MIN-SUM is an online algorithm, and it outputs a MIN-SUM allocation. Hence, $\text{CR}_{\text{MIN-SUM}}(\text{GREEDY MIN-SUM}) = 1$.

However, for the MIN-MAX objective, the lack of information about the future can significantly impact the MIN-MAX value. For the remainder of this section, we focus on the competitive ratio with respect to MIN-MAX.

Similarly to the greedy algorithm for MIN-SUM, we consider the greedy egalitarian algorithm, which, at each timestep k , picks a project so as to minimize the maximum disutility over the first k timesteps. We refer to this algorithm as GREEDY-MIN-MAX. Unfortunately, it turns out that for both GREEDY-MIN-MAX and GREEDY-MIN-MAX the competitive ratio is lower-bounded by $\Omega(n)$.

Proposition 6.2. *GREEDY MIN-SUM and GREEDY MIN-MAX both have a competitive ratio of $\Omega(n)$ with respect to MIN-MAX.*

Next, we present a lower bound for *all* online algorithms. We consider a weak, non-adaptive adversary that does not have access to the randomized results of the algorithm and show that even against such an adversary, any online algorithm has a competitive ratio of at least $\Omega(\log n)$.

Proposition 6.3. *Against a non-adaptive adversary, any online algorithm (deterministic or randomized) has a competitive ratio of $\Omega(\log n)$ with respect to MIN-MAX.*

7 Other Fairness Notions

Throughout this paper, we focused on the MIN-MAX objective—a fairness criterion that has been extensively studied across many topics in social choice. A natural extension of this work would be to explore other well-established fairness concepts, such as envy-freeness, proportionality or equitability. While envy-based measures are commonly studied in the fair division literature, these do not adapt well to the public goods or chores setting. In the remainder of this section, we discuss proportionality and equitability in more detail.

7.1 Proportionality

The first concept we consider is a widely studied notion of fairness in both the fair division of private goods/chores and public goods settings, called *proportionality* [Barman and Krishnamurthy, 2019; Brânzei and Sandomirskiy, 2023; Conitzer *et al.*, 2017]. More specifically, for private chores, proportionality mandates that no agent should receive more than $1/n$ of their pessimal (worst-case) disutility; and for private/public goods, it requires that all agents receive at least $1/n$ of their optimal (best-case) utility. While proportionality may not always be achievable in these settings, mild relaxations of the concept are known to always exist. However, an analogous definition for public chores proves to be unintuitive, as illustrated by the following example.

Example 7.1. Consider an instance with n agents, ℓ timesteps, and set of projects $P = \{p_1, \dots, p_n\}$. At every timestep, each agent $i \in N$ disapproves all projects in $P \setminus \{p_i\}$.

Then, every outcome \mathbf{o} is disapproved by $n - 1$ agents in each round, so by the pigeonhole principle we have $d_i(\mathbf{o}) \geq \ell \cdot \frac{n-1}{n}$ for some $i \in N$. This is despite the fact that i 's disutility from (p_i, \dots, p_i) is 0. Thus, no outcome is close to being proportional for the standard definition of proportionality.

Identifying a suitable notion of proportionality for the public chores setting remains an intriguing open problem.

7.2 Equitability

Another potentially suitable fairness notion is *equitability*, which mandates that agents' disutilities should be equal [Elkind *et al.*, 2022b; Freeman *et al.*, 2019]. While it may not always be achievable in practice (this is also the case for many similar fairness properties in the social choice literature), we may be interested in obtaining an equitable outcome when one exists. Unfortunately, we can show that even determining if an instance admits an equitable outcome is computationally intractable.

Theorem 7.2. *Determining if there exists an equitable outcome is NP-complete.*

Moreover, enforcing equitability comes at a very high cost to (both MIN-SUM and MIN-MAX) welfare. Our definition of *price of equitability* is structurally similar to the definition of the price of temporal fairness, and leads to the following result.

Theorem 7.3. *The price of equitability is $\Omega(n^2)$ with respect to MIN-SUM and at least $\Omega(n)$ with respect to MIN-MAX.*

Exploring more suitable fairness concepts in this setting remains an intriguing direction for future work.

8 Conclusion

In this work, we introduced and studied a model of temporal voting where agents can express disapprovals over candidates. We investigated the computational complexity of two well-studied welfare objectives—MIN-SUM and MIN-MAX—and identified several settings where the problem can be solved efficiently, together with accompanying algorithms. We also quantified the effects of enforcing temporal fairness on social welfare, and analyzed the strategic implications associated with these welfare objectives. Further, we derived bounds on the price of temporal fairness and the competitive ratio of algorithms in the online setting. Finally, we made a case for why proportionality and equitability may not be suitable as fairness measures for the public chores setting.

Directions for future work include defining and studying weaker forms of strategyproofness that may be compatible with MIN-MAX in this setting, or identifying an appropriate (potentially weaker) notion of proportionality in this setting. It would also be interesting to consider a model with both positively- and negatively-valued candidates.