PRELIMINARY PREPRINT VERSION – DO NOT CITE
This is not the version of record. The final version will appear in the Proceedings of IJCAI-ECAI 2026.

The Price of Proportional Representation in Temporal Voting

Nicholas Teh
University of Oxford
nicholas.teh@cs.ox.ac.uk

Abstract

We study proportional representation in the temporal voting model, where collective decisions are made repeatedly over time over a fixed horizon. Prior work has extensively investigated how proportional representation axioms from multiwinner voting (e.g., justified representation (JR) and its variants) can be adapted, satisfied, and verified in this setting. However, much less is understood about their interaction with social welfare. In this work, we quantify the efficiency cost of enforcing proportionality. We formalize the welfare-proportionality tension via the worst-case ratio between the maximum achievable utilitarian welfare and the maximum welfare attainable subject to a proportionality axiom. We show that imposing proportional representation in the temporal setting can incur a growing, yet sublinear, welfare loss as the number of voters or rounds increases. We further identify a clean separation among axioms: for JR, the welfare loss diminishes as the time horizon grows and vanishes asymptotically, whereas for stronger axioms this conflict persists even with many rounds. Moreover, we prove that welfare maximization under each axiom is NP-complete and APX-hard, even under static preferences and bounded-degree approvals, and provide fixed-parameter algorithms under several natural structural parameters.

1 Introduction

Many collective decisions are inherently multi-stage: rather than selecting a single outcome once, an institution commits to a sequence of choices over a horizon. In modern ML-driven systems, this occurs whenever a platform allocates scarce *exposure* repeatedly: for instance, selecting a featured creator each day, a highlighted job posting each week, or a promoted product category each campaign cycle. A large body of work in information retrieval and recommender systems studies exactly this kind of repeated exposure allocation, where fairness is evaluated *across many outputs* (often called *amortized* or *long-run* fairness) and must be balanced against utility or engagement [Biega *et al.*, 2018; Singh and Joachims, 2018; Pitoura *et al.*, 2022]. These settings also need not be online in the algorithmic sense: platforms commonly plan or optimize slates and exposure schedules in batch using historical logs, forecasts, or business constraints.

We study such problems through the lens of the *temporal voting* model, a framework for long-term collective decision-making [Bulteau *et al.*, 2021; Chandak *et al.*, 2024; Elkind *et al.*, 2025b; Phillips *et al.*, 2026]. An instance consists of a finite horizon of ℓ rounds together with an approval profile for each round.[1] In each round the rule selects one alternative, and alternatives may be selected multiple times. This repeated nature of selection creates a fundamental tension: selecting, in each round, an alternative that has maximum support will guarantee high utilitarian welfare. However, repeatedly using such a greedy rule can systematically exclude minorities; for instance, even a sizeable segment of the population may see none of its approved alternatives selected over the entire horizon. Accordingly, the natural fairness question is whether we can achieve some kind of representation *across time*.

The concept of *proportional representation* is particularly relevant here. Recent work adapts classic proportionality axioms from approval-based multiwinner voting—in particular, justified representation (JR), proportional justified representation (PJR), and extended justified representation (EJR)—to the temporal setting [Bulteau *et al.*, 2021; Chandak *et al.*, 2024]. These axioms ensure that any sufficiently large and cohesive group obtains its "fair share" of representation when representation is evaluated over the whole horizon. More recently, strengthened variants such as EJR+ were proposed to further tighten proportionality guarantees over time [Phillips *et al.*, 2026].

However, while these axioms provide compelling group-representation guarantees, they may force the outcome to deviate from the welfare-maximizing optimum. Importantly, this tension arises even under static preferences and offline optimization, because proportionality imposes global, group-based constraints that couple the choices across rounds.

This raises two basic questions that are central both in the-

[1] We study the *offline* planning problem: the entire ℓ-round instance (all round profiles) is given as input, and the rule outputs the full length-ℓ sequence at once.

ory and practice: (1) How costly is enforcing proportional representation in temporal voting? (2) Can we compute a high-welfare outcome that satisfies JR/PJR/EJR/EJR+?

To address the first question, we adopt the "price of fairness" perspective and study the *price of proportional representation*: the worst-case ratio between the maximum utilitarian welfare achievable without constraints and the maximum utilitarian welfare achievable subject to a proportionality axiom. This viewpoint provides a crisp, quantitative answer to the welfare-fairness tension: it identifies when proportionality is essentially "free" and when it provably imposes a substantial welfare penalty.

To address the second question, we study the computational complexity of optimizing welfare subject to these proportionality axioms. While proportionality constraints are conceptually appealing, they implicitly encode global combinatorial structure (i.e., which cohesive groups must be represented and how their representation must be distributed over time) and it is not clear a priori whether this structure admits efficient optimization, even under restrictive preference assumptions.

1.1 Our Contributions

We study four proportionality axioms $\{\text{JR}, \text{PJR}, \text{EJR}, \text{EJR+}\}$ in the temporal voting setting, and analyze their compatibility with utilitarian social welfare.

In Section 3, we *quantify* the welfare-proportionality tension via a *price of proportional representation* framework. We establish a general lower bound showing that enforcing any of JR/PJR/EJR/EJR+ can incur a growing (yet sublinear) loss of $\Omega(\sqrt{\ell})$ in the worst case, where ℓ is the number of rounds. We also prove a universal upper bound of n (the number of agents) on the price. For JR, we go further and obtain a tight, horizon-sensitive guarantee: the price decreases with ℓ and becomes asymptotically negligible when $\ell \gg n$ (approaching 1). In contrast, for the stronger axioms PJR/EJR/EJR+, the price can remain $\Omega(\sqrt{n})$ even when $\ell = \Theta(n)$, giving us a clear separation between JR and its stronger counterparts.

In Section 4, we show that maximizing utilitarian welfare subject to any of JR/PJR/EJR/EJR+ is NP-complete, and that the corresponding optimization problem is APX-hard, even under static preferences and constant approval-degree bounds. This demonstrates that the computational barrier is intrinsic to the group-based proportionality constraints rather than being dependent on time-varying preferences.

Finally, in Section 5, we complement the hardness results by identifying practically meaningful regimes where optimal solutions can be computed efficiently. Under static preferences, we obtain fixed-parameter tractability with respect to the number of candidates m. More generally, we develop algorithms that exploit temporal structure, including FPT results parameterized by the number of voter types (and related temporal-compression parameters), showing that repeated preference patterns across rounds enable tractable welfare maximization under temporal proportionality.

Overall, our results provide the first systematic study of the efficiency cost and computational compatibility of temporal proportional representation axioms with social welfare.

1.2 Related Work

We briefly survey the most relevant strands of the literature and emphasize works that are directly related to our work.

Temporal voting. Proportional representation in temporal voting was first studied by Bulteau et al. [2021]. They adapted two proportional representation axioms (JR and PJR) from multiwinner voting to the temporal setting (with analyses for both static and changing preferences). Chandak et al. [2024] subsequently extended this analysis to the stronger EJR, and showed that EJR exists in the temporal setting. Elkind et al. [2025b] built on this and study the complexity of checking whether a given outcome satisfies JR/PJR/EJR. Recently, Phillips et al. [2026] examined various ways of strengthening these concepts, and showed that EJR+ and FJR (both strengthenings of EJR) are always satisfiable. We refer the reader to the survey by Elkind et al. [2024c] for an overview of this area.

In an adjacent line of work, Elkind et al. [2024b] study temporal elections with a focus on welfare and incentive-compatibility. They investigate the computational complexity of welfare maximization and its compatibility with strategyproofness. They also adapt an individual proportionality notion (primarily studied in fair division and public decision-making), and study its compatibility and tradeoffs with utilitarian and egalitarian welfare objectives. This distinction is important for our results. The individual proportionality notion studied by Elkind et al. [2024b] is vacuous in some of our hardness constructions. Subsequently, Elkind et al. [2025a] consider analogous welfare optimization questions for negatively valued candidates, where proportionality notions are not well-defined. Our work differs in that we focus on *group-based* proportional representation axioms, which capture collective fairness guarantees that are fundamentally different from individual proportionality (where relevant, we will point out the connection and highlight the differences).

Another similar (but fundamentally distinct) body of work studies the *perpetual voting* model [Lackner, 2020; Lackner and Maly, 2023]. This model differs from temporal voting in that decisions explicitly depend on the history of past rounds, giving it a more inherently online flavor.

Proportionality in multiwinner voting. The proportionality axioms we study are inspired by the well-developed literature on proportional representation in approval-based multiwinner voting. In that setting, justified representation (JR) and extended justified representation (EJR) were introduced by Aziz et al. [2017], and proportional justified representation (PJR) was later proposed as an intermediate axiom between JR and EJR by Sanchez-Fernandez et al. [2017]. Subsequently, Brill and Peters [2023] introduced EJR+, a strengthening of EJR that preserves desirable practical properties such as polynomial-time verifiability and satisfaction by polynomial-time computable rules.

Temporal variants of these axioms adapt the same group-protection intuition, but the introduction of rounds fundamentally changes both the structure of the constraints (e.g., how representation is required to accrue over time) and the way proportionality interacts with efficiency. Moreover, unlike in multiwinner voting, the same alternative may be selected repeatedly across rounds, which further alters the behavior of proportionality axioms and the design space of rules. Our work advances this line of research by quantifying the welfare-proportionality trade-off for temporal variants of the JR family, and by showing that the length of the time horizon affects this trade-off in qualitatively different ways depending on the strength of the axiom. Equivalently, even under static preferences, a temporal outcome is a multiset of ℓ candidate occurrences rather than a committee of ℓ distinct or interchangeable winners: multiplicities matter because selecting the same candidate repeatedly consumes rounds and repeatedly satisfies the same voters.

Proportionality has also been studied in the context of *public decision-making* (simultaneous aggregation over multiple issues), both without constraints [Skowron and Górecki, 2022] and with feasibility constraints [Chingoma *et al.*, 2025] through adaptations of proportional representation axioms from multiwinner voting, as well as in even more general models of feasibility constraints [Masařík *et al.*, 2024].

Welfare-proportionality tradeoffs. A complementary line of work studies proportionality constraints through a quantitative lens, often phrased as a "price of fairness" or "price of proportionality" measure. In the context of multiwinner voting, Elkind et al. [2024a] initiate a systematic study of the welfare and coverage losses incurred by imposing JR and EJR constraints. Related quantitative analyses adopt approximation-style guarantees to compare voting rules against welfare- or representation-oriented benchmarks [Lackner and Skowron, 2020]. Similar welfare–representation tradeoffs have also been examined in participatory budgeting [Fairstein *et al.*, 2022], and in designing multiwinner rules that simultaneously provide utilitarian and representation guarantees [Brill and Peters, 2024].

Our work extends this quantitative perspective to the temporal voting setting by defining and bounding the utilitarian price of proportionality.

2 Preliminaries

For every natural number $k \in \mathbb{N}$, we let $[k] = \{1, 2, \dots, k\}$. A *temporal election* is a tuple $E = (P, N, \ell, (\mathbf{s}_i)_{i\in N})$, where P is the set of m *candidates*, N is the set of n *voters*, ℓ is the number of *rounds*, and for each $i \in N$, $\mathbf{s}_i = (s_{i,1}, s_{i,2}, \dots, s_{i,\ell})$, where $s_{i,r} \subseteq P$ is the *approval set* of voter i in round r, which consists of candidates that i approves in round r.[2] We denote the set of all temporal elections by $\mathcal{E}$ and the set of temporal elections whereby in every round every voter approves at least one candidate by $\mathcal{E}_{\geq 1}$ (which we call *complete* elections). We say that voters from a subset $S \subseteq N$ *agree* in a round $r \in [\ell]$ if there exists a candidate that they all approve in this round, i.e., $\bigcap_{i\in S} s_{i,r} \neq \varnothing$.

An *outcome* of a temporal election E is a sequence $\mathbf{o} = (o_1, o_2, \dots, o_\ell) \in P^\ell$ of ℓ candidates such that for every $r \in [\ell]$, candidate $o_r \in P$ is chosen in round r. A candidate may be selected multiple times, i.e., it may be the case that $o_t = o_{r'}$ for some $r \neq r'$. For a subset of rounds $R \subseteq [\ell]$ and an outcome $\mathbf{o}$, we write $\mathbf{o}_R = (o_r)_{r\in R}$ to denote the *suboutcome* with respect to R. The *satisfaction* of a subset of voters $S \subseteq N$ from a suboutcome $\mathbf{o}_R$ is $\mathrm{sat}_S(\mathbf{o}_R) = \left|\{r \in R : \ o_r \in \bigcup_{i\in S} s_{i,r}\}\right|$, i.e., the number of rounds in R in which the selected candidate is approved by at least one voter from S. If $S = \{i\}$ for some $i \in N$, then we simply write $\mathrm{sat}_i(\mathbf{o}_R)$.

2.1 Proportional Representation Axioms

We now provide the definitions of the proportionality axioms we consider, which were originally introduced in prior work on temporal voting [Bulteau *et al.*, 2021; Chandak *et al.*, 2024; Elkind *et al.*, 2025b].

Definition 2.1 (JR/PJR/EJR)**.** Given a temporal election $E = (P, N, \ell, (\mathbf{s}_i)_{i\in N})$ and an outcome $\mathbf{o}$, if, for each $t \in [\ell]$ and every nonempty subset of voters $S \subseteq N$ that agrees in a size-t subset of rounds, it holds that

- $\mathrm{sat}_S(\mathbf{o}) \geq \min(1, \lfloor t \cdot |S|/n \rfloor)$, then $\mathbf{o}$ provides *justified representation* (JR),
- $\mathrm{sat}_S(\mathbf{o}) \geq \lfloor t \cdot |S|/n \rfloor$, then $\mathbf{o}$ provides *proportional justified representation* (PJR),
- $\mathrm{sat}_i(\mathbf{o}) \geq \lfloor t \cdot |S|/n \rfloor$ for some $i \in S$, then $\mathbf{o}$ provides *extended justified representation* (EJR).

While JR, PJR, and EJR impose proportionality guarantees based on the total number of rounds on which a group agrees, they do not ensure that such agreement translates into representation in *any particular round*. In particular, a group may be sufficiently large and cohesive to warrant proportional representation overall, yet still be unrepresented in a round where all its members agree. To capture this stronger notion of temporal proportionality, Phillips et al. [2026] proposed extended justified representation + (EJR+),[3] which requires that sufficiently cohesive groups either obtain proportional satisfaction overall or receive representation in every round on which they unanimously agree. It is defined as follows.

Definition 2.2 (EJR+)**.** Given a temporal election $E = (P, N, \ell, (\mathbf{s}_i)_{i\in N})$ and $\sigma \in [n]$, $\tau \in [\ell]$, we say that a nonempty subset of voters $S \subseteq N$ is (σ, τ)-*cohesive* if there exists a set of τ rounds $R \subseteq [\ell]$ and a suboutcome $\mathbf{o}_R = (o_r)_{r\in R}$ such that for each $r \in R$, at least σ voters in S approve o_r.

We say that an outcome $\mathbf{o}$ provides *extended justified representation + (EJR+)* if for all $\sigma \in [n]$, $\tau \in [\ell]$, every (σ, τ)-cohesive nonempty subset of voters $S \subseteq N$ and every round $r \in [\ell]$ with $\bigcap_{i\in S} s_{i,r} \neq \varnothing$, it holds that (i) $\mathrm{sat}_i(\mathbf{o}) \geq \lfloor \tau \cdot \sigma/n \rfloor$ for some $i \in S$, or (ii) $o_r \in \bigcap_{i\in S} s_{i,r}$.

Phillips et al. [2026] proposed two strengthenings of EJR that are always satisfiable: EJR+ and *full justified representation* (FJR). The latter is inspired by an analogous strengthening studied in the participatory budgeting setting—a generalization of multiwinner voting that allows items to have

[2] Preferences are said to be *static* if for every voter $i \in N$ there exists a fixed approval set $A_i \subseteq P$ such that $s_{i,r} = A_i$ for every round $r \in [\ell]$.

[3] This notion is based on the analogous strengthening of EJR in the multiwinner voting setting [Brill and Peters, 2023].

variable costs [Peters *et al.*, 2021]. In this work, we focus on EJR+ rather than FJR, as EJR+ is better established in the multiwinner voting literature and preserves key practical properties of its multiwinner analogue—most notably polynomial-time verifiability and satisfaction by a polynomial-time computable rule. In contrast, while FJR is always satisfiable, the known constructive approach relies on a procedure whose running time is not polynomial, and Phillips et al. [2026] observe that it remains open whether FJR outcomes can be computed in polynomial time.

2.2 Price of Proportional Representation

In this work, we focus on utilitarian social welfare, which is the canonical efficiency benchmark in approval-based voting and is standard in quantitative welfare-proportionality tradeoff analyses [Lackner and Skowron, 2020; Elkind *et al.*, 2024a; Elkind *et al.*, 2024b]. We view utilitarian welfare as the most informative baseline for isolating the efficiency cost of proportional representation: unlike other (e.g., egalitarian or Nash) objectives, which already encode inequality aversion, utilitarian welfare measures aggregate efficiency directly. Moreover, utilitarian welfare is additively separable across rounds, making the unconstrained optimum efficiently computable in our model and enabling transparent worst-case comparisons.

Given a temporal election $E = (P, N, \ell, (\mathbf{s}_i)_{i \in N})$ and an $\mathbf{o} = (o_1, \dots, o_\ell) \in P^\ell$, the *utilitarian welfare* of $\mathbf{o}$ under E is the total number of approving voter-round pairs:

$$\textsc{Util}_E(\mathbf{o}) := \sum_{i \in N} \mathrm{sat}_i(\mathbf{o}) = \sum_{r=1}^{\ell} \big|\{i \in N : o_r \in s_{i,r}\}\big|.$$

For notational simplicity, we omit the subscript E when it is obvious from context. More generally, one could consider an arbitrary welfare function $W_E : P^\ell \to \mathbb{R}_{\geq 0}$.

Let Φ be a property of the outcome, e.g., $\Phi \in \{\text{JR}, \text{PJR}, \text{EJR}, \text{EJR+}\}$. Fix integers $n, \ell \geq 1$, and let $\mathcal{E}^{n,\ell}$ be the class of temporal elections with n voters and ℓ rounds; define $\mathcal{E}^{n,\ell}_{\geq 1}$ analogously for complete elections. For an election $E \in \mathcal{E}^{n,\ell}$, define the unconstrained optimum by $W^*(E) := \max_{\mathbf{o} \in P^\ell} W_E(\mathbf{o})$, and the Φ-constrained optimum by $W^\Phi(E) := \max\{W_E(\mathbf{o}) : \mathbf{o} \in P^\ell \text{ and } \mathbf{o} \text{ provides } \Phi\}$. For a fixed election E with $W^\Phi(E) > 0$, let $\rho_{W,\Phi}(E) := \frac{W^*(E)}{W^\Phi(E)}$. For a class $\mathcal{C} \subseteq \mathcal{E}^{n,\ell}$, the price of Φ with respect to W over $\mathcal{C}$ is $\mathcal{P}_W(\Phi; \mathcal{C}) := \sup_{E \in \mathcal{C} : W^\Phi(E) > 0} \rho_{W,\Phi}(E)$. In Section 3, we take $\mathcal{C} = \mathcal{E}^{n,\ell}_{\geq 1}$, since the price bounds there are for complete elections. To lighten notation, we let $\mathcal{P}^{n,\ell}_W(\Phi) := \mathcal{P}_W(\Phi; \mathcal{E}^{n,\ell}_{\geq 1})$ when the underlying class is clear.

Then, we will study the following decision problem.

Φ-UTIL

Input: A temporal election E and an integer threshold B.
Question: Does there exist an outcome $\mathbf{o} \in P^\ell$ such that $\mathbf{o}$ provides Φ and $\textsc{Util}(\mathbf{o}) \geq B$?

We also consider the associated optimization problem.

Φ-MAXUTIL

Input: A temporal election E.
Task: Compute an outcome $\mathbf{o} \in P^\ell$ such that $\mathbf{o}$ provides Φ and $\textsc{Util}(\mathbf{o})$ is maximized.

3 Price of Proportional Representation

We begin by quantifying the welfare cost of enforcing proportional representation in temporal elections.

Here, we make two simple assumptions that are needed to ensure that the price measure is well-defined and meaningful. First, we restrict attention to complete elections $E \in \mathcal{E}_{\geq 1}$, i.e., $s_{i,r} \neq \varnothing$ for all $i \in N$ and all $r \in [\ell]$. Without this, the ratio can be inflated by trivial "empty approval" rounds in which no candidate is approved, and the measure would be uninformative. Second, we focus on the setting where $\ell \geq n$. To see why, our proportionality axioms "assign representation" in indivisible units of rounds via terms $\lfloor t \cdot |S|/n \rfloor$. When $\ell < n$, even a voter who has a nonempty approval set at every round has $\lfloor \ell/n \rfloor = 0$, and JR/PJR/EJR impose no per-voter guarantee and can be satisfied while ignoring some voters in every round. Moreover, this is unavoidable: in complete elections, one can have disjoint approvals so that each round can satisfy at most one voter, making "everyone is represented at least once" infeasible unless $\ell \geq n$. This is also the minimal non-degenerate setting in which the proportionality axioms impose positive individual entitlements and have some bite; and thus, the resulting price isolates the meaningful welfare-proportionality tension.

We begin with a general lower bound.

Proposition 3.1. *Fix any $\ell \geq 1$ and $\Phi \in \{JR, PJR, EJR, EJR+\}$. There exists a complete temporal election $E \in \mathcal{E}^{\ell,\ell}_{\geq 1}$ such that $\rho_{\textsc{Util},\Phi}(E) \geq \frac{1}{2}\sqrt{\ell}$. Consequently, $\mathcal{P}^{\ell,\ell}_{\textsc{Util}}(\Phi) \geq \frac{1}{2}\sqrt{\ell} = \frac{1}{2}\sqrt{n}$.*

Intuitively, the construction creates a "core" block of about $\sqrt{\ell}$ voters who share a consistently popular candidate, alongside many voters whose approvals are essentially private. The welfare-maximizing outcome repeatedly satisfies the popular block, but any Φ-feasible outcome must spend many rounds on low-support candidates to ensure some baseline representation for the private voters, producing a $\Theta(\sqrt{\ell})$ welfare gap.

Next, we provide a universal upper bound, which serves more as a coarse sanity check showing imposing proportionality cannot blow up welfare by more than a factor n.

Proposition 3.2. *Fix any $n, \ell \geq 1$, any $\Phi \in \{JR, PJR, EJR, EJR+\}$, and any complete temporal election $E \in \mathcal{E}^{n,\ell}_{\geq 1}$ with $\textsc{Util}^\Phi(E) > 0$. Then $\rho_{\textsc{Util},\Phi}(E) \leq n$. Consequently, $\mathcal{P}^{n,\ell}_{\textsc{Util}}(\Phi) \leq n$.*

The proof combines two simple facts: (i) welfare in any round is at most n, and $\textsc{Util}^*(E) \leq n\ell$; (ii) for $E \in \mathcal{E}_{\geq 1}$, any Φ-feasible outcome can be modified (without affecting feasibility) so that every round selects a candidate approved by at least one voter, guaranteeing welfare at least ℓ.

For JR we can substantially strengthen Proposition 3.2; crucially, the bound improves with number of rounds.

Theorem 3.3. *Fix n, ℓ with $\ell \geq n$. Then, for every complete temporal election $E \in \mathcal{E}^{n,\ell}_{\geq 1}$, $\rho_{\text{UTIL},JR}(E) \leq \frac{\ell}{\ell - n + 2\sqrt{n} - 1}$. Consequently, $\mathcal{P}^{n,\ell}_{\text{UTIL}}(JR) \leq \frac{\ell}{\ell - n + 2\sqrt{n} - 1}$. In particular, for every fixed $\varepsilon > 0$, if $\ell = (1+\varepsilon)n$, then $\mathcal{P}^{n,\ell}_{\text{UTIL}}(JR) \leq \frac{1+\varepsilon}{\varepsilon + o(1)}$, and if $\ell/n \to \infty$, then $\mathcal{P}^{n,\ell}_{\text{UTIL}}(JR) \to 1$.*

At a high level, the proof is constructive: it identifies n "cheap" rounds and uses them to ensure that every voter is represented at least once, while selecting per-round welfare maximizers in the remaining rounds. An averaging argument then shows that even on the reserved n rounds, one can retain a $\Theta(1/\sqrt{n})$ fraction of the local welfare optimum.

We complement Theorem 3.3 with a matching lower bound, showing tightness up to lower-order terms.

Proposition 3.4. *Fix n, ℓ with $\ell \geq n$. There exists a complete temporal election $E \in \mathcal{E}^{n,\ell}_{\geq 1}$ such that $\rho_{\text{UTIL},JR}(E) \geq \frac{\ell}{\ell - n + 2\sqrt{n} - \mathcal{O}(1)}$. Consequently, $\mathcal{P}^{n,\ell}_{\text{UTIL}}(JR) \geq \frac{\ell}{\ell - n + 2\sqrt{n} - \mathcal{O}(1)}$.*

Remark 3.5. Elkind et al. [2024b] study temporal elections under an *individual* proportionality axiom PROP, which requires each voter i to be satisfied in at least $\lfloor \mu_i / n \rfloor$ rounds, where μ_i is the number of rounds in which i approves at least one candidate (so $\mu_i = \ell$ on complete elections). When $\lfloor \ell/n \rfloor = 1$ (in particular, when $\ell = n$), PROP reduces to the requirement that every voter is satisfied at least once; on complete elections, this condition coincides with JR. In this setting, Elkind et al. [Elkind *et al.*, 2024b, Thm. 5.6] determine the worst-case utilitarian price exactly as $n/(2\sqrt{n} - 1)$; their lower-bound instance corresponds to the $\ell = n$ special case of Proposition 3.4. Our focus is different: we analyze group-based axioms and explicitly track how the welfare loss varies with the horizon ℓ; Theorem 3.3 and Proposition 3.4 give a tight ℓ-sensitive bound for JR, showing that its price decreases with ℓ and tends to 1 when $\ell \gg n$.

Finally, we make explicit the qualitative separation suggested by Theorem 3.3: increasing the time horizon can make JR cheap, but it need not reduce the cost of stronger axioms.

Theorem 3.6. *Fix an integer $a \geq 2$ and set $\ell = an$. Then $P^{n,an}_{\text{UTIL}}(JR) \leq \frac{a}{a-1} + o(1)$ as $n \to \infty$. However, for each $\Phi \in \{PJR, EJR, EJR+\}$, $P^{n,an}_{\text{UTIL}}(\Phi) = \Omega(\sqrt{n})$.*

The results in this section show that enforcing proportional representation in temporal elections can incur a provable sublinear welfare loss. For JR, this loss decreases with ℓ and vanishes as ℓ becomes large relative to n (Theorem 3.3), whereas for stronger axioms (PJR/EJR/EJR+) the cost can remain $\Omega(\sqrt{n})$ even when ℓ is a constant-factor multiple of n (Theorem 3.6).

4 Computational Intractability Results

The previous section quantified the welfare loss that can arise from enforcing proportionality in temporal elections. We now turn to the algorithmic question: given a temporal election, can we compute a high welfare outcome that satisfies a chosen proportionality axiom? Formally, we study the decision problem Φ-UTIL and its optimization variant Φ-MAXUTIL (as defined in Section 2) for $\Phi \in \{\text{JR}, \text{PJR}, \text{EJR}, \text{EJR+}\}$.

Our key findings are: for each $\Phi \in \{\text{JR}, \text{PJR}, \text{EJR}, \text{EJR+}\}$, maximizing utilitarian welfare subject to Φ is NP-hard, and even APX-hard, under very strong structural restrictions. In particular, our reductions use static preferences (approval sets do not vary across rounds) and constant approval degree bounds (each candidate is approved by a constant number of voters), so that feasibility verification is straightforward; the hardness comes from the combinatorial property of simultaneously meeting many group-based representation constraints while retaining high welfare.

We begin with two simple structural observations that let us treat JR/PJR/EJR (and, on static instances, also EJR+) in a unified way. We then prove NP-completeness of Φ-UTIL and strengthen this to APX-hardness for Φ-MAXUTIL via a gap-preserving reduction.

Our NP-hardness construction then forces every *positive* proportional entitlement to be exactly 1. Intuitively, whenever a cohesive group is large/frequent enough to trigger a proportionality requirement at all, it is only entitled to *one* unit of representation over the whole horizon. In this setting, the distinctions between JR, PJR, and EJR disappear: all three axioms reduce to the same proportionality condition.

Lemma 4.1. *Fix a temporal election $E = (P, N, \ell, (\mathbf{s}_i)_{i \in N})$. Assume that for every integer $t > 0$ and every nonempty voter group $S \subseteq N$ that agrees in a size-t subset of rounds, we have $\lfloor t \cdot |S|/n \rfloor \leq 1$. Then for any outcome $\mathbf{o}$, the following are equivalent: $\mathbf{o}$ satisfies JR, $\mathbf{o}$ satisfies PJR, and $\mathbf{o}$ satisfies EJR.*

Lemma 4.1 lets a single reduction establish hardness for all of JR/PJR/EJR, provided we ensure the stated condition. Our constructions achieve this in a clean, verifiable way: we enforce a constant bound k on the number of approvers per candidate, which implies that any agreeing group S must be contained in the approver set of some candidate and hence has size at most k. By choosing n sufficiently large relative to ℓ and k, we guarantee $\lfloor t \cdot |S|/n \rfloor \leq \lfloor \ell \cdot k/n \rfloor \leq 1$ for all relevant (S, t).

Next we handle EJR+. While EJR+ is strictly stronger than EJR in general, under static preferences it collapses back to EJR: if a group ever unanimously agrees, it agrees in *every* round, so EJR already forces enough aggregate satisfaction to satisfy the EJR+ disjunction via (i).

Lemma 4.2. *For any temporal election with static preferences, an outcome satisfies EJR+ iff it satisfies EJR.*

Together, Lemmas 4.1 and 4.2 imply that, on the static, "low demand" instances used in our reductions, all axioms $\Phi \in \{\text{JR}, \text{PJR}, \text{EJR}, \text{EJR+}\}$ impose the same feasibility constraints (up to notational differences).

We now show that Φ-UTIL is NP-complete even under very strong restrictions. The reduction is designed to isolate the algorithmic difficulty to the group-based representation constraints: preferences are static, and each candidate is approved by only constantly many voters, so the structure of potentially binding groups is simple; nevertheless, selecting a high welfare, feasible outcome encodes an NP-complete covering problem.

Theorem 4.3. *Fix any $\Phi \in \{JR, PJR, EJR, EJR+\}$. Then, Φ-UTIL is NP-complete, even when voters have static preferences and each candidate is approved by at most 4 voters.*

The above result rules out polynomial-time exact optimization unless $\mathrm{P} = \mathrm{NP}$. One might still hope for a PTAS for Φ-MAXUTIL. We rule this out by proving APX-hardness under similarly strong restrictions.

Theorem 4.4. *Fix any $\Phi \in \{JR, PJR, EJR, EJR+\}$. Then Φ-MAXUTIL is APX-hard, even when voters have static preferences, each voter approves at most 3 candidates and each candidate is approved by at most 8 voters.*

It is worth contrasting our computational results with those of Elkind et al. [2024b]. They show NP-hardness of deciding whether there exists an *individually* proportional (PROP) outcome that also maximizes utilitarian welfare. Our focus is different: we study *group-based* temporal proportionality axioms JR/PJR/EJR/EJR+, which impose fundamentally different constraints (and are in general incomparable with PROP). Moreover, our hardness reductions already go through under static preferences, so the intractability is not due to time-varying approvals. Finally, both reductions operate in settings where $\lfloor \ell/n \rfloor = 0$, so PROP is vacuous, yet JR/PJR/EJR/EJR+ remain nontrivial due to cohesive groups—highlighting that the difficulty is inherent to *group* representation over time.

5 Parameterized Complexity Results

The intractability results in the previous section hold even under static preferences, and rule out polynomial-time algorithms or approximation schemes in the worst case. Nevertheless, temporal elections in practice often exhibit additional structure. In such settings, worst-case hardness need not preclude efficient algorithms. Thus, we study the parameterized complexity of Φ-MAXUTIL (see Section 2). Our goal is to identify natural structural parameters under which the problem becomes tractable, and to make explicit how the algorithmic difficulty depends on these parameters. Throughout, we give exact algorithms (rather than approximation guarantees), thereby complementing the hardness results of Section 4.

As a simple baseline, when the number of rounds ℓ is treated as a parameter, Φ-MAXUTIL can be solved by brute-force enumeration of outcomes: we enumerate all sequences $\mathbf{o} \in P^\ell$, compute $\mathrm{UTIL}(\mathbf{o})$, and retain the best Φ-feasible outcome. Since verifying Φ-feasibility of a given outcome o is in XP with respect to ℓ [Elkind *et al.*, 2025b, Prop. 5.3], enumerating all m^ℓ outcomes gives us an XP algorithm with respect to ℓ: for some computable function f, the running time is $m^\ell \cdot (n+m)^{f(\ell)}$. Thus, Φ-MAXUTIL is in XP with respect to ℓ. While this approach might be impractical for large ℓ, it provides a useful reference point in settings where ℓ is constant/bounded. We then turn to more meaningful structural parameters.

5.1 Fixed m under Static Preferences

Section 4 shows that even under static preferences, Φ-UTIL is NP-hard and Φ-MAXUTIL is APX-hard for all $\Phi \in \{\mathrm{JR}, \mathrm{PJR}, \mathrm{EJR}, \mathrm{EJR+}\}$; thus, tractability cannot be recovered merely by assuming static preferences. However, this structural assumption becomes algorithmically useful when combined with small parameters. In particular, if the number of candidates m is fixed, an outcome is characterized (up to permuting rounds) by its candidate multiplicities (i.e., how many times each candidate is selected). Formally, each voter $i \in N$ approves a fixed set $A_i \subseteq P$ in every round. Any outcome $\mathbf{o} \in P^\ell$ is therefore fully specified (up to permuting rounds) by its *multiplicity vector* $\mathbf{x} \in \mathbb{Z}^P_{\geq 0}$ with $x_p := |\{r \in [\ell] : o_r = p\}|$ and $\sum_{p \in P} x_p = \ell$. Under $\mathbf{x}$, voter i's satisfaction and utilitarian welfare are linear: $\mathrm{sat}_i(\mathbf{x}) = \sum_{p \in A_i} x_p$ and $\mathrm{UTIL}(\mathbf{x}) = \sum_{p \in P} w(p) x_p$, where $w(p) := |\{i \in N : p \in A_i\}|$. To compress the instance, we group identical approval sets into *approval classes* $\mathcal{C} := \{A_i : i \in N\}$, and write $n_C := |\{i \in N : A_i = C\}|$ for each class $C \in \mathcal{C}$. All voters in a class share the same satisfaction value $v_C(\mathbf{x}) := \sum_{p \in C} x_p$.

The remaining work is to encode the proportionality axiom as constraints over $\mathbf{x}$ (and a bounded number of auxiliary integer variables). When the number of candidates m is fixed, these ILPs have a number of integer variables bounded by a function of m only (e.g., via at most $|\mathcal{C}| \leq 2^m$ auxiliary variables), and thus are solvable in time $f(m) \cdot \mathrm{poly}(n, \ell)$ for ILPs in fixed dimension using the result of Lenstra Jr [1983].

Although EJR implies PJR implies JR (and on static instances EJR+ is equivalent to EJR by Lemma 4.2), we state separate optimization theorems because the objective is welfare maximization under constraints: weaker axioms enlarge the feasible set, so an outcome that maximizes welfare subject to EJR may be feasible for PJR and JR but not optimal for them. Moreover, the constraint structure differs qualitatively. This contrasts with the "low demand" setting used in Section 4 (Lemma 4.1), where the three axioms coincide; here demands can exceed 1, so the distinctions matter for optimization.

Theorem 5.1. *Fix a temporal election $E = (P, N, \ell, (A_i)_{i \in N})$ with static preferences, and let $m = |P|$. Then JR-MAXUTIL is solvable in time $f(m) \cdot \mathrm{poly}(n, m, \ell)$ for some computable function f. In particular, it is FPT with respect to m.*

Theorem 5.2. *Fix a temporal election with static preferences $A_1, \ldots, A_n \subseteq P$ and ℓ rounds. Then PJR-MAXUTIL is solvable in time $f(m) \cdot \mathrm{poly}(n, \ell)$ for some computable function f. In particular, it is FPT with respect to m.*

Theorem 5.3. *Fix a temporal election with static preferences $A_1, \ldots, A_n \subseteq P$ and ℓ rounds. Then EJR-MAXUTIL (and thus EJR+-MAXUTIL) is solvable in time $f(m) \cdot \mathrm{poly}(n, \ell)$ for some computable function f. In particular, it is FPT with respect to m.*

5.2 Fixed Voter Types

We now move beyond static preferences while retaining a common form of structure in applications: although approvals may vary over time, the electorate often consists of a small number of *behavioral types* (e.g., user segments) whose members react similarly in every round. We show that this structure gives us fixed-parameter tractability for welfare maximization under temporal proportionality.

Formally, let T be a set of voter types and let $\kappa := |T|$. The voters are partitioned into disjoint sets $(N_\theta)_{\theta \in T}$, where $n_\theta := |N_\theta|$ and $\sum_{\theta \in T} n_\theta = n$. For each round $r \in [\ell]$, voters of type θ share a common approval set $A_{\theta,r} \subseteq P$; i.e., for every $i \in N_\theta$ we have $s_{i,r} = A_{\theta,r}$.

This representation allows us to work at the type level. For an outcome $\mathbf{o} = (o_1, \ldots, o_\ell)$ define the (common) satisfaction of type θ by $\mathrm{sat}_\theta(\mathbf{o}) := |\{r \in [\ell] : o_r \in A_{\theta,r}\}|$. For any voter $i \in N_\theta$, we have $\mathrm{sat}_i(\mathbf{o}) = \mathrm{sat}_\theta(\mathbf{o})$. Moreover, utilitarian welfare can be written compactly as $\textsc{Util}(\mathbf{o}) = \sum_{r=1}^{\ell} \sum_{\theta \in T : o_r \in A_{\theta,r}} n_\theta$. It will be convenient to encode the effect of choosing a candidate in round r by the set of approving types. For each round r and candidate $p \in P$, let $X_r(p) := \{\theta \in T : p \in A_{\theta,r}\} \subseteq T$, and $w(X) := \sum_{\theta \in X} n_\theta$. Then selecting p in round r contributes $w(X_r(p))$ to welfare and increases $\mathrm{sat}_\theta(\cdot)$ by 1 exactly for $\theta \in X_r(p)$. Importantly, for a fixed round r, the welfare and all proportionality-relevant effects depend only on the *type pattern* $X_r(p)$, not on the identity of p. Thus in round r we may compress candidates into the family $\mathcal{X}_r := \{X_r(p) : p \in P\} \subseteq 2^T$, keeping one representative candidate per pattern (so $|\mathcal{X}_r| \leq 2^\kappa$).

We show that this structure leads to efficient algorithms even when preferences can vary across rounds. For JR, we obtain a dynamic program whose running time is exponential only in the number of types κ.

Theorem 5.4. *JR-*$\textsc{MaxUtil}$ *is solvable in time* $f(\kappa) \cdot \mathrm{poly}(n, m, \ell)$ *for some computable function* f*. In particular, it is FPT with respect to* κ*.*

For EJR we need to account for *how many times* each type is satisfied, since demands scale with group size and cohesion frequency. Let $\gamma_U := \left|\left\{r \in [\ell] : \bigcap_{\theta \in U} A_{\theta,r} \neq \emptyset\right\}\right|$ be the number of rounds in which all types in U "agree" (i.e., have a unanimously approved candidate in that round). Then, define for each $U \subseteq T$ the type set U induces an integer demand $d_U := \lfloor \gamma_U \cdot w(U)/n \rfloor$, and $D := \max_{U \subseteq T, U \neq \varnothing} d_U$. Note that $D \leq \ell$ (since $d_U \leq \gamma_U \leq \ell$). Under voter types, EJR can be expressed purely at the type level: for every U with $d_U > 0$, at least one type $\theta \in U$ must achieve satisfaction at least d_U. Then, we get fixed-parameter tractability with respect to κ and this demand parameter D.

Theorem 5.5. *EJR-*$\textsc{MaxUtil}$ *is solvable in time* $f(\kappa, D) \cdot \mathrm{poly}(n, m, \ell)$ *for some computable function* f*. In particular, it is FPT with respect to* $\kappa + D$*.*

The parameter D captures the "largest proportional claim" any cohesive union of types can generate; when no type set can demand many satisfied rounds (e.g., because unanimous agreement is rare), the state space remains small.

Since EJR implies PJR, the DP in Theorem 5.5 also outputs a PJR-feasible outcome. However, because PJR is strictly weaker than EJR in general, this outcome need not be welfare-optimal among all PJR-feasible outcomes.

The above DPs scale linearly with ℓ. In many temporal settings, however, the approval pattern repeats: many rounds share the same "profile" across types (e.g., weekly cycles). We can exploit this further by grouping rounds into a small number of distinct profiles. Two rounds $r, r' \in [\ell]$ have the same profile if $A_{\theta,r} = A_{\theta,r'}$ for every $\theta \in T$. Then, we get the following result.

Proposition 5.6. *On instances with* κ *voter types and* q *round profiles, PJR-*$\textsc{MaxUtil}$ *is solvable in time* $f(\kappa, q) \cdot \mathrm{poly}(n, \ell, m)$ *for a computable function* f*. In particular, it is FPT with respect to* (κ, q)*.*

The same profile-based variable scheme can also give us ILP formulations (and hence FPT algorithms parameterized by (κ, q)) for JR-$\textsc{MaxUtil}$ and EJR-$\textsc{MaxUtil}$ by adapting the constraints to the corresponding feasibility conditions.

Section 4 showed that static preferences alone do not restore tractability. In contrast, bounding the diversity of *temporal behavior*—via a small number of voter types and, optionally, a small number of round profiles—does lead to exact FPT algorithms for welfare maximization under temporal proportionality. This pinpoints temporal structure, rather than time-invariance per se, as a main driver of algorithmic tractability in temporal voting. Together, these results show that although Φ-$\textsc{MaxUtil}$ is intractable in full generality, it becomes efficiently solvable in a range of structured settings that naturally arise in temporal voting settings. This delineates a clear boundary between worst-case hardness and practically relevant tractability.

6 Conclusion

In this work, we investigated proportional representation in temporal voting through a quantitative *price of proportional representation* framework, measuring the worst-case loss in utilitarian welfare incurred by enforcing JR/PJR/EJR/EJR+. Our bounds make the welfare-proportionality tension in this setting precise: enforcing proportionality can cause a growing, yet sublinear, welfare loss, and the magnitude of this loss depends sharply on both the axiom and the time horizon. In particular, the cost of JR vanishes on long horizons: its price approaches 1 as the number of rounds grows large relative to the number of voters. By contrast, stronger notions such as PJR/EJR/EJR+ can still incur an $\Omega(\sqrt{n})$ loss even when ℓ is only a constant-factor multiple of n.

On the algorithmic side, we showed that maximizing utilitarian welfare subject to these proportionality constraints is NP-complete, and APX-hard, even under structural restrictions such as static preferences and bounded approval degrees. At the same time, we identified structured regimes in which optimal outcomes can be computed efficiently via fixed-parameter methods, including parameterizations by the number of candidates, the number of voter types, and the number of distinct round profiles.

Our work naturally gives rise to several promising avenues for further work. One direction is to refine the quantitative bounds, particularly for stronger proportionality notions such as PJR/EJR/EJR+. Such refinements would help clarify whether the observed separations are inherent or driven by worst-case constructions. A second direction is to extend the analysis beyond utilitarian welfare. Studying the price of proportionality with respect to alternative objectives (such as egalitarian or Nash welfare) may uncover qualitatively different trade-offs and lead to new insights into how proportionality interacts with social welfare under richer normative criteria.

References

[Aziz *et al.*, 2017] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48:461–485, 2017.

[Biega *et al.*, 2018] Asia J. Biega, Krishna P. Gummadi, and Gerhard Weikum. Equity of attention: Amortizing individual fairness in rankings. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR)*, pages 405–414, 2018.

[Brill and Peters, 2023] Markus Brill and Jannik Peters. Robust and verifiable proportionality axioms for multiwinner voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, page 301, 2023.

[Brill and Peters, 2024] Markus Brill and Jannik 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.

[Bulteau *et al.*, 2021] Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon. Justified representation for perpetual voting. *IEEE Access*, 9:96598–96612, 2021.

[Chandak *et al.*, 2024] Nikhil Chandak, Shashwat Goel, and Dominik Peters. Proportional aggregation of preferences for sequential decision making. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9573–9581, 2024.

[Chingoma *et al.*, 2025] Julian Chingoma, Umberto Grandi, and Arianna Novaro. Proportionality for constrained public decisions. *arXiv preprint arXiv:2409.02609*, 2025.

[Elkind *et al.*, 2024a] Edith Elkind, Piotr Faliszewski, Ayumi Igarashi, Pasin Manurangsi, Ulrike Schmidt-Kraepelin, and Warut Suksompong. The price of justified representation. *ACM Transactions on Economics and Computation*, 12(3):11:1–11:27, 2024.

[Elkind *et al.*, 2024b] Edith Elkind, Tzeh Yuan Neoh, and Nicholas Teh. Temporal elections: Welfare, strategyproofness, and proportionality. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*, pages 3292–3299, 2024.

[Elkind *et al.*, 2024c] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. Temporal fairness in multiwinner voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 22633–22640, 2024.

[Elkind *et al.*, 2025a] Edith Elkind, Tzeh Yuan Neoh, and Nicholas Teh. Not in my backyard! Temporal voting over public chores. In *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3814–3820, 2025.

[Elkind *et al.*, 2025b] Edith Elkind, Svetlana Obraztsova, Jannik Peters, and Nicholas Teh. Verifying proportionality in temporal voting. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, pages 13805–13813, 2025.

[Fairstein *et al.*, 2022] Roy Fairstein, Dan Vilenchik, Reshef Meir, and Kobi Gal. Welfare vs. representation in participatory budgeting. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 409–417, 2022.

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

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

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

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

[Masařík *et al.*, 2024] Tomáš Masařík, Grzegorz Pierczyński, and Piotr Skowron. A generalised theory of proportionality in collective decision making. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, pages 734–754, 2024.

[Peters *et al.*, 2021] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS)*, pages 12726–12737, 2021.

[Phillips *et al.*, 2026] Bradley Phillips, Edith Elkind, Nicholas Teh, and Tomasz Wąs. Strengthening proportionality in temporal voting. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2026.

[Pitoura *et al.*, 2022] Evaggelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. Fairness in rankings and recommendations: an overview. *The VLDB Journal*, 31:431–458, 2022.

[Sánchez-Fernández *et al.*, 2017] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 670–676, 2017.

[Singh and Joachims, 2018] Ashudeep Singh and Thorsten Joachims. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, pages 2219–2228, 2018.

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