

Fairness in the Multi-Secretary Problem

GEORGIOS PAPASOTIROPOULOS, *University of Warsaw*

ZEIN PISHBIN, *University of Warsaw*

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.

CONTENTS

1	Introduction	2
2	Preliminaries	3
2.1	Fairness Considerations from Social Choice	4
3	Proportional Selection with Certainty	4
3.1	Satisfying and Approximating Extended Justified Representation	5
3.2	Satisfying Weakenings of Extended Justified Representation	7
4	Proportional Selection with Risk	9
4.1	Online Method of Equal Shares and Beyond	9
4.2	Online Nash Product Rule	11
5	Empirical Evaluation	13
5.1	Pabulib Instances	13
5.2	Sushi and MovieLens Datasets	15
5.3	Sampled Datasets	15
5.4	Polarized Instances	17
6	Concluding Discussion	19

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. Consider a committee or hiring board in a Mathematics department, composed equally of theoreticians and applied mathematicians, with each group preferring candidates from their own field. Due to the heavy workload involved in evaluating individual applications, the committee decides not to review all candidates at once but to use a sequential process. Since strong candidates may accept other offers if kept waiting, the committee aims to decide on each candidate immediately upon evaluation. The committee is tasked with selecting $k = 2$ candidates from a pool of six applicants $c_1, c_2, \dots, c_6$ considered in the order of their arrival (i.e., increasing index). The preferences of the two groups within the committee are expressed through cardinal ballots, as shown below; the colors indicate different candidate selections, which will be discussed shortly.

	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

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 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 et al., 2023; Peters et al., 2021; Fain et al., 2016; Bhaskar et al., 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 et al., 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 and is the first:

- *to explicitly adapt 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
- *to apply 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 et al. \[2025\]](#); [Dong and Peters \[2025\]](#); [Halpern et al. \[2023\]](#); [Hossain et al. \[2021\]](#), or the survey by [Elkind et al. \[2024\]](#). 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 et al. \[2024\]](#) 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 the most attention in the social choice literature when it comes to formalizing fairness—especially in the context of multi-winner elections and participatory budgeting (PB): the election settings most closely related to our framework—and its influence extends even beyond traditional voting settings [Kellerhals and Peters, 2024; Papsotiropoulos 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.

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 the reader to the work of Peters et al. [2021]. Apart from EJR (and variants of it), we also explore an additional fairness concept in Section 4.2, namely Nash Welfare Maximization.

3 Proportional Selection with Certainty

As mentioned, EJR has become the most 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 Extended Justified Representation

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 et al., 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 $u_i(\mathcal{R}(E)) \geq 1/\beta \sum_{c \in T} \alpha(c)$.
- *Extended Justified Representation up to γ candidates ($EJR\text{-}\gamma$)*, 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. Fix an arbitrary β , some k and an arbitrarily small positive value ε . Consider an instance on $2k$ candidates, namely $\{a_1, a_2, \dots, a_k, b_1, \dots, b_k\}$ and k voters, the cardinal ballots of which appear in the table below and are upper bounded by $B = \beta$.

	a_1	a_2	$\dots$	a_k	b_1	$\dots$	b_k
v_1	$1 - \varepsilon$				β/k	$\dots$	β/k
v_2		$1 - \varepsilon$			β/k	$\dots$	β/k
$\dots$					$\dots$		
v_k				$1 - \varepsilon$	β/k	$\dots$	β/k

Say that the candidates are arriving in the order presented, i.e. first the a -candidates and then the b -candidates, all in increasing index order. Candidate a_1 appears first. Say that $T = \{a_1\}$ and $S = \{v_1\}$. For these groups it holds that $|S| \geq |T|^{n/k}$, and $\alpha(c) = 1 - \varepsilon, \forall c \in T$. According to β -EJR, at least one voter in S should receive a satisfaction of at least $\frac{1-\varepsilon}{\beta}$ from the outcome. The decision on whether to include a_1 in the winning committee or not should be done at this point and any online multi-winner election rule that rejects a_1 can't revoke this decision in case the candidates that will come later do not satisfy v_1 to an extent of $1 - \varepsilon$. Therefore, 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 at this point. By the same argument, all candidates in the set $\{a_1, a_2, \dots, a_k\}$ will

be selected one after the other. Hence, there is no place to include any candidate from $\{b_1, \dots, b_k\}$ in the solution. It holds that any rule that does not elect any candidate from $\{b_1, \dots, b_k\}$ does not satisfy β -EJR. This is because if we set $T = \{b_1, \dots, b_k\}$ and $S = \{v_1, v_2, \dots, v_k\}$, then S is $((\frac{\beta}{k})_{i \in [k]}, T)$ -cohesive. Hence, in order to satisfy β -EJR, there should be a voter i in S such that $u_i(\mathcal{R}(E)) \geq \frac{1}{\beta} \sum_{c \in T} \frac{\beta}{k} = 1$. But, under a rule that does not elect any candidate from $\{b_1, \dots, b_k\}$, each voter in S gets a satisfaction of at most $1 - \varepsilon$. $\square$

The axiom of EJR up to γ candidates is also unsatisfiable. The proof follows the same rationale as that of [Theorem 1](#), though it employs a somewhat more involved construction.

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

Proof. Fix an arbitrary γ , some k and an arbitrarily small positive value ε . Consider an instance on $k^2 + k$ candidates, namely $\{a_1^1, a_1^2, \dots, a_1^k, a_2^1, a_2^2, \dots, a_2^k, \dots, a_k^1, a_k^2, \dots, a_k^k, b_1, \dots, b_k\}$ and k voters, the cardinal ballots of which appear in the table below and are upper bounded by $B = 2^{k+1}\varepsilon$.

	a_1^1	a_1^2	...	a_1^k	a_2^1	a_2^2	...	a_2^k	...	a_k^1	a_k^2	...	a_k^k	b_1	...	b_k
v_1	ε	2ε	...	$2^k\varepsilon$										$2^{k+1}\varepsilon$		$2^{k+1}\varepsilon$
v_2					ε	2ε	...	$2^k\varepsilon$						$2^{k+1}\varepsilon$		$2^{k+1}\varepsilon$
...																
v_k										ε	2ε	...	$2^k\varepsilon$	$2^{k+1}\varepsilon$		$2^{k+1}\varepsilon$

Say that the candidates are arriving in the order presented, i.e. first all the a_1 -candidates, then all the a_2 -candidates, etc, then the all a_k -candidates and then the b -candidates, all in increasing order of superscript. First, we note that at least one candidate from $\{a_1^1, \dots, a_1^k\}$ will be included in the solution by the time that a_1^k has been examined. This is because if at that point we consider $T = \{a_1^k\}$ and $S = \{v_1\}$, it holds that $|S| \geq |T|^{n/k}$, and $\alpha(c) = 2^k\varepsilon, \forall c \in T$. According to EJR- γ , at least one voter in S should receive a satisfaction of at least $2^k\varepsilon$ from $\mathcal{R}(E) \cup X$, for some X where $|X| \leq k - 1$. This is because any online multi-winner election rule that rejects a_1^k can't revoke this decision in case the candidates that will come are not approved by v_1 , and then, if $X = \{a_1^1, a_1^2, \dots, a_1^{k-1}\}$ it holds $u_i(\mathcal{R}(E) \cup X) = \varepsilon \sum_{i=0}^{k-1} 2^i = \varepsilon(2^k - 1) = 2^k\varepsilon - \varepsilon$. Therefore, a_1^k should be selected at the time it arrives. By similar arguments, after examining all a -candidates, it should hold that at least one candidate approved by each voter should be selected. Hence, there is no place to include any candidate from $\{b_1, \dots, b_k\}$ in the solution.

It holds that any rule that does not elect any candidate from $\{b_1, \dots, b_k\}$ does not satisfy EJR- γ . This is because if we set $T = \{b_1, \dots, b_k\}$ and $S = \{v_1, v_2, \dots, v_k\}$, then S is $((2^{k+1}\varepsilon)_{i \in [k]}, T)$ -cohesive. Hence, in order to satisfy EJR- γ , there should be a voter i in S such that $u_i(\mathcal{R}(E) \cup X) \geq \sum_{c \in T} \alpha(c) = \varepsilon k 2^{k+1}$. However, for each voter i in S , using the set X that consists of $(k - 1)$ b -candidates, it holds that $u_i(\mathcal{R}(E) \cup X) \leq 2^k\varepsilon + \sum_{i \in [k-1]} \varepsilon 2^{k+1} = 2^k\varepsilon + \varepsilon(k - 1)2^{k+1} = 2^k\varepsilon(1 + 2(k - 1)) = 2^k\varepsilon(2k - 1) < 2^k\varepsilon 2k$. $\square$

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.*

Proof. Fix an arbitrary δ , some k and an arbitrarily small positive value ε . Consider an instance on $2k$ candidates, namely $\{a_1, a_2, \dots, a_k, b_1, \dots, b_k\}$ and 1 voter, the cardinal ballot of which appears in the table below and are upper bounded by $B = 1 + k\varepsilon$.

	a_1	a_2	$\dots$	a_k	b_1	$\dots$	b_k
v_1	$1 + \varepsilon$	$1 + 2\varepsilon$		$1 + k\varepsilon$	$1 + \varepsilon(\frac{k+1}{2} + \frac{1}{k})$	$\dots$	$1 + \varepsilon(\frac{k+1}{2} + \frac{1}{k})$

Say that the candidates are arriving in the order presented, i.e. first the a -candidates and then the b -candidates, all in increasing index order. Candidate a_1 appears first. Say that $T = \{a_1\}$ and $S = \{v_1\}$. For these groups it holds that $|S| \geq |T|^{n/k}$, and $\alpha(c) = 1 + \varepsilon, \forall c \in T$. According to EJR for δ -cohesive groups, at least one voter in S should receive a satisfaction of at least $1 + \varepsilon$ from the outcome. The decision on whether to include a_1 in the winning committee or not should be done at this point and any online multi-winner election rule that rejects a_1 can't revoke this decision in case the candidates that will come do not satisfy v_1 to an extent of $1 + \varepsilon$. Therefore, 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 at this point. Similarly, when a_2 arrives, we have that v_1 should now receive a satisfaction of at least $1 + 2\varepsilon$ from the outcome. Currently, having included a_1 , the satisfaction of v_1 is $1 + \varepsilon$, so not sufficiently large to satisfy EJR for δ -cohesive groups in case the candidates that will come next do not satisfy v_1 , and, hence a_2 should be included in the committee as well. By similar arguments, all candidates in the set $\{a_1, a_2, \dots, a_k\}$ will be selected one after the other. Hence, there is no place to include any candidate from $\{b_1, \dots, b_k\}$ in the solution.

It holds that any rule that does not elect any candidate from $\{b_1, \dots, b_k\}$ does not satisfy EJR for δ -cohesive groups. This is because if we set $T = \{b_1, \dots, b_k\}$ and $S = \{v_1\}$, then S is $((1 + \varepsilon(\frac{k+1}{2} + \frac{1}{k}))_{i \in [k]}, T)$ -cohesive. Hence, in order to satisfy EJR for δ -cohesive groups, there should be a voter i in S such that $u_i(\mathcal{R}(E)) \geq \sum_{c \in T} (1 + \varepsilon(\frac{k+1}{2} + \frac{1}{k})) = k + \varepsilon \frac{k(k+1)}{2} + \varepsilon$. But, under a rule that does not elect any candidate from $\{b_1, \dots, b_k\}$, each voter in S gets a satisfaction of at most $\sum_{i \in [k]} 1 + i\varepsilon = k + \varepsilon \frac{k(k+1)}{2}$. $\square$

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 Extended Justified Representation

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 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.*

Proof. Regarding feasibility, we first note that the initially distributed amount of money equals $n \cdot k/n = k$. Each candidate costs 1 and is being bought only if their supporters have enough remaining funds to cover this cost. Therefore, no more than k candidates can be bought by the voters. The final step of the rule ensures that it will always return a set of exactly k candidates.

We will now show that the Greedy Budgeting rule satisfies JR. For contradiction, suppose that it does not. Then, there is a set of voters S witnessing the violation of the axiom. So the voters of this set are at least n/k and all value positively one candidate, say c . Moreover, their satisfaction from the outcome is zero, but then, they have not contributed any amount of money towards including any of their approved candidates in the outcome. Therefore, at the moment of consideration of c , they all had their initial budget which equals $|S| \frac{k}{n} \geq \frac{n}{k} \frac{k}{n} = 1$. Consequently, c should have been hired at that point by the rule. $\square$

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](#). Below we define a stronger version of JR, motivated by the observation that [Definition 3](#) only guarantees some positive utility from $\mathcal{R}(E)$ regardless of the value of $\alpha(c)$ when $T = \{c\}$. Instead, one could aim to adapt JR staying closer to the spirit of EJRP. Such a strengthening of [Definition 3](#) follows.

Definition 4. A rule $\mathcal{R}$ satisfies *Strong Justified Representation (strong-JR)* if for each election instance E , each $\alpha : C \mapsto \mathbb{R}$, $T = \{c\}$ for some $c \in C$, and each (α, T) -cohesive group of voters S , there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) \geq \alpha(c)$.

As also supported by the results in [Section 3.1](#), there is little hope when moving toward EJRP. Indeed, for the stronger adaptation of JR defined in [Definition 4](#), we also get a negative result.

Theorem 5. *There is no online voting rule that satisfies strong-JR.*

Proof. Consider the simple instance where there are three candidates $\{a, b, c\}$ and two voters, v_1, v_2 submitting the ballots $(1, 0, 0)$ and $(0, 1, 2)$ respectively. Say that $k = 2$ and that the candidate a comes before b who comes before c . Say that $T = \{a\}$ and $S = \{v_1\}$. For these groups it holds that $|S| \geq |T|n/k$, and $\alpha(c) = 1, \forall c \in T$. According to strong-JR, at least one voter in S should receive a satisfaction of at least 1 from the outcome. The decision on whether to include a in the winning committee or not should be done at this point and any online multi-winner election rule that rejects a can't revoke this decision in case the candidates that will come do not satisfy v_1 to an extent of 1. Therefore, 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 in the committee at this point. Using the same argument for $T = \{b\}$ and $S = \{v_2\}$, we can show that any such rule will also include b in the committee. Hence, then, there is no place to include candidate c in the solution. But, it holds

that any rule that does not elect c does not satisfy strong-JR. This is because if we set $T = \{c\}$ and $S = \{v_2\}$, then S is $(2, T)$ -cohesive. Hence, in order to satisfy strong-JR, there should be a voter i in S such that $u_i(\mathcal{R}(E)) \geq 2$. But, each voter in S gets a satisfaction of at most 1 if c is not hired. $\square$

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, the very voters who support c might 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 selection 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](#) (and [Theorem 5](#)). 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 hereinafter, we adopt the commonly used *random arrival model* from the online voting literature [[Gupta and Singla, 2021](#)] which assumes that elements arrive one by one in a uniformly random order. Other models, such as the adversarial one, are clear candidates 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 probability as m increases. This probability depends exponentially

¹We assume wlog that $t \geq k$; otherwise, we can add dummy candidates with no support.

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 for instance Experiment 2 in [Section 5](#)). We note that 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 empirically show in [Section 5](#), the proposed method also succeeds in producing fair solutions for general instances.

Theorem 6. *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.*

Proof. Regarding feasibility, we note that a candidate c is hired only if, at the time of consideration, it manages to replace a candidate $\hat{c} \in C_R$ currently in C_R^r . Since C_R is a feasible outcome of Offline MES, we have that $|C_R| \leq k$. Therefore, at most k candidates can be removed from C_R^r during the execution of the algorithm, and thus no more than k candidates can be hired. The final step of the rule ensures that the output always contains exactly k candidates.

We now turn to the proportionality guarantee of the rule. As a first step, we show that each candidate that would have been selected by Offline MES is hired by Online MES with probability at least $1/e$. Fix an arbitrary t and say that c is a candidate that is hired under Offline MES and arrives at time $i > t$ in Online MES. Moreover say that C_R^r is the running sample at the moment of appearance of c .

First, we need to show that c will be included in the winning outcome of MES when the available candidates are $C_R^r \cup \{c\}$. Towards a contradiction, suppose that c is not included in the winning outcome. This could happen either because the supporters of c do not have enough budget to afford c , or because there are k other candidates, say a set C^* , that are ρ -affordable at a lower value of ρ . The first case cannot occur because the supporters of c can only spend money towards candidate c , due to the single approval assumption. In the second case, all candidates in C^* must have been ρ -affordable at the same values of ρ even when considering the full set of candidates. This is because, under the single approval assumption, each voter can only contribute to the single candidate they approve. Therefore, the outcome of MES would have been C^* or would consist of candidates that are ρ -affordable for even lower values of ρ than the candidates in C^* , contradicting the fact that c is a MES-winning candidate.

Under online-MES, c will be selected only if the candidate $\hat{c}$ of C_R^r that is being kicked out of the MES-winning committee by c at the moment of consideration of c was in the reference set C_R found during the sampling phase. Since we have assumed uniformly random permutation of candidates, this happens with probability $\frac{t}{i-1}$, since $i-1$ spots of appearance of $\hat{c}$ exist before the appearance of c and only t out of them are leading to the selection of c . Moreover, each step i is equally likely for c to be spotted then (with probability of $1/m$), so the probability of selecting c is

$$\sum_{t+1}^m \frac{1}{m} \frac{t}{i-1} = \frac{t}{m} \sum_{t+1}^m \frac{1}{i-1} > \frac{t}{m} \ln \frac{m}{t}.$$

The function $\frac{t}{m} \ln \frac{m}{t}$ is being maximized for $t = m/e$. Hence, for $t = \lfloor m/e \rfloor$, with a probability of at least $1/e$, the candidate c is hired by Online MES. If all the k candidates of the MES-winning committee are in the winning committee of Online MES, then its outcome satisfies EJR. This is a consequence of the fact that MES satisfies EJR for multi-winner elections [[Peters et al., 2021](#)] and happens with a probability of $(1/e)^k$. Moreover, to satisfy EJR up to p candidates, it suffices to elect only $k-p$ candidates from the outcome of MES, and this happens with a probability of $(1/e)^{k-p}$. $\square$

The proof of [Theorem 6](#) 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. Hence, Online MES is the first rule providing such bounded guarantees; whether analyses improving the one in [Theorem 6](#) 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 BOS 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 [Papasotiropoulos 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, [Papasotiropoulos et al. \[2025a\]](#) introduced the Method of Equal Shares with Bounded Overspending (BOS). BOS is considered better suited for practical applications than MES, with the council of Pruszków, Poland, already adopting it for an election featuring the largest budget to date allocated using a voting rule with theoretical fairness guarantees. 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 $\Pi(W) := \sum_{i \in V} \log(1 + u_i(W))$.

The Nash Product rule is well-studied in computational social choice [[Caragiannis et al., 2019](#); [Freeman et al., 2017](#); [Conitzer et al., 2017](#); [Fain et al., 2018](#); [Ebadian et al., 2024](#)]. To the best of our knowledge, our work is the first to 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 it, 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.

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 $g_i(c) = \Pi(T_{i-1} \cup \{c\})$, 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 7. *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.*

Proof. [Bateni et al. \[2013\]](#) proposed an online algorithm that, given a monotone nonnegative submodular function f , returns a solution whose expected value under f is within a multiplicative factor of $\frac{1-1/e}{7}$ of the optimum. The Online Nash Rule is simply an application of the rule by [Bateni et al. \[2013\]](#) with the Nash Welfare function Π as the maximizing function f . Therefore, to prove the claimed guarantee it suffices to show that the Nash Welfare function is a monotone nonnegative submodular one. The first two properties are trivial by definition. It remains to prove submodularity.

We want to show that the function $\Pi(W) = \sum_{i \in V} \log(1 + u_i(W))$ is submodular. Recall that a function $f : 2^C \mapsto \mathbb{R}_{\geq 0}$ is submodular if for all $A \subseteq B \subseteq C$ and every $c \in C \setminus B$, it holds $f(A \cup \{c\}) - f(A) \geq f(B \cup \{c\}) - f(B)$. We want to show that this holds for Π . Fix any sets of candidates A, B such that $A \subseteq B$ and a candidate $c \notin B$. Then, for both $X = A$ and $X = B$ it holds that

$$\begin{aligned} \Pi(X \cup \{c\}) - \Pi(X) &= \sum_{i \in V} [\log(1 + u_i(X \cup \{c\})) - \log(1 + u_i(X))] = \\ \sum_{i \in V} [\log(1 + u_i(X) + u_i(c)) - \log(1 + u_i(X))] &= \sum_{i \in V} \log \frac{1 + u_i(X) + u_i(c)}{1 + u_i(X)} = \sum_{i \in V} \log \left(1 + \frac{u_i(c)}{1 + u_i(X)} \right). \end{aligned}$$

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

Fix an arbitrary voter $i \in V$. Since $A \subseteq B$ it holds that $u_i(A) \leq u_i(B)$. Therefore, $\frac{1}{u_i(A)} \geq \frac{1}{u_i(B)}$, and hence we get $\Pi(A \cup \{c\}) - \Pi(A) \geq \Pi(B \cup \{c\}) - \Pi(B)$. The feasibility of the outcome follows directly from the mechanics of the rule. $\square$

5 Empirical Evaluation

Finally, we experimentally study Greedy Budgeting, Online MES,³ 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.

5.1 Experiment 1: Pabulib Instances

We computed the outcomes of the proposed rules across all the 773 participatory budgeting instances that were available in the PABULIB library as of 22-08-2025. To allow the experiment to run in reasonable time, we only focused on instances of manageable size and constructed the dataset as follows. We used all approval-based instances (excluding those with $m \leq 2$) where the number of voters was not excessively large ($n \leq 50,000$). Among those, we selected instances where either both n and m are small ($n < 10,000$ and $m < 100$), or where n is large ($10,000 < n < 50,000$) but m is small ($m < 80$), or where m is large ($m > 100$), but n is small ($n < 5,000$). We then partitioned these instances into three groups based on the number of candidates m : instances with $m < 10$ were considered small (202 instances), those with $m \geq 30$ were considered large (396 instances), and the rest were classified as medium (175 instances). For each instance, we consider 5 different values of committee size, chosen as functions of m to allow consistent comparison across datasets. Additionally, for each setting, we run 5 iterations with candidate arrival orders drawn uniformly at random, in line with our theoretical analyses. To assess proportionality, we use EJR+ as its satisfaction (and degree of violation) can be verified efficiently. In [Tables 1](#) and [2](#) we report the (average) share of voters who are not adequately represented and the number of candidates witnessing the violation, respectively, for various values of k , with instances broken down by the number of candidates.

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 2.5% 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. The objective of this experiment is not to draw comparisons between the rules; the subsequent experiments in this section are devoted to comparative analyses. Finally, in [Table 3](#), we report the average running time⁴ per instance for each rule, showcasing a noticeable difference among them—we briefly return to it in [Section 6](#).

³Offline MES is typically paired with a completion method to ensure the selection of exactly k candidates [Faliszewski et al., 2023; Brill and Peters, 2024]. Several such methods have been proposed in the literature. Since Online MES relies on Offline MES as a subroutine, we had to choose a specific variant for our experiments—though this choice does not affect the theoretical guarantees of our paper. We tested both the Utilitarian and Add1U completion methods on small instances and found their outcomes to be largely similar. Given the already high running time of Online MES without completion, we focused on the Utilitarian variant for efficiency.

⁴The experimental part of our work was implemented in Python and run on an ASUS VivoBook laptop with an Intel i7-11370H processor, 16GB of RAM, and a 512GB PCIe SSD.

	$k = m/20$	$k = m/15$	$k = m/10$	$k = m/7$	$k = m/4$
Greedy Budgeting					
Small instances	0.0000	0.0000	0.0000	0.0137	0.0117
Medium instances	0.0242	0.0200	0.0188	0.0181	0.0156
Large instances	0.0176	0.0171	0.0164	0.0155	0.0135
All instances	0.0206	0.0188	0.0181	0.0167	0.0141
Online MES					
Small instances	0.0000	0.0000	0.0000	0.0137	0.0116
Medium instances	0.0236	0.0196	0.0188	0.0181	0.0157
Large instances	0.0177	0.0172	0.0165	0.0157	0.0135
All instances	0.0203	0.0186	0.0181	0.0168	0.0142
Online BOS					
Small instances	0.0000	0.0000	0.0000	0.0137	0.0116
Medium instances	0.0236	0.0196	0.0185	0.0180	0.0155
Large instances	0.0176	0.0171	0.0163	0.0155	0.0133
All instances	0.0203	0.0186	0.0179	0.0167	0.0140
Online Nash Rule					
Small instances	0.0000	0.0000	0.0000	0.0138	0.0117
Medium instances	0.0226	0.0190	0.0194	0.0187	0.0164
Large instances	0.0181	0.0177	0.0170	0.0161	0.0140
All instances	0.0201	0.0185	0.0186	0.0172	0.0147

Table 1: Average EJR+ Violations by Rule and Committee Size (Experiment 1). Values represent percentages of voters witnessing a villation, where 0 corresponds to 0% and 1 to 100%

	$k = m/20$	$k = m/15$	$k = m/10$	$k = m/7$	$k = m/4$
Greedy Budgeting					
Small instances	0.0000	0.0000	0.0000	0.3656	0.5192
Medium instances	0.2257	0.2423	0.3388	0.4915	0.9038
Large instances	0.3380	0.4883	0.7399	1.0469	1.7986
All instances	0.2877	0.3436	0.4617	0.6138	1.0142
Online MES					
Small instances	0.0000	0.0000	0.0000	0.3566	0.5115
Medium instances	0.2228	0.2375	0.3305	0.4800	0.9079
Large instances	0.3313	0.4800	0.7377	1.0623	1.8746
All instances	0.2827	0.3374	0.4553	0.6097	1.0319
Online BOS					
Small instances	0.0000	0.0000	0.0000	0.3566	0.5088
Medium instances	0.2228	0.2375	0.3298	0.4781	0.8939
Large instances	0.3306	0.4775	0.7254	1.0241	1.7721
All instances	0.2823	0.3363	0.4511	0.5988	1.0003
Online Nash Rule					
Small instances	0.0000	0.0000	0.0000	0.3709	0.5408
Medium instances	0.2296	0.2444	0.3424	0.5180	1.0003
Large instances	0.3478	0.5104	0.7995	1.1666	2.0922
All instances	0.2948	0.3540	0.4825	0.6609	1.1375

Table 2: Average Shortfall of EJR+ Violations by Rule and Committee Size (Experiment 1).

Rule	Time (seconds)
Greedy Budgeting Rule	0.0181
Online Method of Equal Shares	1.1999
Online Method of Equal Shares with Bounded Overspending	6.7802
Online Nash Rule	0.3237

Table 3: Average running time per instance from PABULIB, for each rule (Experiment 1).

5.2 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 in computational social choice and beyond.

The Sushi dataset [Kamishima, 2003; Kamishima et al., 2010] contains 5,000 individual rankings of 100 items grouped into 5 preference tiers with some items left unranked, which we interpret as scores from 0 to 5. We evaluate 9 different values of k ; specifically $k \in \{2, 5, 8, 15, 20, 25, 35, 40, 45\}$. For each, we run 20 iterations of each algorithm over random arrival orders of candidates.

The MovieLens dataset we used [Harper and Konstan, 2015] contains 100,000 movie ratings where 1000 users are giving scores from 0 to 5 on 1700 movies. The same values of k were considered for this dataset. We run 50 iterations of each algorithm, each with a random arrival order. The two datasets differ significantly in the number of candidates, a deliberate choice to provide insights into rules’ behavior under different candidates’ set sizes.

The goal here 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 a satisfaction of 0, average satisfaction of the least satisfied quartile of voters, and the Gini coefficient capturing inequality in satisfaction. In Table 4 we report, for each metric, the percentage that a rule ranks first, or among the two best or last among the examined rules, for different committee sizes, for the Sushi dataset. Table 5 report the results for the MovieLens dataset.

For the Sushi dataset, Greedy Budgeting is most often the top-performing rule across all the examined metrics, but there are also multiple instances where it fails to rank among the top two. Online BOS, on the other hand, appears more consistently reliable: for three metrics it ranks among the two best significantly more frequently than Greedy, and even for the 25th percentile metric—where it is slightly behind—it still does so in the vast majority of instances. The MovieLens dataset is considerably more favorable to Online BOS, which ranks first in far more instances than any other rule. Online MES also sees a notable 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.

5.3 Experiment 3: Sampled Datasets

According to Boehmer et al. [2024b], the two most commonly used models for generating synthetic data in computational social choice are Impartial Culture (IC) and Mallows model. 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 et al., 2023].

Best Top-2 Worst	Avg Sat	Excl Ratio	25th Perc	Gini
Greedy Budgeting	47.78% 72.22% 3.33%	56.11% 74.44% 3.33%	64.44% 91.67% 0.00%	48.89% 73.89% 3.33%
Online MES	8.33% 32.78% 3.89%	8.89% 33.33% 5.00%	4.44% 27.78% 2.78%	8.89% 30.56% 5.00%
Online BOS	43.33% 93.89% 1.11%	34.44% 90.56% 1.11%	31.11% 80.56% 0.00%	41.67% 94.44% 1.11%
Online Nash Rule	0.56% 1.11% 91.67%	0.56% 1.67% 90.56%	0.00% 0.00% 97.22%	0.56% 1.11% 90.56%

Table 4: Percentage where each rule performed best, among the two best and worst, according to different metrics, for the Sushi dataset (Experiment 2).

Best Top-2 Worst	Avg Sat	Excl Ratio	25th Perc	Gini
Greedy Budgeting	4.00% 22.89% 0.00%	11.78% 54.67% 0.00%	12.89% 56.67% 0.00%	14.67% 58.00% 0.00%
Online MES	8.67% 80.00% 0.67%	7.78% 48.22% 0.67%	3.56% 53.56% 0.00%	6.67% 44.67% 0.67%
Online BOS	87.33% 96.67% 0.89%	80.44% 96.67% 0.89%	83.56% 89.78% 0.00%	78.67% 96.89% 0.89%
Online Nash Rule	0.00% 0.44% 98.44%	0.00% 0.44% 98.44%	0.00% 0.00% 100.00%	0.00% 0.44% 98.44%

Table 5: Percentage where each rule performed best, among the two best and worst, according to different metrics, for the MovieLens dataset (Experiment 2).

For IC, we select 5 fixed values of a parameter $p \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ and we say that each voter approves a number of candidates drawn from a binomial distribution centered around the fraction p of m . Having fixed their number, the actual approved candidates are chosen at random. Then, positive utilities capped at 200 are assigned following a normal distribution with mean value 150 and standard deviation 140. We examine the values of n from the set $\{5, 10, 20, 30, 50, 70, 100, 150\}$ and we also fix 32 pairs of (m, k) values selected so as for each m to have both small, medium and large fraction of it as k . For each combination, we run 10 iterations with different candidate arrival orders.

For Mallows model, we first generate ranked ballots in the usual way [Mallows, 1957; Boehmer et al., 2022, 2021], and then we assign utilities evenly spaced between 0 and 200 based on the ranking. We also add some utility noise across voters for achieving larger variance of scores in ballots. Again, we fix n , m , and k as in the IC model and evaluate our rules on profiles with 5 different dispersion parameters $\phi \in \{0.2, 0.4, 0.6, 0.8, 1\}$ and, again, 10 different candidate arrival orders. Regarding the parameterization of the experiment for Normalized Mallows, we set it up in the exact same way as for the classic Mallows model and we follow the implementation from Python PrefSampling library [Boehmer et al., 2024c].

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. In this experiment, we assess how much our rules deviate from the Offline MES outcomes based on the four statistics also used in Experiment 2. Specifically, we examined the average satisfaction of all voters and of the least satisfied quarter of voters as well as the Gini coefficient and exclusion ratio as measures of fairness. Naturally, given their online nature, our rules fall short compared to Offline MES, but the key question is by how much, and which rule performs best in minimizing this gap.

We first pay particular attention to the exclusion ratio metric for the Impartial Culture model against the approval probability p (Figure 1, bottom-right). When voters approve many candidates, each with varying scores, all rules exhibit an exceptionally low exclusion ratio on average, and as the approval probability decreases, differences between the rules become apparent. Across all values of p , the exclusion ratio of all four rules remain within an additive gap of at most 2.5 candidates from that of Offline MES. For low approval probabilities, BOS and Greedy perform nearly as well, being the top performers, with Online MES staying slightly behind. For $p = 0.6$, Greedy becomes the worst rule among the examined, while BOS achieves again a guarantee particularly close to the one provided by Offline MES.

In the same sampling model (IC), the Gini coefficient (Figure 1, top-left) of BOS is consistently the smallest of our proposed rules and the gap between each rule and Offline MES narrows as n increases. In Figures 2 and 3 (top-left), we also report the Gini coefficient of each rule’s outcome for the Mallows cultures, the values of which remain consistently particularly low for all rules, with their ranking following the same order as in the case of IC.

For the two satisfaction-based metrics (average and worst quartile), the conclusions are broadly similar: Online BOS consistently ranks highest, followed by Online MES with a clear margin over the remaining rules, and, as n increases, the performance of all rules approaches that of Offline MES. These patterns hold across all three cultures (see also top-right and bottom-left of Figures 1 to 3) and can be observed also from the plots showing the 25th percentile utility in relation to the dispersion parameter of the Mallows models (bottom-right of Figures 2 and 3).




Figure 1: Gini coefficient against the number of voters (top-left), average satisfaction (top-right) and 25th percentile (bottom-left) relative to Offline MES against the number of voters, and Exclusion ratio against the approval probability p (bottom-right) in the IC model (Experiment 3).

5.4 Experiment 4: Polarized Instances

In this 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. In this setup, 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 created 300 profiles by randomly selecting values for n , m , k , x and the approval rate of voters from group B, running 10 different random candidate orders for each profile. We measure the frequency with which group A receives less representation than they deserve and we also quantify the extent of this underrepresentation. Since JR guarantees no shortfall here, Greedy Budgeting produces outcomes with no violations by design. Online MES fails in about a fifth of the cases, Online BOS in slightly more, and the Nash Rule only in a very few. 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. We report results on percentage and deficit of underperformance in Table 6.




Figure 2: Gini coefficient against the number of voters (top-left), average satisfaction (top-right) and 25th percentile (bottom-left) relative to Offline MES against the number of voters, and 25th percentile relative to Offline MES against the dispersion parameter ϕ (bottom-right) in the Mallovs model (Experiment 3).




Figure 3: Gini coefficient against the number of voters (top-left), average satisfaction (top-right) and 25th percentile (bottom-left) relative to Offline MES against the number of voters, and 25th percentile relative to Offline MES against the dispersion parameter ϕ (bottom-right) in the Normalized Mallovs model (Experiment 3).

Algorithm	Underperformance	Average Deficit	Maximum Deficit
Greedy Budgeting	0.00%	0.0000	0.0000
Online MES	19.67%	1.8576	8.9000
Online BOS	27.00%	0.8556	4.1000
Online Nash Rule	2.67%	3.9375	8.1000

Table 6: Underperformance percentages, average deficits across all instances, and maximum deficits per instance, computed as the average of maxima over multiple runs of the same instance under different candidate arrival orders (Experiment 4).

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. 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. The PB setting brings new challenges because the number of winners is not fixed, which calls for significantly different insights or techniques.

Acknowledgments. We thank Tomasz Wąs, Piotr Faliszewski and Piotr Skowron 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.



European Research Council
Established by the European Commission

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*, pages 461–485.
- Babai, 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*, pages 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*.
- Batoni, M., Hajiaghayi, M., and Zadimoghaddam, M. (2013). Submodular secretary problem and extensions. *ACM Transactions on Algorithms*, pages 1–23.
- 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., Bredereck, R., Elkind, E., Faliszewski, P., and Szufa, S. (2022). Expected frequency matrices of elections: Computation, geometry, and preference learning. In *Proceedings of the Conference on Neural Information Processing Systems*.
- Boehmer, N., Bredereck, R., Faliszewski, P., Niedermeier, R., and Szufa, S. (2021). Putting a compass on the map of elections. In *Proceedings of the International Joint 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., Janeczko, Ł., Kaczmarczyk, A., Lisowski, G., Pierczyński, G., Rey, S., Stolicki, D., Szufa, S., and Wąs, T. (2024c). 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 Kraiczy, 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*.
- Brill, M. and Peters, J. (2024). Completing priceable committees: Utilitarian and representation guarantees for proportional multiwinner voting. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- 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*, pages 1–32.
- Conitzer, V., Freeman, R., and Shah, N. (2017). Fair public decision making. In *Proceedings of the ACM Conference on Economics and Computation*, pages 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*, pages 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*, pages 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*, pages 1–39.
- Elkind, E., Obraztsova, 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*, pages 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*, pages 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*, pages 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*, pages 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*, pages 630–631.
- Lackner, M. and Skowron, P. (2023). *Multi-Winner Voting with Approval Preferences*. Springer Nature. Chapter 4: “Proportionality”.
- Mallows, C. L. (1957). Non-null ranking models. i. *Biometrika*, pages 114–130.
- Nash, J. F. (1950). The bargaining problem. *Econometrica*, pages 155–162.
- Papastotiropoulos, G., Pishbin, S. Z., Skibski, O., Skowron, P., and Wąs, T. (2025a). Method of equal shares with bounded overspending. In *Proceedings of the ACM Conference on Economics and Computation*.
- Papastotiropoulos, G., Skibski, O., Skowron, P., and Wąs, 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*, pages 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*.