

Fairness in the Multi-Secretary Problem

Georgios Papasotiropoulos¹, Zein Pishbin^{1,2}

¹University of Warsaw, Warsaw, Poland

²Université Paris Dauphine–PSL, Paris, France

gpapasotiropoulos@gmail.com, zeinab.s.pishbin@gmail.com

Abstract

This paper bridges two perspectives: it studies the multi-secretary problem through the fairness lens of social choice, and examines multi-winner elections from the viewpoint of online decision making. After identifying the limitations of the prominent proportionality notion of Extended Justified Representation (EJR) in the online domain, the work proposes a set of mechanisms that merge techniques from online algorithms with rules from social choice—such as the Method of Equal Shares and the Nash Rule—and supports them through both theoretical analysis and extensive experimental evaluation.

1 Introduction

The *secretary problem* appeared in the 1950s and quickly became a central topic in probability theory, optimization, and algorithms (Freeman 1983; Ferguson 1989; Samuels 1991; Dynkin 1963). It is a simple and elegant model that led to the development of an entire field around online and sequential decision-making processes. In its base form, a known number of candidates arrive one by one in random order. A single decision maker must either accept or reject a candidate immediately upon their arrival. Once rejected, a candidate cannot be recalled, and once accepted, they cannot later be rejected or exchanged for another one. At the moment of arrival of a candidate, the decision maker gathers enough information to evaluate them with respect to those seen so far, but remains unaware of the quality of the yet unseen ones. The goal is to pick the best candidate overall.

Often, real-world scenarios depart from this setting. It is typical to *select multiple candidates* instead of one, and for *decisions to be made by electorates*, rather than individuals, whose members may have varying and conflicting preferences. Then, *fairness among decision-makers* becomes crucial to ensure that diverse interests are adequately represented. In our work, we extend the base secretary model to reflect these aspects, still under the constraint that decisions must be made sequentially and irrevocably as candidates arrive. The following example illustrates our fairness considerations.

Example 1. *A hiring board in a Mathematics department is composed equally of theoreticians and applied mathematicians. Each group prefers candidates from their own field.*

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

	c_1	c_2	c_3	c_4	c_5	c_6
Theory	0	1	2	0	0	0
Applied	2	0	0	3	1	3

Table 1: Preferences of the two electorate groups in Example 1. Colors mark the candidates hired by the different rules proposed in the paper.

Due to the heavy workload involved in evaluating individual applications, the members of the board decide not to review all candidates at once but to use a sequential process. Since strong candidates may accept other offers if kept waiting, the board aims to decide on each candidate immediately upon evaluation. The task is to select $k = 2$ candidates from a pool of six applicants $\{c_1, c_2, \dots, c_6\}$ considered in arrival order (i.e., increasing index). The two groups in the hiring committee express their preferences through cardinal ballots; see Table 1. The solution that consists of the candidates highlighted in black, namely $\{c_4, c_6\}$, clearly achieves the highest total utility. However, it leaves half the electorate (the theoretical mathematicians) completely unsatisfied. A fairer solution is to hire one candidate from the applied field (c_1, c_4, c_5, c_6) and one from theory (c_2, c_3). The three rules proposed in our work, namely Greedy Budgeting, Online Method of Equal Shares and Online Nash Product, would take the online element into account and select the (blue) solution $\{c_1, c_2\}$, the (green) solution $\{c_3, c_4\}$, and the (red) solution $\{c_3, c_6\}$, respectively. While these sacrifice some overall utility, they ensure that no field is left unrepresented.

Our work is shaped by three core dimensions: (i) *online arrival of candidates*, (ii) *fairness goal*, (iii) *cardinal ballots of voters*. This is the first paper to bring all these elements together. *Fairness under cardinal ballots* in offline settings has been extensively explored in computational social choice (Benadè, Gal, and Fairstein 2023; Peters, Pierczyński, and Skowron 2021; Fain, Goel, and Munagala 2016; Bhaskar, Dani, and Ghosh 2018); when fairness is not a goal and total utility is considered in an online fashion, the problem reduces to the *multi-secretary problem* (Kleinberg 2005; Besbes, Kanoria, and Kumar 2022); and when ballots are binary (i.e., approval-based), the problem corresponds to *online multi-winner elections* (Do et al. 2022). Our study lies at the intersection of online algorithms and social choice.

- It explicitly adapts concepts and rules from voting theory (such as the Extended Justified Representation axiom, the Method of Equal Shares, the Method of Equal Shares with Bounded Overspending, the Nash Product rule) to address fairness considerations in classical online selection problems (i.e., in the (multi-)secretary problem), and
- It applies algorithmic techniques from the online algorithms literature (specifically, two distinct optimal stopping policies) to address foundational problems in social choice theory (namely, multi-winner/committee elections with cardinal ballots) when extended to dynamic settings.

Therefore, it aligns with two lines of work. *Dynamic social decision frameworks* have recently attracted considerable attention in computational social choice, with fairness considerations often being the central focus; see, indicatively, the works by Kehne, Schmidt-Kraepelin, and Sornat (2025); Dong and Peters (2025); Halpern et al. (2023); Hossain, Micha, and Shah (2021), or the survey by Elkind, Obraztsova, and Teh (2024). Yet, only Do et al. (2022) consider an online arrival of candidates requiring immediate decisions, akin to secretary problems. The study of *fairness in online algorithms* has also developed recently. Balkanski, Ma, and Maggiori (2024) study fairness for variants of the secretary problem in learning-augmented settings. Fairness in the single-candidate selection model is the focus of Arsenis and Kleinberg (2022). Another notable work is that of Correa et al. (2021), which is conceptually close to ours but also limited to selecting one candidate; it explicitly states the multi-selection setting as an open problem—which we address.

Contributions. Our work can be seen as extending the study of Do et al. (2022) by considering cardinal ballots. Do et al. showed that moving from the offline to the online setting entails only a modest loss in fairness: although the prominent notion of Extended Justified Representation (EJR)—which is achievable offline—cannot be fully satisfied online, it can still be well-approximated. This was somewhat surprising, as one might reasonably expect a more significant drop in quality of the outcome when the selection mechanism lacks knowledge of the entire input. We show that this counterintuitive conclusion no longer holds when voters submit more expressive ballots than simple approvals. Specifically, we analyze three natural relaxations of EJR and prove that none of them can be reasonably approximated in the setting we examine. In response to these negative observations, we explore alternative paths to achieving fairness guarantees online. We propose three distinct rules and we evaluate their performance both theoretically and experimentally. One of the rules we define, namely Online Method of Equal Shares, can be seen as fitting a general framework we introduce that converts offline multi-winner election rules into online ones. Based on this observation we additionally formalize an online variant of the recently introduced Method of Equal Shares with Bounded Overspending (Papasotiropoulos et al. 2025a), and show that it consistently stands out due to its strong performance. Finally, in the course of formalizing fairness, we also introduce the angle of probabilistic satisfiability of known proportionality notions: a conceptual contribution that may be of interest to social choice beyond the online domain.

2 Preliminaries

We adopt the terminology and notation of social choice theory as the fairness concepts and selection methods we explore are mostly inspired by it. An *election* is a tuple $E = (C, V, k, \{u_i\}_{i \in [n]})$:

- $C = \{1, 2, \dots, m\}$ and $V = \{1, 2, \dots, n\}$ are the sets of *candidates* and *voters* respectively,
- k is the upper bound on the *size of the committee* to be selected, and it is assumed to satisfy $2 \leq k < m$,
- for each voter $i \in V$, $u_i : C \mapsto \mathbb{R}_{\geq 0}$ is a function that corresponds to the *cardinal ballot* that voter i casts indicating their *additive utility*. If a set $W \subseteq C$ is elected, the overall utility of voter i is $u_i(W) = \sum_{j \in W} u_i(j)$.

If C is an ordered set, meaning the candidates are presented in a specific sequence that determines their appearance order, such that for each $i \in [n]$ the value of $u_i(j)$ is revealed at the moment when the candidate j appears, then E is an *online election*. Otherwise, it is an *offline election*. In case $u_i(j) \in \{0, 1\}, \forall (i, j) \in [n] \times [m]$, we say that we are in the approval setting and voters are expressing *approval preferences*. The definition of cardinal ballots above allows for arbitrary scores. A more realistic variant, which already generalizes approval preferences, involves *range ballots*. There, voters rate each candidate on a fixed scale—such as with integers from 0 to 5, 10 or 20 (as done in online reviews, grading, employee evaluations, or satisfaction surveys). Then, we say that the considered election E comes with an additional integer parameter B and voters’ ballots satisfy $u_i(j) \leq B$.

A subset of candidates $W \subseteq C$ is a *feasible solution* if $|W| \leq k$. We denote the set of all feasible solutions by \mathcal{W} . Our goal is to choose a feasible subset of candidates, which we call an *outcome*, based on voters’ ballots. A *multi-winner election rule* \mathcal{R} , in short, a rule, is a function that takes as input an election $E = (C, V, k, \{u_i\}_{i \in [n]})$ and returns an outcome $\mathcal{R}(E)$ called the *winning committee*. In particular, an *online multi-winner election rule* is an algorithm that processes candidates of an online election in the predefined sequence given by C , making an irrevocable decision at each step whether to include the currently examined candidate in the winning committee or not. If included, the candidate is said to be *hired/elected*; otherwise, they are *rejected*. When making such a decision on candidate $j \in [m]$, the rule can only use information about voters’ preferences for candidates that have already appeared, i.e., $u_i(c), \forall c \leq j$ and $i \in [n]$.

2.1 Fairness Considerations from Social Choice

A prominent concept in achieving fair representation is the axiom of Extended Justified Representation (EJR) (Aziz et al. 2017). It has received significant attention in the literature when it comes to formalizing fairness—especially in the context of multi-winner elections and participatory budgeting (PB): the settings most closely related to our framework—and its influence extends even beyond traditional voting (Kellerhals and Peters 2024; Papasotiropoulos et al. 2025b; Boehmer et al. 2024a). As such, EJR plays a central role in our work. It applies directly to both offline and online elections. Further proportionality axioms, which are implied by EJR, will also be presented and examined later, in Section 3.2.

Definition 1. A group of voters S is (α, T) -cohesive, where $\alpha : C \mapsto \mathbb{R}_{\geq 0}$ and $T \subseteq C$, if $\frac{|S|}{n} \geq \frac{|T|}{k}$ and $u_i(c) \geq \alpha(c)$ for all $i \in S$ and $c \in T$. A rule \mathcal{R} satisfies Extended Justified Representation (EJR) if for each election instance E , and each (α, T) -cohesive group of voters S there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) \geq \sum_{c \in T} \alpha(c)$.

Among the election rules examined in computational social choice from the perspective of proportionality, the *Method of Equal Shares (MES)*, which satisfies EJR, has received particular attention and has also been applied in real-life (offline) elections. For its formal definition we refer to the work of (Peters and Skowron 2020). Apart from EJR (and variants of it), we also explore an additional fairness concept in Section 4.2, namely Nash Welfare Maximization.

Complete proofs, additional experimental findings, and further material are provided in the full version of the paper.

3 Proportional Selection with Certainty

As mentioned, EJR has become a very well-studied fairness notion in social choice. Therefore, we first focus on EJR, along with its approximations and relaxations, to explore what is impossible and what can be guaranteed in this regard, in the online setting. In parallel, we aim to understand how these results differ from those by Do et al. (2022), where voters were limited to expressing only approval preferences.

3.1 Satisfying and Approximating EJR

It is already known that EJR cannot be satisfied by an online rule (Do et al. 2022). Consequently, in hopes of positive results, we turn to approximate notions of EJR instead, by exploring three distinct ways to get around this limitation. The first variant we introduce aims to guarantee that each group of voters receives at least a fraction of the utility they would deserve according to EJR and it is simply inspired by multiplicative guarantees of approximation algorithms. The second notion closely relates to an established relaxation of EJR, ensuring that every group of voters would be appropriately represented if additional candidates were allowed in the outcome (Peters, Pierczyński, and Skowron 2021). The third variant ensures fairness for sufficiently large groups of voters, specifically those that exceed the size required by EJR by a fixed factor—this is the one explored by Do et al. (2022).

Definition 2. A rule \mathcal{R} satisfies:

- β -approximation of Extended Justified Representation (β -EJR), for $\beta \geq 1$, if for each election instance E , and each (α, T) -cohesive group of voters S there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) \geq \frac{1}{\beta} \sum_{c \in T} \alpha(c)$.
- Extended Justified Representation up to γ candidates (EJR- γ), for $\gamma \leq k - 1$, if for each election instance E , and each (α, T) -cohesive group of voters S there exists a voter $i \in S$ and a set of candidates X where $|X| \leq \gamma$, such that $u_i(\mathcal{R}(E) \cup X) \geq \sum_{c \in T} \alpha(c)$.
- Extended Justified Representation for δ -cohesive groups, for $\delta \leq k$, if for each election instance E , and each (α, T) -cohesive group of voters S satisfying $\frac{|S|}{n} \geq \delta \frac{|T|}{k}$, there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) \geq \sum_{c \in T} \alpha(c)$.

Below we show that in the examined framework, not only EJR cannot be satisfied, but it also cannot be reasonably approximated under any of the three approximate notions defined. The constructed instances highlight that the impossibility comes exactly because of the online nature of the examined problem combined with cardinal ballots and they show a notable contrast to the approval setting, where EJR could be well approximated (Do et al. 2022).

Theorem 1. *There is no online voting rule that satisfies β -EJR, for any finite positive value of β , even for range ballots.*

Proof Sketch. Fix an arbitrarily β , some k and an arbitrary small positive value ε . Consider an instance on $2k$ candidates, namely $\{a_1, a_2, \dots, a_k, b_1, b_2, \dots, b_k\}$ and k voters, the cardinal ballots of which are described below and are upper bounded by $B = \beta$. For each $i \in [k]$, voter v_i casts $u_i(a_i) = 1 - \varepsilon$ and $u_i(a_j) = 0$, for each $j \neq i$. Moreover, $v_i(b_j) = \beta/k$, for every $j \in [k]$. At the point when a_1 arrives, any rule that does not have a knowledge of the preferences of v_1 for the candidates that are yet to come, should include a_1 in the committee to satisfy β -EJR with certainty. Similarly, all candidates in $\{a_2, \dots, a_k\}$ also need to be selected one after another. It also holds that any rule that does not elect any candidate from $\{b_1, \dots, b_k\}$ does not satisfy β -EJR. \square

The axiom of EJR up to γ candidates is also unsatisfiable. The proof follows the same rationale as that of Theorem 1.

Theorem 2. *There is no online voting rule that satisfies EJR- γ , for any integer $\gamma \leq k - 1$, even for range ballots*

Finally, akin to the previous relaxations of EJR, the third one introduced in Definition 2 is also impossible to satisfy.

Theorem 3. *There is no online voting rule that satisfies EJR for δ -cohesive groups, for any $\delta \leq k$, even for range ballots.*

Observe that the negative results of our work do not rely on complexity assumptions, demonstrating that the guarantees we aim for as per Definition 2 are impossible to achieve under any rule regardless of its running time. On the other hand, naturally, when presenting positive results we will focus exclusively on polynomial-time rules.

While one could explore combinations of the approximate notions we defined due to the strongly negative results of the preceding theorems it seems more prudent to move away from EJR. This is what we do in Section 3.2 and Section 4.

3.2 Satisfying Weakenings of EJR

Proportional Justified Representation (PJR) and Justified Representation (JR) are two well studied relaxations of EJR. For the definition of PJR, we refer to the work of Sánchez-Fernández et al. (2017), and we note that natural approximate notions of it can be defined in accordance to Definition 2. Importantly, all the negative results presented in Theorems 1 to 3 continue to hold for PJR as well—for the latter, this follows directly from its proof; for the others, it follows after dividing voters' utilities from all but the last k candidates by k in the corresponding proofs. Hence we turn our attention to JR. This is currently defined in the literature only for the approval setting (Aziz et al. 2017) so we first adapt it to the setting with cardinal ballots.

Definition 3. A rule \mathcal{R} satisfies Justified Representation (JR) if for each election instance E , and each (α, T) -cohesive group of voters S , with $|T| = 1$, there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) > 0$.

Towards designing a first method with provable fairness guarantees in the examined setting it is natural to revisit the Greedy Budgeting rule from the work of Do et al. (2022), which was proven to satisfy PJR under approval preferences. We begin by extending its definition to the cardinal setting.

The *Greedy Budgeting rule* works as follows. Consider an online election $E = (C, V, k, \{u_i\}_{i \in [n]})$ given as input. Each voter is initially given k/n dollars. When a candidate $j \in C$ arrives we look if the voters in the set $A_j := \{i \in V : u_i(j) > 0\}$ have at least 1 dollar in total. If not, we reject j , otherwise, we hire j and ask the voters in A_j to pay 1; each voter pays the same, or all of their remaining budget. We reduce their available budget accordingly and continue with the next candidate. If at some point the algorithm has hired $k - x$ candidates, for some $x > 0$, and there are exactly x candidates that remain to come, we hire all of the remaining candidates.

We show that, when applied to the cardinal setting, the rule loses some of its theoretical appeal (recall that PJR is unsatisfiable in the setting), but still manages to achieve certain proportionality guarantees, despite its simplicity.

Theorem 4. *Greedy Budgeting rule returns a feasible solution and satisfies JR.*

While JR is a relatively weak axiom, it nonetheless guarantees a basic form of fairness. Hence, Greedy Budgeting is arguably more fair than an approach focused on maximizing total utility, which could completely overlook minority preferences and overrepresent certain groups—recall Example 1. Moreover, Theorem 4 can also be strengthened to an analog of PJR, which must again refer to the number of approved candidates in the outcome and not accounting for utilities.

4 Proportional Selection with Risk

All the negative results we unfolded regarding the satisfiability of approximate and weakened notions of EJR, while differing in their technical details, are driven by the same underlying idea: towards satisfying a strong fairness guarantee in the online setting a rule must select a candidate c at the moment of their arrival, provided that c receives sufficient support. Under this principle, voters who support c may later encounter a candidate they value significantly more; the online nature of the problem prevents any knowledge of future arrivals. This tension is inherent in the examined setting and any rule aiming to always act safely towards satisfying EJR must adopt such a conservative strategy.

In this section we propose an alternative perspective; one that accepts risk as a tool for bypassing the impossibility results of Theorems 1 to 3. Rather than insisting on achieving proportionality in every instance, we propose accepting occasional shortfalls. More concretely, we will propose rules that do not commit to hiring candidates early solely because their supporters deserve some representation, because, by doing so, we risk exhausting the committee capacity too early, before more valuable candidates have had the chance to arrive. For the theoretical guarantees of the rules we propose, we adopt

the commonly used *random arrival model* from the online voting literature (Gupta and Singla 2021). Other models, such as the adversarial one, are natural options for follow-up work.

4.1 Online Method of Equal Shares and Beyond

Our approach draws inspiration from both the literature on secretary problems (Babaioff et al. 2007) and the most prominent voting rule for proportional representation (Peters and Skowron 2020), combining ideas from both.

The algorithm we propose, to be called *Online Method of Equal Shares*, operates in two phases: an exploration phase of length t , followed by a selection phase over the remaining $m - t$ candidates. The parameter t , referred to as the threshold time, will be fixed later based on the analysis. Let $C_S \subseteq C$ denote the set of the first t candidates of C , i.e., those observed during the exploration phase.¹ At the end of the first phase, we construct a reference committee $C_R \subseteq C_S$ by running MES on the election induced by the first t candidates. Importantly, candidates in C_R will not be hired, they are only used as samples for evaluation of the quality of subsequent candidates from $C \setminus C_S$. We next proceed to the selection phase. There we maintain a running committee C_R^r , initialized to C_R . For each arriving candidate c , we consider the set of candidates in $C_R^r \cup \{c\}$, which are $k + 1$ many. We compute the MES outcome for this candidate set. If MES returns the set C_R^r , then c is rejected. If MES returns a new winning committee, and there is a candidate $\hat{c} \in C_R^r$ that hasn't been selected by MES (one could see it as if c took the place of \hat{c}) we proceed as follows: We check whether $\hat{c} \in C_R$. If so, then c is hired; otherwise is rejected. In other words, c is hired not only if they are preferred over \hat{c} but also if \hat{c} was part of the initial trusted set C_R . Regardless of whether c is hired or not, as long as c was included in the outcome of MES, we update C_R^r to $C_R^r \cup \{c\} \setminus \{\hat{c}\}$ and proceed with the next candidate. To ensure the algorithm selects exactly k candidates, we include a final safeguard, the same one used in the Greedy Budgeting rule.

In what follows, we present a theoretical guarantee for Online MES. Specifically, we show that it satisfies EJR with a bounded from below probability as m increases. This probability depends exponentially on k , but also improves exponentially when aiming to satisfy the approximate notion of EJR up to p candidates (see Definition 2). Most importantly, this probability is independent of the total number of candidates m , unlike what a naive random selection would suggest. It only depends on the target committee size k , making the rule particularly well-suited for applications where the goal is to elect small committees— we validate this experimentally (see Section 5). Our theoretical guarantee holds for the single-approval case, i.e., for instances where for every voter i there is at most one candidate j such that $u_i(j) > 0$, yet, as we will empirically show later, the proposed method also succeeds in producing fair solutions for general instances.

Theorem 5. *Online Method of Equal Shares returns a feasible solution and satisfies EJR up to p candidates with probability $(1/e)^{k-p}$ as m tends to infinity, for $t = \lfloor m/e \rfloor$, for the single-approval case.*

¹We assume without loss of generality that $t \geq k$; otherwise, we can add dummy candidates with no support.

Proof Sketch. The proof relies on computing the probability that the outcome of Online MES matches that of Offline MES. The main component of the proof is showing that each candidate selected by Offline MES is hired by Online MES with probability at least $1/e$, which is attainable and maximized for $t = \lfloor m/e \rfloor$. This follows from the rule’s mechanics in the single-approval setting and the assumption of a uniformly random arrival order, using basic combinatorial arguments. As a result, with probability at least $(1/e)^k$, all k MES winners appear in the online outcome, which then satisfies EJR. \square

The proof of Theorem 5 directly extends to further guarantees satisfied by Offline MES, including EJR+ (Brill and Peters 2023) and other notions of representation (see the book by Lackner and Skowron (2023) for more). EJR+ is one of the strongest satisfiable proportionality axioms known in the literature, and strictly stronger than EJR. It is violated by a candidate not in the winning set if they are approved by $\ell \frac{n}{k}$ voters, each of whom approves less than ℓ members of the outcome.

Probabilistically guaranteeing proportionality axioms (via deterministic rules) is a conceptually appealing and novel approach. Online MES is the first rule providing such bounded guarantees; whether analyses improving the one in Theorem 5 are possible remains open. It is also the first classic offline voting rule adapted to the online setting. Yet, the main appeal of the method is explained in the next paragraph.

A Framework for Bringing Voting Rules Online and the Case of MES with Bounded Overspending The principle behind Online Method of Equal Shares has a broader applicability. We highlight that its mechanics are not specific to MES, and can be adapted to work with any voting rule. In that sense, our method can be seen as an instance of a *general framework that converts any rule for offline elections into one that works in the online setting*. MES is the most widely known and used voting method for achieving proportional outcomes. However, it is not without drawbacks. As discussed in the very recent work by Papatotiropoulos et al. (2025a), there are natural cases where MES behaves in counterintuitive ways. One such case arises precisely due to the use of cardinal ballots—which is exactly the focus of our work. A simple but insightful example, known as the “weakness of tail utilities,” captures this issue well. To address limitations of the Method of Equal Shares, Papatotiropoulos et al. (2025a) introduced the Method of Equal Shares with Bounded Overspending (BOS). BOS is considered better suited for practical applications than MES. The approach we proposed in Section 4.1 is fully compatible with BOS as well and it can be applied using BOS as the underlying rule instead of MES. Therefore, in our experiments, (Section 5) we also consider *Online BOS* and we show that it exhibits exceptionally good behavior.

4.2 Online Nash Product Rule

Given an offline election, a natural first objective is to maximize the total voter utility, hence elect the candidates in $\arg \max_{W \in \mathcal{W}} \sum_{i \in V} u_i(W)$. As observed by Do et al. (2022), the approval-based variant of this problem essentially matches the multi-secretary problem studied by Kleinberg (2005). The online algorithm proposed there applies directly to the setting with cardinal ballots as well, yielding a constant

competitive ratio against the offline optimum. However, this objective entirely overlooks fairness among voters. It values all gains equally, without considering who benefits. To address this, a more fairness-aware objective is to *maximize the product of utilities*. This approach promotes balanced distribution of voter satisfaction and, here, fairness is understood as ensuring equitable utility among voters, rather than proportional representation as captured by EJR and related axioms discussed in Section 3 and Section 4.1. This objective is known as the *Nash Product rule*. Formally, it is defined as $\arg \max_{W \in \mathcal{W}} \prod_{i \in V} u_i(W)$ (Nash 1950), or, equivalently in terms of maximizers, $\arg \max_{W \in \mathcal{W}} \sum_{i \in V} \log(1 + u_i(W))$, where the addition of 1 ensures that all terms are well-defined, even when some voters have a utility of zero (Fluschnik et al. 2019). We call *Nash Welfare* the quantity

$$\Pi(W) := \sum_{i \in V} \log(1 + u_i(W)).$$

This and related formulations of the Nash Product rule are well-studied in computational social choice (Caragiannis et al. 2019; Freeman, Zahedi, and Conitzer 2017; Conitzer, Freeman, and Shah 2017; Fain, Munagala, and Shah 2018; Ebadian et al. 2024).

We introduce this rule to a new application domain: the fair selection of candidates in an online secretary hiring framework. From the original axiomatic characterization (Nash 1950), we can derive that it is the only one satisfying several important axioms—each natural and desirable in our setting, even when fairness is not the main objective. According to this characterization, a more equal distribution of satisfaction among voters is better, increasing a voter’s satisfaction without decreasing that of others is also an improvement, and the optimal outcome remains unchanged under scaling of valuations; we refer to the work of Kilgour and Eden (2010) for details and to the papers by Ramezani and Endriss (2009) and Lu et al. (2024) for further properties.

The Nash Product rule can be seen as an analog of PAV (Fluschnik et al. 2019), a rule with strong guarantees of proportionality in the approval setting (Lackner and Skowron 2023) which does not extend directly to settings with cardinal utilities. On the negative side, likewise PAV, computing a committee that maximizes the Nash Welfare is NP-hard in the offline setting (Ramezani and Endriss 2009). The committee returned by our method, which runs in polynomial time, achieves a Nash Welfare within a multiplicative factor of roughly $1/11$ compared to the offline optimal.

The algorithm we propose, to be called *Online Nash Rule* works as follows. We divide the input stream into k consecutive and equally-sized² segments, with the goal of selecting one candidate per segment. Intuitively, in each segment we aim to select the best available candidate based on the contributions to the overall Nash Welfare. Let $T_0 = \emptyset$ and, for $i \in \{1, 2, \dots, k\}$ let T_i denote the set of candidates selected by the time a decision has been made for the last candidate of the i -th segment. For each segment $i \in \{1, \dots, k\}$, we define the marginal gain function $g_i(c) = \Pi(T_{i-1} \cup \{c\})$,

²We assume wlog that m is a multiple of k ; otherwise, we can add dummy candidates with no support arriving in random order.

which measures the value of adding the candidate c to the set of already selected candidates with respect to the function Π that denotes the Nash Welfare. Our objective in segment i is to select the candidate maximizing g_i . To do so, we apply the classical algorithm for the secretary problem within each segment (Dynkin 1963). Specifically, for the i -th segment, we observe the first $\lfloor 1/e \rfloor$ -fraction of candidates of the segment, say R_i , without making any selection. Let $\hat{c}_i := \arg \max_{c \in R_i} g_i(c)$. Then, we include in T_i the first candidate that comes after this observation phase and has $g_i(c) \geq g_i(\hat{c}_i)$. If no such candidate arrives, we select the last candidate of the segment. The final selected committee is T_k .

Theorem 6. *Online Nash Rule returns a feasible solution and the expected value of its outcome is at least $\frac{1-1/e}{7}$ of the offline maximum Nash Welfare.*

5 Empirical Evaluation

Finally, we experimentally study Greedy Budgeting, Online MES, Online BOS, and Online Nash rules on both real-world (Experiments 1 and 2) and synthetic datasets (Experiments 3 and 4), designing a set of simulations with distinct goals and evaluation metrics. The key takeaways are that (i) all the proposed rules perform well in terms of fairness across a range of setups and metrics, (ii) when comparing them directly, Online BOS stands out for its consistency: it is the only rule among those we examined that is consistently either the best or close to it, (iii) Online BOS outperforms Online MES across all experiments and setups.

Experiment 1: Pabulib Instances. We computed the outcomes of the proposed rules across a set of 700 participatory budgeting instances (Faliszewski et al. 2023). To keep running times reasonable, we focused on instances where at least one of the two main complexity-driving parameters—the number of voters n or the number of candidates m —is not excessively large. To assess proportionality, we use EJR+ as its satisfaction (and degree of violation) can be verified efficiently. Table 2 reports the (average) share of voters who are not adequately represented (“viol. voters” in the corresponding table) and the number of candidates witnessing the violation (“shortfall”), for three values of k .

Overall, the experiment showcases that all of our rules perform very well with respect to EJR+ violations: the violating voters percentage is only marginally different across rules and there are around 1% to 2% of voters who are underrepresented on average; the average shortfall indicates that all rules either satisfy EJR+ up to one candidate (see Definition 2) or come very close to doing so. As k increases, the percentage of violating voters decreases, while the shortfall increases.

Experiment 2: Sushi and MovieLens Datasets. A limitation of the datasets used in Experiment 1 is that they come from processes where candidates are project proposals with associated costs—which we had to ignore since our rules are not designed for cost-based settings. Additionally, since Experiment 1 focused on the EJR+ metric which is only defined for approval preferences, the datasets used involve only binary preferences. Our second experiment utilizes two datasets which avoid these limitations and are standard benchmarks

viol. voters shortfall	$k = m/20$	$k = m/10$	$k = m/4$
Greedy Budg.	2.11% 0.296	1.87% 0.476	1.45% 1.032
Online MES	2.08% 0.291	1.87% 0.470	1.46% 1.049
Online BOS	2.07% 0.290	1.85% 0.465	1.44% 1.016
Nash Rule	2.05% 0.303	1.93% 0.498	1.51% 1.155

Table 2: Share of voters not adequately represented and number of candidates witnessing a violation, under EJR+, shown for each rule and committee size (Experiment 1).

Best Top-2	Excl. Ratio (S)	Excl. Ratio (M)	Gini (S)	Gini (M)
Greedy Budg.	56.1% 74.4%	11.8% 54.7%	48.9% 73.9%	14.7% 58.0%
Online MES	8.9% 33.3%	7.8% 48.2%	8.9% 30.6%	6.7% 44.7%
Online BOS	34.4% 90.6%	80.4% 96.7%	41.7% 94.4%	78.7% 96.9%
Nash Rule	0.6% 1.7%	0.0% 0.4%	0.6% 1.1%	0.0% 0.4%

Table 3: Percentage of instances in which each rule achieves the best performance among all evaluated rules, and percentage of instances in which it ranks within the top two, reported for the exclusion ratio and Gini coefficient metrics, for the two examined datasets (Experiment 2).

in computational social choice and beyond. The goal is to analyze the relative performance of the rules under varying committee sizes. We evaluate performance using four metrics: average voter satisfaction, percentage of voters with satisfaction 0 (exclusion ratio), average satisfaction of the least satisfied quartile of voters, and the Gini coefficient capturing inequality in satisfaction; we report two of them below and we note that the conclusions for the rest are similar. Table 3 shows the percentage of instances where each rule achieves the best performance, as well as the percentage where it ranks either first or second. For the Sushi dataset (Kamishima 2003; Kamishima, Kazawa, and Akaho 2010)—denoted by “S” in the table—while Greedy Budgeting is most often the top-performing rule, there are multiple instances where it fails to rank among the top two; unlike Online BOS, which appears to be a more robust overall choice. The MovieLens dataset (Harper and Konstan 2015)—denoted by “M”—is even more favorable to Online BOS. Online MES also sees a significant increase in the percentage of instances where it ranks among the top two. These shifts can be explained by the increase in the number of available candidates (from 100 in the Sushi dataset to 1700 in MovieLens), which, with k fixed, inherently favors Online MES and BOS by design.

Experiment 3: Sampled Datasets. The two most commonly used models for generating synthetic data in computational social choice are Impartial Culture (IC) and Mallows model (Boehmer et al. 2024b). We evaluate our rules using these and we also explore a recently proposed culture for synthetic preference generation that is well suited for experiments involving varying numbers of candidates: the Normalized Mallows model (Boehmer, Faliszewski, and Kraiczky 2023). Even in the well-studied offline setting of multi-winner elections, there is no single widely accepted metric for measuring proportionality. However, the Method of Equal Shares is highly regarded for its proportionality guarantees. We assess how much our rules deviate from the Offline MES outcomes based on the four statistics used in Experiment 2. Naturally, given their

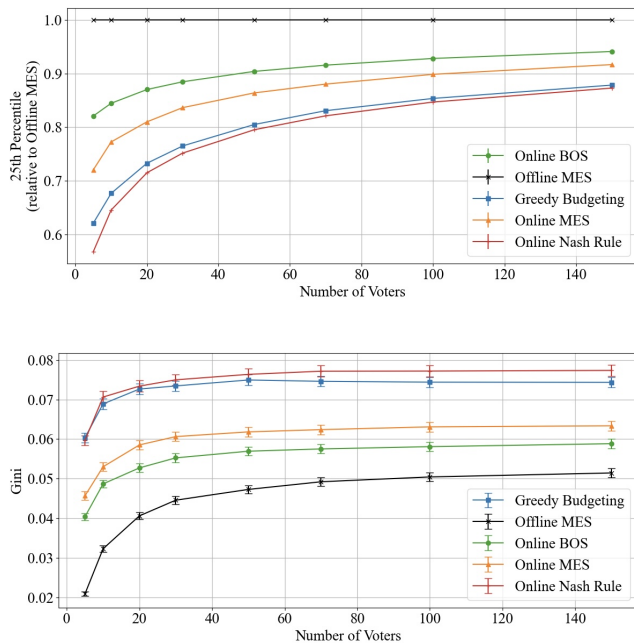


Figure 1: (Top) 25th percentile of voter utilities under the IC model relative to the analogous value for Offline MES rule; (Bottom) Gini coefficient of voter utilities under the Normalized Mallows model (Experiment 3).

online nature, our rules fall short compared to MES—but the key question is by how much, and which rule performs best in minimizing this gap. Figure 1 presents two indicative results. For the IC model, the average utility of the least satisfied quartile under Online BOS is clearly better than the rest of the rules, consistently reaching over $4/5$ of what Offline MES achieves. Under the Normalized Mallows model, measuring utility inequality again favors Online BOS, followed by Online MES, with Greedy Budgeting and Nash rules trailing closely together. Averages of Gini coefficients across all rules remain notably low, indicating very limited inequality.

Experiment 4: Polarized Instances. In this final experiment, we generate profiles where proportionality is intuitively easy to capture. Specifically, there are two voter groups: group A approves the first half of candidates, and group B approves some from the second half. For a given value x , group A constitutes an x -percentage of the voters. Proportionality can be interpreted as requiring at least an x -percentage of the winning committee to come from the first half of the candidates. We measure the frequency with which group A receives less representation than they deserve and quantify the extent of this underrepresentation. Since JR guarantees no shortfall here, Greedy Budgeting produces outcomes with no violations. Online MES fails in 19.67% of the cases, Online BOS in 27%, and the Nash Rule in just 2.67%. However, in terms of average shortfall, the Nash Rule shows a deficit of around 4 candidates, while Online BOS stays below 1; Online MES falls in between. BOS also performs twice as well in terms of the maxima of deficits observed per instance.

6 Concluding Discussion

It is worth emphasizing that the proposed rules follow fundamentally different strategies, each with its own selection pattern. The Greedy Budgeting rule prioritizes candidates who appear early. This principle does not apply to Online MES (and similarly, BOS), which reserve approximately the first 37% of candidates as an observation phase and only begin selecting afterwards. The Nash Rule, in contrast, ensures a more distributed selection by dividing candidates into segments and selecting one from each. These behaviors are clearly illustrated in Example 1, where the blue committee reflects the outcome of the Greedy rule, the green committee corresponds to Online MES (and BOS), and the red one to the Nash Rule. These differences showcase the importance of the arrival order of candidates. If some additional information were available—say, from past elections or recruitment cycles—it could guide the choice of which rule to apply. This points to a promising direction for future work: designing fair online multi-winner election rules that make use of predictions such as forecasts of future preferences. In fact, the negative results from Section 3 might have been avoided if such predictive information had been available. Another way to potentially avoid such strongly negative findings is to significantly restrict the allowed scores in range voting.

Among the proposed rules, Greedy Budgeting has a remarkable advantage: it is the only one that does not require knowing the total number of candidates in advance (except during its final safeguard step ensuring exactly k selections). The rest rely on the value of m to determine the transition point between the observation and selection phases. While assuming knowledge of the number of candidates is standard in the literature of online algorithms, designing fair rules that operate without this assumption remains an important direction for future work. The simplicity of Greedy Budgeting is also a practical benefit and this relates to a further drawback of Online MES and BOS: their running time. Even though they run in polynomial time, they need to call the MES or BOS subroutine $O(m)$ times, which can be impractical in some scenarios. Indicatively, in the real-world datasets of Experiment 1, Greedy Budgeting rule took on average 0.01s per instance, making it significantly faster than Online MES, which needed around 1.1s. Online MES was itself about 6 times faster than Online BOS. Nash Product rule, with an average of 0.3s per instance, was also sufficiently fast.

As mentioned, our work can be seen as a follow-up to the work by Do et al. (2022). There are two clear directions to generalize their work: one is what we pursued, i.e., generalizing ballots from approval to cardinal (a setting perfectly suited for scenarios such as secretary hiring that motivated our study), and the other is moving beyond multi-winner elections. The natural, yet non-trivial, next step is to explore proportionality in settings like PB, where candidates have costs and must be selected within a budget, still in an online manner. This can, in principle, be applied in online blockchain-based public funding platforms, such as Bitcoin, the Web3 Foundation, or Project Catalyst. Participatory Budgeting brings new challenges because the number of winners is not fixed, which calls for significantly different insights or techniques.

Acknowledgements

We thank Piotr Faliszewski, Piotr Skowron, and Tomasz Wąs for their assistance and their feedback on earlier versions of this work. The authors were supported by the European Union (ERC, PRO-DEMOCRATIC, 101076570). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.



References

- Arsenis, M.; and Kleinberg, R. 2022. Individual Fairness in Prophet Inequalities. In *Proceedings of the ACM Conference on Economics and Computation*.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare*, 461–485.
- Babaioff, M.; Immorlica, N.; Kempe, D.; and Kleinberg, R. 2007. A Knapsack Secretary Problem with Applications. In *Proceedings of the International Workshop on Approximation Algorithms for Combinatorial Optimization*, 16–28.
- Balkanski, E.; Ma, W.; and Maggiori, A. 2024. Fair Secretaries with Unfair Predictions. In *Proceedings of the Conference on Neural Information Processing Systems*.
- Benadè, G.; Gal, K.; and Fairstein, R. 2023. Participatory Budgeting Design for the Real World. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Besbes, O.; Kanoria, Y.; and Kumar, A. 2022. The Multi-Secretary Problem With Many Types. In *Proceedings of the ACM Conference on Economics and Computation*.
- Bhaskar, U.; Dani, V.; and Ghosh, A. 2018. Truthful and Near-Optimal Mechanisms for Welfare Maximization in Multi-Winner Elections. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Boehmer, N.; Brill, M.; Cevallos, A.; Gehrlein, J.; Sánchez-Fernández, L.; and Schmidt-Kraepelin, U. 2024a. Approval-Based Committee Voting in Practice: A Case Study of (Over-)Representation in the Polkadot Blockchain. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Boehmer, N.; Faliszewski, P.; Janeczko, Ł.; Kaczmarczyk, A.; Lisowski, G.; Pierczyński, G.; Rey, S.; Stolicki, D.; Szufa, S.; and Wąs, T. 2024b. Guide to Numerical Experiments on Elections in Computational Social Choice. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Boehmer, N.; Faliszewski, P.; and Kraicz, S. 2023. Properties of the Mallows Model Depending on the Number of Alternatives: A Warning for an Experimentalist. In *Proceedings of the International Conference on Machine Learning*.
- Brill, M.; and Peters, J. 2023. Robust and Verifiable Proportionality Axioms for Multiwinner Voting. In *Proceedings of the ACM Conference on Economics and Computation*.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 1–32.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair Public Decision Making. In *Proceedings of the ACM Conference on Economics and Computation*, 629–646.
- Correa, J.; Cristi, A.; Duetting, P.; and Norouzi-Fard, A. 2021. Fairness and Bias in Online Selection. In *Proceedings of the International Conference on Machine Learning*, 2112–2121.
- Do, V.; Hervouin, M.; Lang, J.; and Skowron, P. 2022. Online Approval Committee Elections. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Dong, C.; and Peters, J. 2025. Proportional Multiwinner Voting with Dynamic Candidate Sets. In *Proceedings of the International Conference on Machine Learning*.
- Dynkin, E. B. 1963. The Optimum Choice of the Instant for Stopping a Markov Process. *Soviet Mathematics*, 627–629.
- Ebadian, S.; Kahng, A.; Peters, D.; and Shah, N. 2024. Optimized Distortion and Proportional Fairness in Voting. *ACM Transactions on Economics and Computation*, 1–39.
- Elkind, E.; Obratzsova, S.; and Teh, N. 2024. Temporal Fairness in Multiwinner Voting. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Fain, B.; Goel, A.; and Munagala, K. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the International Conference on Web and Internet Economics*.
- Fain, B.; Munagala, K.; and Shah, N. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the ACM Conference on Economics and Computation*.
- Faliszewski, P.; Flis, J.; Peters, D.; Pierczyński, G.; Skowron, P.; Stolicki, D.; Szufa, S.; and Talmon, N. 2023. Participatory Budgeting: Data, Tools, and Analysis. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Ferguson, T. S. 1989. Who Solved the Secretary Problem? *Statistical Science*, 282–289.
- Fluschnik, T.; Skowron, P.; Triphaus, M.; and Wilker, K. 2019. Fair Knapsack. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Freeman, P. R. 1983. The Secretary Problem and Its Extensions: A Review. *International Statistical Review/Revue Internationale de Statistique*, 189–206.
- Freeman, R.; Zahedi, S. M.; and Conitzer, V. 2017. Fair Social Choice in Dynamic Settings. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Gupta, A.; and Singla, S. 2021. Random-Order Models. In *Beyond the Worst-Case Analysis of Algorithms*, 234–258. Cambridge University Press.
- Halpern, D.; Kehne, G.; Procaccia, A. D.; Tucker-Foltz, J.; and Wüthrich, M. 2023. Representation with Incomplete Votes. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Harper, F. M.; and Konstan, J. A. 2015. The MovieLens Datasets: History and Context. In *ACM Transactions on Interactive Intelligent Systems*, 1–19.

- Hossain, S.; Micha, E.; and Shah, N. 2021. Fair algorithms for multi-agent multi-armed bandits. *Proceedings of the Conference on Neural Information Processing Systems*, 34: 24005–24017.
- Kamishima, T. 2003. Nantonac Collaborative Filtering: Recommendation Based on Order Responses. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining*.
- Kamishima, T.; Kazawa, H.; and Akaho, S. 2010. A Survey and Empirical Comparison of Object Ranking Methods. In *Preference Learning*.
- Kehne, G.; Schmidt-Kraepelin, U.; and Sornat, K. 2025. Robust Committee Voting, or The Other Side of Representation. In *Proceedings of the ACM Conference on Economics and Computation*.
- Kellerhals, L.; and Peters, J. 2024. Proportional Fairness in Clustering: A Social Choice Perspective. In *Proceedings of the Conference on Neural Information Processing Systems*.
- Kilgour, D. M.; and Eden, C. 2010. *Handbook of Group Decision and Negotiation*. Springer.
- Kleinberg, R. D. 2005. A Multiple-Choice Secretary Algorithm with Applications to Online Auctions. In *Proceedings of the Symposium on Discrete Algorithms*, 630–631.
- Lackner, M.; and Skowron, P. 2023. *Multi-Winner Voting with Approval Preferences*. Springer Nature. Chapter 4: “Proportionality”.
- Lu, X.; Peters, J.; Aziz, H.; Bei, X.; and Suksompong, W. 2024. Approval-based voting with mixed goods. *Social Choice and Welfare*, 62(4): 643–677.
- Nash, J. F. 1950. The Bargaining Problem. *Econometrica*, 155–162.
- Papasotiropoulos, G.; Pishbin, S. Z.; Skibski, O.; Skowron, P.; and Waş, T. 2025a. Method of Equal Shares with Bounded Overspending. In *Proceedings of the ACM Conference on Economics and Computation*.
- Papasotiropoulos, G.; Skibski, O.; Skowron, P.; and Waş, T. 2025b. Proportional Selection in Networks. *arXiv preprint*. 2502.03545.
- Peters, D.; Pierczyński, G.; and Skowron, P. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the Conference on Neural Information Processing Systems*.
- Peters, D.; and Skowron, P. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the ACM Conference on Economics and Computation*.
- Ramezani, S.; and Endriss, U. 2009. Nash Social Welfare in Multiagent Resource Allocation. In *Proceedings of the International Workshop on Agent-Mediated Electronic Commerce*.
- Samuels, S. M. 1991. Secretary Problems. *Handbook of Sequential Analysis*, 381–405.
- Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J.; Val, P. B.; and Skowron, P. 2017. Proportional Justified Representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.