ECAI 2024
U. Endriss et al. (Eds.)

doi:10.3233/FAIA240877

Temporal Elections: Welfare, Strategyproofness, and Proportionality

Edith Elkind[1,2], Tzeh Yuan Neoh[3] and Nicholas Teh[1]

[1]University of Oxford, UK
[2]Alan Turing Insitute, UK
[3]Agency for Science, Technology and Research, Singapore

Abstract. We investigate a model of sequential decision-making where a single alternative is chosen at each round. We focus on two objectives—utilitarian welfare (UTIL) and egalitarian welfare (EGAL)—and consider the computational complexity of the associated maximization problems, as well as their compatibility with strategyproofness and proportionality. We observe that maximizing UTIL is easy, but the corresponding decision problem for EGAL is NP-complete even in restricted cases. We complement this hardness result for EGAL with parameterized complexity analysis and an approximation algorithm. Additionally, we show that, while a mechanism that outputs a UTIL outcome is strategyproof, all deterministic mechanisms for computing EGAL outcomes fail a very weak variant of strategyproofness, called non-obvious manipulability (NOM). However, we show that when agents have non-empty approval sets at each timestep, choosing an EGAL-maximizing outcome while breaking ties lexicographically satisfies NOM. Regarding proportionality, we prove that a proportional (PROP) outcome can be computed efficiently, but finding an outcome that maximizes UTIL while guaranteeing PROP is NP-hard. We also derive upper and lower bounds on the price of proportionality with respect to UTIL and EGAL.

1 Introduction

Consider a group of friends planning their itinerary for a two-week post-graduation trip across Europe. They have selected their activities, but still need to decide on their choice of meals for each day. As popular restaurants typically require reservations, everyone is asked to declare their preferences upfront before the trip commences.

Suppose that 55% of them prefer Asian cuisine, 25% prefer European cuisine, 10% prefer Oceanic cuisine, and the remaining 10% prefer South American cuisine. Given that there is a maximum of three meals in any given day, it may be impossible to satisfy everyone on any single day. However, over multiple days, it may be feasible to eventually satisfy everyone. Still, adopting a day-to-day majority voting may not lead to a desirable outcome, as the Asian cuisine would be chosen for every meal, and, as a result, 45% of the group will be perpetually unhappy. A natural question is then: what would be an appropriate notion of *satisfaction*, and can we (efficiently) obtain an outcome that offers high satisfaction?

As the group moves from city to city, the set of available restaurants changes. Even within the same town, one may have different preferences for lunch and dinner, opting for a heavier meal option at lunch and a lighter one at dinner. As both preferences and the set of alternatives may evolve with time, traditional multiwinner voting models [16, 23, 30] do not fully capture this setting.

This problem fits within the *temporal elections* framework, a model where a sequence of decisions is made, and outcomes are evaluated with respect to agents' temporal preferences. It was first introduced as *perpetual voting* by Lackner [27]; see the survey of Elkind et al. [20] for a discussion of subsequent work. We consider the *offline* variant of this model where preferences are provided upfront. That is, at each timestep, each agent has a set of approved alternatives, and the goal is to select a single alternative per timestep.

While this model has been considered in prior work [11, 12, 21], earlier papers focus on fairness and proportionality notions, with only a few (if at all) looking into welfare objectives and strategic considerations. Against this background, in this work we focus on the algorithmics of maximizing two classic welfare objectives: the utilitarian welfare (the sum of agents' utilities) and the egalitarian welfare (the minimum utility of any agent), both in isolation and in combination with strategyproofness and proportionality axioms.

Our Contributions In this paper, we investigate the utilitarian (UTIL) and egalitarian (EGAL) welfare-maximization objectives from three perspectives: the computational complexity of welfare maximization, compatibility with strategyproofness, and trade-offs with proportionality.

In Section 3, we show that UTIL is solvable in polynomial-time, and the decision problem associated with EGAL is NP-complete even in very restricted instances. Given this, we analyze the parameterized complexity with respect to several natural parameters, and provide an ILP-based approximation.

In Section 4, we show that while UTIL is strategyproof, any deterministic mechanism for EGAL may fail non-obviously manipulability (NOM), a relaxation of strategyproofness. In the case where each agent has a non-empty approval set at each timestep, we show that EGAL admits no deterministic strategyproof mechanism in general, but is NOM when using leximin tie-breaking.

Finally, in Section 5, we show that while a simple greedy algorithm can return a proportional (PROP) outcome, it is NP-hard to determine if there exists a PROP outcome that maximizes utilitarian welfare when each agent has a non-empty approval set at each timestep. We also provide upper and lower bounds for the (strong) price of proportionality with respect to both UTIL and EGAL. To the best of our knowledge, our work is the first to investigate the price of proportionality in temporal elections.

Related Work Perhaps the most popular line of work in this area is that of *perpetual voting* by Lackner [27] and Lackner and Maly [28]. These works focus on temporal extensions of traditional multi-winner voting rules, and analyze them with respect to novel axioms developed for the temporal setting. Bulteau et al. [11] built upon this framework and proposed temporal extensions of popular *justified representation* axioms. Chandak et al. [12] extended this work by studying efficient algorithms that achieve these axioms. Elkind et al. [21] studied the computational problems associated with verifying whether outcomes satisfy these axioms.

An adjacent line of work by Bredereck et al. [7, 8] and Zech et al. [36] look into sequential committee elections where an entire committee is elected at each timestep. They consider constraints on the extent to which a committee can change, while ensuring a certain level of support remains for the committee. Deltl et al. [15] looked into a similar model, but with the restriction that agents can only approve at most one project per timestep.

A line of work also worth mentioning is that of apportionment with approval preferences [10, 14]. In this setting, the goal is to allocate the seats of a fixed-size committee to parties based on voters' (approval) preferences over the parties. This is equivalent to a restricted setting of temporal voting where voters have static preferences that do not change over time.

Yet another related model is that of fair scheduling, whereby both agents' preferences and the outcome are permutations of projects. Elkind et al. [18] studied the computational issues of associated with maximizing various welfare objectives, along with several other fairness properties. Our work differs from this setting in that we allow projects to be chosen more than once (both in agents' preferences and the outcome).

One could also view temporal elections as allocating public goods or decision-making on public issues, making the fair public decision-making model [13, 22, 34] particularly relevant.

Other models of issue-by-issue decision-making can also be considered as a restricted setting of temporal voting. In particular, Alouf-Heffetz et al. [2] considered such a model with uncertainty in voters' preferences, and the goal is to recover the majority-supported outcome for each issue.

Finally, the *price of fairness* concept for proportionality, which captures the welfare loss from mandating proportionality, has also been studied for more demanding proportionality guarantees in the single-round multiwinner voting literature [9, 17, 29].

2 Preliminaries

For each positive integer k, let $[k] := \{1, \dots, k\}$. Let $N = [n]$ be a set of n *agents*, let $P = \{p_1, \dots, p_m\}$ be a set of m *projects* (or *candidates*), and let $T = [\ell]$ be a set of ℓ *timesteps*. For each $k \in [\ell]$, the set of projects approved by agent $i \in N$ at timestep k is captured by her *approval set* $S_{ik} \subseteq P$. The approval sets of agent i are collected in her *approval vector* $\mathbf{S}_i = (S_{i1}, \dots, S_{i\ell})$. Thus, an *instance* of our problem is a tuple $(N, P, \ell, (\mathbf{S}_i)_{i \in N})$.

An *outcome* is a vector $\mathbf{o} = (o_1, \dots, o_\ell)$, where $o_k \in P$ for each $k \in [\ell]$. The utility of an agent $i \in N$ for an outcome $\mathbf{o}$ is given by $u_i(\mathbf{o}) = |\{k \in [\ell] : o_k \in S_{ik}\}|$. Let $\Pi(\mathcal{I})$ denote the space of all possible outcomes for an instance $\mathcal{I}$. A *mechanism* maps an instance $\mathcal{I} = (N, P, \ell, (\mathbf{S}_i)_{i \in N})$ to an outcome in $\Pi(\mathcal{I})$.

We do not require each agent to approve at least one project at each timestep; however, we do require that each agent approves at least one project at some timestep, i.e., for each $i \in N$ there exists a $k \in [\ell]$ with $S_{ik} \neq \varnothing$; indeed, if this condition is failed for some $i \in N$, we can simply delete i, as it can never be satisfied. If $S_{ik} \neq \varnothing$ for all $i \in N$ and $k \in [\ell]$, we say that we are in the *complete preference (CP) setting*. We believe that the CP setting captures many real-life applications of our model: for instance, in our motivating example, having no particular opinion on any cuisine would be more reasonably interpreted as approving all options rather than having an empty approval set.

We assume that the reader is familiar with basic notions of classic complexity theory [33] and parameterized complexity [24, 32].

All missing proofs can be found in the full version of our paper.

3 Welfare Maximization

We first focus on welfare maximization, without combining it with other desiderata. The two objectives we consider are defined as follows.

Definition 1 (Social Welfare)**.** *Given an outcome* $\mathbf{o}$*, its* utilitarian social welfare *is defined as* $\sum_{i \in N} u_i(\mathbf{o})$ *and its* egalitarian social welfare *is defined as* $\min_{i \in N} u_i(\mathbf{o})$*. We refer to outcomes that maximize the utilitarian/egalitarian welfare as* UTIL/EGAL *outcomes, respectively.*

A UTIL outcome can be found in polynomial time: at each timestep, one can simply select the project that has the highest number of approvals. However, observe that if 51% of the population approves project p at each timestep, while 49% of the population approves project q at each timestep (and there are no other approvals), if we select the UTIL outcome, close to half of the population will not get a single project approved at any timestep. This underscores the need to consider other criteria for selecting outcomes, such as, e.g., egalitarian welfare. However, while computing a UTIL outcome is computationally feasible, this is not the case for EGAL outcomes.

The decision problem associated with computing EGAL outcomes, which we denote by EGAL-DEC, is defined as follows.

> MAXIMIZING EGALITARIAN WELFARE (EGAL-DEC):
> **Input**: A problem instance $\mathcal{I} = (N, P, \ell, (\mathbf{S}_i)_{i \in N})$ and a parameter $\lambda \in \mathbb{Z}^+$.
> **Question**: Is there an outcome $\mathbf{o}$ such that $\min_{i \in N} u_i(\mathbf{o}) \geq \lambda$?

The following result shows that, perhaps surprisingly, EGAL-DEC is NP-complete even when the goal is to guarantee each agent a utility of 1, and when there are only two projects.[1]

Theorem 3.1. EGAL-DEC *is* NP*-complete, even if* $m = 2$*,* $\lambda = 1$*.*

The above negative result effectively rules out the possibility of maximizing the egalitarian welfare even in simple settings.

Nevertheless, in what follows, we show that when the number of agents or timesteps is constant, we are able to efficiently find a solution to this problem. More precisely, we show that EGAL-DEC is fixed-parameter tractable (FPT) with respect to the number of agents (n). Our approach is based on integer linear programming; we show how to encode EGAL-DEC as an integer linear program (ILP) whose number of variables depends on n only; our claim then follows from Lenstra's classic result [31]. To accomplish this, we classify the projects and timesteps into 'types', so that the number of types is exponential in n, but does not depend on m or ℓ respectively.

[1] This result is equivalent to Theorem 2 of Deltl et al. [15]. Nevertheless, for completeness, we include a proof in the full version of this paper.

Theorem 3.2. EGAL-DEC *is* FPT *with respect to* n.

Proof. As a preprocessing step, we create ℓ copies of each project. That is, we replace a project p with projects $p^1, \dots, p^\ell$ and modify the approval vectors: for each $i \in N$, $k \in [\ell]$, $p \in S_{ik}$ we place p^k in S_{ik} and remove p. This does not change the nature of our problem, since in our setting, there is no dependence between timesteps, and a project's label can be re-used arbitrarily between timesteps. For the modified instance, it holds that for each project p there is at most one timestep $k \in [\ell]$ such that $p \in \cup_{i \in N} S_{ik}$. Then, we define the *type* of a project as the set of voters who approve it: the type of p is $\tau(p) = \{i \in N : p \in S_{ik}$ for some $k \in [\ell]\}$. Because of the preprocessing step, for each $p \in P$ there is a timestep $k \in [\ell]$ such that $p \in S_{ik}$ for all $i \in \tau(p)$, and $p \notin S_{ik'}$ for all $i \in N$, $k' \in [\ell] \setminus \{k\}$. Note that there are at most 2^n different project types.

In the same way, we can classify timesteps by the types of projects present in them, giving us at most 2^{2^n} different *timestep types*. Let $Q \subseteq 2^N$ be the set of all project types and let $\mathcal{R} \subseteq 2^Q$ be the set of all timestep types.

Now, we construct our ILP. For each $R \in \mathcal{R}$, let z_R be the number of timesteps of type R. For each $i \in N$, let Q_i be the set of project types that agent i approves of. For each $R \in \mathcal{R}$ and $\tau \in R$, we introduce an integer variable $x_{R,\tau}$ representing the number of timesteps of type R in which a project of type τ was chosen.

The ILP is defined as follows, with objective function

$$\text{maximize } \lambda$$

subject to the following constraints:
(1) $\sum_{\tau \in R} x_{R,\tau} \le z_R$ for each $R \in \mathcal{R}$;
(2) $\sum_{R \in \mathcal{R}} \sum_{\tau \in Q_i} x_{R,\tau} \ge \lambda$ for each $i \in N$;
(3) $x_{R,\tau} \ge 0$ for each $R \in \mathcal{R}$ and $\tau \in R$.

The first constraint is equivalent to requiring that we select at most one project per timestep, whereas the second constraint ensures that each agent has utility of λ from the outcome.

There are at most $\mathcal{O}(2^{n+2^n})$ variables in the ILP, so the classic result of Lenstra Jr [31] implies that our problem is FPT in n. □

Next, we show that when the number of timesteps (ℓ) is constant, EGAL can be solved in polynomial time, i.e., EGAL-DEC is slicewise polynomial (XP) with respect to the number of timesteps.

Theorem 3.3. EGAL-DEC *is* XP *with respect to* ℓ.

Proof. Observe that there are m^ℓ possible outcomes. Thus, when ℓ is constant, the number of outcomes is bounded by a polynomial. We can thus iterate through all outcomes; we output 'yes' if we find an outcome that provides utility γ to all agents, and 'no' otherwise. By combining this approach with binary search over γ, we can also find an EGAL outcome. □

We complement the above result by showing that EGAL-DEC is W[2]-hard with respect to the number of timesteps. This indicates that an FPT (in ℓ) algorithm does not exist unless FPT = W[2], and hence the XP result of Theorem 3.3 is tight.

Theorem 3.4. EGAL-DEC *is* W[2]*-hard with respect to* ℓ.

Proof. We reduce from the DOMINATING SET (DS) problem. An instance of DS consists of a graph $G = (V, E)$ and an integer κ; it is a yes-instance if there exists a subset $D \subseteq V$ such that $|D| \le \kappa$ and every vertex of G is either in D or has a neighbor in D, and a no-instance otherwise. DS is known to be W[2]-complete with respect to the parameter κ [32].

Given an instance (G, κ) of DOMINATING SET with $G = (V, E)$, $V = \{v_1, \dots, v_n\}$, set $N = [n]$, $P = \{p_1, \dots, p_n\}$, $\ell = \kappa$. Then for each $i \in N$ and $k \in [\ell]$ let $S_{ik} = \{p_j : i = j \text{ or } \{v_i, v_j\} \in E\}$. We claim that G admits a dominating set D with $|D| \le \kappa$ if and only if there exists an outcome $\mathbf{o}$ such that $u_i(\mathbf{o}) \ge 1$ for all agents $i \in N$.

For the 'if' direction, consider an outcome $\mathbf{o} = (p_{j_1}, \dots, p_{j_\kappa})$ that provides positive utility to all agents, and set $D = \{v_{j_1}, \dots, v_{j_\kappa}\}$. Then D is a dominating set of size at most κ. Indeed, consider a vertex $v_i \in V$. Since voter i approves p_{j_k} for some $k \in [\ell]$, we have $v_{j_k} \in D$ and $i = j_k$ or $\{v_i, v_{j_k}\} \in E$. Note that if there are projects chosen more than once, we simply have $|D| < \kappa$.

For the 'only if' direction, observe that a dominating set $D = \{v_{j_1}, \dots, v_{j_s}\}$ with $s \le \kappa$ can be mapped to an outcome $\mathbf{o} = (p_{j_1}, \dots, p_{j_s}, p_1, \dots, p_1)$ (with p_1 selected in the last $\kappa - s$ timesteps). As any vertex of G is either in D, or has a neighbor in D, we have $u_i(\mathbf{o}) \ge 1$ for each agent $i \in N$. □

The construction in the proof of Theorem 3.4 can be used to derive the following corollary

Corollary 3.5. EGAL-DEC *is* NP*-complete, even for* $\lambda = 1$ *and in the* CP *setting.*

As a special case, we further show that when agents have non-empty approval sets at all timesteps, and there are two projects, then EGAL-DEC is XP with respect to λ.

Theorem 3.6. EGAL-DEC *is* XP *with respect to* λ *in the* CP *setting with* $m = 2$.

Theorem 3.1 shows that EGAL-DEC is NP-complete even when $\lambda = 1$. This implies that EGAL-DEC is inapproximable: an approximation algorithm would be able to detect whether a given instance admits an outcome with positive egalitarian social welfare. However, suppose we redefine each agent's utility function as $u'_i(\mathbf{o}) = 1 + u_i(\mathbf{o})$; this captures, e.g., settings where there is a timestep in which all agents approve the same project. We will now show that we can obtain an $\frac{1}{4 \log n}$-approximation to the optimal egalitarian welfare with respect to the utility profile $(u'_1, \dots, u'_n)$.

Theorem 3.7. *There is a polynomial-time algorithm that, for any* $\varepsilon > 0$*, given an instance* $(N, P, \ell, (\mathbf{S}_i)_{i \in N})$*, with probability* $1 - \varepsilon$ *outputs an outcome* $\mathbf{o}$ *whose egalitarian social welfare is at least* $\frac{1}{4 \log n}$ *of the optimal egalitarian social welfare with respect to the modified utility functions* $(u'_1, \dots, u'_n)$.

Proof. First, we construct a polynomial-size integer program for finding outcomes whose egalitarian welfare with respect to modified utilities is at least a given quantity η. For each $p \in P$ and $k \in [\ell]$, we define a variable $c_{(p,k)} \in \{0, 1\}$: $c_{(p,k)} = 1$ encodes that p is selected at time k. Our constraints require that (1) for each $k \in [\ell]$, at most one project can be chosen in timestep k: $\sum_{p \in P} c_{(p,k)} \le 1$, and (2) each agent $i \in N$ approves the outcome in at least $\eta - 1$ timesteps, so her modified utility is at least η: $\sum_{k=1}^{\ell} \sum_{p \in S_{ik}} c_{(p,k)} + 1 \ge \eta$. By relaxing the 0-1 variables $c_{(p,k)}$ to take values in $\mathbb{R}_+$, we obtain the following LP relaxation:

$$\begin{aligned} \text{LP}(\eta): & \sum_{p\in P} c_{(p,k)} \le 1 \quad \text{for all } k\in[\ell] \\ & \sum_{k=1}^{\ell}\sum_{p\in S_{ik}} c_{(p,k)} \ge \eta-1 \quad \text{for all } i\in N \\ & c_{(p,k)} \ge 0 \quad \text{for all } p\in P, k\in[\ell]. \end{aligned}$$

We can use binary search over η to find the largest value of η for which our LP is feasible; denote this value by η^*. Let η' be the optimal egalitarian welfare with respect to $u'_1,\dots,u'_n$ in our instance; then η' together with an encoding of the outcome that provides this welfare forms a feasible solution to our LP, and hence $\eta'\le\eta^*$.

When $\eta^*\le 4\log n$, for every outcome $\mathbf{o}$ we have $u'_i(\mathbf{o})\ge 1$ for all $i\in N$, and hence $\mathbf{o}$ is a $\frac{1}{4\log n}$-approximation. Thus, we can output an arbitrary outcome in this case. Hence, from now we assume that $\eta^*>4\log n$.

We can find an optimal (fractional) solution to LP(η^*) in polynomial time; let $\{c^*_{(p,k)}\}$ be some such solution. To transform it into a feasible integer solution, we use randomized rounding: for each $k\in[\ell]$ we select p to be implemented at timestep k with probability $c^*_{(p,k)}$. These choices are independent across timesteps. For each $i\in N$ and $k\in[\ell]$, define a Bernoulli random variable Z^i_k to indicate if agent i approves the project randomly selected at timestep t. Then, for each agent $i\in N$, we define a random variable $Z_i=\sum_{k=1}^{\ell} Z^i_k$. Note that the utility of an agent $i\in N$ is given by $u'_i = Z_i+1$. Then, the expected value of Z^i_k is

$$\mathbb{E}[Z^i_k]=\sum_{p\in S_{ik}} c^*_{(p,k)}.$$

By linearity of expectation, we get

$$\mathbb{E}[Z_i]=\sum_{k=1}^{\ell}\mathbb{E}[Z^i_k]=\sum_{k=1}^{\ell}\sum_{p\in S_{ik}} c^*_{(p,k)}\ge\eta^*-1.$$

Applying the multiplicative Chernoff bound [1], we obtain

$$\mathbb{P}\{u'_i\le\eta^*(1-\delta)\}\le\exp\left(\frac{-\eta^*\delta^2}{2}\right) \text{ for any } \delta>0.$$

Recall that $\eta^*>4\log n$. Thus, by letting $\delta=\frac{4}{5}$, we have

$$\mathbb{P}\left\{u'_i\le\frac{\eta^*}{5}\right\}\le\exp\left(-\frac{32\log n}{25}\right)=n^{-\frac{32}{25}}.$$

Finally, by applying the union bound, we get

$$\mathbb{P}\left\{u'_i\ge\frac{\eta^*}{5}\text{for all } i\in N\right\}\ge 1-n\cdot n^{-\frac{32}{25}} \quad =1-n^{-\frac{7}{25}}>0.$$

Consequently, there exists a $\frac{1}{5}$-OPT integer solution; using probability amplification techniques, we can obtain it with probability $1-\varepsilon$. It remains to observe that $\frac{1}{5}>\frac{1}{4\log n}$ when $n>3$. □

4 Strategyproofness and Non-Obvious Manipulability

An important consideration in the context of collective decision-making is *strategyproofness*: no agent should be able to increase their utility by misreporting their preferences. It is formally defined as follows. Note that agent i's utility function u_i is computed with respect to his (truthful) approval vector $\mathbf{S}_i$.

Definition 2 (Strategyproofness)**.** *For each $i\in N$, let $\mathcal{S}_{-i}$ denote the list of all approval vectors except that of agent i: $\mathcal{S}_{-i}=(\mathbf{S}_1,\dots,\mathbf{S}_{i-1},\mathbf{S}_{i+1},\dots,\mathbf{S}_n)$. A mechanism $\mathcal{M}$ is* strategyproof *(SP) if for each instance $(N,P,\ell,(\mathbf{S}_i)_{i\in N})$, each agent $i\in N$ and each approval vector $\mathbf{B}_i$ it holds that $u_i(\mathcal{M}(\mathbf{S}_i,\mathcal{S}_{-i}))\ge u_i(\mathcal{M}(\mathbf{B}_i,\mathcal{S}_{-i}))$.*

We first show that the algorithm that obtains a UTIL outcome by choosing a project with the highest number of approvals at each timestep (breaking ties lexicographically) satisfies this property. We will refer to this algorithm as GREEDYUTIL.

Theorem 4.1. GREEDYUTIL *is strategyproof.*

In contrast, no deterministic mechanism that maximizes egalitarian welfare can be strategyproof. Intuitively, this is because agents have an incentive to not report their approval for already-popular projects.

Proposition 4.2. *Let $\mathcal{M}$ be a deterministic mechanism that always outputs an* EGAL *outcome. Then $\mathcal{M}$ is not strategyproof, even in the* CP *setting.*

Proof. Consider an instance with $P=\{p_1,p_2,p_3\}$, $n=3$, $\ell=2$, and the approval sets $\mathbf{S}_1,\mathbf{S}_2,\mathbf{S}_3$ such that $S_{11}=S_{21}=S_{31}=\{p_1\}$ and $S_{i2}=\{p_i\}$ for each $i\in\{1,2,3\}$.

If p_1 is not selected at the first timestep, at most one agent receives positive utility, so the egalitarian welfare is 0. Thus, for every EGAL outcome $\mathbf{o}=(o_1,o_2)$ we have $o_1=p_1$ and $o_2\in\{p_1,p_2,p_3\}$. This ensures that the egalitarian welfare in 1.

Assume without loss of generality that $\mathcal{M}$ selects $\mathbf{o}=(p_1,p_2)$ when the agents report truthfully. Then $u_1(\mathbf{o})<2$. Then, agent 1 can misreport their approval vector as $\mathbf{S}'_1=(S_{11},S_{12})$, where $S_{11}=\{p_2,p_3\}$, $S_{12}=\{p_1\}$. In this case, the only EGAL outcome is $\mathbf{o}'=(p_1,p_1)$, so $\mathcal{M}$ is forced to output $\mathbf{o}'$. Moreover, agent 1's utility (with respect to his true preference) from $\mathbf{o}'$ is $u_1(\mathbf{o}')=2>u_1(\mathbf{o})$, i.e., agent 1 has an incentive to misreport. □

Having ruled out the compatibility of EGAL and strategyproofness, we consider a relaxation of strategyproofness known as *non-obvious manipulability*. It was introduced by Troyan and Morrill [35], and has been studied in the single-round multiwinner voting literature [3, 4]. It is formally defined as follows.

Definition 3 (Non-Obvious Manipulability)**.** *A mechanism $\mathcal{M}$ is* not obviously manipulable (NOM) *if for every instance $(N,P,\ell,(\mathbf{S}_i)_{i\in N})$, each agent $i\in N$, and each approval vector $\mathbf{B}_i$ that i may report, the following conditions hold:*

$$\begin{aligned} \min_{\mathcal{S}_{-i}\in\Sigma^{n-1}_{P,\ell}} u_i(\mathcal{M}(\mathbf{S}_i,\mathcal{S}_{-i})) &\ge \min_{\mathcal{S}_{-i}\in\Sigma^{n-1}_{P,\ell}} u_i(\mathcal{M}(\mathbf{B}_i,\mathcal{S}_{-i})) \\ \max_{\mathcal{S}_{-i}\in\Sigma^{n-1}_{P,\ell}} u_i(\mathcal{M}(\mathbf{S}_i,\mathcal{S}_{-i})) &\ge \max_{\mathcal{S}_{-i}\in\Sigma^{n-1}_{P,\ell}} u_i(\mathcal{M}(\mathbf{B}_i,\mathcal{S}_{-i})), \end{aligned}$$

where $\Sigma^{n-1}_{P,\ell}$ denotes the space of all $(n-1)$-voter profiles where each voter expresses her approvals of projects in P over ℓ steps.

Intuitively, a mechanism is NOM if an agent cannot increase her worst-case utility or her best-case utility (with respect to her true utility function) by misreporting. Clearly, strategyproofness implies NOM: if a mechanism is strategyproof, no agent can increase her utility in *any* case by misreporting.

While NOM is a much weaker condition than strategyproofness, it turns out that it is still incompatible with EGAL.

Proposition 4.3. *Let $\mathcal{M}$ be a deterministic mechanism that always outputs an* EGAL *outcome. Then $\mathcal{M}$ is not NOM.*

Proof. We will prove that an agent can increase her worst-case utility, and hence the mechanism fails NOM.

Fix $P = \{p_1, p_2\}$, $n = \ell = 2$. Consider first an instance $\mathcal{I} = (N, P, \ell, (\mathbf{S}_i)_{i\in N})$ with $S_{11} = S_{21} = \varnothing$, $S_{12} = \{p_1\}$, and $S_{22} = \{p_2\}$. Let $\mathbf{o} = (o_1, o_2)$ be the output of $\mathcal{M}$ on this instance; assume without loss of generality that $o_2 = p_1$.

Now, consider an instance $\mathcal{I}' = (N, P, \ell, (\mathbf{S}'_i)_{i\in N})$ with $S'_{11} = \{p_1, p_2\}$, $S'_{21} = \varnothing$, $S'_{12} = \{p_1\}$, and $S'_{22} = \{p_2\}$. For $\mathbf{o}' = (o'_1, o'_2)$ to be an EGAL outcome for this instance, it has to provide positive utility to both agents; this is only possible if $o'_2 = p_2$. Thus, $\mathcal{M}$ has to select p_2 at timestep 2 (and one of p_1, p_2 at timestep 1), so the utility of agent 1 from the outcome selected by $\mathcal{M}$ is 1.

However, we will now argue that if the first agent misreports her approval vector as $(\varnothing, \{p_1\})$, she is guaranteed utility 2 no matter what the second agent reports, i.e., her worst-case utility is 2.

Indeed, by our assumption on $\mathbf{o}$, if agent 2 reports $(\varnothing, \{p_2\})$, $\mathcal{M}$ selects p_1 at timestep 2 (and one of p_1, p_2 at timestep 1). Further, if agent 2 reports $(\varnothing, \{p_1\})$ or $(\varnothing, \{p_1, p_2\})$, then $\mathcal{M}$ selects p_1 at timestep 2 (and one of p_1, p_2 at timestep 1), as this is the only way to guarantee positive utility to both agents. Finally, if $S_{21} \neq \varnothing$, to guarantee positive utility to both agents, $\mathcal{M}$ would have to select a project from S_{21} at the first timestep, and p_1 at the second timestep. That is, if agent 1 reports $(\varnothing, \{p_1\})$, $\mathcal{M}$ selects an outcome (o^*_1, o^*_2) with $o^*_2 = p_1$, and this outcome provide utility 2 to agent 1 according to his true preferences (i.e., $S'_{11} = \{p_1, p_2\}$, $S'_{12} = \{p_1\}$). □

However, we obtain a positive result for the CP setting. Let $\mathcal{M}_{\text{lex}}$ be the mechanism that outputs an EGAL outcome, breaking ties in favor of agents with lower indices. Formally, we define an order $\succ$ on the set $\Pi(\mathcal{I})$ of possible outcomes for a given instance as follows: (1) if $\min_{i\in N} u_i(\mathbf{o}) > \min_{i\in N} u_i(\mathbf{o}')$, then $\mathbf{o} \succ \mathbf{o}'$; (2) if $\min_{i\in N} u_i(\mathbf{o}) = \min_{i\in N} u_i(\mathbf{o}')$ and there is an $i \in N$ such that $u_{i'}(\mathbf{o}) = u_{i'}(\mathbf{o}')$ for $i' < i$ and $u_i(\mathbf{o}) > u_i(\mathbf{o}')$ then $\mathbf{o} \succ \mathbf{o}'$. We then complete $\succ$ to a total order on $\Pi(\mathcal{I})$ arbitrarily. $\mathcal{M}_{\text{lex}}$ outputs an outcome $\mathbf{o}$ with $\mathbf{o} \succ \mathbf{o}'$ for all $\mathbf{o}' \in \Pi(\mathcal{I}) \setminus \{\mathbf{o}\}$.

Theorem 4.4. *$\mathcal{M}_{lex}$ is* NOM *in the* CP *setting.*

Proof. In the CP setting, the best-case utility of each agent is ℓ when they report truthfully: this is, e.g., achieved if all other agents have the same preferences. Thus, it remains to establish that under $\mathcal{M}_{\text{lex}}$ no agent can improve their worst-case utility by misreporting.

Let $\mathbf{S}_i$ be the true approval vector of agent i, and let $\mathbf{B}_i$ be another approval vector that i may report. Consider a minimum-length sequence of elementary operations that transforms $\mathbf{S}_i$ into $\mathbf{B}_i$ by first adding approvals in $B_{ik} \setminus S_{ik}$, $k \in [\ell]$, one by one, and then removing approvals in $S_{ik} \setminus B_{ik}$, $k \in [\ell]$, one by one. Let $\mathbf{X}_0, \mathbf{X}_1, \dots, \mathbf{X}_k, \dots, \mathbf{X}_{\gamma+1}$ be the resulting sequence of approval vectors, with $\mathbf{X}_0 = \mathbf{S}_i$, $\mathbf{X}_{\gamma+1} = \mathbf{B}_i$. Note that all approval vectors in this sequence consist of non-empty approval sets, i.e., we remain in the CP setting. Suppose this sequence starts with t additions, so that $\mathbf{X}_s$ is obtained from $\mathbf{X}_{s-1}$ by adding a single approval is $s \leq t$ and by deleting a single approval if $s > t$.

We will first show that reporting $\mathbf{X}_t$ instead of $\mathbf{X}_0 = \mathbf{S}_i$ does not increase i's worst-case utility. Then, we will show that for all $s = t+1, \dots, \gamma+1$, reporting $\mathbf{X}_s$ instead of $\mathbf{X}_{s-1}$ does not increase i's worst-case utility. This implies that reporting $\mathbf{B}_i$ instead of $\mathbf{S}_i$ does not increase her worst-case utility either.

Fix a list $\mathcal{S}_{-i}$ of other agents' approval vectors, and let $\mathcal{S} = (\mathcal{S}_{-i}, \mathbf{S}_i)$, $\mathcal{S}' = (\mathcal{S}_{-i}, \mathbf{X}_t)$. Suppose $\mathcal{M}_{\text{lex}}(\mathcal{S}') = \mathbf{o}$. Let η be the egalitarian welfare of $\mathbf{o}$ with respect to the reported utilities at $\mathcal{S}'$, and let η' be the utility of i at $\mathbf{o}$ according to $\mathbf{X}_t$; note that $\eta' \geq \eta$ and $\eta' \geq u_i(\mathbf{o})$. By choosing $\mathbf{o}$ at $\mathcal{S}$, the mechanism can guarantee utility η to all agents other than i, and $u_i(\mathbf{o}) \leq \eta'$ to i. If $u_i(\mathbf{o}) \leq \eta$, the egalitarian welfare of choosing $\mathbf{o}$ at $\mathcal{S}$ is $u_i(\mathbf{o})$, so under any EGAL outcome at $\mathcal{S}$ the utility of i is at least $u_i(\mathbf{o})$. In this case we are done, as $\mathcal{M}_{\text{lex}}$ always chooses an EGAL outcome.

Now, suppose $u_i(\mathbf{o}) > \eta$, and let $\mathbf{o}' = \mathcal{M}_{\text{lex}}(\mathcal{S})$. Note that the egalitarian welfare of $\mathbf{o}$ at $\mathcal{S}$ is η, so the egalitarian welfare of $\mathbf{o}'$ at $\mathcal{S}$ is at least η. Moreover, it cannot be strictly higher than η, because then the egalitarian welfare of $\mathbf{o}'$ at $\mathcal{S}'$ according to the reported utilities would be strictly higher than η as well, a contradiction with $\mathcal{M}_{\text{lex}}$ outputting $\mathbf{o}$ on $\mathcal{S}'$. Thus, $\mathbf{o}'$ and $\mathbf{o}$ provide the same egalitarian welfare at $\mathcal{S}$, and $\mathcal{M}_{\text{lex}}$ favors $\mathbf{o}'$ over $\mathbf{o}$ at $\mathcal{S}$ due to lexicographic tie-breaking. Let i' be the smallest index such that $u_{i'}(\mathbf{o}') > u_{i'}(\mathbf{o})$. If $i' \geq i$, we are done, as this means that $u_i(\mathbf{o}') = u_i(\mathbf{o})$, so i does not benefit from reporting $\mathbf{X}_s$ instead of $\mathbf{S}_i$. Otherwise, $\mathcal{M}_{\text{lex}}$ should favor $\mathbf{o}'$ over $\mathbf{o}$ at $\mathcal{S}'$. Indeed, i's utility from $\mathbf{o}'$ according to $\mathbf{X}_s$ is at least $u_i(\mathbf{o}') \geq \eta$, so $\mathbf{o}'$ and $\mathbf{o}$ provide the same egalitarian welfare. As $i' < i$ and agents $1, \dots, i'$ have the same preferences in $\mathcal{S}'$ and $\mathcal{S}$, $\mathcal{M}_{\text{lex}}$ should choose $\mathbf{o}'$ over $\mathbf{o}$, a contradiction with the choice of $\mathbf{o}$.

Now, for each $s > t$ we will argue that for every $\mathcal{S}_{-i}$ there is an $\mathcal{S}'_{-i}$ such that $u_i(\mathcal{M}_{\text{lex}}(\mathcal{S}_{-i}, \mathbf{X}_{s-1})) \geq u_i(\mathcal{M}_{\text{lex}}(\mathcal{S}'_{-i}, \mathbf{X}_s))$. Suppose that $\mathbf{X}_s$ is obtained from $\mathbf{X}_{s-1}$ by deleting a project p at timestep k. Let $\mathbf{o} = \mathcal{M}_{\text{lex}}(\mathcal{S}_{-i}, \mathbf{X}_{s-1})$. If $o_k \neq p$ then $\mathbf{o} = \mathcal{M}_{\text{lex}}(\mathcal{S}_{-i}, \mathbf{X}_s)$. Otherwise, consider a project $p' \neq p$ approved by agent i at timestep k according to $\mathbf{X}_s$. We construct $\mathcal{S}'_{-i}$ by swapping all other agents' approvals for p and p' at timestep k. Let $\mathbf{o}' = \mathcal{M}_{\text{lex}}(\mathcal{S}'_{-i}, \mathbf{X}_s)$. Then $\mathbf{o}$ and $\mathbf{o}'$ will only differ at timestep k, choosing project p' instead of p. As agent i approves p, we have $u_i(\mathbf{o}) \geq u_i(\mathbf{o}')$ Hence, for each $\mathcal{S}_{-i}$, there is an $\mathcal{S}'_{-i}$ such that $u_i(\mathcal{M}_{\text{lex}}(\mathcal{S}_{-i}, \mathbf{X}_{s-1})) \geq u_i(\mathcal{M}_{\text{lex}}(\mathcal{S}'_{-i}, \mathbf{X}_s))$, i.e., reporting $\mathbf{X}_s$ instead of $\mathbf{X}_{s-1}$ does not increase i's worst-case utility. As this holds for all $s > t$, the proof is complete. □

5 Proportionality

Another property that may be desirable in the context of temporal voting (and has been considered by others in similar settings [13, 18]) is *proportionality* (PROP).

Definition 4 (Proportionality)**.** *Given an instance $\mathcal{I} = (N, P, \ell, (\mathbf{S}_i)_{i\in N})$, for each $i \in N$ let $\mu_i = |\{k \in [\ell] : S_{ik} \neq \varnothing\}|$. We say that an outcome* $\mathbf{o}$ *is* proportional (PROP) *for $\mathcal{I}$ if for all $i \in N$ it holds that $u_i(\mathbf{o}) \geq \lfloor \frac{\mu_i}{n} \rfloor$.*

We note that proportionality is often understood as guaranteeing each agent at least $1/n$-th of her maximum utility, which would correspond to using $\frac{\mu_i}{n}$ instead of $\lfloor \frac{\mu_i}{n} \rfloor$ in the right-hand side of our definition [13, 18, 26]. However, the requirement that $u_i(\mathbf{o}) \geq \frac{\mu_i}{n}$ may be impossible to satisfy: e.g. if $N = \{1, 2\}$, $\ell = 3$, $P = \{p_1, p_2\}$ and for $i = 1, 2$ agent i approves project p_i at each timestep, we cannot simultaneously guarantee utility $3/2$ to both agents. Moreover, the proof of Theorem 3.4 shows that the associated decision problem is NP-complete. In contrast, our definition can be satisfied by a simple greedy algorithm. This follows from a similar result obtained by Conitzer et al. [13] in the setting of public decision-making; we provide a proof in the full version of this paper for completeness.

Theorem 5.1. *A* PROP *outcome always exists and can be computed by a polynomial-time greedy algorithm.*

We note that PROP can be seen as a specialization of the Strong PJR axiom for temporal voting [12, 21] to voter groups of size 1 (this offers additional justification for our definition), and hence the existence of PROP outcomes also follows from Theorem 4.1 in the work of Chandak et al. [12]; we omit a formal proof of this connection due to space constraints.

As finding *some* PROP outcome is not hard, one may wish to select the "best" PROP outcome. A natural criterion would be to pick a PROP outcome with the maximum utilitarian or egalitarian welfare (i.e., a PROP outcome that is also UTIL or EGAL).

However, the proof of Theorem 3.1 implies that selecting the PROP outcome with maximum egalitarian welfare is computationally intractable. Our next result shows that combining proportionality with utilitarian welfare is hard, too, even though both finding a PROP outcome and finding a UTIL outcome is easy. It also implies that finding a utilitarian welfare-maximizing outcome among all PROP outcomes is NP-hard.

Theorem 5.2. *Determining if there exists a* PROP *outcome that is* UTIL *is* NP*-complete, even in the* CP *setting.*

Proof. To see that this problem is is NP, recall that we can compute the maximum utilitarian welfare for a given instance; thus, we can check if a given outcome is UTIL and PROP.

From Corollary 3.5, in the CP setting, it is NP-complete even to determine if there is an outcome $\mathbf{o}$ such that $u_i(\mathbf{o}) \geq 1$ for all agents $i \in N$. We construct a new instance $\mathcal{I}'$ with $2n$ agents and timesteps such that there is an outcome that is proportional and maximizes utilitarian welfare if and only if there is an outcome in the original instance $\mathcal{I}$ that offers positive utility to all agents.

Note that if $\ell \geq n$ in the CP setting, then there always exists an outcome $\mathbf{o}$ such that $u_i(\mathbf{o}) \geq 1$ for all $i \in N$. Hence, assume that $\ell < n$. Let A_{pk} be the set of agents that approve project p at timestep k in $\mathcal{I}$. We construct an instance $\mathcal{I}'$ with $2n$ agents and timesteps and $m+n$ projects. Let $P'_{ik} = \{p \in P : i-n \leq n-|A_{pk}|\}$. We define the approval sets S'_{ik} for $\mathcal{I}'$ as follows:

$$S'_{ik} = \begin{cases} \{p_i\} & \text{if } i \leq n \text{ and } k \geq \ell \\ \{p_{n+1}\} & \text{if } i \geq n \text{ and } k \geq \ell \\ S_{ik} & \text{if } i \leq n \text{ and } k \leq \ell \\ \{p_{m+i-n}\} \cup P'_{ik} & \text{if } i \geq n \text{ and } k \leq \ell \end{cases}$$

Note that this is an instance in the CP setting. At timestep $\ell < k \leq 2n$, each agent prefers exactly one project. At timesteps $k \leq \ell$, agents $i \leq n$ prefer at least one project as $S_{ik} \neq \varnothing$ and agents $i > n$ prefer project p_{m+i-n}.

Also note that an outcome for $\mathcal{I}'$ maximizes utilitarian welfare if and only if a project in P is chosen for timesteps $k \leq \ell$ and the project p_{n+1} is chosen at timesteps $k > \ell$. For timesteps $k \leq \ell$, all projects $p \in P$ have exactly n agents approving the project and all projects $p \notin P$ has at most one agent approving the project. For timesteps $k > \ell$, project p_{n+1} has exactly n agents approving the project and all other projects has at most one agent approving the project. There is an outcome that is proportional and maximizes utilitarian welfare for $\mathcal{I}'$ if and only if there is an outcome $\mathbf{o}$ for $\mathcal{I}$ such that $u_i(\mathbf{o}) \geq 1$ for all agents $i \in N$.

For the 'if' direction, suppose there were an outcome $\mathbf{o}$ for $\mathcal{I}$ such that $u_i(\mathbf{o}) \geq 1$ for all agents $i \in N$. We construct outcome $\mathbf{o}'$ by picking the same project as $\mathbf{o}$ for the first ℓ timesteps and picking project p_{n+1} for the remaining timesteps. For agent $i \leq n$, $u_i(\mathbf{o}') \geq 1$ as we select the same first ℓ projects as $\mathbf{o}$ and the approval set of these agents are identical to the original instance for the first ℓ timesteps. For agents $i > n$, $u_i(\mathbf{o}') \geq 2n - \ell$ as these agents prefer the project p_{n+1} for the last $2n - \ell$ timesteps. Hence, $\mathbf{o}'$ achieves proportionality and maximizes utilitarian welfare.

For the 'only if' direction, suppose there were an outcome $\mathbf{o}'$ for $\mathcal{I}'$ that achieves proportionality and maximizes utilitarian welfare. As the outcome achieves proportionality, $u_i(\mathbf{o}') \geq 1$ for all agents i. Furthermore, as the outcome maximizes utilitarian welfare, project $n+1$ was chosen for the last $2n-\ell$ timesteps and no agent in N prefers any project chosen in the last $2n - \ell$ timesteps. Hence, each agent $i \in N$ prefers a project chosen in the first ℓ timesteps. We construct outcome $\mathbf{o}$ by picking the same project as $\mathbf{o}'$ for the first ℓ timesteps. As the approval set of agents $i \in N$ are identical to the original instance for the first ℓ timesteps, $u_i(\mathbf{o}) \geq 1$ for all agents $i \in N$. □

Beyond the computational aspects of picking the 'best' (according to our welfare objective) PROP outcome, another interesting question is to consider the trade-off between proportionality and one of our welfare objectives.

The concepts of the *price of fairness* and the *strong price of fairness* [5] formalize this trade-off, and have been studied for several proportionality guarantees in the single-round multiwinner voting literature [9, 17, 29]. As we are only considering PROP, we instantiate the definition of the (strong) price of fairness accordingly, as follows.

For any problem instance $\mathcal{I}$, let $\Pi_{\text{PROP}}(\mathcal{I}) \subseteq \Pi(\mathcal{I})$ denote the set of all outcomes that are proportional for $\mathcal{I}$. Also, given a welfare objective function W, let $W(\mathbf{o})$ denote the welfare of an outcome $\mathbf{o}$. For instance, $\text{UTIL}(\mathbf{o}) = \sum_{i \in N} u_i(\mathbf{o})$, and $\text{EGAL}(\mathbf{o}) = \min_{i \in N} u_i(\mathbf{o})$. Furthermore, given a welfare objective W, let W-OPT($\mathcal{I}$) denote the maximum W-welfare over all outcomes in $\Pi(\mathcal{I})$.

Definition 5 (Price of Proportionality)**.** *For a welfare objective W, the* price of proportionality (PoPROP) *is the supremum over all instances $\mathcal{I}$ of the ratio between the maximum W welfare of an outcome for $\mathcal{I}$ and the maximum W welfare of an outcome for $\mathcal{I}$ that satisfies PROP:*

$$\text{PoPROP} = \sup_{\mathcal{I}} \frac{W\text{-}OPT(\mathcal{I})}{\max_{\mathbf{o} \in \Pi_{PROP}(\mathcal{I})} W(\mathbf{o})}.$$

Definition 6 (Strong Price of Proportionality)**.** *For a welfare objective W, the* strong price of proportionality (s-PoPROP) *is the supremum over all instances $\mathcal{I}$ of the ratio between the the maximum W-welfare of an outcome for $\mathcal{I}$ and the minimum W welfare of an outcome for $\mathcal{I}$ that satisfies PROP:*

$$\text{s-PoPROP} = \sup_{\mathcal{I}} \frac{W\text{-}OPT(\mathcal{I})}{\min_{\mathbf{o} \in \Pi_{PROP}(\mathcal{I})} W(\mathbf{o})}.$$

We first observe that requiring proportionality has no impact on egalitarian welfare: any outcome $\mathbf{o}$ can be transformed (in polynomial time) into a proportional outcome $\mathbf{o}'$ so that the egalitarian welfare of $\mathbf{o}'$ is at least as high as that of $\mathbf{o}$.

Proposition 5.3. *Given an outcome $\mathbf{o}$, we can construct in polynomial time another outcome $\mathbf{o}'$ so that $\mathbf{o}'$ is proportional and its egalitarian welfare is at least as high as that of $\mathbf{o}$.*

By applying Proposition 5.3 to an outcome $\mathbf{o}$ that maximizes the egalitarian welfare, we obtain the following corollary.

Corollary 5.4. *The* PoPROP *with respect to* EGAL *is* 1.

In contrast, for utilitarian welfare, the price of proportionality scales as $\sqrt{n}$, even in the CP setting.

Theorem 5.5. *In the* CP *setting, the* PoPROP *with respect to* UTIL *is* $\frac{n}{2\sqrt{n}-1}$ *or* $\Theta(\sqrt{n})$.

Proof. We first prove the lower bound. Let k be an integer such that $1 \leq k \leq n$. Consider the instance with $P = \{p_1, \ldots, p_{n-k+1}\}$, n agents, and $\ell = n$. The agents have static preferences: their approval sets, for all timesteps $k \in [\ell]$, are defined as follows:

$$S_{ik} = \begin{cases} \{p_i\} & \text{if } i \leq n-k \\ \{p_{n-k+1}\} & \text{otherwise} \end{cases}$$

Under a UTIL outcome, project p_{n-k+1} will be chosen for all timesteps. As k agents approve p_{n-k+1}, the maximum utilitarian welfare in this case is $n \cdot k$. However, in order for an outcome to be proportional, each project must be chosen at least once. Hence, the maximum utilitarian welfare for a proportional outcome is $n-k+k^2$. Then, we get that

$$\text{PoPROP} = \frac{k \cdot n}{n-k+k^2} = \frac{n \cdot \sqrt{n}}{2n-\sqrt{n}} = \frac{n}{2\sqrt{n}-1},$$

where the second equality is obtained by letting $k = \sqrt{n}$.

Next, we prove the upper bound. We provide a constructive proof by selecting an outcome that has a utilitarian welfare of at least $\frac{2}{\sqrt{n}} - \frac{1}{n}$ of the maximum utilitarian welfare attainable. First, we split up the ℓ timesteps into *groups* of n timesteps. For the last group, if there are less than n timesteps (i.e., the number of timesteps is not a multiple of n), we simply choose projects that maximize utilitarian welfare. Next, we ensure that each agent approves of at least one of the projects chosen in each group of n timesteps. This would satisfy proportionality, as it ensures that $u_i(\mathbf{o}) \geq \lfloor \frac{\ell}{n} \rfloor$ for each agent $i \in N$. Furthermore, this also guarantees that for each group, the outcome selected in those timesteps n has utilitarian welfare at least $\frac{2}{\sqrt{n}} - \frac{1}{n}$ of the maximum utilitarian welfare attainable.

Next, we analyze a group of n timesteps. First, let p be the project in any timestep with the highest number of agents approving it. Let q be the number of agents that approve p. First, we select p at that timestep, and we pick $q-1$ more projects greedily by targeting projects at unselected timesteps with the most agents approving those projects. Note that as project p has q agents approving it, after our first q selections, there are at most $n-q$ agents that do not approve any project in the initial selection. For the remaining $n-q$ timesteps, we select projects to ensure that each agent approves of at least one project chosen in this group of timesteps. Let Q be the utilitarian welfare from our initial q selection and r be the maximum number of agents that approve a project in the remaining $n-q$ unselected timesteps. The maximum utilitarian welfare in this group of n time steps is at most $Q + (n-q) \cdot r$ and the minimum utilitarian welfare from our selection is at least $Q + n - q$, giving us:

$$\begin{aligned} \frac{1}{\text{PoPROP}} &\geq \frac{Q+n-q}{Q+(n-q)\cdot r} \geq \frac{r \cdot q + n - q}{n \cdot r} \quad (\text{as } Q \geq r \cdot q) \\ &\geq \frac{(r-1)\cdot r + n}{n \cdot r} \quad (\text{as } r \leq q) \\ &= \frac{r-1}{n} + \frac{1}{r} \end{aligned}$$

Using elementary calculus, the term $\frac{1}{\text{PoPROP}}$ is minimized when $r = \sqrt{n}$, giving us PoPROP $\leq \frac{n}{2\sqrt{n}-1} < \frac{n}{\sqrt{n}} = \Theta(\sqrt{n})$. □

For s-PoPROP, we obtain the following bounds.

Theorem 5.6. *The* s-PoPROP *with respect to* UTIL *or* EGAL *is* ∞. *However, if for all agents* $i \in N$, $\lfloor \frac{\mu_i}{n} \rfloor \geq 1$, *then* s-PoPROP *with respect to* UTIL *or* EGAL *is* $2n-1$ *or* $\Theta(n)$.

6 Conclusion and Future Work

We investigated the problem of maximizing utilitarian and egalitarian welfare for temporal elections. We showed that, while UTIL outcomes can be computed in polynomial time and can be achieved in a strategyproof manner, EGAL is NP-complete and obviously manipulable. To circumvent the NP-hardness of EGAL, we analyzed its parameterized complexity with respect to n, m and ℓ, and provided an approximation algorithm that is based on randomized rounding. We also established the existence of a NOM mechanism for EGAL under a mild constraint on agents' preferences. Finally, we considered proportionality and showed that it is computationally hard to select the 'best' proportional outcome. We also gave upper and lower bounds for the (strong) price of proportionality with respect to both UTIL and EGAL.

We discuss several possible directions for future work.

Maximizing the p-mean welfare objective One possible extension to our work is analyzing the more general *p-mean welfare* objective. The associated decision problem is defined as follows.

> MAXIMIZING p-MEAN WELFARE:
> **Input**: A problem instance $\mathcal{I} = (N, P, \ell, (\mathbf{S}_i)_{i \in N})$ and a parameter $\lambda \in \mathbb{Z}^+$.
> **Question**: Is there an outcome $\mathbf{o}$ such that $\left(\frac{1}{n}\sum_{i \in N} u_i(\mathbf{o})^p\right)^{1/p} \geq \lambda$?

Note that setting $p = 1$ (respectively, $p = -\infty$) would correspond to the utilitarian (respectively, egalitarian) welfare. Setting $p \to 0$ corresponds to maximizing the geometric mean, or Nash welfare (we denote the corresponding decision problem of maximizing the Nash welfare by NASH). We also note that many of the computational hardness and impossibility results for EGAL directly translate to similar results for NASH: NASH is obviously manipulable in the general setting and not strategyproof even in the CP setting. NASH is also NP-complete even when $m = 2$ and is W[2]-hard with respect to ℓ. The XP algorithm with respect to ℓ also works for NASH. While our FPT algorithm (with respect to n) for EGAL relies on an ILP that does not extend to NASH, a randomized XP algorithm has been proposed for a more demanding setting [18]. Now, for our setting, we can show that there is a deterministic XP algorithm (with respect to n) for any *p-mean welfare* objective, which includes NASH.

Theorem 6.1. *There exists a deterministic* XP *algorithm (with respect to n) for maximizing the p-*mean welfare.

However, the question of whether NASH admits an FPT algorithm is an open problem. Beyond NASH, it may be interesting to identify values of p for which maximizing *p-mean welfare* is tractable.

Cardinal preferences Another possible direction for building upon this work is looking into cardinal preferences, which has been studied in various social choice settings recently [6, 13, 19, 22, 25]. In our model with approval preferences, agents can also be thought of as having *binary utilities* over projects. One can extend this model by allowing each agent $i \in N$ to have a valuation function $v_i : P \times [\ell] \to \mathbb{R}_{[0,1]}$ instead of an approval set. Then, it would be interesting to investigate whether the positive results in our setting extend to the setting with cardinal preferences.

Acknowledgements

Edith Elkind was supported by the AI Programme of The Alan Turing Institute and an EPSRC Grant EP/X038548/1.

References

[1] N. Alon and J. H. Spencer. *The probabilistic method*. John Wiley & Sons, 2016.

[2] S. Alouf-Heffetz, L. Bulteau, E. Elkind, N. Talmon, and N. Teh. Better collective decisions via uncertainty reduction. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 24–30, 2022.

[3] R. P. Arribillaga and A. G. Bonifacio. Obvious manipulations of tops-only voting rules. *Games and Economic Behavior*, 143:12–24, 2024.

[4] H. Aziz and A. Lam. Obvious manipulability of voting rules. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*, pages 179–193, 2021.

[5] X. Bei, X. Lu, P. Manurangsi, and W. Suksompong. The price of fairness for indivisible goods. *Theory of Computing Systems*, 65:1069–1093, 2021.

[6] N. Boehmer, R. Bredereck, P. Faliszewski, A. Kaczmarczyk, and R. Niedermeier. Line-up elections: Parallel voting with shared candidate pool. In *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT)*, pages 275–290, 2020.

[7] R. Bredereck, A. Kaczmarczyk, and R. Niedermeier. Electing successive committees: Complexity and algorithms. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1846–1853, 2020.

[8] R. Bredereck, T. Fluschnik, and A. Kaczmarczyk. When votes change and committees should (not). In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 144–150, 2022.

[9] M. Brill and J. Peters. Completing priceable committees: Utilitarian and representation guarantees for proportional multiwinner voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9528–9536, 2024.

[10] M. Brill, P. Gölz, D. Peters, U. Schmidt-Kraepelin, and K. Wilker. Approval-based apportionment. *Mathematical Programming*, 203:77–105, 2024.

[11] L. Bulteau, N. Hazon, R. Page, A. Rosenfeld, and N. Talmon. Justified representation for perpetual voting. *IEEE Access*, 9:96598–96612, 2021.

[12] N. Chandak, S. Goel, and D. Peters. Proportional aggregation of preferences for sequential decision making. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9573–9581, 2024.

[13] V. Conitzer, R. Freeman, and N. Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.

[14] T. Delemazure, T. Demeulemeester, M. Eberl, J. Israel, and P. Lederer. Strategyproofness and proportionality in party-approval multiwinner elections. *arXiv preprint arXiv:2211.13567*, 2022.

[15] E. M. Deltl, T. Fluschnik, and R. Bredereck. Algorithmics of egalitarian versus equitable sequences of committees. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2651–2658, 2023.

[16] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48:599–632, 2017.

[17] E. Elkind, P. Faliszewski, A. Igarashi, P. Manurangsi, U. Schmidt-Kraepelin, and W. Suksompong. The price of justified representation. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 4983–4990, 2022.

[18] E. Elkind, S. Kraiczy, and N. Teh. Fairness in temporal slot assignment. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory (SAGT)*, pages 490–507, 2022.

[19] E. Elkind, W. Suksompong, and N. Teh. Settling the score: Portioning with cardinal preferences. In *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI)*, pages 621–628, 2022.

[20] E. Elkind, S. Obraztsova, and N. Teh. Temporal Fairness in Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 22633–22640, 2024.

[21] E. Elkind, S. Obraztsova, and N. Teh. Verifying proportionality in temporal multiwinner voting. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 2246–2248, 2024.

[22] B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018.

[23] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In *Trends in Computational Social Choice*, chapter 2, pages 27–47. 2017.

[24] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006.

[25] R. Freeman, S. M. Zahedi, and V. Conitzer. Fair and efficient social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4580–4587, 2017.

[26] A. Igarashi, M. Lackner, O. Nardi, and A. Novaro. Repeated fair allocation of indivisible items. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9781–9789, 2024.

[27] M. Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2103–2110, 2020.

[28] M. Lackner and J. Maly. Proportional decisions in perpetual voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5722–5729, 2023.

[29] M. Lackner and P. Skowron. Utilitarian welfare and representation guarantees of approval-based multiwinner rules. *Artificial Intelligence*, 288: 103366, 2020.

[30] M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2023.

[31] H. W. Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.

[32] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.

[33] C. H. Papadimitriou. *Computational complexity*. Academic Internet Publ., 2007. ISBN 978-1-4288-1409-7.

[34] P. Skowron and A. Górecki. Proportional public decisions. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5191–5198, 2022.

[35] P. Troyan and T. Morrill. Obvious manipulations. *Journal of Economic Theory*, 185:104970, 2020.

[36] V. Zech, N. Boehmer, E. Elkind, and N. Teh. Multiwinner temporal voting with aversion to change. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*, 2024.