
Planning to Fairly Allocate: Probabilistic Fairness in the Restless Bandit Setting

Christine Herlihy^{*1} Aviva Prins^{*1} Aravind Srinivasan¹ John P. Dickerson¹

Abstract

Restless and collapsing bandits are often used to model budget-constrained resource allocation in settings where receiving the resource increases the probability that an arm will transition to, or remain in, a desirable state. However, SOTA Whittle-index-based approaches to this planning problem either do not consider fairness among arms, or incentivize fairness without guaranteeing it. We introduce PROBFAIR, an algorithm which finds the best (reward-maximizing) policy that: (a) satisfies the budget constraint; and (b) enforces bounds $[\ell, u]$ on the probability of being pulled at each timestep. We evaluate our algorithm on a real-world application, where interventions support continuous positive airway pressure (CPAP) therapy adherence among patients, as well as on a broader class of synthetic transition matrices. PROBFAIR preserves utility while providing fairness guarantees.

1. Introduction

Restless multi-armed bandits (RMABs) are used to model budget-constrained resource allocation tasks in which a decision-maker must select a subset of arms (e.g., projects, patients, assets) to receive a beneficial intervention at each timestep, while the state of each arm evolves over time in an action-dependent, Markovian fashion. Such problems are common in healthcare, where clinicians may be tasked with monitoring large, distributed patient populations and determining which individuals to expend scarce resources on to maximize total welfare. RMABs have been proposed to prioritize which inmates should receive hepatitis C treatment in U.S. prisons (Ayer et al., 2019), and which tuberculosis patients should receive medication adherence support in India (Mate et al., 2020).

^{*}Equal contribution ¹Department of Computer Science, University of Maryland, College Park, MD, USA. Correspondence to: Christine Herlihy <cherlihy@cs.umd.edu>, Aviva Prins <aviva@cs.umd.edu>.

A full version of this paper is available at [arXiv:2106.07677](https://arxiv.org/abs/2106.07677).
ICML 2022 Workshop on Responsible Decision Making in Dynamic Environments, Baltimore, Maryland, USA, 2022. Copyright 2022 by the author(s).

Current SOTA approaches to solving RMAB problems rely on the indexing work introduced by Whittle (1988). While the Whittle index solves an otherwise PSPACE-complete problem in an asymptotically optimal fashion by decoupling arms (Weber & Weiss, 1990), it fails to provide any guarantees about how pulls will be distributed *among arms*.

Though the intervention is canonically assumed to be beneficial for *every* arm, the marginal benefit (i.e., relative increase in the probability of a favorable state transition) varies in accordance with each arm’s underlying state transition function. Consequently, Whittle index-based maximization of total expected reward *without regard for distributive fairness* empirically allocates all available interventions to a small subset of arms, ignoring the rest (Prins et al., 2020). We observe that this distributive outcome may be perceived as unfair by arms who fail to receive any pulls, as well as ethically unacceptable by decision-makers who must balance utilitarian welfare maximization with arm-level considerations such as beneficence. Relatedly, while purely egalitarian policies such as ROUND-ROBIN *do* guarantee individual fairness with respect to receipt of the intervention, they do not maximize the probability of beneficial outcomes and therefore incur a large loss of total welfare.

We find that it is intractable to modify the Whittle index in order to guarantee distributive fairness. We thus introduce PROBFAIR, a state-agnostic policy that maps each arm to a fairness-constraint satisfying, stationary probability distribution over actions that takes the arm’s transition matrix into account. At each timestep, we use a dependent rounding algorithm (Srinivasan, 2001) to sample from this distribution to produce a budget-constraint satisfying discrete action vector. We present empirical results in Appendix E.

2. Restless Multi-Armed Bandit Model

A restless multi-armed bandit consists of $N \in \mathbb{N}$ independent arms, each of which evolves over a finite time horizon $T \in \mathbb{N}$, according to an associated Markov Decision Process (MDP). Each arm’s MDP is characterized by a 4-tuple $(\mathcal{S}, \mathcal{A}, P, r)$ where \mathcal{S} represents the state space, \mathcal{A} represents the action space, P represents an $|\mathcal{S}| \times |\mathcal{A}| \times |\mathcal{S}|$ transition matrix, and $r : \mathcal{S} \rightarrow \mathbb{R}$ represents a local reward function that maps states to real-valued rewards. Appendix A summarizes notation; note that $[N]$ denotes the set $\{1, 2, \dots, N\}$.

States, actions, and observability: We specifically consider a discrete two-state system $\mathcal{S} := \{0, 1\}$ where 1 (0) represents being in the “good” (“bad”) state, and a set of two possible actions $\mathcal{A} := \{0, 1\}$ where 1 represents the decision to select (“pull”) arm $i \in [N]$ at time $t \in [T]$, and 0 represents the choice to be passive (not pull). In the general RMAB setting, each arm’s state s_t^i is observable. We consider the partially-observable extension introduced by Mate et al. (2020), where each arm’s state is replaced with the probabilistic *belief* $b_t^i \in [0, 1]$ that it is in state 1. Partial observability captures uncertainty regarding patient status and treatment efficacy associated with outpatient or remotely-administered interventions.

Transition matrices: Each arm i is characterized by a set of transition matrices P , where $P_{s,s'}^a$ denotes the probability of transitioning from state s to state s' when action a is taken. We assume P to be (a) static and (b) known by the agent at planning time. Assumptions (a) and (b) are likely to be violated in practice; however, they provide a useful foundation, and can be modified to incorporate additional uncertainty, such as the requirement that transition matrices must be learned (Ortner et al., 2012; Jung & Tewari, 2019; Jung et al., 2019). Researchers often use longitudinal data to construct risk-adjusted transition matrices that encode cohort-specific transition probabilities to guide patient-level decision-making (Steimle & Denton, 2017).

Consistent with previous literature, we assume strictly positive transition matrix entries, and impose four *structural constraints*: (a) $P_{0,1}^0 < P_{1,1}^0$; (b) $P_{0,1}^1 < P_{1,1}^1$; (c) $P_{0,1}^0 < P_{0,1}^1$; (d) $P_{1,1}^0 < P_{1,1}^1$ (Liu & Zhao, 2010; Mate et al., 2020; 2021; Ou et al., 2022). These constraints are application-motivated, and imply that arms are more likely to remain in a “good” state than change from a bad state to a good one, and that a pull is helpful when received. In the absence of such constraints, receiving the intervention would be either superfluous or harmful, rather than desirable.

Objective and constraints: In the canonical RMAB setting, the agent’s goal is to find a policy π^* that maximizes total expected reward—i.e., $\arg \max_{\pi} \mathbb{E}_{\pi}[R(r(s))]$, while satisfying a *budget constraint*, $k \ll N \in \mathbb{N}$, which allows the agent to select at most k arms at each timestep. Our objective function, which is applied to the entire cohort, rewards arms in the “good” state. Specifically, we consider a cumulative reward function $R(\cdot) := \sum_{i \in [N]} \sum_{t \in [T]} \beta^{t-1} r(s_t^i)$, for discount rate $\beta \in [0, 1]$, and non-decreasing $r(s)$.

We extend this model by introducing a Boolean-valued, distributive fairness-motivated constraint, which may take one of two general forms:

1. *Time-indexed:* A function $g(\cup_{t \in [T]} \{\vec{a}_t\})$ which is satisfied if each arm is pulled at least once within each user-defined time interval $\nu \leq T$ (e.g., at least once every seven days), or a minimum fraction $\psi \in (0, 1)$

of times over the entire time horizon (Li et al., 2019).

2. *Probabilistic:* A function $g'(\vec{p}^i | \vec{a}_t \sim \vec{p}^i \forall t)$ which operates on the stationary probability vector \vec{p}^i , from which discrete actions are drawn, by requiring the probability that each arm receives a pull at any given t to fall within an interval $[\ell, u]$ where $0 < \ell \leq \frac{k}{N} \leq u \leq 1$.

3. Background, Motivation, and Related Work

3.1. Background: Whittle Index-based Policies

Pre-computing the optimal policy for a given set of restless or collapsing arms is PSPACE-hard in the general case (Papadimitriou & Tsitsiklis, 1994). However, as established by Whittle (1988) and formalized by Weber & Weiss (1990), if a set of arms are *indexable*, we can decouple the arms and efficiently solve the problem using an asymptotically-optimal heuristic index policy.

Mechanics: The Whittle index represents the subsidy, m , required to make the agent indifferent between *pulling* and *not pulling* arm i at time t . (Per Section 2, b denotes the probabilistic belief that an arm is in state $s = 1$; for restless arms, $b_t^i = s_t^i \in \{0, 1\}$.)

$$W(b_t^i) = \inf_m \left\{ m \mid V_m(b_t^i, a_t^i = 0) \geq V_m(b_t^i, a_t^i = 1) \right\} \quad (1)$$

The value function $V_m(b)$ represents the maximum expected discounted reward under passive subsidy m and discount rate β for arm i with belief state $b_t^i \in [0, 1]$ at time t :

$$V_m(b_t^i) = \max \begin{cases} m + r(b_t^i) + \beta V_m(b_{t+1}^i) & \text{passive} \\ r(b_t^i) + \beta [b_t^i V_m(P_{1,1}^1) + (1 - b_t^i) V_m(P_{0,1}^1)] & \text{active} \end{cases} \quad (2)$$

Once the Whittle index has been computed for each arm, the agent sorts the indices, and the k arms with the greatest index values receive a pull at time t , while the remaining $N - k$ arms are passive. Weber & Weiss (1990) give sufficient conditions for *indexability*:

Definition 3.1. An arm is indexable if the set of beliefs for which it is optimal to be passive for a given m , $\mathcal{B}^*(m) = \{b \mid \forall \pi \in \Pi_m^*, \pi(b) = 0\}$, monotonically increases from \emptyset to the entire belief space as m increases from $-\infty$ to $+\infty$. An RMAB is indexable if every arm is indexable.

Indexability is often difficult to establish in practice, and computing the Whittle index can be complex (Liu & Zhao, 2010). Prevailing approaches rely on proving the optimality of a *threshold policy* for a subset of transition matrices (Niño-Mora, 2020). A *forward* threshold policy pulls an arm when its state is at or below a given threshold, and makes the arm passive otherwise; the converse is true for a *reverse* threshold policy. Mate et al. (2020) give such conditions for this RMAB setting, when $r(b) = b$, and provide

an algorithm, THRESHOLD WHITTLE, that is asymptotically optimal for forward threshold-optimal arms. Mate et al. (2021) expand on this for any non-decreasing $r(b)$ and present the RISK-AWARE WHITTLE algorithm. As far as we are aware, no algorithm exists for reverse threshold-optimal arms.

3.2. Motivation: Individual Welfare Under Whittle

Bimodal allocation: Existing theory does not offer any guarantees about how the sequence of actions will be distributed over arms under Whittle index-based policies, nor about the probability with which a given arm can expect to be pulled at any particular timestep. Prins et al. (2020) demonstrate that Whittle-based policies tend to allocate all pulls to a small number of arms, neglecting most of the population. We present similar findings in Appendix B.

Ethical implications: This zero-valued lower bound on the number of pulls an arm can receive aligns with a *utilitarian* approach to distributive justice, in which the decision-maker seeks to allocate resources so as to maximize total expected utility (Bentham, 1781; Marseille & Kahn, 2019). This may be incompatible with competing pragmatic and ethical desiderata, including *egalitarian* and *prioritarian* notions of distributive fairness, in which the decision-maker seeks to allocate resources equally among arms (e.g., ROUND-ROBIN), or prioritize arms considered to be worst-off under the status quo, for some quantifiable notion of *worst-off* that induces a partial ordering over arms (Rawls, 1971; Scheunemann & White, 2011). We consider the *worst off* arms to be those who would be deprived of algorithmic attention (e.g., not receive any pulls), or, from a probabilistic perspective, would have a zero-valued lower bound on the probability of receiving a pull at any given timestep.

Why algorithmic attention? This choice is motivated by our desire to improve *equality of opportunity* (i.e., access to the beneficial intervention) rather than *equality of outcomes* (i.e., observed adherence). The agent directly controls who receives the intervention, but has only indirect control (via actions) over the sequence of state transitions an arm experiences. Additionally, proclivity towards adherence may vary widely in the absence of restrictive assumptions about cohort homogeneity, and focusing on equality of outcomes could thus entail a significant loss of total welfare.

Distributive fairness and algorithmic acceptability: To realize the benefits associated with an algorithmically-derived resource allocation policy, practitioners tasked with implementation must find the policy to be acceptable and potential beneficiaries must find participation to be rational (i.e., to yield an increase in expected time in the adherent state, relative to non-participation). Providing fairness-aware decision support can improve acceptability (Rajkomar et al., 2018; Kelly et al., 2019) and minimize the loss of reward associ-

ated with ethically-motivated deviation to a sub-optimal but equitable approach such as ROUND-ROBIN (De-Arteaga et al., 2020; Dietvorst et al., 2015).

3.3. Time-indexed Fairness and Indexability

In Appendix C.1, we demonstrate that it is not possible to modify the Whittle index to guarantee time-indexed fairness while preserving our ability to decouple arms.

3.4. Additional Related Work

Multi-armed bandit optimization problems are canonically framed from the perspective of the decision-maker; interest in individual and group fairness in this setting is a relatively recent phenomenon (Joseph et al., 2016). To the extent that fairness among arms has been considered, existing approaches either: (1) do not consider restless arms (Chen et al., 2020; Li et al., 2019); (2) redistribute pulls without providing arm-level guarantees (Mate et al., 2021; Killian et al., 2021); or (3) provide time-indexed fairness guarantees without optimality guarantees (Prins et al., 2020).

4. PROBFAIR: a Probabilistically Fair Policy

We now introduce PROBFAIR, a state-agnostic policy that maximizes reward subject to the satisfaction of both budget and probabilistic fairness constraints. PROBFAIR maps each arm i to an arm-specific, stationary probability distribution over atomic actions, such that for each timestep t , $P[a_t^i = 1] = p_i$ and $P[a_t^i = 0] = 1 - p_i$, where $p_i \in [\ell, u]$ for all $i \in [N]$ and $\sum_i p_i = k$. Here, ℓ and u are user-defined fairness parameters satisfying $0 < \ell \leq \frac{k}{N} \leq u \leq 1$, per Section 2. In Section 4.1, we describe how to construct the p_i 's so as to efficiently approximate our constrained reward-maximization objective within a multiplicative factor of $(1 - \epsilon)$, for any given constant $\epsilon > 0$.

To build intuition as to why the state-agnostic mapping from arms to probability distributions over atomic actions is consistent with constrained reward maximization, observe that when we take the union of each arm's stationary probability vector, we obtain a system-level policy, $\pi_{PF} : \{i \mid i \in N\} \rightarrow [1 - p_i, p_i]^N$. Regardless of the system's initial state, repeated application of this policy will result in convergence to a steady-state distribution in which (WLOG) arm i is in the adherent state (i.e., state 1) with probability $x_i \in [\ell, u]$, and the non-adherent state (i.e., state 0) with probability $1 - x_i \in [0, 1]$. By definition, for any arm i , x_i will satisfy the equation

$$x_i [(1 - p_i)P_{1,1}^0 + p_i P_{1,1}^1] + (1 - x_i)[(1 - p_i)P_{0,1}^0 + p_i P_{0,1}^1] = x_i. \quad (3)$$

Thus, $x_i = f_i(p_i)$, where

$$f_i(p_i) = \frac{(1 - p_i)P_{0,1}^0 + p_i P_{0,1}^1}{1 - (1 - p_i)P_{1,1}^0 - p_i P_{1,1}^1 + (1 - p_i)P_{0,1}^0 + p_i P_{0,1}^1}. \quad (4)$$

We seek the policy which maximizes total expected reward.

Thus, PROBFair is defined as:

$$\pi_{PF} = \arg \max_{p^t \in [\ell, u]} \sum_i f_i(p_i) \text{ s.t. } \sum_i p_i \leq k \quad (5)$$

Solving this constrained maximization problem is thus consistent with maximizing the expected number of timesteps each arm will spend in the adherent state, subject to satisfying the budget and fairness constraints, since each f_i takes into account arm i 's underlying transition probabilities.

4.1. Computing the p_i 's: Algorithmic Approach

To construct π_{PF} (Eq. 5), we: (1) partition the arms based on the shapes of their respective f_i functions (Eq. 4); (2) perform a grid search over possible ways to allocate the budget, k , between the two subsets of arms; (2a) solve each sub-problem to produce a probabilistic policy for the arms in that subset; (2b) compute the total expected reward of the policy; (3) take the argmax over this set of grid search values to determine the approximately optimal budget allocation; and (4) form π_{PF} by taking the union over the policies produced by evaluating each sub-problem at its approximately optimal share of the budget. Figure 5 visualizes. We begin by introducing two theorems (see Appendix D for proofs):

Theorem 4.1. *For every arm $i \in [N]$, $f_i(p_i)$ is either concave or strictly convex in all of $p_i \in [0, 1]$.*

Theorem 4.2. *For each arm $i \in [N]$, the structural constraints introduced in Section 2 ensure that $f_i(p_i)$ is monotonically non-decreasing in p_i over the interval $[0, 1]$.*

Our algorithm takes these monotonicity and shape properties of the f_i 's into account and splits the problem into two disjoint subproblems based on the concavity of the arms (Theorem 4.1). Since $f_i'(p_i) \geq 0$ for all arms i per Theorem 4.2, our budgetary constraint $\sum_i p_i \leq k$ becomes an equality, $\sum_i p_i = k$. Let $\mathcal{X} = \{i \mid f_i \text{ is concave}\}$ and $\mathcal{Y} = N \setminus \mathcal{X} = \{i \mid f_i \text{ is strictly convex}\}$, and let:

- (P1) maximize $\sum_{i \in \mathcal{X}} f_i(p_i)$ subject to: $p_i \in [\ell, u]$ for all $i \in \mathcal{X}$, and $\sum_{i \in \mathcal{X}} p_i = z$
- (P2) maximize $\sum_{i \in \mathcal{Y}} f_i(p_i)$ subject to: $p_i \in [\ell, u]$ for all $i \in \mathcal{Y}$, and $\sum_{i \in \mathcal{Y}} p_i = k - z$

Then, π_{PF} is the union of the solutions to **P1** and **P2** at the optimal grid search value $z^* = \arg \max_z \sum_{i \in \mathcal{X}} f_i(p_i) + \sum_{i \in \mathcal{Y}} f_i(p_i)$. Algorithm 1 provides pseudocode.

P1 is a concave-maximization problem that can be solved efficiently via gradient descent. The computational complexity is $O\left(\frac{|\mathcal{X}|}{\epsilon^2}\right)$ (Nesterov et al., 2018). To solve **P2**, we begin by introducing a lemma that we prove in Appendix D:

Lemma 4.3. *P2 has an optimal solution in which $p_i \in (\ell, u)$ for at most one $i \in \mathcal{Y}$.*

Algorithm 1 PROBFair

```

1: procedure PROBFair( $[N], k, \epsilon, \ell, u$ )
2:    $\mathcal{X} \leftarrow \{i \mid f_i \text{ is concave in all of } p_i \in [0, 1]\}$ 
3:    $\mathcal{Y} \leftarrow \{i \mid f_i \text{ is strictly convex in all of } p_i \in [0, 1]\}$ 
4:    $\text{grid\_search\_vals} \leftarrow \{\epsilon^j \mid j \in [0, k - \ell|\mathcal{Y}|\}\}$ 
5:   for  $z \in \text{grid\_search\_vals}$  do
6:      $\pi_{\mathcal{X}, z} \leftarrow \text{SOLVEP1}(\mathcal{X}, k, z, \ell, u)$ 
7:      $\pi_{\mathcal{Y}, z} \leftarrow \text{SOLVEP2}(\mathcal{Y}, k, z, \ell, u)$ 
8:      $\mathcal{V}_z \leftarrow \sum_{p_i \in \pi_{\mathcal{X}, z}} f_i(p_i) + \sum_{p_i \in \pi_{\mathcal{Y}, z}} f_i(p_i)$ 
9:    $z^* \leftarrow \arg \max_z \mathcal{V}_z$ 
10:   $\pi_{PF} \leftarrow \pi_{\mathcal{X}, z^*} \cup \pi_{\mathcal{Y}, z^*}$ 
11:  return  $\pi_{PF}$ 
    
```

Given this structure, an optimal solution $\{p_i^* \mid i \in \mathcal{Y}\}$ will set some number of arms $\gamma \in \mathbb{Z}^+$ to ℓ , at most one arm to $p' \in (\ell, u]$, and the remaining $|\mathcal{Y}| - \gamma - 1$ arms to u . We represent these subsets by $\mathcal{Y}_1, \mathcal{Y}_2$, and \mathcal{Y}_3 , respectively. Let $\gamma = \left\lfloor \frac{|\mathcal{Y}|u - (k - z)}{u - \ell} \right\rfloor$, and $p' = k - z - |\mathcal{Y}_1|\ell - |\mathcal{Y}_3|u \in (\ell, u]$. With the cardinality of each subset thus established, per Theorem 4.4 (see below), we use Algorithm 2 to optimally partition the arms in \mathcal{Y} .

Algorithm 2 SOLVEP2

Note: all sorts are ascending; arrays are zero-indexed.

```

1: procedure SOLVEP2( $\mathcal{Y} \subseteq N, k, z, \ell, u$ )
2:    $\gamma \leftarrow \left\lfloor \frac{|\mathcal{Y}|u - (k - z)}{u - \ell} \right\rfloor$ 
3:    $p' \leftarrow k - z - \gamma\ell - (|\mathcal{Y}| - 1 - \gamma)u$ 
4:   if  $|\mathcal{Y}| - \gamma - 1 > 0$  then
5:      $\Delta_{\mathcal{Y}} = \text{sort}(\{f_i(u) - f_i(\ell) \mid i \in \mathcal{Y}\})$ 
6:      $\mathcal{Y}_3 \leftarrow \{(\Delta_{\mathcal{Y}})[(|\mathcal{Y}| - \gamma - 1) : ]\}$ 
7:   else  $\mathcal{Y}_3 \leftarrow \emptyset$ 
8:    $\Delta_{\mathcal{Y} \setminus \mathcal{Y}_3} = \text{sort}(\{f_i(p') - f_i(\ell) \mid i \in \mathcal{Y} \setminus \mathcal{Y}_3\})$ 
9:    $\mathcal{Y}_1 \leftarrow \{(\Delta_{\mathcal{Y} \setminus \mathcal{Y}_3})[ : \gamma]\}$ 
10:   $\mathcal{Y}_2 \leftarrow \{(\Delta_{\mathcal{Y} \setminus \mathcal{Y}_3})[\gamma]\}$ 
11:   $\pi_{\mathcal{Y}} := i \mapsto \ell$  for  $i \in \mathcal{Y}_1$ ;  $p'$ , for  $i \in \mathcal{Y}_2$ ;  $u$ , for  $i \in \mathcal{Y}_3$ 
12:  return  $\pi_{\mathcal{Y}}$ 
    
```

Theorem 4.4. *Alg. 2 yields the mapping from arms in \mathcal{Y} to subsets in $\{\mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Y}_3\}$ which maximizes $\sum_{i \in \mathcal{Y}} f_i(p_i)$ s.t. $\sum_{i \in \mathcal{Y}} p_i = k - z$. (See Appendix D for the complete proof).*

Corollary 4.5. *Alg. 2 has time complexity $O(|\mathcal{Y}| \log |\mathcal{Y}|)$.*

With our solutions to **P1** and **P2** so defined, the cost of finding our in this way is $O\left(\frac{(k - \ell|\mathcal{Y}|)}{\epsilon} \left(\frac{|\mathcal{X}|}{\epsilon^2} + |\mathcal{Y}| \log |\mathcal{Y}|\right)\right)$, which is at worst $O\left(\frac{kN}{\epsilon^3}\right)$ when all N arms are in \mathcal{X} .

4.2. Sampling Approach

We use a linear-time algorithm introduced by Srinivasan (2001) and detailed in Appendix D.2 to sample from π_{PF} at each timestep, such that: (1) with probability one, we satisfy the budget constraint by pulling exactly k arms; and (2) any given arm i is pulled with probability p_i . Formally, each time we draw a vector of binary random variables $(X_1, X_2 \dots X_N) \sim \pi_{PF}$, $\Pr[|i : X_i = 1| = k] = 1$ and $\forall i, \Pr[X_i = 1] = p_i$.

5. Empirical Evaluation and Future Work

We present **empirical results** in Appendix E. Next steps include: (1) extending PROBFair for larger state or action spaces; (2) considering the online version of this problem where transitions are unknown, and (3) relaxing the stationarity condition in the construction of π_{PF} .

Acknowledgements

Christine Herlihy was supported by the National Institute of Standards and Technology’s (NIST) Professional Research Experience Program (PREP). John Dickerson and Aviva Prins were supported in part by NSF CAREER Award IIS-1846237, NSF D-ISN Award #2039862, NSF Award CCF-1852352, NIH R01 Award NLM-013039-01, NIST MSE Award #20126334, DARPA GARD #HR00112020007, DoD WHS Award #HQ003420F0035, and a Google Faculty Research award. Aravind Srinivasan was supported in part by NSF awards CCF-1749864 and CCF-1918749, as well as research awards from Adobe, Amazon, and Google. The views and conclusions contained in this publication are those of the authors and should not be interpreted as representing official policies or endorsements of U.S. government or funding agencies. We thank Samuel Dooley, Dr. Furong Huang, Naveen Raman, and Daniel Smolyak for helpful input and feedback.

References

- Askland, K., Wright, L., Wozniak, D. R., Emmanuel, T., Caston, J., and Smith, I. Educational, Supportive and Behavioural Interventions to Improve Usage of Continuous Positive Airway Pressure Machines in Adults with Obstructive Sleep Apnoea. *Cochrane Database of Systematic Reviews*, (4), 2020.
- Ayer, T., Zhang, C., Bonifonte, A., Spaulding, A. C., and Chhatwal, J. Prioritizing Hepatitis C Treatment in US Prisons. *Operations Research*, 67(3):853–873, 2019.
- Bentham, J. An Introduction to the Principles of Morals and Legislation. Technical report, McMaster University Archive for the History of Economic Thought, 1781.
- Bertsimas, D., Farias, V. F., and Trichakis, N. The Price of Fairness. *Operations research*, 59(1):17–31, 2011.
- Chen, Y., Cuellar, A., Luo, H., Modi, J., Nemlekar, H., and Nikolaidis, S. Fair Contextual Multi-Armed Bandits: Theory and Experiments. volume 124 of *Proceedings of Machine Learning Research*, pp. 181–190, Virtual, 03–06 Aug 2020. PMLR. URL <http://proceedings.mlr.press/v124/chen20a.html>.
- De-Arteaga, M., Fogliato, R., and Chouldechova, A. A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2020.
- Dietvorst, B. J., Simmons, J. P., and Massey, C. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1):114, 2015.
- Ghouila-Houri, A. Caractérisation des matrices totalement unimodulaires. *Comptes Rendus Hebdomadaires des Séances de l’Académie des Sciences (Paris)*, 254: 1192–1194, 1962.
- Hirschman, A. O. *National Power and the Structure of Foreign Trade*, volume 105. Univ of California Press, 1980.
- Joseph, M., Kearns, M., Morgenstern, J. H., and Roth, A. Fairness in Learning: Classic and Contextual Bandits. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 325–333. Curran Associates, Inc., 2016. URL <http://papers.nips.cc/paper/6355-fairness-in-learning-classic-and-contextual-bandits.pdf>.
- Jung, Y. H. and Tewari, A. Regret Bounds for Thompson Sampling in Episodic Restless Bandit Problems. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 9007–9016. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9102-regret-bounds-for-thompson-sampling-in-episodic-restless-bandit-problems.pdf>.
- Jung, Y. H., Abeille, M., and Tewari, A. Thompson sampling in non-episodic restless bandits. *CoRR*, abs/1910.05654, 2019. URL <http://arxiv.org/abs/1910.05654>.
- Kang, Y., Prabhu, V. V., Sawyer, A. M., and Griffin, P. M. Markov models for treatment adherence in obstructive sleep apnea. *Age*, 49:11–6, 2013.
- Kang, Y., Sawyer, A. M., Griffin, P. M., and Prabhu, V. V. Modelling Adherence Behaviour for the Treatment of Obstructive Sleep Apnoea. *European journal of operational research*, 249(3):1005–1013, 2016.
- Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., and King, D. Key Challenges for Delivering Clinical Impact with Artificial Intelligence. *BMC medicine*, 17(1): 195, 2019.

- Killian, J. A., Biswas, A., Shah, S., and Tambe, M. Q-Learning Lagrange Policies for Multi-Action Restless Bandits. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD)*, pp. 871–881, 2021.
- Li, F., Liu, J., and Ji, B. Combinatorial Sleeping Bandits with Fairness Constraints. *CoRR*, abs/1901.04891, 2019. URL <http://arxiv.org/abs/1901.04891>.
- Liu, K. and Zhao, Q. Indexability of Restless Bandit Problems and Optimality of Whittle Index for Dynamic Multichannel Access. *IEEE Transactions on Information Theory*, 56(11):5547–5567, Nov 2010. ISSN 1557-9654. doi: 10.1109/tit.2010.2068950. URL <http://dx.doi.org/10.1109/TIT.2010.2068950>.
- Marseille, E. and Kahn, J. G. Utilitarianism and the Ethical Foundations of Cost-Effectiveness Analysis in Resource Allocation for Global Health. *Philosophy, Ethics, and Humanities in Medicine*, 14(1):1–7, 2019. doi: 10.1186/s13010-019-0074-7.
- Mate, A., Killian, J., Xu, H., Perrault, A., and Tambe, M. Collapsing Bandits and Their Application to Public Health Intervention. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 2020.
- Mate, A., Perrault, A., and Tambe, M. Risk-Aware Interventions in Public Health: Planning with Restless Multi-Armed Bandits. In *20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, London, UK, 2021.
- Nesterov, Y. et al. *Lectures on convex optimization*, volume 137. Springer, 2018.
- Niño-Mora, J. A Verification Theorem for Threshold-Indexability of Real-State Discounted Restless Bandits. *Mathematics of Operations Research*, 45(2):465–496, 2020.
- Ortner, R., Ryabko, D., Auer, P., and Munos, R. Regret bounds for restless markov bandits. In *International conference on algorithmic learning theory*, pp. 214–228. Springer, 2012.
- Ou, H.-C., Siebenbrunner, C., Killian, J., Brooks, M. B., Kempe, D., Vorobeychik, Y., and Tambe, M. Networked restless multi-armed bandits for mobile interventions. *arXiv preprint arXiv:2201.12408*, 2022.
- Papadimitriou, C. H. and Tsitsiklis, J. N. The Complexity of Optimal Queueing Network Control. In *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory*, pp. 318–322. IEEE, 1994.
- Prins, A., Mate, A., Killian, J. A., Abebe, R., and Tambe, M. Incorporating Healthcare Motivated Constraints in Restless Bandit Based Resource Allocation. *preprint*, 2020. URL <https://teamcore.seas.harvard.edu/publications/incorporating-healthcare-motivated-constraints-restless-bandit-based-resource>.
- Punjabi, N. M. The epidemiology of adult obstructive sleep apnea. *Proceedings of the American Thoracic Society*, 5(2):136–143, 2008.
- Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., and Chin, M. H. Ensuring Fairness in Machine Learning to Advance Health Equity. *Annals of internal medicine*, 169(12):866–872, 2018.
- Rawls, J. *A Theory of Justice*. Belknap Press of Harvard University Press, Cambridge, Massachusetts, 1 edition, 1971. ISBN 0-674-88014-5.
- Rhoades, S. A. The Herfindahl-Hirschman Index. *Fed. Res. Bull.*, 79:188, 1993.
- Rotenberg, B. W., Murariu, D., and Pang, K. P. Trends in CPAP adherence over twenty years of data collection: a flattened curve. *Journal of Otolaryngology-Head & Neck Surgery*, 45(1):1–9, 2016.
- Rubner, Y., Tomasi, C., and Guibas, L. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40:99–121, 11 2000. doi: 10.1023/A:1026543900054.
- Sawyer, A. M., Gooneratne, N. S., Marcus, C. L., Ofer, D., Richards, K. C., and Weaver, T. E. A systematic review of CPAP adherence across age groups: clinical and empiric insights for developing CPAP adherence interventions. *Sleep medicine reviews*, 15(6):343–356, 2011.
- Scheunemann, L. and White, D. The Ethics and Reality of Rationing in Medicine. *Chest*, 140:1625–32, 12 2011. doi: 10.1378/chest.11-0622.
- Srinivasan, A. Distributions on Level-Sets with Applications to Approximation Algorithms. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, FOCS ’01*, pp. 588, USA, 2001. IEEE Computer Society. ISBN 0769513905.
- Steimle, L. N. and Denton, B. T. *Markov Decision Processes for Screening and Treatment of Chronic Diseases*, pp. 189–222. Springer International Publishing, Cham, 2017. ISBN 978-3-319-47766-4. doi: 10.1007/978-3-319-47766-4_6. URL https://doi.org/10.1007/978-3-319-47766-4_6.

Weber, R. R. and Weiss, G. On an Index Policy for Restless Bandits. *Journal of Applied Probability*, pp. 637–648, 1990.

Whittle, P. Restless Bandits: Activity Allocation in a Changing World. *Journal of applied probability*, 25(A):287–298, 1988.

A. Notation

In Table 1, we present an overview of the notation used in the paper. $[N]$ denotes the set $\{1, 2, \dots, N\}$.

Table 1: Notation used in our model and notes on their interpretation.

MDP Variables Here, timestep $t \in [T] = \{1, 2, \dots, T\}$ (subscript) and arm index $i \in [N] = \{1, 2, \dots, N\}$ (superscript) are implied.		
State space	$s \in \mathcal{S} = \{0, 1\}$	$s = \begin{cases} 1 & \text{arm is in the 'good' state.} \\ 0 & \text{else} \end{cases}$
Belief space	$b \in \mathcal{B} = [0, 1]$ $b_{t+1} = \begin{cases} s_{t+1} & \text{if known} \\ b_t P_{1,1}^0 + (1 - b_t) P_{0,1}^0 & \text{else} \end{cases}$	If an arm's true state is unknown, the recursively computed belief state approximates it.
Action space	$a \in \mathcal{A} = \{0, 1\}$	$a = \begin{cases} 1 & \text{pull arm (i.e., provide intervention)} \\ 0 & \text{else, don't pull} \end{cases}$
MDP Functions		
Transition function	$P: \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ $s_t, a_t, s_{t+1} \mapsto \Pr(s_{t+1} \mid s_t, a_t)$	The probability of an arm going from state s_t to s_{t+1} , given action a_t . Equivalent (matrix) notation: $P_{s_t, s_{t+1}}^{a_t}$.
Reward function	$r: \mathcal{S} \text{ or } \mathcal{B} \rightarrow \mathbb{R}$	$r(b)$ is used in computing the Whittle index.
Policy function	$\pi: \mathcal{S} \rightarrow \mathcal{A}$	A policy for actions. The set of optimal policies is $\pi^* \in \Pi^*$.
RMAB Variables		
Timestep	$\{t \in \mathbb{N} \mid t \leq T\}$	This timestep is implicit in the MDP.
Arm index	$i \in \{1, 2, \dots, N\}$	Each arm can represent a patient. k arms can be pulled at any timestep t .
Objective Functions The objective is to find a policy $\pi^* = \max_{\pi} \mathbb{E}_{\pi}[R(\cdot)]$.		
Discounted reward function	$R_{\beta}^{\pi}: \mathcal{S}^N \rightarrow \mathbb{R}$ $s_0^1, s_0^2, \dots, s_0^N \mapsto \sum_{i \in [N]} \sum_{t \in [T]} \beta^{t-1} r(s_t^i)$	$\beta \in [0, 1]$ is some <i>discount parameter</i> .
Fairness-motivated Constraint Functions		
Integer periodicity	$\bigwedge_{j=0}^{\lceil \frac{T}{\nu} \rceil} \sum_{t=j\nu+1}^{(j+1)\nu} a_t^i \geq 1$	A form of time-indexed fairness. Guarantees arm i is pulled at least once within each period of ν timesteps.
Minimum selection fraction	$\bigwedge_{i \in [N]} \frac{1}{T} \sum_{t=1}^T a_t^i \geq \psi$	A form of time-indexed fairness. Arm i should be pulled at least some minimum fraction $\psi \in (0, 1)$ of timesteps.
Probabilistic	$\bigwedge_{i \in [N]} \bigwedge_{t \in [T]} \Pr(a_t^i = 1 \mid i, t) \in [\ell, u]$	Pull each arm with probability $p_i \in [\ell, u]$, where $\ell \in (0, \frac{k}{N}]$ and $u \in [\frac{k}{N}, 1]$.

B. Empirical Inequity in the Distribution of Actions under Whittle Index Policies

Here, we present numerical results confirming Prins et al. (2020)’s findings that THRESHOLD WHITTLE tends to allocate pulls according to a bimodal distribution: a small subset of arms are pulled frequently, while others are largely ignored.

This bimodal distribution is a consequence of how the Whittle index prioritizes arms. Whittle favors arms for whom a pull is most beneficial to achieving sustained occupancy in the “good” state, regardless of whether this results in the same subset of arms repeatedly receiving pulls. While the structural constraints in Sec. 2 ensure that a pull is beneficial for every arm, marginal benefit varies. Since reward is a function of each arm’s underlying state, arms whose trajectories are characterized by a relative—but *not absolute*—indifference to receipt of the intervention are likely to be ignored.

Experimental Setup: For each iteration, we generate $N = 2$ forward threshold-optimal arms and run THRESHOLD WHITTLE for a $T = 365$ horizon simulation, where the budget constraint $k = 1$. We run 1,000 such iterations.

Results: In 515 out of 1,000 (51.5%) simulations, the arms’ Whittle indices never overlap, meaning that for *any* combination of initial states, state transitions, and pulls, THRESHOLD WHITTLE would pull one arm for all timesteps $t \in T$ and completely ignore the second arm.

We visualize one such case in Figure 1. Recall that THRESHOLD WHITTLE precomputes the infimum subsidy m per arm and belief combination. Since belief is a function of last known state $s \in \{0, 1\}$ and time-since-seen $u \in [T]$ (using the notation of Mate et al. (2020)), we plot the infimum subsidy of each arm-state combination with time-since-seen, u , on the x -axis. There exists a horizontal line that divides the two arms, so arm $i = 1$ will be pulled for every timestep $t \in T$ and arm $i = 2$ will never be pulled.

In order to modify the Whittle index to guarantee time-indexed fairness constraint satisfaction, one would need to ensure that no such horizontal line exists. Additionally, if we consider a specific form of time-indexed fairness known as an *integer periodicity* constraint, which allows a decision-maker to guarantee that arm i is pulled at least once within each period of ν days, the lines associated with the arms in Figure 1 must cross before ν timesteps elapse to guarantee fairness constraint satisfaction.

Another perspective we can take is to ask: what’s the smallest interval ν_i for each arm i we could have specified such that THRESHOLD WHITTLE would have satisfied the integer periodicity constraint? Note that this is retrospective, as there is no way to enforce this constraint at planning time.

We visualize the minimum such ν_i in Figure 2. On the far right, we see the 515 cases where (without loss of generality) the second arm is never pulled—that is, the minimum ν_i such that THRESHOLD WHITTLE satisfies the hard integer periodicity constraint must be *larger* than the horizon, $T = 365$. There is one case where arm $i = 2$ is pulled exactly once. In a majority of the remaining simulations, THRESHOLD WHITTLE pulls each arm with approximately equal frequency.




Figure 1: The Whittle index values for Arm 1 and 2 can be separated by a horizontal line, meaning that (WLOG) Arm 1 will always be chosen over Arm 2 because its index value dominates.

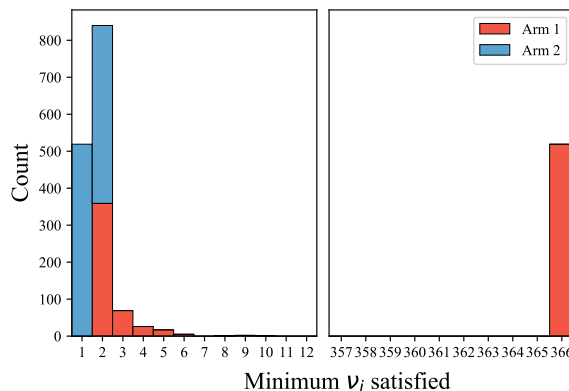


Figure 2: The smallest interval ν_i such that THRESHOLD WHITTLE satisfies an integer periodicity definition of time-indexed fairness, given $N = 2$ random arms. In over 50% of iterations, no such fairness constraint satisfaction is possible (i.e., $\exists i$ s.t. $\nu_i > T$).

C. Intractability of Alternative Approaches

In this section, we motivate the algorithmic design choices we have made when constructing PROBFair by discussing the feasibility of possible alternatives, including: (1) direct modification of the Whittle index to guarantee time-indexed fairness constraint satisfaction, and (2) a math programming-based approach.

C.1. Why not modify the Whittle index to guarantee time-indexed fairness constraint satisfaction?

In Section 3.3, we demonstrate that it is not possible to guarantee time-indexed fairness when arms are decoupled. If arms cannot be decoupled, the tractability of a Whittle index-based approach breaks down. Here, we discuss this topic in greater detail, and provide specific examples of possible Whittle index modifications. We also take this opportunity to emphasize that our focus is on *guaranteeing* fairness rather than incentivizing it. Mate et al. (2021) provide an example of the latter, which we discuss and include as a comparison algorithm in Section E.2.

To begin, recall that the efficiency of Whittle index-based policies stems from our ability to decouple arms when we are only concerned with maximizing total expected reward (Whittle, 1988; Weber & Weiss, 1990). However, guaranteeing time-indexed fairness (as defined in Section 2) in the planning setting requires time-stamped record keeping. It is no longer sufficient to compute each arm’s infimum subsidy in isolation and order the resulting set of values. Instead, for an optimal index policy to be efficiently computable, it must be possible to modify the value function (Equation 2) so as to ensure that the infimum subsidy each arm would require in the absence of fairness constraints is minimally perturbed via augmentation or “donation”, so as to maximize total expected reward while ensuring its own fairness constraint satisfaction or the constraint satisfaction of other arms, respectively, *without* requiring input from other arms.

Plausible modifications include altering the conditions under which an arm receives the subsidy associated with passivity, m , or introducing a modified reward function, $r'(b)$ that is capable of accounting for an arm’s fairness constraint satisfaction status in addition to its state at time t . For example, we might use an indicator function to “turn off” the subsidy until arm i has been pulled at least once within the interval in question, or increase reward as an arm’s time-since-pulled value approaches the interval cut-off, so as to incentivize a constraint-satisfying pull. When these modifications are viewed from the perspective of a single arm, they *appear* to have the desired effect: if no subsidy is received, it will be optimal to pull for all belief states; similarly, for a fixed m , as reward increases it will be optimal to pull for an increasingly large subset of the belief state space.

Recall, however, that structural constraints ensure that when an arm is considered in isolation, the optimal action will *always* be to pull. Whether or not arm i is *actually* pulled at time t depends on how the infimum subsidy, m , it requires to accept passivity at time t compares to the infimum subsidies required by other arms. Thus, any modification intended to *guarantee* time-indexed fairness constraint satisfaction must be able to alter the ordering *among* arms, such that any arm i which would otherwise have a subsidy with rank $> k$ when sorted in descending order will now be in the top- k arms. Even if we were able to construct such a modification for a single arm without requiring system state, if *every* arm had this same capability, then a new challenge would arise: we would be unable to distinguish among arms, and arbitrary tie-breaking could again jeopardize fairness constraint satisfaction.

If it is not possible to decouple arms, then we must consider them in tandem. Papadimitriou & Tsitsiklis (1994) prove that the RMAB problem is PSPACE-complete even when transition rules are action-dependent but deterministic, via reduction from the halting problem.

C.2. Why not use a math-programming approach?

Our constrained maximization problem can be readily formulated as an integer program (IP) with a totally unimodular (TU) constraint matrix. However, this approach is intractable because the objective function coefficients of this IP cannot be efficiently enumerated. To support this intractability claim, we begin by presenting an integer program (IP) that maximizes total expected reward under both budget *and* time-indexed fairness constraints, for problem instances with feasible hyperparameters. We then prove that any problem instance with feasible hyperparameters yields a totally unimodular (TU) constraint matrix, which ensures that the linear program (LP) relaxation of our IP will yield an integral solution. We proceed to demonstrate that tractability issues arise because we incur an exponential dependency on the time horizon, T , when we construct the IP’s objective function coefficients. We conclude by comparing PROBFair to the IP for small values of N and T .

C.2.1. INTEGER PROGRAM FORMULATION

To leverage a math programming approach for our constrained reward maximization task, we seek to construct an integer program (IP) whose solution is the policy $\vec{x} \in \{0, 1\}^{N|\mathcal{A}|T}$. We require this policy to be reward-maximizing, subject to the guaranteed satisfaction of both budget and time-indexed fairness constraints. To begin, let each decision variable $x_{i,a,t} \in \{0, 1\}$ represent whether or not we take action $a \in \mathcal{A} = \{0, 1\}$ for arm $i \in [N]$ at time $t \in [T]$. Then, let each objective function coefficient $c_{i,a,t}$ represent the expected reward associated with an arm-action-timestep combination.

To formalize the objective function, recall that the agent seeks to maximize total expected reward, $\mathbb{E}_\pi[R(\cdot)]$. For clarity of exposition, we specifically consider the linear global reward function $R(r(s)) = \sum_{i=1}^N \sum_{t=0}^T s_t^i$. Note that this implies the discount rate, $\beta = 1$; however, the approach outlined here can be extended in a straightforward manner for $\beta \in (0, 1)$. In order to compute the expected reward associated with taking action a for arm i at time t , we must consider: (1) what state is the arm currently in (i.e., what is the realized value of $s_t^i \in \{0, 1\}$)? (2) when the arm transitions from s_t to s_{t+1} by virtue of taking action a , what reward, $r(\cdot)$, should we expect to earn?

Because we define $r(s) = s$, (2) can be reframed as: what is the probability $p(s_{t+1} = 1 | s_t^i, a_t^i)$ that action a causes a transition from s_t to the adherent state? Because each arm's state at time t is stochastic, depending not only on the sequence of actions taken in previous timesteps, but the associated set of stochastic transitions informed by the arm's underlying MDP, each coefficient of our objective function must be computed as the expectation over the possible values of $s_t \in \mathcal{S}$:

$$\vec{c} = \mathbb{E}_s[p(s_{t+1} = 1 | x_{i,a,t}, s_t)] \quad \forall i, a, t \in [N] \times \mathcal{A} \times [T] \quad (6)$$

$$= \frac{1}{2^t} \sum_{s \in \mathcal{S}} p(s_t = s) \sum_{s' \in \mathcal{S}} p(s_{t+1} = s' | x_{i,a,t}, s_t = s) r(s') \quad \forall i, a, t \in [N] \times \mathcal{A} \times [T] \quad (7)$$

$$= \frac{1}{2^t} \sum_{s \in \mathcal{S}} p(s_t = s) p(r(s_{t+1}) = 1 | x_{i,a,t}, s_t = s) \quad \forall i, a, t \in [N] \times \mathcal{A} \times [T] \quad (8)$$

$$= \frac{1}{2^t} \sum_{s \in \mathcal{S}} p(s_t = s) p(s_{t+1} = 1 | x_{i,a,t}, s_t = s) \quad \forall i, a, t \in [N] \times \mathcal{A} \times [T] \quad (9)$$

Within the context of this IP, the time-indexed fairness constraint we introduce in Section 2 can be more specifically defined as either an *integer periodicity* or *minimum selection fraction* constraint. We formalize each of these below:

The **integer periodicity constraint** allows a decision-maker to guarantee that arm i is pulled at least once within each period of ν days. We define this constraint as a function g , over the vector of actions, \vec{a}^i associated with arm i , and user-defined interval length $\nu \in [1, T]$:

$$g(\vec{a}^i) = \sum_{t=j\nu+1}^{(j+1)\nu} a_t^i \geq 1 \quad (10)$$

$$\forall j \in \left\{ 0, 1, 2, \dots, \left\lceil \frac{T}{\nu} \right\rceil \right\}; \forall i \in \{1, 2, \dots, N\}$$

The **minimum selection fraction constraint** introduced by Li et al. (2019) forces the agent to pull arm i at least a minimum fraction, $\psi \in (0, 1)$, of the total number of steps, but is agnostic to how these pulls are distributed over time. We define this constraint, g' , as a function over the vector of actions, \vec{a}^i associated with arm i and user-defined ψ :

$$g'(\vec{a}^i) = \frac{1}{T} \sum_{t=1}^T a_t^i \geq \psi \quad \forall i \in \{1, 2, \dots, N\} \quad (11)$$

The resulting integer program is given by:

$$\begin{aligned}
 \max \quad & c^T x & (12) \\
 \text{s.t.} \quad & \sum_{a=1}^{|\mathcal{A}|} x_{i,a,t} = 1 & \forall i \in [N], t \in [T] \quad \text{(a) Select exactly one action per arm } i \text{ at each } t \\
 \text{if int. per:} \quad & \sum_{t \in I_j} x_{i,1,t} \geq 1 & \forall j \in \left\{ 0, 1, \dots, \frac{T-\nu}{\nu} \right\} \quad \text{(b.i) Pull each arm } i \text{ at least once during each interval of length } \nu \\
 \text{if min. sel:} \quad & \frac{1}{T} \sum_{t=1}^T x_{i,1,t} \geq \psi & \psi \in (0, 1), \forall i \in N \quad \text{(b.ii) Pull each arm } i \text{ at least a minimum fraction } \psi \text{ of } T \text{ rounds} \\
 & \sum_{i=1}^N x_{i,1,t} = k & \forall t \in [T] \quad \text{(c) Pull exactly } k \text{ arms at each } t \\
 & x_{i,a,t} \in \{0, 1\} & \forall i \in [N], a \in \mathcal{A}, t \in [T] \quad \text{(d) Each arm-action-timestep choice is a binary decision variable}
 \end{aligned}$$

C.2.2. LP RELAXATION AND INTEGRALITY OF SOLUTION

We now prove that the IP we have formulated in Section C.2.1 has an attractive property: namely, any feasible problem instance will produce a totally unimodular constraint matrix. Our proof leverages a theorem introduced by Ghouila-Houri and restated below for convenience, which can be used to determine whether a matrix, $A \in \mathbb{R}^{m \times n}$ is totally unimodular:

Lemma C.1. (*Ghouila-Houri (1962)*) *A matrix $A \in \mathbb{Z}^{m \times n}$ is totally unimodular (TU) if and only if for every subset of the rows $R \subseteq [m]$, there is a partition $R = R_1 \cup R_2$ such that for every $j \in [n]$,*

$$\sum_{i \in R_1} A_{ij} - \sum_{i \in R_2} A_{ij} \in \{-1, 0, 1\} \quad (13)$$

Theorem C.2. *Within the context of the integer program outlined in Appendix C.2.1, any feasible problem instance will produce a constraint matrix that is totally unimodular (TU).*

Proof. To begin, we establish the dimensions of any such constraint matrix \mathbf{A} and note the maximum possible column-wise sum that each of its component submatrices may contribute. Note that the minimum selection fraction constraint (b.ii in Equation 12), which requires the agent to pull each arm i at least a minimum fraction, $\psi \in (0, 1)$, of T rounds, can be thought of as a special case of the integer periodicity constraint, (b.i in Equation 12), where $\nu = T$ and each arm must be pulled at least $\lceil T\psi \rceil$ times. As such, we assume that at most one of the time-indexed fairness constraints can be selected, and focus on the more general of the two, which is the integer periodicity constraint. For notational convenience, we refer to constraints by their alphabetic identifiers. Let (b) represent the integer periodicity constraint, and define a function $\varphi : r \in \mathbf{R} \subseteq \mathbf{A} \mapsto e \in \{a, b, c\}$ that maps each row to its corresponding constraint type.

First, recall that each $x_{i,a,t}$ represents a single binary decision variable, and corresponds to a column in \mathbf{A} . There are $N \times |\mathcal{A}| \times T$ such columns. Next, note that constraint (a) enforces the requirement that we select *exactly* one action per arm per timestep. Formally, $\forall i, t \in N \times T, \exists! a \in \mathcal{A} \text{ s.t. } x_{i,a,t} = 1$. Correspondingly, $\forall a' \in \mathcal{A} \setminus a, x_{i,a',t} = 0$. The column vectors of the associated sub-matrix, $\mathbf{A}_a \in \mathbb{Z}^{N \times T \times |\mathcal{A}|}$, are indexed by disjoint $(i, a, t) \in N \times |\mathcal{A}| \times T$; thus, each column vector contains a single non-zero entry and for $\mathbf{R}_a \subseteq \mathbf{A}_a$, taking the column-wise sum will yield a vector $\vec{v} \in \mathbb{Z}^{N \times |\mathcal{A}| \times T}$ with every entry equal to 1.

In a similar vein, equity constraint (b) enforces the requirement that we must pull each arm, i at least once during each interval I_j of length ν_i . Within the associated sub-matrix, $\mathbf{A}_b \in \mathbb{Z}^{N \times \lceil \frac{T}{\nu_i} \rceil \times N \times |\mathcal{A}| \times T}$, each column that corresponds to a passive action (e.g., $x_{i,a=0,t}$) will have *only* zero-valued entries, since passive action decision variables are not impacted by constraint (b). Conversely, each column that corresponds to an active action (e.g., $x_{i,a=1,t}$) will have a single non-zero entry. Each active action column corresponding to a specific arm-timestep can be mapped to exactly one interval. Thus, for $\mathbf{R}_b \subseteq \mathbf{A}_b$, taking the column-wise sum will yield a vector $\vec{v} \in \mathbb{Z}^{N \times |\mathcal{A}| \times T}$ with every entry taking a value $\in \{0, 1\}$.

The budget constraint (c) enforces the requirement that we must pull exactly k of the N arms at each timestep. Much like equity constraint (b), only columns corresponding to active actions are impacted. Thus, within the associated sub-matrix, $\mathbf{A}_c \in \mathbb{Z}^{T \times N|\mathcal{A}|T}$, each column that corresponds to a passive action (e.g., $x_{i,a=0,t}$) will have *only* zero-valued entries, while each column that corresponds to an active action (e.g., $x_{i,a=1,t}$) can be mapped to a single timestep, and will have a single non-zero entry. Thus, for $\mathbf{R}_c \subseteq \mathbf{A}_c$, taking the column-wise sum also yield a vector $\vec{v} \in \mathbb{Z}^{N|\mathcal{A}|T}$ with every entry taking a value $\in \{0, 1\}$.

The complete constraint matrix \mathbf{A} thus contains $NT + N\lceil \frac{T}{\nu_i} \rceil + T$ rows. Three possible cases arise when we consider every subset of these rows: (1) $\mathbf{R} \subsetneq \mathbf{A} = \emptyset$; (2) $\mathbf{R} \subsetneq \mathbf{A}$; $\mathbf{R} \cap \mathbf{A} \neq \emptyset$; (3) $\mathbf{R} \subseteq \mathbf{A}$.

Case 1. $\mathbf{R} \subsetneq \mathbf{A} = \emptyset$. To satisfy Lemma C.1, partition \mathbf{R} such that $\mathbf{R} = \mathbf{R}_1 \cup \mathbf{R}_2 = \emptyset \cup \emptyset$. Then, for every $j \in [n]$,

$$\sum_{i \in \mathbf{R}_1} \mathbf{A}_{ij} - \sum_{i \in \mathbf{R}_2} \mathbf{A}_{ij} = 0 - 0; 0 \in \{-1, 0, 1\} \square$$

Case 2. $\mathbf{R} \subsetneq \mathbf{A}$; $\mathbf{R} \cap \mathbf{A} \neq \emptyset$. If we consider $\cup_{r \in \mathbf{R}} \varphi(r)$, there are $\sum_{k=1}^3 \binom{3}{k}$ possible sets of observed constraint types: $\{a\} \vee \{b\} \vee \{c\} \vee \{a, b\} \vee \{a, c\} \vee \{b, c\} \vee \{a, b, c\}$.

1. If $|\cup_{r \in \mathbf{R}} \varphi(r)| = 1$, then any partition of \mathbf{R} will satisfy Lemma C.1. Without loss of generality, let each row $r \in \mathbf{R}$ belong to \mathbf{R}_1 and $\mathbf{R}_2 = \emptyset$.

i. If $\cup_{r \in \mathbf{R}} \varphi(r) = \{a\}$, taking the column-wise sum of \mathbf{R}_1 will yield a vector $\vec{v} \in \mathbb{Z}^{N|\mathcal{A}|T}$ with every entry $\in \{1\}$ if $\mathbf{R} \subseteq \mathbf{A}_a$, and $\in \{0, 1\}$ otherwise. Thus, $\forall j \in [N|\mathcal{A}|T]$,

$$\sum_{i \in \mathbf{R}_1} \mathbf{A}_{ij} - \sum_{i \in \mathbf{R}_2} \mathbf{A}_{ij} = 0 - 0 \vee 1 - 0; \{0, 1\} \subsetneq \{-1, 0, 1\} \square$$

ii. If $\cup_{r \in \mathbf{R}} \varphi(r) = \{b\} \vee \{c\}$, taking the column-wise sum of \mathbf{R}_1 will yield a vector $\vec{v} \in \mathbb{Z}^{N|\mathcal{A}|T}$ with every entry corresponding to a passive action column $\in \{0\}$ and every entry corresponding to an active action column $\in \{1\}$ if $\mathbf{R} \subseteq \mathbf{A}_{b \vee c}$, and $\in \{0, 1\}$ otherwise. Thus, $\forall j \in [N|\mathcal{A}|T]$,

$$\sum_{i \in \mathbf{R}_1} \mathbf{A}_{ij} - \sum_{i \in \mathbf{R}_2} \mathbf{A}_{ij} = 0 - 0 \vee 1 - 0; \{0, 1\} \subsetneq \{-1, 0, 1\} \square$$

2. If $|\cup_{r \in \mathbf{R}} \varphi(r)| = 2$, without loss of generality, partition as follows: sort the elements $\in \cup_{r \in \mathbf{R}} \varphi(r)$ lexicographically, and let $\mathbf{R}_1 = \{r | \varphi(r) = \min \cup_{r \in \mathbf{R}} \varphi(r)\}$ and $\mathbf{R}_2 = \mathbf{R} \setminus \mathbf{R}_1$. Per Case 2.1 (i) and (ii), taking the column-wise sums of \mathbf{R}_1 and \mathbf{R}_2 will yield two vectors, $\vec{v}_1, \vec{v}_2 \in \mathbb{Z}^{N|\mathcal{A}|T}$, each of which will contain only entries $\in \{0, 1\}$. Thus, $\forall j \in [N|\mathcal{A}|T]$,

$$\sum_{i \in \mathbf{R}_1} \mathbf{A}_{ij} - \sum_{i \in \mathbf{R}_2} \mathbf{A}_{ij} = 0 - 0 \vee 0 - 1 \vee 1 - 0 \vee 1 - 1; \{-1, 0, 1\} \subseteq \{-1, 0, 1\} \square$$

3. If $|\cup_{r \in \mathbf{R}} \varphi(r)| = 3$, partition according to Algorithm 3:

Algorithm 3 Partition for Case 2.3

```

1: procedure PARTITION( $N, \mathcal{A}, T, \mathbf{R}$ )
2:    $C_{R_1}; C_{R_2} \leftarrow \{0\}^{N|\mathcal{A}|T}$  ▷ Initialize two sets of counters
3:    $\mathbf{R}_1; \mathbf{R}_2 \leftarrow \emptyset$ 
4:   for element  $e \in \{a, b, c\}$  do
5:      $\mathbf{R}_e \leftarrow \{r \mid \varphi(r) = e\}$  ▷ Use  $\varphi$  to partition the rows of  $\mathbf{R}$  by constraint type
6:     for  $r \in \mathbf{R}_a$  do
7:        $\mathbf{R}_1 \leftarrow \mathbf{R}_1 \cup r$  ▷ Let  $r \in \mathbf{R}_1$ 
8:       for  $i \in 0 : N|\mathcal{A}|T$  do
9:         if  $r[i] > 0$  then ▷ For each non-zero entry  $\in r$ 
10:           $C_{R_1}[i] \leftarrow C_{R_1}[i] + 1$  ▷ Increment corresponding  $C_{R_1}$  counter
11:       for  $r \in \mathbf{R}_b$  do
12:         flag  $\leftarrow$  false
13:         for  $i \in 0 : N|\mathcal{A}|T$  do
14:           if  $r[i] > 0 \wedge C_{R_1}[i] > 0$  then ▷ If any non-zero element  $r_i$  has  $C_{R_1} > 0$ 
15:             flag  $\leftarrow$  true ▷ Set flag to true
16:         if flag then
17:            $\mathbf{R}_2 \leftarrow \mathbf{R}_2 \cup r$  ▷ Let  $r \in \mathbf{R}_2$ 
18:           for  $i \in 0 : N|\mathcal{A}|T$  do
19:             if  $r[i] > 0$  then ▷ For each non-zero entry  $\in r$ 
20:                $C_{R_2}[i] \leftarrow C_{R_2}[i] + 1$  ▷ Increment corresponding  $C_{R_2}$  counter
21:         else
22:            $\mathbf{R}_1 \leftarrow \mathbf{R}_1 \cup r$  ▷ Let  $r \in \mathbf{R}_1$ 
23:           for  $i \in 0 : N|\mathcal{A}|T$  do
24:             if  $r[i] > 0$  then ▷ For each non-zero entry  $\in r$ 
25:                $C_{R_1}[i] \leftarrow C_{R_1}[i] + 1$  ▷ Increment corresponding  $C_{R_1}$  counter
26:           for  $r \in \mathbf{R}_c$  do
27:              $\mathbf{R}_2 \leftarrow \mathbf{R}_2 \cup r$  ▷ Let  $r \in \mathbf{R}_2$ 
28:             for  $i \in 0 : N|\mathcal{A}|T$  do
29:               if  $r[i] > 0$  then ▷ For each non-zero entry  $\in r$ 
30:                  $C_{R_2}[i] \leftarrow C_{R_2}[i] + 1$  ▷ Increment corresponding  $C_{R_2}$  counter
return  $\mathbf{R}_1; \mathbf{R}_2$ 

```

Taking the column-wise sums of the resulting \mathbf{R}_1 and \mathbf{R}_2 will yield two vectors $\vec{v}_1, \vec{v}_2 \in \mathbb{Z}^{N|\mathcal{A}|T}$, which can contain entries $\in \{0, 1\}$ and $\{0, 1, 2\}$, respectively. Note that since \vec{v}_2 is constructed by taking only rows with constraint types $\in \{b, c\}$, only entries corresponding to active action columns can take values > 1 . Moreover, $\forall j \in [N|\mathcal{A}|T], \sum_{i \in \mathbf{R}_2} \mathbf{A}_{ij} = 2 \rightarrow \sum_{i \in \mathbf{R}_1} \mathbf{A}_{ij} = 1$. Thus, $\forall j \in [N|\mathcal{A}|T]$,

$$\sum_{i \in \mathbf{R}_1} \mathbf{A}_{ij} - \sum_{i \in \mathbf{R}_2} \mathbf{A}_{ij} = 0 - 0 \vee 0 - 1 \vee 1 - 0 \vee 1 - 1 \vee 1 - 2; \{-1, 0, 1\} \subseteq \{-1, 0, 1\} \square$$

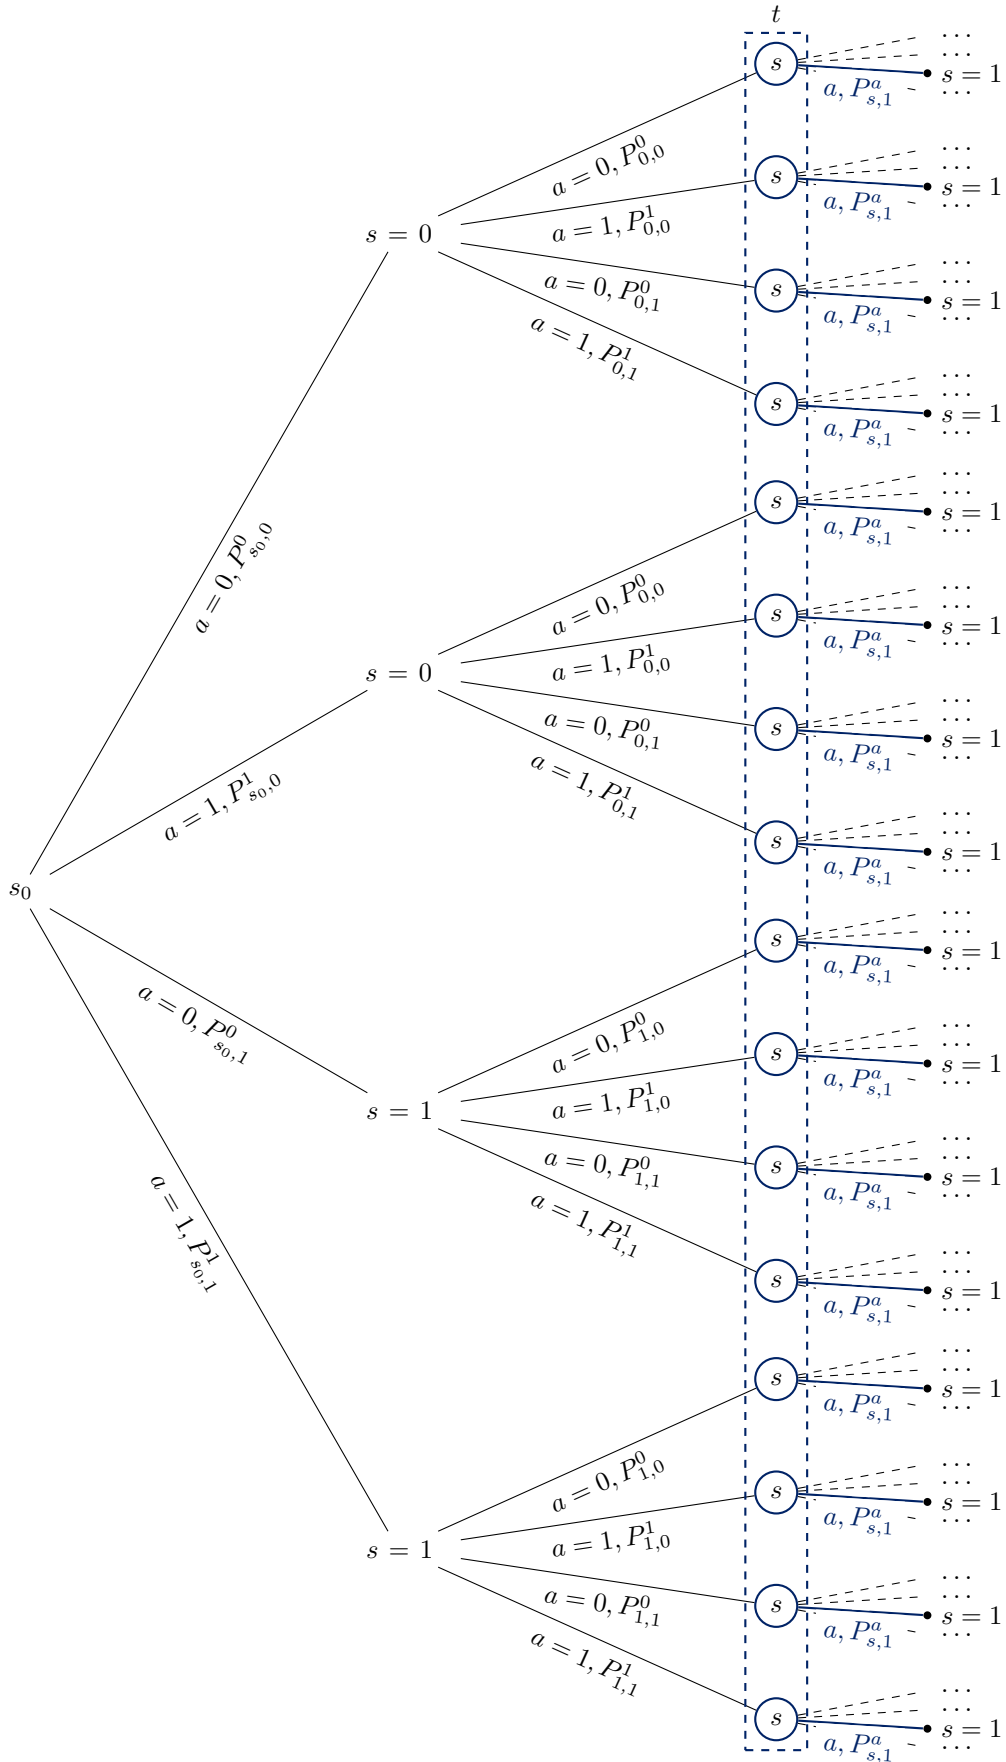
Case 3. $R \subseteq A$. Since $|\cup_{r \in \mathbf{R}} \varphi(r)| = 3$, proceed as outlined in Case 2.3. Only a slight modification is required: since \mathbf{R} is now equal to \mathbf{A} , taking the column-wise sums of the resulting \mathbf{R}_1 and \mathbf{R}_2 will yield two vectors $\vec{v}_1, \vec{v}_2 \in \mathbb{Z}^{N|\mathcal{A}|T}$, which can contain entries $\in \{1\}$ and $\{0, 2\}$, respectively. Thus, $\forall j \in [N|\mathcal{A}|T]$,

$$\sum_{i \in \mathbf{R}_1} \mathbf{A}_{ij} - \sum_{i \in \mathbf{R}_2} \mathbf{A}_{ij} = 1 - 0 \vee 1 - 2; \{-1, 1\} \not\subseteq \{-1, 0, 1\} \square$$

C.2.3. ENUMERATION OF OBJECTIVE FUNCTION COEFFICIENTS

The key challenge we encounter when we seek to enumerate the IP outlined in Section C.2.1 is that exact computation of the objective function coefficients, $\vec{c} \in \mathbb{R}^{N|\mathcal{A}|T}$ is intractable. Each arm contributes $|\mathcal{A}| \times T$ coefficients, and while calculation is trivially parallelizable over arms, we must consider a probability tree like the one in Figure 3 for each arm.

Figure 3: Illustration of the probability tree for finding the coefficient corresponding to $x_{i,a=0,t=2}$.



The number of decision variables required to enumerate each arm’s game tree is of order $O(|\mathcal{A}||\mathcal{S}|^T)$ and there are N such trees, so even a linear program (LP) relaxation is not tractable for larger values of T and N , which motivates us to propose PROBFair (Section 4) as an efficient alternative.

Example 1. Suppose we wish to find the coefficient c' corresponding to $x_{i,a=0,t=2}$. From Equation 9, we have $c' = \frac{1}{2^2} \sum_{s \in \mathcal{S}} p(s_2 = s)p(s_3 = 1|x_{i,a=0,t=2}, s_2 = s)$. Equivalently, we sum the weight of each path from the root node to the highlighted end nodes in Figure 3 and normalize by $\frac{1}{2^2}$:

$$\begin{aligned} c' = \frac{1}{4} & (P_{s_0,0}^0 P_{0,0}^0 P_{0,1}^0 + P_{s_0,0}^0 P_{0,0}^1 P_{0,1}^0 + P_{s_0,0}^0 P_{0,1}^0 P_{1,1}^0 + P_{s_0,0}^0 P_{0,1}^1 P_{1,1}^0 \\ & + P_{s_0,0}^1 P_{0,0}^0 P_{0,1}^0 + P_{s_0,0}^1 P_{0,0}^1 P_{0,1}^0 + P_{s_0,0}^1 P_{0,1}^0 P_{1,1}^0 + P_{s_0,0}^1 P_{0,1}^1 P_{1,1}^0 \\ & + \dots) \quad \text{For each of the } (|\mathcal{A}||\mathcal{S}|)^t = 16 \text{ paths to a blue node in Figure 3.} \end{aligned} \quad (14)$$

C.2.4. COMPARISON OF PROBFair WITH THE TRUE OPTIMAL POLICY

In Section E, we normalize intervention benefit with THRESHOLD WHITTLE, which is asymptotically optimal for forward threshold-optimal transition matrices under a budget constraint k (Mate et al., 2020). However, with the integer program (IP) we formulate in Section C.2.1, we can find the *optimal* policy for any set of transition matrices under budget *and* fairness constraints as long as N and T are small.




Figure 4: Adherences of PROBFair, compared to the IP formulation

We generate $N = 2$ random arms such that the structural constraints outlined in Section 2 are satisfied. We set $k = 1$ and $T = 6$. Though the variance in reward is large due to the small T , Figure 4 shows that PROBFair obtains 100% of the intervention benefit when no fairness constraints are applied. Similarly, PROBFair with $\ell = 0.33$ obtains the same adherence behavior as the IP policy with under hard fairness constraint $\nu = 3$ or minimum selection fraction constraint $\psi = 0.33$. (within 95% confidence interval shown). All results shown are bootstrapped over 500 iterations.

Minimum Selection Fraction Constraints. As we discuss in Appendix B, the optimal policy is often to pull the same k arms at every timestep and ignore all other arms. Under minimum selection fraction constraints (Equation 11), each arm must be pulled at least a minimum fraction ψ of T rounds, with no conditions on when these pulls should take place. We confirm with the optimal IP implementation our intuition that these additional pulls are allocated at the beginning or end of the simulation. That is, the optimal policy under minimum selection fraction constraints is to take advantage of the finite time horizon, which is not suitable for the applications we consider.

D. PROBF AIR: a Probabilistically Fair Policy

Our main contribution is the novel probabilistic policy PROBF AIR (see Section 4.1). In contrast to prior work, we seek to *guarantee* rather than incentivize fairness without incurring an exponential dependency on the time horizon or sacrificing optimality guarantees. We thus seek an efficient policy which is reward-maximizing, subject to the satisfaction of both budget and probabilistic fairness constraints.

Recall to construct π_{PF} , we: (1) partition the arms based on the shapes of their respective f_i functions (Eq. 4); (2) perform a grid search over possible ways to allocate the budget, k , between the two subsets of arms; (2a) solve each sub-problem to produce a probabilistic policy for the arms in that subset; (2b) compute the total expected reward of the policy; (3) take the argmax over this set of grid search values to determine the approximately optimal budget allocation; and (4) form π_{PF} by taking the union over the policies produced by evaluating each sub-problem at its approximately optimal share of the budget.

Here, in Section D.1, we present complete proofs to Theorems 4.1-4.4. Then, in Section D.2, we provide additional details about how we sample from our probabilistic policy to select discrete actions at each timestep.

D.1. Proofs

In this section, we provide proofs for the theorems introduced in Section 4.1. When relevant, we begin by restating the theorem for convenience.

Theorem 4.1. *For every arm $i \in [N]$, $f_i(p_i)$ is either **concave** or **strictly convex** in all of $p_i \in [0, 1]$.*

Proof. For notational convenience, let:

$$\begin{aligned} c_1 &= P_{0,1}^0; \\ c_2 &= P_{0,1}^1 - P_{0,1}^0; \\ c_3 &= 1 - P_{1,1}^0 + P_{0,1}^0; \\ c_4 &= P_{1,1}^0 - P_{1,1}^1 - P_{0,1}^0 + P_{0,1}^1. \end{aligned}$$

Then, $f_i(p_i) = \frac{c_1 + c_2 p_i}{c_3 + c_4 p_i}$. We observe that $\forall i \in [N]$, $f_i(p_i)$ is a valid probability since the term $1 - (1 - p_i)P_{1,1}^0 - p_i P_{1,1}^1$ in the denominator is at least $1 - (1 - p_i) - p_i = 0$ for all $p_i \in [0, 1]$. Then, there are three cases which describe the possible shapes of $f_i(p_i)$:

Case 1. $c_4 = 0$. Here, $f_i(p_i)$ is **linear** and hence, **concave**.

Case 2. $c_4 \neq 0; c_2 = 0$. Here, $f_i''(p_i) = \frac{2c_1 c_3^2}{(c_3 + c_4 p_i)^3} \geq 0$, so $f_i(p_i)$ is linear (hence **concave**) if $c_1 = 0$ or **strictly convex** (if $c_1 > 0$) in the domain $p_i \in [0, 1]$.

Case 3. $c_4 \neq 0; c_2 \neq 0$. Here,

$$f_i(p_i) = \frac{\frac{c_2}{c_4} \left(\frac{c_1 c_4}{c_2} + c_4 p_i \right)}{c_3 + c_4 p_i} = \frac{c_2}{c_4} + \frac{\left(c_1 - \frac{c_2 c_3}{c_4} \right)}{c_3 + c_4 p_i}. \quad (15)$$

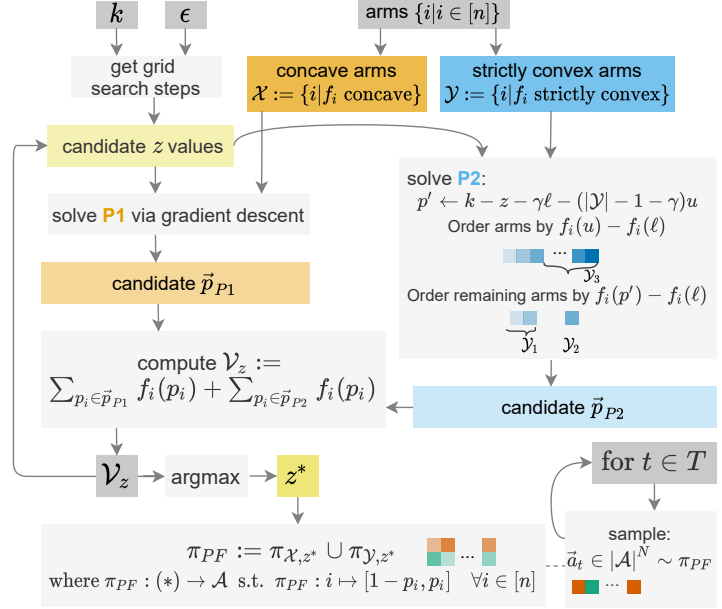


Figure 5: A visual representation of PROBF AIR's probabilistic policy construction and discrete action selection.

Thus,

$$f_i''(p_i) = \frac{2c_4^2 \left(c_1 - \frac{c_2 c_3}{c_4} \right)}{(c_3 + c_4 p_i)^3} \quad (16)$$

The sign of $f_i''(p_i)$ is the same as the sign of $d = c_1 - \frac{c_2 c_3}{c_4}$. It follows that $f_i(p_i)$ is **strictly convex** if $c_1 > \frac{c_2 c_3}{c_4}$, and **concave** otherwise for $p_i \in [0, 1]$. □

Theorem 4.2. *For each arm $i \in [N]$, the structural constraints introduced in Section 2 ensure that $f_i(p_i)$ is monotonically non-decreasing in p_i over the interval $[0, 1]$.*

Proof. For notational convenience, let:

$$\begin{aligned} c_1 &= P_{0,1}^0; \\ c_2 &= P_{0,1}^1 - P_{0,1}^0; \\ c_3 &= 1 - P_{1,1}^0 + P_{0,1}^0; \\ c_4 &= P_{1,1}^0 - P_{1,1}^1 - P_{0,1}^0 + P_{0,1}^1. \end{aligned}$$

Then $f_i(p_i) = \frac{c_1 + c_2 p_i}{c_3 + c_4 p_i}$ and $f_i'(p_i) = \frac{c_2 c_3 - c_1 c_4}{(c_3 + c_4 p_i)^2}$.

Observe $c_2 c_3 - c_1 c_4 \geq 0$ implies $f_i'(p_i) \geq 0$.

$$\begin{aligned} c_2 c_3 - c_1 c_4 &\geq 0 \\ c_2 c_3 &\geq c_1 c_4 \\ (P_{0,1}^1 - P_{0,1}^0)(1 - P_{1,1}^0 + P_{0,1}^0) &\geq P_{0,1}^0 (P_{1,1}^0 - P_{1,1}^1 - P_{0,1}^0 + P_{0,1}^1) \\ (1 - P_{1,1}^0) P_{0,1}^1 &\geq (1 - P_{1,1}^1) P_{0,1}^0 \end{aligned}$$

Per our structural constraints, $P_{1,1}^0 < P_{1,1}^1$ and $P_{0,1}^1 > P_{0,1}^0$. □

Lemma 4.3. *P_2 has an optimal solution in which $p_i \in (\ell, u)$ for at most one $i \in \mathcal{Y}$.*

Proof. Note by compactness that P_2 has not just a supremum, but an actual maximum solution. Suppose for contradiction there is some optimal solution \vec{p} with distinct indices $i, j \in \mathcal{Y}$ such that $p_i, p_j \in (\ell, u)$. Now, suppose we perturb by an infinitesimal ϵ (of arbitrary sign but tiny positive absolute value) such that $p_i := p_i + \epsilon$ and $p_j := p_j - \epsilon$. This satisfies all our constraints for small-enough $|\epsilon|$. The change in the objective $\sum_{i \in \mathcal{Y}} f_i(p_i)$ is now $\epsilon \cdot (f_i'(p_i) - f_j'(p_j)) + O(\epsilon^2)$; hence, if $f_i'(p_i) - f_j'(p_j)$ is nonzero, then we can take a tiny ϵ of the appropriate sign to increase the objective, a contradiction. Therefore, $f_i'(p_i) - f_j'(p_j) = 0$, and so, we now focus on lower-order terms: the change in the objective $\sum_{i \in \mathcal{Y}} f_i(p_i)$ is now $(\epsilon^2/2) \cdot (f_i''(p_i) + f_j''(p_j)) + O(\epsilon^3)$. However, since f_i and f_j are strictly convex, we have that $f_i''(p_i) + f_j''(p_j) > 0$, and hence the objective increases regardless of the sign of (the tiny) ϵ , again a contradiction. Thus we have our structural result. □

Theorem 4.4. *Alg. 2 yields the mapping from arms in \mathcal{Y} to subsets in $\{\mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Y}_3\}$ which maximizes $\sum_{i \in \mathcal{Y}} f_i(p_i)$ s.t. $\sum_{i \in \mathcal{Y}} p_i = k - z$. (See Appendix D for the complete proof).*

We begin by introducing Lemma D.1, which we use in our proof of Theorem 4.4:

Lemma D.1. *For a given $\gamma, p' \in (\ell, u)$.*

Proof. To begin, observe that $z - k = \gamma\ell + (|\mathcal{Y}| - 1 - \gamma)u + p'$. Then, to prove the lower bound, observe that:

$$\begin{aligned}\gamma &= \left\lfloor \frac{|\mathcal{Y}|u - (k - z)}{u - \ell} \right\rfloor \\ \gamma &> \frac{|\mathcal{Y}|u - (k - z)}{u - \ell} - 1 \\ \rightarrow \gamma &> \frac{|\mathcal{Y}|u - (\gamma\ell + (|\mathcal{Y}| - 1 - \gamma)u + p')}{u - \ell} - 1 \\ \gamma &> \frac{|\mathcal{Y}|u - \gamma\ell - |\mathcal{Y}|u + u + \gamma u - p' - u + \ell}{u - \ell} \\ 0 &> \ell - p' \implies p' > \ell\end{aligned}$$

To prove the upper bound, observe that:

$$\begin{aligned}\gamma &= \left\lfloor \frac{|\mathcal{Y}|u - (k - z)}{u - \ell} \right\rfloor \\ \gamma &\leq \frac{|\mathcal{Y}|u - (k - z)}{u - \ell} \\ \rightarrow \gamma &\leq \frac{|\mathcal{Y}|u - (\gamma\ell + (|\mathcal{Y}| - 1 - \gamma)u + p')}{u - \ell} \\ \gamma &\leq \frac{\gamma(u - \ell) + u - p'}{u - \ell} \\ 0 &\leq u - p' \implies p' \leq u\end{aligned}$$

Thus, $\ell < p' \leq u$. □

Now, we proceed with our proof of **Theorem 4.4**.

Proof. Recall **P2**: maximize $\sum_{i \in \mathcal{Y}} f_i(p_i)$ such that $p_i \in [\ell, u]$ for all $i \in \mathcal{Y}$ and $\sum_{i \in \mathcal{Y}} p_i = k - z$. By **Lemma 4.3**, there exists *at most one arm* with optimal value $p_i^* \in (\ell, u)$.

First, we discuss an edge case. If $k - z = |\mathcal{Y}|\ell$, Line 2 of Algorithm 2 assigns $\gamma = |\mathcal{Y}|$, so Line 11 assigns

$$\pi_{\mathcal{Y}} := i \mapsto \begin{cases} \ell, & \text{for } i \in \mathcal{Y}_1 = \{\mathcal{Y}\} \\ p', & \text{for } i \in \mathcal{Y}_2 = \emptyset \\ u, & \text{for } i \in \mathcal{Y}_3 = \emptyset \end{cases} \quad (17)$$

Thus Algorithm 2 returns the only valid solution to **P2** in this case, which is to set $p_i = \ell$ for all arms $i \in \mathcal{Y}$.

For all other cases $k - z > |\mathcal{Y}|\ell$, we introduce the following notation: let \mathcal{Y}_1 be the set of arms for which $p_i = \ell$, \mathcal{Y}_2 be a set containing exactly one arm (WLOG j) where $p_j = p' \in (\ell, u]$, and \mathcal{Y}_3 be the remaining set of arms for which $p_i = u$, with $\bigcap_{x=1}^3 \mathcal{Y}_x = \emptyset$. Then by **Lemma D.1**, $\gamma = |\mathcal{Y}_1| = \left\lfloor \frac{|\mathcal{Y}|u - (k - z)}{u - \ell} \right\rfloor$ and $p' = k - z - \gamma\ell - (|\mathcal{Y}| - 1 - \gamma)u \in (\ell, u]$.

P2 is then equivalent to finding a partition $\mathcal{Y} \rightarrow \mathcal{Y}_1 \cup \mathcal{Y}_2 \cup \mathcal{Y}_3$ which maximizes the following:

$$\begin{aligned}\arg \max_{\{\mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Y}_3\}} & \sum_{i \in \mathcal{Y}_1} f_i(\ell) + f_j(p') + \sum_{i'' \in \mathcal{Y}_3} f_{i''}(u) \\ \text{s.t. } & |\mathcal{Y}_1| = \gamma, \mathcal{Y}_2 = \{j\}, \\ & \bigcap_{x=1}^3 \mathcal{Y}_x = \emptyset, \text{ and } \bigcup_{x=1}^3 \mathcal{Y}_x = \mathcal{Y}\end{aligned} \quad (18)$$

Subtracting the constant $\sum_{i \in \mathcal{Y}} f_i(\ell)$ and simplifying yields:

$$\begin{aligned} & \arg \max_{\{\mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Y}_3\}} f_j(p') - f_j(\ell) + \sum_{i'' \in \mathcal{Y}_3} f_{i''}(u) - f_{i''}(\ell) \\ & \text{s.t. } |\mathcal{Y}_1| = \gamma, \mathcal{Y}_2 = \{j\}, \\ & \quad \bigcap_{x=1}^3 \mathcal{Y}_x = \emptyset, \text{ and } \bigcup_{x=1}^3 \mathcal{Y}_x = \mathcal{Y} \end{aligned} \quad (19)$$

Suppose we sort arms in ascending order by $f_i(u) - f_i(\ell)$. Let us create the set \mathcal{Y}'_3 from the last $|\mathcal{Y}| - \gamma - 1$ arms.

By monotonicity, for all $i \in \mathcal{Y}'_3$ and $j \notin \mathcal{Y}'_3$,

$$f_j(p') - f_j(\ell) \leq f_i(u) - f_i(\ell) \quad (20)$$

Thus, setting $\mathcal{Y}_3^* = \mathcal{Y}'_3$ reduces the optimization problem in Eq. 18 to finding a partition over the remaining sets \mathcal{Y}_1 and \mathcal{Y}_2 .

$$\begin{aligned} & \arg \max_{\{\mathcal{Y}_1, \mathcal{Y}_2\}} f_j(p') - f_j(\ell) \\ & \text{s.t. } |\mathcal{Y}_1| = \gamma, \mathcal{Y}_2 = \{j\}, \\ & \quad \bigcap_{x=1}^2 \mathcal{Y}_x = \emptyset, \text{ and } \bigcup_{x=1}^2 \mathcal{Y}_x = \mathcal{Y} \setminus \mathcal{Y}_3^* \end{aligned} \quad (21)$$

Finally, we solve Equation 21 by finding the arm j with maximal value $f_j(p') - f_j(\ell)$. Then $\mathcal{Y}_2^* = j$, $\mathcal{Y}_1^* = \mathcal{Y} \setminus (\mathcal{Y}_2 \cup \mathcal{Y}_3^*)$, and we are done. \square

Corollary 4.5. *Alg. 2 has time complexity $O(|\mathcal{Y}| \log |\mathcal{Y}|)$.*

At worst, Algorithm 2 requires two sorts: once on Line 5, and a second time on Line 8, for a total computational cost of $O(2|\mathcal{Y}| \log |\mathcal{Y}|)$. In total, the computational cost of Algorithm 1 is at worst $O\left(\frac{kN}{\epsilon^3}\right)$ when all N arms are in \mathcal{X} .

D.2. Dependent Rounding Sampling Approach

Here we provide pseudocode for the sampling algorithm introduced in Section 4.2, along with its associated SIMPLIFY subroutine (Srinivasan, 2001).

Algorithm 4 Sampling Subroutine (adapted from Srinivasan (2001))

```

1: procedure SIMPLIFY( $\alpha \in [0, 1], \beta \in [0, 1]$ )
2:   if  $\alpha = \beta = 0$  then
3:      $p_i, p_j \leftarrow [0, 0]$ 
4:   else if  $\alpha = \beta = 1$  then
5:      $p_i, p_j \leftarrow [1, 1]$ 
6:   else if  $\alpha + \beta = 1$  then
7:      $\text{flag} \leftarrow X \sim B(n = 1, p = \alpha)$ 
8:      $p_i, p_j \leftarrow [1, 0]$  if  $\text{flag}$  else  $[0, 1]$ 
9:   else if  $0 < \alpha + \beta < 1$  then
10:     $\text{flag} \leftarrow X \sim B\left(n = 1, p = \frac{\alpha}{\alpha + \beta}\right)$ 
11:     $p_i, p_j \leftarrow [\alpha + \beta, 0]$  if  $\text{flag}$  else  $[0, \alpha + \beta]$ 
12:   else if  $1 < \alpha + \beta < 2$  then
13:     $\text{flag} \leftarrow X \sim B\left(n = 1, p = \frac{1 - \beta}{2 - \alpha - \beta}\right)$ 
14:     $p_i, p_j \leftarrow [1, \alpha + \beta - 1]$  if  $\text{flag}$  else  $[\alpha + \beta - 1, 1]$ 
   return  $p_i, p_j$ 

```

Algorithm 5 Sampling Algorithm (adapted from Srinivasan (2001))

```

1: procedure SAMPLE( $G = (V, E)$ )
2:    $H \leftarrow G \setminus \{v \mid \exists e \in G \text{ s.t. } e_{\text{dst}} = v\}$ 
3:   if  $|H| = 1$  then return  $G$ 
4:   else if  $|H| \geq 2$  then
5:      $A \subseteq G \in \binom{H}{\lfloor \frac{|H|}{2} \rfloor}$ 
6:      $B \leftarrow G \setminus A$ 
7:      $\text{pairs} \leftarrow \{(a_i, b_i) \in A \times B \mid i \in \mathcal{I}\}$ 
8:      $H' \leftarrow (V = \emptyset, E = \emptyset)$ 
9:     for  $(v_i, v_j) \in \text{pairs}$  do
10:       $H' \leftarrow H' \cup \{v_i, v_j\}$ 
11:       $H' \leftarrow H' \cup \{v' \mid e_{v', v_\alpha} \mid \alpha \in \{i, j\}\}$ 
12:       $X_i, X_j \leftarrow \text{SIMPLIFY}(p_{v_i}, p_{v_j})$ 
13:       $z_{v_i} \leftarrow X_i$ 
14:       $z_{v_j} \leftarrow X_j$ 
15:      if  $z_{v_i} \in \{0, 1\}$  then
16:         $p_{v'} \leftarrow X_j$ 
17:      else  $p_{v'} \leftarrow X_i$ 
18:    $F \leftarrow G \cup H'$ 
19:   return SAMPLE( $F$ )

```

\triangleright Select subgraph containing nodes without a parent
 $\triangleright z_v \in \{0, 1\} \forall v \in G; \sum_v z_v = k$

\triangleright Defined in Algorithm 4
 \triangleright If X_i was fixed, $z_{v_i} \in \{0, 1\}$
 \triangleright If X_j was fixed, $z_{v_j} \in \{0, 1\}$

$\triangleright \forall v \in G \cap H',$ update attribute values per H'

E. Experimental Evaluation

In this section, we empirically demonstrate that PROBFair enforces the probabilistic fairness constraint introduced in Section 2 with minimal loss in total expected reward, relative to fairness-aware alternatives. We present results from three experiments: (1) PROBFair versus fairness-inducing alternative policies, holding the cohort fixed and considering fairness-aligned sets of hyperparameters; (2) PROBFair evaluated on a breadth of cohorts representing different types of patient populations; and (3) PROBFair when fairness is *not* enforced (i.e., $\ell = 0$), to examine the cost of state agnosticism.

E.1. Experimental Design

Policies: In each of our experiments, we compare PROBFair against the following baseline[§] and fairness- $\{\text{inducing}^\dagger, \text{guaranteeing}^\ddagger, \text{and agnostic}^*\}$ policies:

RANDOM [§]	Select k arms uniformly at random at each t .
ROUND-ROBIN ^{§, ‡}	Select k arms at each t in fixed, sequential order.
TW-BASED HEURISTICS [‡]	Select top- k arms based on Whittle index values. Available arms vary based on time-indexed fairness constraint satisfaction (Prins et al., 2020).
RISK-AWARE TW (RA-TW) [†]	Select top- k arms based on Whittle index values. Incentivizes fairness via concave reward function (Mate et al., 2021).
THRESHOLD WHITTLE (TW) [*]	Select top- k arms based on Whittle index values (Whittle, 1988; Mate et al., 2020).

We specifically consider three **THRESHOLD WHITTLE**-based heuristics: **FIRST**, **LAST**, and **RANDOM**. These heuristics partition the k pulls available at each timestep into (un)constrained subsets, where a pull is *constrained* if it is executed to satisfy a time-indexed fairness constraint. During constrained pulls, only arms that have not yet been pulled the required number of times within a ν -length interval are available; other arms are excluded from consideration, unless *all* arms have already satisfied their constraints. **FIRST**, **LAST**, and **RANDOM** position constrained pulls at the beginning, end, or randomly within each interval of length ν , respectively. Appendix F.1 provides pseudocode.

Objective: In all experiments, we assign equal value to the adherence of a given arm over time. Thus, we set our objective to reward occupancy in the “good” state: a simple local reward $r_t(s_t^i) := s_t^i \in \{0, 1\}$ and undiscounted cumulative reward function, $R(r(s)) := \sum_{i \in [N]} \sum_{t \in [T]} r(s_t^i)$.

Evaluation metrics: We are interested in comparing policies along two dimensions: reward maximization and fairness (i.e., with respect to the distribution of algorithmic attention). To this end, we rely on two performance metrics: (a) intervention benefit and (b) earth mover’s distance.

Intervention benefit (IB) is the total expected reward of an

algorithm, normalized between the reward obtained with no interventions (0% intervention benefit) and the asymptotically optimal but fairness-agnostic **THRESHOLD WHITTLE** algorithm (100%) (Mate et al., 2020). Formally,

$$\text{IB}_{\text{NoAct}, \text{TW}}(\text{ALG}) := \frac{\mathbb{E}_{\text{ALG}}[R_{\text{ALG}}(\cdot)] - \mathbb{E}_{\text{NoAct}}[R(\cdot)]}{\mathbb{E}_{\text{TW}}[R(\cdot)] - \mathbb{E}_{\text{NoAct}}[R(\cdot)]} \quad (22)$$

Earth mover’s distance (EMD) is a metric that allows us to compute the minimum cost required to transform one probability distribution into another (Rubner et al., 2000). We use it to compare algorithms with respect to fairness—i.e., how evenly pulls are allocated among arms. (Other metrics that may measure individual distributive fairness are discussed in Appendix F.2.)

For each algorithm, we consider a discrete distribution F of observed pull counts, where each bucket, $j \in \{0 \dots T\}$, corresponds to a feasible number of total pulls that an arm could receive, and $F[j] \in \{0 \dots N\}$ corresponds to the number of arms whose observed pull count is equal to j . All of the distributions we consider have the same total mass because each algorithm produces kT total pulls.

We use **ROUND-ROBIN** as a fair reference algorithm since it distributes pulls evenly among arms. We then compute the minimum cost required to transform each algorithm’s distribution, F_{ALG} , into that of **ROUND-ROBIN**’s, F_{RR} .

For our application this is equivalent to:

$$\text{EMD}_{\text{RR}}(\text{ALG}) := \left| \sum_{h=0}^T \sum_{j=0}^h F_{\text{ALG}}[j] - F_{\text{RR}}[j] \right| \quad (23)$$

Unless otherwise noted, we normalize EMD with respect to the distance obtained by **TW**, $\text{EMD}_{\text{RR}}(\text{ALG})/\text{EMD}_{\text{RR}}(\text{TW})$.

Datasets: We evaluate performance on two datasets: (a) a general set of randomly generated synthetic transition matrices and (b) a realistic patient adherence behavior model.

Synthetic. We construct a synthetic dataset of randomly generated arms such that the structural constraints outlined in Section 2 are preserved.

CPAP Adherence. Continuous positive airway pressure therapy (CPAP) is an effective treatment for obstructive sleep apnea (OSA). However, poor adherence behavior in using CPAP reduces its beneficial outcomes. Non-adherence to CPAP therapy affects an estimated 30-40% of patients (Rotenberg et al., 2016). We adapt Kang et al. (2016; 2013)’s Markov model of CPAP adherence behavior (a three-state system of hours of nightly CPAP usage) to a two-state system with the clinical standard for adherence—at least four hours of CPAP machine usage per night (Sawyer et al., 2011). We consider an intervention effect that broadly characterizes supportive interventions such as telemonitoring and phone support, which are associated with a moder-

ate 0.70 hours (95% CI ± 0.35) increase in device usage per night (Askland et al., 2020). For full details, see Appendix F.4.

E.2. PROBFAIR versus Fairness-aware Alternatives

In this experiment, we compare PROBFAIR to alternative policies which either *induce* or *guarantee* fairness. The former includes RISK-AWARE WHITTLE, which modifies the objective function via concave reward $r(b)$ (Mate et al., 2021). We consider the authors’ suggested reward function $r(b) = -e^{\lambda(1-b)}$, $\lambda = 20$, which imposes a large negative utility on lower belief values. However, RA-TW does not *guarantee* fair resource allocation. The latter includes ROUND-ROBIN and the heuristics, which guarantee time-indexed fairness but do not provide optimality guarantees.

In Table 2, we report average results for each policy, along with margins of error for 95% confidence intervals, computed over 100 simulation seeds for a synthetic cohort of 100 collapsing arms, with $k = 20$ and $T = 180$. To facilitate meaningful comparisons between PROBFAIR and the heuristics, we consider combinations of values for ℓ and ν that produce equivalent, integer-valued lower bounds on the number of pulls any arm can expect to receive—i.e., $\min_i \mathbb{E}[\sum_t \mathbb{1}(a_t^i = 1)] = \ell \times T = T/\nu$.

$\min_i \mathbb{E}[\# \text{ pulls}]$	Policy		$\mathbb{E}[\text{IB}]$ (%)	$\mathbb{E}[\text{EMD}]$ (%)
10 $\ell = 0.056$ $\nu = 18$	PF	ℓ	88.45 ± 0.27	81.11 ± 0.18
	FIRST	ν	88.75 ± 0.27	68.19 ± 0.14
	LAST	ν	89.32 ± 0.26	69.17 ± 0.11
	RANDOM	ν	92.02 ± 0.18	71.24 ± 0.13
18 $\ell = 0.1$ $\nu = 10$	PF	ℓ	81.57 ± 0.29	60.04 ± 0.22
	FIRST	ν	81.07 ± 0.31	47.44 ± 0.09
	LAST	ν	81.30 ± 0.29	48.47 ± 0.08
	RANDOM	ν	84.33 ± 0.26	51.67 ± 0.10
30 $\ell = 0.167$ $\nu = 6$	PF	ℓ	68.22 ± 0.33	22.66 ± 0.17
	FIRST	ν	70.22 ± 0.30	19.10 ± 0.03
	LAST	ν	69.41 ± 0.33	19.70 ± 0.03
	RANDOM	ν	70.52 ± 0.34	19.96 ± 0.04
comparison	TW		100.00 ± 0.00	100.00 ± 0.00
	RA-TW		72.73 ± 0.38	115.14 ± 0.26
baseline	RANDOM		54.66 ± 0.35	10.44 ± 0.11
	NOACT		0.00 ± 0.00	76.08 ± 0.11
	RR		62.96 ± 0.33	0.00 ± 0.00

Table 2: Expected intervention benefit and normalized earth mover’s distance by policy and fairness bracket.

Key findings include:


- *Fairer* hyperparameters ($\ell \uparrow$, $\nu \downarrow$), yield decreased $\mathbb{E}[\text{IB}]$ and $\mathbb{E}[\text{EMD}]$, reflecting improved individual fairness at the expense of total reward.
- RA-TW appears to *discourage* distributive fairness ($\mathbb{E}[\text{EMD}] > 100\%$), and yields reduced expected reward.
- For each (ℓ, ν) , PROBFAIR performs competitively

with respect to the best-performing heuristic (which, like TW, are state-aware, see Appendix E.4).


E.3. PROBFAIR Evaluated on a Breadth of Cohorts

Here we conduct sensitivity analysis with respect to cohort composition. For each dataset, we identify a transition matrix characteristic that can be modified during the generation process to produce a subset of arms that will exhibit less favorable transition dynamics than their peers. For the synthetic dataset, this characteristic is *strict convexity*. For the CPAP dataset, it is *non-adherence*, a mnemonic coined by Kang et al. (2013) to characterize a cluster of study participants, and contrast this to a model fit on the general patient population (see Appendix F.4 for details).

For each dataset, we generate 10 different cohorts, each of which is characterized by the percentage of unfavorable arms that it contains. We use a seed to control the generation process such that each cohort contains 100 collapsing arms in total, and a sliding window of the unfavorable (favorable) arms we can generate with this seed are included (excluded) as we increase the cardinality of the unfavorable subset. For ease of interpretation, we present unnormalized results over 100 simulation seeds with $k = 20$ and $T = 180$ in Figure 6, and then proceed to summarize normalized performance.



(a) Synthetic results



(b) CPAP results

Figure 6: Expected total reward (left) and unnormalized earth mover’s distance (right) on a breadth of cohorts.

Key findings from this experiment include:

- Per Figure 6, for each dataset, $\mathbb{E}[R]$ predictably declines for all policies as the % of unfavorable arms increases, while unnormalized $\mathbb{E}[EMD]$ rises for TW and PROBFair.
 - *Synthetic (6(a))*: As the proportion of strictly convex arms increases, PROBFair’s allocation of pulls tends towards the bimodality of TW.
 - *CPAP (6(b))*: As the proportion of non-adherent arms increases, the level of intervention required to improve trajectories rises, but the available budget remains static.
- PROBFair’s *normalized* performance remains stable even as cohort composition is varied:
 - *Synthetic*: With respect to IB (EMD), PROBFair achieves an average (over all cohorts) of averages (over 100 simulations per cohort) of $80.45\% \pm 1.30\%$ ($60.42\% \pm 2.03\%$).
 - *CPAP*: The corresponding values for IB (EMD): $80.44\% \pm 0.57\%$ ($60.54\% \pm 1.26\%$).

E.4. PROBFair: Price of State Agnosticism

Here, we investigate the cost associated with PROBFair’s state agnosticism, relative to state-aware THRESHOLD WHITTLE (TW). To ensure a fair comparison, we set $\ell = 0$ and $u = 1$, effectively constructing a version of PROBFair in which probabilistic fairness is *not* enforced. (Recall that TW is fairness-agnostic; in the previous results, we do not expect PROBFair to obtain the same total reward as TW). While PROBFair incorporates each arm’s structural information (i.e., transition matrices), it produces a set of *stationary* probability distributions over actions from which all discrete actions are subsequently drawn. TW, in contrast, ingests each arm’s current state at each timestep, and is thus able to exploit *realized* sequences of state transitions. While we thus expect PROBFair to incur some loss in IB, our results (computed over 100 simulation seeds, with $k = 20$, $N = 100$, and $T = 180$) indicate that this loss is acceptable rather than catastrophic. Relative to TW, $\text{PROBFair}_{\ell=0}$ obtains $96.04\% \pm 0.22$ of $\mathbb{E}[IB]$ and incurs an increase of only $4.61\% \pm 0.16$ with respect to $\mathbb{E}[EMD]$.

F. Additional Experimental Details

In this section, we discuss additional details of our empirical study in Appendix E. We provide a description and pseudocode of the heuristic policies (F.1), discuss our choice of fairness metric (F.2), and provide additional details of our Synthetic and CPAP datasets (F.3, F.4).

All results presented in this paper are bootstrapped over 100 simulation iterations, with a time horizon $T = 180$, cohort size $N = 100$, and budget $k = 20$, unless otherwise noted. We utilize seeds to ensure reproducible variation for each randomized parameter, including actualized transitions in each simulation. We have run simulations on an Intel(R) Core i7 CPU with 16Gb of RAM. Simulations are configurable via configuration files; runs are trivially parallelizable via these configuration files.

F.1. Heuristic Algorithms

In Section E, three heuristics based on the THRESHOLD WHITTLE algorithm are introduced: FIRST, LAST, and RANDOM. Here, we go into more detail and provide pseudocode.

Definition F.1. Within the context of Algorithm 6, we define a *constrained pull* to be one that is executed to satisfy an integer periodicity constraint. Only arms that have not yet been pulled the required number of times within the ν -length interval are available; other arms are excluded from consideration, unless *all* arms have already satisfied their constraints. In this case, all arms are available to be pulled.

If a pull is not constrained, we say it is *unconstrained* or *residual*.

The FIRST heuristic requires that all constrained pulls must occur at the start of the interval. This implies that the first N/k timesteps in each interval are dedicated to pulling all N arms.

The LAST heuristic requires that all constrained pulls must occur at the end of the interval. Unlike the FIRST heuristic, not all arms will necessarily be pulled in the last N/k timesteps, as some arms will have already satisfied their constraint earlier in the interval via unconstrained pull(s). These leftover constrained pulls function as unconstrained pulls, per Definition F.1.

The RANDOM heuristic chooses random positions within the interval for constrained pulls to occur. Similarly to the LAST heuristic, some of the later constrained pulls may become unconstrained pulls if all arms have already satisfied their constraint earlier in the interval.

Algorithm 6 Periodicity Constraint-Enforcing Heuristic Based on THRESHOLD WHITTLE

```

1: procedure SIMULATION( $A, T, \nu, k$ )
2:   for interval  $\in [0, T]$  with step size  $\nu$  do
3:      $C_{\text{interval}} \leftarrow \emptyset$   $\triangleright C_{\text{interval}} := \text{arm(s) with constraint satisfied during the interval}$ 
4:
5:   for  $a \in A$  do
6:      $a.\text{last observed state} \leftarrow 1$ 
7:      $a.\text{time since pull} \leftarrow 1$ 
8:
9:   for  $t \in T$  do
10:     $i \leftarrow \text{GetInterval}(t)$ 
11:
12:    if  $t$  is a constrained pull  $\wedge C_i \subsetneq A$  then  $\triangleright$  Consider arms with constraint not yet satisfied in interval
13:       $A' \leftarrow \{a | a \in A \setminus C_i\}$ 
14:    else if  $t$  is a residual pull  $\vee C_i = A$  then  $\triangleright$  Consider all arms
15:       $A' \leftarrow A$ 
16:
17:       $A'_k \leftarrow \text{SelectTopK}(A', k, t)$   $\triangleright$  Select  $k$  arms with highest Whittle index
18:       $C_i \leftarrow C_i \cup A'_k$ 
19:
20:    for  $a \in A$  do
21:       $s_{t+1}(a) \leftarrow \text{UpdateState}(a)$   $\triangleright$  Update each arm's state using belief
    return

```

F.2. Fairness Metric Choices

It is not immediately obvious which evaluation metric(s) best indicate whether we have improved distributive fairness. While constraint satisfaction itself is a logical candidate, it is Boolean-valued at the arm level, and thus does not reflect *to what extent* a policy fairly allocates pulls. Even if we were to report population-level constraint satisfaction (i.e., by noting the proportion of arms for which a given fairness constraint is satisfied, either over the course of a single simulation, or in expectation over a set of simulation iterations), this would be tautologically biased in favor of PROBFair and the THRESHOLD WHITTLE-based heuristics, which explicitly encode constraint satisfaction. This observation motivates us to consider proxy metrics, including the price of fairness (PoF), the Herfindahl–Hirschman Index (HHI), and the earth mover’s distance (EMD).

Price of Fairness. Consider *price of fairness*, defined formally as:

$$\text{PoF}_{\text{TW}}(\text{ALG}) := \frac{\mathbb{E}_{\text{TW}}[R(\cdot)] - \mathbb{E}_{\text{ALG}}[R(\cdot)]}{\mathbb{E}_{\text{TW}}[R(\cdot)]} \quad (24)$$

Price of fairness is the relative loss in total expected reward associated with following a distributive fairness-enforcing policy, as compared to THRESHOLD WHITTLE (Bertsimas et al., 2011). A small loss ($\sim 0\%$) indicates that fairness has a small impact on total expected reward; conversely, a large loss means total expected reward is sacrificed in order to satisfy the fairness constraints.

Lemma F.2. *Price of fairness is inversely proportional to intervention benefit.*

Proof. The statement in Lemma F.2 is equivalent to the statement “Given $y, z > 0$, there exists $\alpha \in \mathbb{R}$ such that $\frac{x-z}{y-z} = \alpha \frac{z-x}{z}$ for all $x > 0$ ”. Here $x = \mathbb{E}_{\text{ALG}}[R(\cdot)]$, $y = \mathbb{E}_{\text{NoAct}}[R(\cdot)]$, and $z = \mathbb{E}_{\text{TW}}[R(\cdot)]$. Consider $\alpha = \frac{-z}{y-z}$. Then $\alpha \frac{z-x}{z} = \frac{x-z}{y-z}$. Thus, for any algorithm ALG, $\text{PoF}_{\text{TW}}(\text{ALG}) \propto \text{IB}_{\text{NoAct, TW}}(\text{ALG})^{-1}$. \square

Herfindahl–Hirschman Index (HHI). The Herfindahl–Hirschman Index (HHI) (Rhoades, 1993), is a statistical measure of concentration useful for measuring the extent to which a small set of arms receive a large proportion of attention due to

an unequal distribution of scarce pulls (Hirschman, 1980). It is defined as:

$$\text{HHI}(\text{ALG}) := \sum_{i=1}^N \left(\frac{1}{kT} \sum_{t=1}^T a_t^i \right)^2 \quad (25)$$

HHI ranges from $1/N$ to 1; higher values indicate that pulls are concentrated on a small subset of arms. However, HHI is an imperfect evaluation metric for addressing our prioritarian concern for arms that would be *deprived* of algorithmic attention (i.e., fail to receive any pulls) under THRESHOLD WHITTLE (see Appendix B). Since entries are squared, reducing u offers a more direct path to lowering HHI than increasing ℓ . However, reducing u will not accomplish our stated goal of guaranteeing each arm a strictly positive lower bound on the probability that it will receive a pull at any given timestep.

Earth Mover’s Distance The earth mover’s distance (EMD), or Wasserstein metric, is a measure of distance between two distributions. Specifically, we measure the distance of an algorithm’s distribution of cumulative pull allocations to a fair reference distribution, ROUND-ROBIN. Though differences in distances are meaningful, EMD does not directly map to our fairness desiderata. That is, a given level of fairness enforcement (e.g., as characterized by the hyperparameters ℓ or ν) is not associated with a specific range of EMD values. Hence, our discussion of (normalized) earth mover’s distances in Section E focuses on relative differences between policies.

E.3. Synthetic Dataset

Conjecture F.3. The set of forward (reverse) threshold-optimal arms are a subset of the set of concave (strictly convex) arms for the local reward function we consider, $r(s) = s$.

Mate et al. (2021) provide conditions for threshold optimality. First, the arm must satisfy the structural constraints (Section 2) and the *indexability* condition $P_{1,1}^0 - P_{0,1}^0 + P_{1,1}^1 - P_{0,1}^1 \leq 1$. Then, the following inequalities determine forward (reverse) threshold optimality:

$$\begin{cases} P_{1,1}^0 - P_{0,1}^0 \geq P_{1,1}^1 - P_{0,1}^1 & \text{forward threshold-optimal} \\ P_{1,1}^0 - P_{0,1}^0 \leq P_{1,1}^1 - P_{0,1}^1 & \text{reverse threshold-optimal} \end{cases} \quad (26)$$

We conjecture that these conditions necessarily imply the conditions for concavity, repeated here for convenience:

$$\begin{cases} P_{0,1}^0 \leq \frac{(P_{0,1}^1 - P_{0,1}^0)(1 - P_{1,1}^0 + P_{0,1}^0)}{P_{1,1}^0 - P_{1,1}^1 - P_{0,1}^0 + P_{0,1}^1} & \text{concave} \\ P_{0,1}^0 > \frac{(P_{0,1}^1 - P_{0,1}^0)(1 - P_{1,1}^0 + P_{0,1}^0)}{P_{1,1}^0 - P_{1,1}^1 - P_{0,1}^0 + P_{0,1}^1} & \text{strictly convex} \end{cases} \quad (27)$$

E.4. CPAP Dataset

Obstructive sleep apnea (OSA) is a common condition that causes interrupted breathing during sleep (Punjabi, 2008); when used throughout the entirety of sleep, continuous positive airway pressure therapy (CPAP) eliminates nearly 100% of obstructive apneas for the majority of treated patients (Sawyer et al., 2011). Our CPAP dataset is provided by Kang et al. (2013; 2016), who model the dynamics and patterns of patient adherence behaviour as a basis for designing effective and economical interventions. We include this dataset in the supplement. Kang et al. (2013) specifically divide patients into two clusters using expectation-maximization on CPAP usage patterns. Patients in the first cluster exhibit ‘adherent’ behavior—though they occasionally miss a night, these patients utilize a CPAP machine for more than four hours every night without assistance. The ‘non-adherent’ patient type we consider in Section E.3 is the second cluster; otherwise, we consider data that encapsulates all study participants.

Kang et al. (2016) consider many intervention effects; in this paper, we specifically consider $\alpha_{\text{interv}} = 1.1$ for all patients, i.e. the impact of supportive interventions (Askland et al., 2020). Additionally, we add random $\sigma = 1$ logistic noise to the transition matrices so that there is some variance in individual arm dynamics. For non-adherent arms, added noise can only further hinder the probability of adherence. The initial state of each arm (before the simulation starts) is randomly assigned between $s = 0$ (did not adhere), and $s = 1$ (adhered).