
From Restless to Contextual: A Thresholding Bandit Reformulation for Finite-horizon Improvement

Jiamin Xu
Cornell Tech

Ivan Nazarov
Causal Foundry

Aditya Rastogi
Causal Foundry

África Perriñez
Causal Foundry

Kyra Gan
Cornell Tech

Abstract

This paper addresses the poor finite-horizon performance of existing online *restless bandit* (RB) algorithms, which stems from the prohibitive sample complexity of learning a full *Markov decision process* (MDP) for each agent. We argue that superior finite-horizon performance requires *rapid convergence* to a *high-quality* policy. Thus motivated, we introduce a reformulation of online RBs as a *budgeted thresholding contextual bandit*, which simplifies the learning problem by encoding long-term state transitions into a scalar reward. We prove the first non-asymptotic optimality of an oracle policy for a simplified finite-horizon setting. We propose a practical learning policy under a heterogeneous-agent, multi-state setting, and show that it achieves a sublinear regret, achieving *faster convergence* than existing methods. This directly translates to higher cumulative reward, as empirically validated by significant gains over state-of-the-art algorithms in large-scale heterogeneous environments. The code is provided in [github](#). Our work provides a new pathway for achieving practical, sample-efficient learning in finite-horizon RBs.

1 INTRODUCTION

Restless bandits extend classic contextual bandits by modeling state transitions and incorporating per-step budget constraints. In a RB, a set of independent arms (agents) evolves over discrete time steps as MDP, with the goal of dynamically allocating a fixed number of active plays to maximize cumulative rewards.

However, this modeling power comes with significant learning challenges. Even in the *offline* setting with known model parameters, RB is PSPACE-hard (Papadimitriou & Tsitsiklis, 1999). This computational intractability has directed much of the research toward policies that provide *asymptotic guarantees*, such as Whittle’s Index Policy (WIP) (Whittle, 1988; Weber & Weiss, 1990; Hsu, 2018; Kadota et al., 2016) and linear programming-based methods (Verloop, 2016; Zhang & Frazier, 2021; Brown & Zhang, 2023; Hong et al., 2024; Gast et al., 2024; Xiong et al., 2022a), which are proven to approach optimal as the number of agents approaches infinity.

This complexity carries over into the *online* setting, where transition dynamics are unknown. Prior online RB algorithms largely seek to learn the optimal policy by first estimating the full underlying MDP for each agent. Consequently, online RB methods are largely inspired by, and mirror the standard frameworks of, *reinforcement learning* (RL).

They predominantly fall into two canonical RL categories. *Model-based* approaches explicitly learn transition probabilities, typically within either *episodic finite-horizon* (Jung & Tewari, 2019; Wang et al., 2023b; Liang et al., 2024) or *infinite-horizon* (Ortner et al., 2012; Wang et al., 2020; Xiong et al., 2022b; Akbarzadeh & Mahajan, 2023; Xiong et al., 2022a) frameworks. *Model-free* methods focus on learning the Q-function specifically in the infinite-horizon setting (Killian et al., 2021; Xiong & Li, 2023; Xiong et al., 2024; Kakarapalli et al., 2024). A common strategy in both categories is to then integrate these learned models with WIP for decision-making.¹ These online methods are designed to converge to their asymptotically optimal offline counterparts when agents are *homogeneous*, where the *convergence rate* is measured by the notion of *regret* (formalized in Sec. 2.2).

This reliance on standard RL frameworks means that

Proceedings of the 29th International Conference on Artificial Intelligence and Statistics (AISTATS) 2026, Tangier, Morocco. PMLR: Volume 300. Copyright 2026 by the author(s).

¹We note that while LP-based methods are powerful in the offline setting, they are generally more difficult to adapt for online learning.

the performance guarantees for these online algorithms are inherently *asymptotic* in the number of episodes or time steps. However, this focus constitutes a fundamental limitation in practical applications,² where horizons are finite due to budget cycles, program durations, or policy deadlines (del Río et al., 2024a,b). In domains like targeted patient interventions (Villar et al., 2015; Biswas et al., 2021; Mate et al., 2021) or resource allocation (Borkar et al., 2017; Brown & Smith, 2020), systems evolve continuously and cannot be reset, making the standard episodic RL framework impractical.

Consequently, an algorithm’s performance is judged by its cumulative reward within a *single, limited-time episode*. An asymptotically optimal policy that requires a long time to learn offers no practical value *if it fails to achieve a high reward within the operational horizon*. This critical limitation of prior work motivates a fundamental shift away from the objective of learning complex MDPs, and towards a reformulation designed for sample-efficient learning in the non-episodic, finite-horizon setting. Our contributions are fourfold:

- **A Novel Reformulation:** we interpret regret as convergence rate (Sec. 2.2) and show that online algorithms may achieve higher cumulative reward in a finite horizon by targeting alternative design objectives, which leads to accelerated convergence (Sec. 2.3). This motivates us to introduce a *budgeted thresholding contextual multi-armed bandit* (BT-CMAB), which simplifies the learning objective of RBs to a contextual bandit problem through carefully designed rewards (Sec. 2.4), enabling sample-efficient learning in finite horizons.
- **First Non-asymptotic Finite-Horizon Optimality Guarantee:** we establish the first finite-horizon non-asymptotic (valid for an arbitrary number of agents) optimality guarantee for a tractable RB setting, proving our oracle greedy policy is optimal for a homogeneous 2-state model (Theorem 3.1). This provides a theoretical foundation absent in prior work.
- **A Fast-Converging Learning Algorithm:** we propose ϵ -GT (*ϵ -Greedy Thresholding*, Algorithm 1, Sec. 4), a novel learning algorithm for RB with *heterogeneous agents*. We establish that it achieves sublinear convergence rate (Theorem 4.4), enabled by the careful choice of decaying ϵ .
- **Empirical Validation:** we numerically demonstrate (in Section 5) that our algorithm consistently outperform existing algorithms in various finite-horizon online RB settings across a wide range of instances *with heterogeneous agents*.

²Table A.1 provides key application areas for RBs.

1.1 Additional Related Work

Thresholding Bandit/MDP Our paper relates to regret-minimizing thresholding bandit problems, where regret is defined by the gap between the arm’s mean and a predefined threshold when the mean falls below the threshold. Under this setting, minimax optimal (max is taken over all possible bandit instances and min taken over all possible learning algorithms) constant regret is achievable (Tamatsukuri & Takahashi, 2019; Michel et al., 2022; Feng et al., 2024). Our problem setup differs from the traditional thresholding bandit framework by incorporating context, where the arrival of contexts follows an MDP. In contrast, Hajiabolhasan & Ortner (2023) study threshold MDPs under an infinite horizon setting, seeking policies with average rewards exceeding a predefined threshold. We, however, focus on the finite-horizon setting.

Episodic MDP RB can be viewed as a single, massive MDP with a state space that is the combinatorial product of all agent states. Thus, our work also relates to episodic MDP, which focuses on maximizing cumulative rewards across a finite horizon. Let K be the number of episodes, each containing T horizon. Azar et al. (2017) establishes a minimax regret of $\mathcal{O}(\sqrt{T^2 K})$ for stationary transitions, and Jin et al. (2018) derives a regret of $\mathcal{O}(\sqrt{T^4 K})$ for nonstationary transitions.

MAB Experimental Design Our incremental reward explicitly incorporates the trade-off between reward maximization and treatment effect estimation (i.e., estimating the difference in rewards under $A = 0$ and $A = 1$), previously studied in MAB with Pareto-optimal algorithms provided (Simchi-Levi & Wang, 2025; Zhong et al., 2021; Zhang & Wang, 2024). Unlike prior work that treats these aspects separately, ours inherently combines both aspects through *reward* that depends on intervention (i.e., $A = 1$) and non-intervention (i.e., $A = 0$) *outcomes*.

2 FINITE-HORIZON RESTLESS BANDITS AND OUR REFORMULATION

We introduce restless bandits in Sec. 2.1 and reinterpret regret as convergence in Sec. 2.2. Sec. 2.3 argues that better finite-sample performance requires new design objectives, while maintaining reward-based evaluation. We present our reformulation in Sec. 2.4.

2.1 Restless Bandits

In a finite-horizon RB problem, the learner is provided with M *independent* arms, each representing one independent agent. Each agent m is modeled as an

MDP represented by the tuple $(\mathcal{S}, \mathcal{A}, P_a^m, R^m)$, where $P_a^m(s, s') : \mathcal{S} \times \mathcal{S} \rightarrow [0, 1]$ represents the transition probability from state s to s' under action a for agent m , and $R^m(s) : \mathcal{S} \rightarrow \mathbb{R}$ represents the *expected reward* that agent m receives at state s . We focus on binary action, i.e., $\mathcal{A} = \{0, 1\}$. We highlight that in classical RBs, the *expected reward*, $R^m(s)$, is independent of the action, and the learner observes the reward *immediately after arriving at* the state.

At each time t , given the state s_t , the learner’s decision \mathbf{a}_{t+1} follows a history-dependent random policy, π_{t+1} . Let $\mathcal{H}_t := \{(\mathbf{a}_i, \mathbf{s}_i, \mathbf{r}_i)\}_{i=1}^t \cup \{\mathbf{s}_0, \mathbf{r}_0\}$ be the history observed up to time t , and let $\mathcal{H}_t^+ := \mathcal{H}_t \cup \{\mathbf{a}_{t+1}\}$ be the augmented history up to taking action \mathbf{a}_{t+1} . Similarly, denote \mathcal{F}_t as the σ -algebra generated by \mathcal{H}_t and \mathcal{F}_t^+ as the σ -algebra generated by \mathcal{H}_t^+ . Then, $\pi_{t+1} : \mathcal{H}_t \rightarrow \Delta(\{0, 1\}^M)$, where $\Delta(\{0, 1\}^M)$ denotes the probability simplex over all possible actions. The learner is provided with a budget constraint, allowing B arms to be pulled at each time step, i.e., $\sum_{m \in [M]} a_{t+1}^m = B$. After pulling B arms, the learner observes the next state $s_{t+1}^m \sim P_{a_{t+1}^m}^m(s_t^m, s_{t+1}^m)$, and receives a feedback \mathbf{r}_{t+1} . Suppose that the reward \mathbf{r}_{t+1} is generated according to a distribution \mathcal{P}_t , we will assume stationarity of the *expected reward* received at time $t + 1$, \mathbf{r}_{t+1} , when conditioned on history \mathcal{H}_t^+ and the current state.

Assumption 2.1. $\mathbb{E}[r_{t+1}^m | s_{t+1}^m, \mathcal{H}_t^+] = R^m(s_{t+1}^m)$.

We permit non-stationarity in the rewards, provided that Assumption 2.1 is satisfied, offering a relaxation of the standard stationarity assumption typically imposed in literature.

The goal of the learner is to find the policy $\boldsymbol{\pi} := \{\pi_t\}_{t \in [T]}$ that maximizes the expected cumulative reward, where the expectation is further taken over the randomized policy ($\mathbf{a}_t \sim \pi_t$), in addition to the states ($s_{t+1}^m \sim P_{a_{t+1}^m}^m(s_t^m, s_{t+1}^m)$), subject to the full budget utilization constraint:³

$$\begin{aligned} \max_{\boldsymbol{\pi}} \mathbb{E}_{\boldsymbol{\pi}} \left[\sum_{m=1}^M \sum_{t=0}^T R^m(s_t^m) \mid \mathbf{s}_0 \right] \\ \text{s.t. } \sum_{m=1}^M a_t^m = B, \forall t \in [T]. \end{aligned} \quad (1)$$

We will use $\boldsymbol{\pi}^*$ to denote the optimal policy corresponding to Problem (1). Since $\boldsymbol{\pi}^*$ is intractable and unknown (Papadimitriou & Tsitsiklis, 1999), WIP (Whittle, 1988) and LP-based policies (Hong et al., 2024; Xiong et al., 2022a) are proposed as approximations.

³Our model incorporates the standard full budget utilization constraint, consistent with the existing literature (Whittle, 1988; Zhang & Frazier, 2021; Hong et al., 2024).

These policies are proven to be asymptotically optimal as $M \rightarrow \infty$ for infinite-horizon rewards with homogeneous agents, a restrictive setting that often does not occur in practice.

2.2 Regret As Convergence Rate

A critical shift occurs in online learning: *the homogeneous agent assumption is often dropped to model more realistic, heterogeneous environments*. Under heterogeneity, these different approximation policies are no longer equivalent and will achieve different asymptotic performance levels. Consequently, comparing the regret magnitudes of different online algorithms is not meaningful when they target different reference policies (e.g., one converging to WIP versus another converging to an LP-based policy), as they are measuring convergence to fundamentally different benchmarks. This reality reframes the meaning of regret in RBs. It is not an absolute measure of sub-optimality, but a relative measure of the *convergence rate to a chosen reference policy* $\boldsymbol{\pi}^e$. We formalize this idea below:

Definition 2.2 (Regret as Convergence Rate). Let $\boldsymbol{\pi}^e$ be the reference policy that an online algorithm ω aims to converge to. We define the regret of ω relative to this reference policy over horizon T as $\mathcal{R}_l^{\boldsymbol{\pi}^e}(\omega)$:⁴

$$\begin{aligned} \mathcal{R}_l^{\boldsymbol{\pi}^e}(\omega) := \mathbb{E}_{\boldsymbol{\pi}^e} \left[\sum_{m=1}^M \sum_{t=0}^T l^m(s_t^m, a_{t+1}^m) \mid \mathbf{s}_0 \right] \\ - \mathbb{E}_{\omega} \left[\sum_{m=1}^M \sum_{t=0}^T l^m(s_t^m, a_{t+1}^m) \mid \mathbf{s}_0 \right], \end{aligned}$$

where $l^m(s_t^m, a_{t+1}^m)$ serves as an instantaneous performance metric where the choice of l^m depends on the objective of $\boldsymbol{\pi}^e$, which we elaborate below. For heterogeneous agents with single-episode finite horizons, existing RB oracle policies are not optimal and hence not comparable. Thus, the regret $\mathcal{R}_l^{\boldsymbol{\pi}^e}(\omega)$ measures the *convergence rate of the learning algorithm ω to its corresponding oracle (or reference) policy $\boldsymbol{\pi}^e$* , where sublinear regret implies that ω converges to $\boldsymbol{\pi}^e$ within sublinear time.

The choice of the loss function l^m is directly linked to the objective of $\boldsymbol{\pi}^e$. For instance, in MAB, if $\boldsymbol{\pi}^e$ targets reward maximization, l is immediate reward (Auer et al., 2002); if $\boldsymbol{\pi}^e$ targets thresholding, l penalizes action on arms below threshold (Feng et al., 2024). Thus, l always measures how well online algorithms are achieving the corresponding objectives of $\boldsymbol{\pi}^e$.

Remark 2.3. The interpretation of regret as a convergence rate appears implicitly in prior work. For

⁴Appendix B.1 extends this definition to episodic settings, which is common in online RBs.

example, Wang et al. (2023b) use Def. 2.2 with π^e as WIP, while Xiong et al. (2022a) use it with π^e as the infinite-horizon average-reward optimal policy.

2.3 The Need for Indirect Reward Maximization in Online RB

Building on our reinterpretation of regret as convergence rate, we reconsider Problem (1)’s reward maximization objective. In finite-horizon settings, algorithms targeting alternative objectives with faster convergence to high-quality reference policies may achieve superior empirical performance. This approach addresses the fundamental limitation of direct methods for online RB—including those learning π^* , WIP (Jung & Tewari, 2019; Wang et al., 2023b), or LP-based policies (Xiong et al., 2022a)—which typically suffer from slow convergence in practice.⁵

This slow convergence is well-characterized theoretically. For algorithms directly targeting reward maximization, Auer et al. (2008) establishes a $\Omega(T)$ lower bound when $\pi^e = \pi^*$ and $l^m = R^m$, while the tightest upper bound for learning WIP or LP-based policies remains $\mathcal{O}(T)$ (Wang et al., 2023b). This linear rate stems from the need to construct confidence intervals for *transition probabilities*, which requires extensive exploration of state-action pairs (s_t, a_t, s_{t+1}) to achieve sufficient visitation counts.

Given these fundamental constraints on estimating transition probabilities, we pursue an alternative approach: proposing an oracle (or reference) policy that enables faster convergence by circumventing the need for direct transition probability estimation. We develop this in Sec. 2.4, focusing on reference policies that avoid this bottleneck while maintaining strong *empirical performance*.

2.4 BT-CMAB Reformulation

Thus motivated, we reformulate restless bandit learning via two key modifications:

1. *Contextual bandit framework*: Transition probabilities need not be explicitly estimated. It embeds non-stationarity into rewards, yielding policies that depend only on the current context. This reduces sample complexity when compared with MDPs (Simchi-Levi & Xu, 2022).
2. *Thresholding objective*: At each step, we identify agents whose intervention benefit exceeds a thresh-

⁵Although Xiong et al. (2022a) claim sublinear convergence for finite-horizon RB, their analysis relies on a generative model that permits restarting the MDP from arbitrary states—an assumption rarely satisfied in practice. Without this, their algorithm would incur linear convergence.

old (vs. optimizing cumulative rewards), further lowering sample complexity (Feng et al., 2024) to achieve fast finite-horizon convergence.

Under suitable assumptions, our approach with known transitions/rewards recovers the optimal policy of *offline* RB for any M (Sec. 3), unlike WIP’s asymptotic optimality requiring $M, T \rightarrow \infty$.

Regret-Minimizing BT-CMAB Reduction In the reduced contextual bandit problem, the context space is the state space. At time t , the reward ϕ^m that agent m receives at state s_t^m under action a_{t+1}^m is defined as the *expected reward* that the agent will receive in the next time period:

$$\begin{aligned} \phi^m(s_t^m, a_{t+1}^m) &:= \sum_{s_{t+1}^m \in \mathcal{S}} P_{a_{t+1}^m}^m(s_t^m, s_{t+1}^m) R^m(s_{t+1}^m) \\ &= \sum_{s_{t+1}^m \in \mathcal{S}} \mathbb{P}(s_{t+1}^m | \mathcal{H}_t^+) \mathbb{E}[r_{t+1}^m | s_{t+1}^m, \mathcal{H}_t^+] \\ &= \mathbb{E}[r_{t+1}^m | \mathcal{H}_t^+]. \end{aligned} \quad (2)$$

The second equality follows from conditional stationarity (Assump. 2.1) and the Markovian property. Eq. (2) shows that the immediate reward, $\phi^m(s_t^m, a_{t+1}^m)$, can be estimated without modeling the transitions, consistent with the standard bandit reward formulation.

As in online RBs, at each time t , the learner observes s_t and the corresponding reward r_t . The learner then updates the estimated uncertainty associated with the estimated $\phi^m(s_{t-1}^m, a_t^m) \forall m \in [M]$ according to Eq. (2), and uses it to guide the action a_{t+1} , with $\sum_{m \in [M]} a_{t+1}^m = B$. The goal of the learner is to learn a policy that identifies and selects B agents whose expected gain from taking the action (i.e., $a_{t+1}^m = 1$) exceeds a predefined threshold, γ , in their current state s_t^m . To quantify this expected gain in state $s \in \mathcal{S}$, we define

$$I^m(s) := \phi^m(s, 1) - \phi^m(s, 0),$$

representing the *incremental reward* an agent receives by taking the action versus not at state s . At time t , we say that an agent m is in a good state if $I^m(s_t^m) \geq \gamma$.

At each time t , let $G_t := \{m \in [M] : I^m(s_t^m) \geq \gamma\}$ denote the set of agents currently in a good state. In this BT-CMAB formulation, the optimal policy, π^b , involves randomly selecting B agents from G_t for intervention when $|G_t| \geq B$ and selecting B agents according to the ordering of incremental reward otherwise.

In our reformulation, we define regret (Def. 2.2) for any policy π by setting the reference policy $\pi^e = \pi^b$ and using the loss function $l^m(s_t^m, a_{t+1}^m) := -c_t^m(a_{t+1}^m)$, where the cost function is:

$$c_t^m(a_{t+1}^m) := (\gamma - I^m(s_t^m)) \mathbb{I}\{I^m(s_t^m) < \gamma, a_{t+1}^m = 1\}. \quad (3)$$

This cost function has an intuitive interpretation: it penalizes allocations to agents in “bad” states (where $I^m(s_t^m) < \gamma$), with the penalty proportional to how far the state index falls below the threshold γ . The baseline policy π^b , by construction, never incurs these costs since it only allocates to agents in good states.

We simplify notation by writing $\mathcal{R}_c^{\pi^b}(\pi)$ as $\mathcal{R}(\pi)$. This regret directly measures how frequently π deviates from π^b by allocating to suboptimal states, thus quantifying the convergence rate to this reference policy. The learner’s objective is to minimize cumulative regret $\mathcal{R}(\pi)$ over horizon T , which equivalently maximizes the convergence rate to π^b .

3 OPTIMALITY OF BT-CMAB REDUCTION

Before establishing the sublinear convergence rate of our proposed algorithm for BT-CMAB with *heterogeneous* agent, we first establish the optimality of our BT-CMAB formulation for 2-state homogeneous agents—the first non-asymptotic optimality result for finite-horizon RB. This demonstrates that the contextual bandit provides an exact reformulation of the offline RB problem under these conditions. *We note that Sections 4 and 5 lift these assumptions.*

We assume *binary states* and we further assume that all agents share a good state with higher R^m . Without loss of generality (WLOG), we will use 1 to denote the good state and 0 for the bad state. We first describe the greedy policy. WLOG, suppose that at each time step t , the incremental rewards are ordered decreasingly as $I^1(s_t^1) \geq I^2(s_t^2) \geq \dots \geq I^M(s_t^M)$. The corresponding greedy policy $\pi_t^{\text{greedy}} : \mathcal{S} \rightarrow \{0, 1\}^M$ is to give actions to the first B agents according to the order. Formally, with a slight abuse of notation, let $\pi_t^{\text{greedy}}(s_t^i)$ be the i^{th} entry of π_t^{greedy} , defined as:

$$\pi_t^{\text{greedy}}(s_t^i) = \begin{cases} 1 & i \leq B \\ 0 & i > B \end{cases}.$$

Let $\pi^{\text{greedy}} := \{\pi_t^{\text{greedy}}\}_{t \in [T]}$. We note that $\pi^{\text{greedy}} = \pi^b$ with suitable choices of γ . A trivial choice of γ would be $\max_{m,s} I^m(s)$. Now we are ready to show that the greedy policy is optimal for any M .

Theorem 3.1. *Assuming all agents are homogeneous, i.e., $P^m = P^{m'}, R^m = R^{m'}, \forall m, m' \in [M]$ and $\mathcal{S} = \{0, 1\}$, $\pi^{\text{greedy}} = \pi^*$ for any $0 \leq B \leq M < \infty$.*

Theorem 3.1 establishes that our reference policy π^b —which corresponds to the greedy policy π^{greedy} with appropriate threshold selection—achieves non-asymptotic optimality for homogeneous, 2-state finite-horizon restless bandits. This provides stronger guaran-

tees than prior reference policies like WIP or LP-based methods (Hong et al., 2024; Whittle, 1988; Zhang & Frazier, 2021), which only offer asymptotic optimality for infinite-horizon settings under homogeneous agents.

While our analysis builds on the practical two-state assumption commonly used in recent work (Liang et al., 2024; Killian et al., 2023; Wang et al., 2023a), proving Theorem 3.1 required developing novel analytical techniques. Existing tools are designed for analyzing value function-based policies (Eq. (E.11)) in infinite-horizon settings (Whittle, 1988; Hong et al., 2024), making them unsuitable for our problem setup.

Remark 3.2 (Difficulty of Relaxation). Theorem 3.1 does not extend to heterogeneous agents (Appendix C) or multi-state systems (Appendix D). Our proof technique—which reduces Bellman optimality to conditions on value functions—relies critically on the two-state property that transition dynamics can be fully characterized by the probability of transitioning to state 1. This characterization breaks down in multi-state systems, where transitions require tracking probabilities to multiple states. While we conjecture that appropriately designed greedy policies may achieve constant-factor approximations even in these more general settings, such extensions are beyond the scope of this paper.

Although π^b is only guaranteed to be optimal for homogeneous agents and a high threshold, our BT-CMAB reduction is valid for heterogeneous agents and an arbitrary threshold γ . As mentioned in Section 2.4, the threshold is associated with fast convergence. We will show in Section 4.1 that a higher threshold will lead to slower convergence for heterogeneous agents. The detailed discussion on how to choose the threshold is deferred to Remark 4.5.

Theorem 3.1 Proof Sketch The proof consists of two steps: 1) We reformulate the problem as an MDP where the state space is the Cartesian product of the individual state space, i.e., \mathcal{S}^M ; the action space is the subset of $\{0, 1\}^M$ that satisfies the budget constraint, and the transition matrix is the product of P^m as all agents are independent. 2) We show that the proposed greedy policy π_t^{greedy} satisfies the Bellman equation (Eq. (E.21)) for all $t \leq T$ by using backward induction. We prove the induction hypothesis by carefully controlling the difference of the value function (Eq. (E.11)).

4 LEARNING IN BT-CMAB

We now present our learning algorithm (Algorithm 1) for BT-CMAB (as reformulated in Sec. 2.4) in a general setting, and establish that it achieves sublinear regret for this new problem. The sublinear regret bound explicitly characterizes the algorithm’s convergence rate, demonstrating its fast convergence in finite sample

when compared with those of existing restless bandit algorithms. We note that *the results hold without agent homogeneity or a binary state space*.

Learning in BT-CMAB framework is challenging because it requires estimating the incremental reward—the outcome difference between acting and not acting—which necessitates observing both actions for each agent-state pair (paired exploration). However, to minimize regret, the policy must also exploit this knowledge by taking the optimal action (which could be either 0 or 1, depending on the state and agent). The tension between these two requirements—thorough exploration for accurate estimation and greedy exploitation for regret minimization—defines the core difficulty of the problem.

ϵ -Greedy Thresholding Algorithm To resolve this, we propose an ϵ -Greedy Thresholding algorithm, Algorithm 1, that explicitly balances this trade-off. By using a decaying exploration rate ϵ_t , it ensures we sample both actions for each state-agent pair enough for learning while progressively prioritizing the optimal action to minimize regret, as formalized in our proof.

At each time step t , Algorithm 1 proceeds as follows: With probability ϵ_t , the algorithm explores by uniformly selecting B agents to act upon, regardless of their state or estimated reward. The exploration rate ϵ_t decays with the inverse of the minimum time an agent has spent in its current state, encouraging more exploration in under-sampled contexts. With probability $1 - \epsilon_t$, the algorithm exploits its knowledge: it computes the empirical mean incremental reward for each agent in its current state, forms a candidate set \tilde{G}_{t+1} of agents exceeding the threshold γ . The algorithm then proceeds as follows: if $|\tilde{G}_{t+1}| \geq B$, it gives actions to the top- B agents within \tilde{G}_{t+1} (e.g., those with the highest empirical means); if $|\tilde{G}_{t+1}| < B$, it gives actions to all agents in \tilde{G}_{t+1} and then allocates the remaining $B - |\tilde{G}_{t+1}|$ actions by selecting the agents with the highest *upper confidence bound* (UCB) on the incremental reward from the rest agents. After the actions are taken, the algorithm observes the new state vector \mathbf{s}_{t+1} and the corresponding reward vector \mathbf{r}_{t+1} .

Next, the algorithm updates its internal estimates using the newly observed rewards and states. For each agent, we update the cumulative rewards and visitation counts for the state-action pairs that were just observed. These are used to compute a new empirical mean for the incremental reward, $\hat{I}_{t+1}^m(s)$, for each agent-state pair. Finally, the upper confidence bound is recalculated. The new $\text{UCB}_{t+1}^m(s)$ for the next round adds an optimism bonus to the empirical incremental reward. This bonus is derived from an anytime-valid confidence bound (Howard et al., 2021) and is inversely

proportional to the square root of the state-action visit counts, ensuring that estimates for less-explored states have higher uncertainty. The specific function used is $f(t, T) = \log \log(2t + 0.72 \log(10.4T))$.

As we address a thresholding contextual bandit with MDP-governed contexts instead of MAB, Alg. 1 differs from Auer et al. (2002). We tailor ϵ_t to decay with the minimum state occupancy time rather than the current timestep, ensuring sufficient exploration across states.

4.1 Theoretical Guarantee

We establish that Algorithm 1 achieves a $\mathcal{O}(\sqrt{T})$ regret upper bound under two assumptions. Assump. 4.1 ensures strictly positive transition probabilities (even if arbitrarily small), while Assump 4.2 guarantees the existence of a policy with sufficient good agent-state pairs at each timestep.

Assumption 4.1 (Positive Transition). $P_a^m(s, s') > 0$ for all $s, s' \in \mathcal{S}$, $a \in \mathcal{A}$, and $m \in [M]$.

Assump. 4.1 guarantees that any state $s' \in \mathcal{S}^M$ is reachable from any given state s in a single step. Importantly, we only require these one-step transition probabilities to be positive; we do not assume they are bounded away from zero, and our result *does not* depend on the minimum transition probability.

Assumption 4.2 (Exists A Good Policy). There exists a policy achieving zero cumulative cost.

Assump. 4.2 serves primarily to simplify our theoretical analysis by ensuring problem instances are well-behaved. This condition guarantees that an oracle policy exists which can spend its full budget of B actions per round without being forced to assign actions with negative expected returns. In practice, this assumption can be satisfied for any problem instance by carefully choosing the threshold parameter γ ; for example, setting $\gamma = -\infty$ guarantees the condition is met trivially. Thus, γ serves as a *tunable parameter* to be optimized for overall performance, a point we discuss further in Remark 4.5.

When combined with Assump. 4.1, Assump. 4.2 yields a stronger guarantee (Lemma 4.3) that is inherently assumed by any algorithm required to spend the full budget B at each round. This is common in prior RB algorithms (Wang et al., 2020, 2023b).

Let $\mathcal{G} := \{(m, s) : m \in [M], s \in \mathcal{S}, I^m(s) \geq \gamma\}$ be the set of good agent-state pairs. Lemma 4.3 ensures that for any state realization $\mathbf{s} \in \mathcal{S}^M$, there are always at least B good agent-state pairs.

Lemma 4.3 (Existence of Sufficient Good Agent-State Pairs). *Under Assumptions 4.2 and 4.1, for any joint state $\mathbf{s} \in \mathcal{S}^M$, the number of good agent-*

Algorithm 1 ϵ -Greedy Thresholding (ϵ -GT)

- 1: **Input:** initial state vector \mathbf{s} , round budget B , exploration function $f(t)$, threshold γ , instance parameter η , scaling constant D , exploration decay rate p
 - 2: Initialize: initial states $\mathbf{s}_0 \leftarrow \mathbf{s}$; $\forall (s, a, m) \in (\mathcal{S} \times \mathcal{A} \times [M])$: agent state visit count $N_{s,m}^{(0)} \leftarrow 0$, agent state-action visit count $N_{s,m,a}^{(0)} \leftarrow 0$, incremental reward $\hat{I}_0^m(s) \leftarrow 0$, cumulative immediate rewards $\hat{\phi}^m(s, a) \leftarrow 0$, UCB of incremental reward $\text{UCB}_0^m(s) \leftarrow \infty$
 - 3: **for** $t = 0, 1, \dots, T$ **do**
 - 4: Retrieve \mathbf{s}_t ; identify good agents above the threshold: $\tilde{G}_{t+1} \leftarrow \{m \in [M]: \hat{I}_t^m(s_t^m) \geq \gamma\}$.
 - 5: Initialize the candidate set $C \leftarrow \tilde{G}_{t+1}$;
 - 6: Sort the remaining agents: $[M] \setminus \tilde{G}_{t+1}$, in descending order of $\text{UCB}_t^m(s_t^m)$; append top $\max(0, B - |\tilde{G}_{t+1}|)$ agents from sorted list to C .
 - 7: $\tilde{C} \leftarrow B$ agents in C with highest $\hat{I}_t^m(s_t^m)$.
 - 8: With probability $\epsilon_t = \min \left\{ 1, \min_m \frac{D}{(N_{s_t^m, m}^{(t)} + 1)^p \min\{\eta^2, 1/2\}} \right\}$:
 - 9: Random exploration: uniformly choose B agents from $[M]$ and set their $a_{t+1}^m \leftarrow 1$.
 - 10: Otherwise greedy: set $a_{t+1}^m = 1$ for $m \in \tilde{C}$, 0 otherwise
 - 11: Observe $(\mathbf{s}_{t+1}, \mathbf{r}_{t+1})$
 - 12: For each agent-state pair $(m, s) \in [M] \times \mathcal{S}$, update the UCB of the incremental rewards as:
 - 13: $\forall a \in \mathcal{A}: \hat{\phi}^m(s_t^m, a) = \hat{\phi}^m(s_t^m, a) + r_{t+1}^m \mathbb{I}\{a_{t+1}^m = a\}$.
 - 14: $\forall a \in \mathcal{A}: N_{s,m,a}^{(t+1)} = N_{s,m,a}^{(t)} + \mathbb{I}\{a_{t+1}^m = a\} \mathbb{I}\{s = s_t^m\}$
 - 15: $N_{s,m}^{(t+1)} = N_{s,m}^{(t)} + \mathbb{I}\{s = s_t^m\}$
 - 16: $\hat{I}_{t+1}^m(s) = \frac{\hat{\phi}^m(s, 1)}{N_{s,m,1}^{(t+1)}} - \frac{\hat{\phi}^m(s, 0)}{N_{s,m,0}^{(t+1)}}$
 - 17: $\text{UCB}_{t+1}^m(s) = \hat{I}_{t+1}^m(s) + \sqrt{\frac{f(N_{s,m,1}^{(t+1)}, T)}{N_{s,m,1}^{(t+1)} + 2}} + \sqrt{\frac{f(N_{s,m,0}^{(t+1)}, T)}{N_{s,m,0}^{(t+1)} + 2}}$
 - 18: **end for**
-

state pairs observed in \mathbf{s} is at least B . That is, $|\{m \in [M]: (m, s^m) \in \mathcal{G}\}| \geq B$.

Despite its strong guarantee, Lemma 4.3 (proof in Appendix F.1) is derived from two relatively mild assumptions. This result is critical for two reasons. First, it ensures that any algorithm required to spend a full budget of B actions is never forced to assign actions whose incremental reward is below the threshold. Second, it guarantees that while choosing a suboptimal agent may affect future state distributions, this effect does not accumulate: sufficient good agents are always available. Thus, a policy achieving zero per-step regret is attainable, which forms the foundation for our algorithm’s performance.

Let $\Delta_s^m := |\gamma - I^m(s)|$. Define (m_{\min}, s_{\min}) as the minimizer of $I^m(s) - \gamma$ over $(m, s) \in \mathcal{G}$, and set $\Delta := \Delta_{s_{\min}}^{m_{\min}}$. Next, we present our main theorem on the regret upper bound, which characterizes the finite-horizon convergence rate of Algorithm 1:

Theorem 4.4 (Regret Upperbound). *Assume r_t^m is 1-subgaussian random variable for all $t \in [T]$. Let $D := \max \left\{ \frac{8M}{B}, \frac{8M}{M-B}, 2 \right\}$, $p = 1, \eta = \Delta$. Under Assump.s 4.2 and 4.1, the regret, $\mathcal{R}(\pi)$, defined by Def. 2.2*

with the loss function l^m given by Eq. (3), satisfies:

$$\mathcal{R}(\pi) = \min \left\{ \mathcal{O} \left(\frac{|\mathcal{S}| (M|\mathcal{S}| - |\mathcal{G}|) MD \log T}{\Delta} \right), \mathcal{O} \left(\sqrt{MDB (|\mathcal{S}| (M|\mathcal{S}| - |\mathcal{G}|)) T \log T} \right) \right\}. \quad (4)$$

Eq. (4) achieves a sublinear convergence rate to π^b , contrasting the linear rate to WIP in Wang et al. (2023b). This bound is independent of γ and the smallest transition probability, with the constant formalized in Appendix F. Although Theorem 4.4 sets $\eta = \Delta$, empirical results show algorithm stability across η values in terms of average cumulative reward (Appendix G.3).

Theorem 4.4 Proof Sketch Our analysis deviates from standard CMAB analysis because estimating the incremental reward requires pulling suboptimal arms; consequently, we must *lower* bound the number of times suboptimal arms are pulled, in addition to the usual upper bound. The argument proceeds in three steps. 1) Regret equivalence: we show that regret is equivalent to the expected number of interventions given to bad agents. 2) Bounding the interventions: we prove that this expected number can be bounded either by the

overestimation of the incremental reward of bad agents or by the underestimation of that of good agents. We further establish that at each time step, both the number of interventions and non-interventions assigned to agents are lower bounded. 3) Concentration: we apply concentration inequality (Lemma F.2) to obtain the desired result. Importantly, stationarity of the reward is not required here, since the concentration inequality we employ holds without such an assumption.

Remark 4.5 (How to Select the Threshold γ ?). In our problem, there is an inherent trade-off between satisfying Assump. 4.2 and maximizing cumulative reward. While our reformulation focuses on the threshold identification problem, maximizing cumulative reward remains our ultimate goal. Specifically, a smaller threshold γ makes it easier to satisfy Assump. 4.2, as more state-agent pairs qualify as good. However, a low γ also makes any arm above the threshold optimal, potentially leading to lower rewards. Empirically, we observe that there exists a universal choice of γ which satisfies the Assump. 4.2 that achieves the highest reward as shown in Appendix G.3.

5 EXPERIMENTS

We numerically evaluate the performance of our algorithm, showing: 1) sublinear regret under the BT-CMAB reformulation (results in Appendix G.1), and 2) the achieved average cumulative reward, as defined in the finite-horizon RB setting, outperforms that of existing online RB algorithms. These results underscore the practicality of the proposed reformulation and algorithms in finite-horizon restless bandit problems.

Benchmarks We compare the average cumulative reward obtained by our algorithm with the following benchmarks in online RB settings: WIP based: 1) WIQL (Biswas et al., 2021), 2) UCWhittlePv (Wang et al., 2023b); LP-based: 3) R(MA)²B (Xiong et al., 2022a); and 4) *Random Policy* (uniformly pick B agents), which follows the experimental setting of Wang et al. (2023b); Xiong et al. (2022a).⁶

With heterogeneous agents, no algorithm has been established as optimal. We incorporate three benchmark algorithms with full MDP knowledge: 5) *Oracle Whittle* (Whittle, 1988), 6) *Oracle Greedy* and an LP-based policy 7) *OMR* (Xiong et al., 2022a). Baseline details and implementation are in Appendix G.2.

Setup We set $M = 50, \mathcal{S} = \{0, 1\}$, and compare the results under three different budget settings: $B = 5, 10, 20$. We use: 1) **Synthetic Instances**: The

expected reward is set to be $R^m(s) = s$ and randomly generate 50 instances with different transition matrices, 2) **Wireless Communication Dataset (Prieto-Cerdeira et al., 2010)**: We use the same wireless communication dataset as Wang et al. (2020). Details will be deferred to Section 5.2. We also consider the setting where $\mathcal{S} = [10]$ and randomly generate 50 instances. The results are deferred to Appendix G.2.

In Figure 1, each line is averaged over 50 instances, each with 13 repetitions. We consider two reward settings: 1) a noiseless case, and 2) a noisy setting where the rewards at state $s \in \mathcal{S}$ follow a normal distribution with mean $R^m(s)$ and variance 1 for the synthetic dataset and 0.1 for the real-world dataset. We note that we did not fine-tune our hyperparameters. Specifically, we universally choose a threshold that satisfies Assump. 4.2. We have also shown that our algorithm is robust across different η and we have done ablations on the choice of threshold (Appendix G.3).

5.1 Synthetic Instances

Results Under Noiseless Rewards In Figure 1 (top row), we observe that Alg. 1 consistently outperforms online RB algorithms within a horizon of 1000 in all budget settings. As shown in Figure 1, Alg. 1 quickly converges to a stable yet suboptimal policy, whereas WIQL converges more slowly. However, given a sufficiently long horizon, WIQL will eventually converge to the *Oracle Whittle* policy. Nevertheless, Alg. 1 outperforms WIQL when the horizon is short, even when the budget is low, highlighting the practicality of our algorithm. UCWhittlePv requires extensive exploration to converge, and R(MA)²B performs similarly to random policy.

Results Under Noisy Rewards In Figure 1 (bottom row), we excluded UCWhittlePv and R(MA)²B since they underperformed WIQL in the noiseless setting. In all three budget settings, Alg. 1 consistently outperforms WIQL.

Lastly, we compare *Oracle Whittle*, *Oracle Greedy*, and *OMR*. *Oracle Whittle* consistently outperforms *Oracle Greedy* across all budgets, with the gap narrowing as budget increases. Despite this, Alg. 1 surpasses all online benchmarks, demonstrating the benefits of rapid convergence. *OMR* performed worst, likely because its LP-based construction leverages homogeneity and falters in heterogeneous settings.

5.2 Wireless Communication Dataset

This problem (Prieto-Cerdeira et al., 2010) involves digital video broadcasting satellite services to handheld devices via land mobile satellite. In this problem, the

⁶Our results differ from those of Wang et al. (2023b) due to detected bugs in their code; a detailed discussion is provided in Appendix G.2.




Figure 1: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 50 instances with 13 repetitions each. Top: noiseless reward; bottom: noisy reward. Exclude $R(MA)^2B$ and UCWhittlePV in the noisy setting since they require known reward.




Figure 2: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 100 repetitions of the real-world dataset. Top: noiseless reward; bottom: noisy reward. Exclude $R(MA)^2B$ and UCWhittlePV in a noisy setting since they require known reward.

environment is modeled as a two-state system (“good” and “bad”), with five different transition models corresponding to different satellite locations. In their original problem, each agent is provided with multiple actions which is modeled as the angle of the satellite. Similar to Wang et al. (2020), we choose two angles as the action: 40 degrees as action 0 and 60 degrees as action 1. The transition probability is chosen according to Table IV in Prieto-Cerdeira et al. (2010). The original reward is dependent on both action and state, but the reward difference between different actions are low (Table III in Prieto-Cerdeira et al. (2010)). Further, the reward difference between the good state and bad state is also small. Therefore, we reconstruct the reward and set the reward for good state as 1 and bad state as 0.9 to mimic the real-world situations where signal is weak. We randomly assign 50 agents with one of the five transition probabilities.

In Figure 2, we observe that in this problem, Our method consistently outperforms all benchmarks within a horizon of 1000 in all budget settings under noiseless and noisy rewards. As shown in Figure 2, Alg. 1 con-

verges to its oracle policy (*OracleGreedy*) quickly while other online RB algorithms converge to their oracle policies more slowly. This translates to the superior performance of Alg. 1, demonstrating our observation that quick convergence translates to better finite sample performance.

Discussions We establish greedy policy optimality in finite-horizon RBs and reformulate online RB as a BT-CMAB with a learning algorithm. Our approach leverages the insight that fast convergence to a sub-optimal policy can outperform slow convergence to a better one in finite horizons. As the BT-CMAB formulation is novel, our work appropriately focuses on establishing its theoretical and empirical value. Deriving a problem-specific lower bound is a natural and important direction for future work. Further, we plan to analyze the optimality of the oracle policy in multi-state setting.

References

- Akbarzadeh, N. and Mahajan, A. On learning whittle index policy for restless bandits with scalable regret. *IEEE Transactions on Control of Network Systems*, 2023.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- Auer, P., Jaksch, T., and Ortner, R. Near-optimal regret bounds for reinforcement learning. *Advances in neural information processing systems*, 21, 2008.
- Azar, M. G., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. In *International conference on machine learning*, pp. 263–272. PMLR, 2017.
- Biswas, A., Aggarwal, G., Varakantham, P., and Tambe, M. Learn to intervene: An adaptive learning policy for restless bandits in application to preventive healthcare. *arXiv preprint arXiv:2105.07965*, 2021.
- Borkar, V. S., Kasbekar, G. S., Pattathil, S., and Shetty, P. Y. Opportunistic scheduling as restless bandits. *IEEE Transactions on Control of Network Systems*, 5:1952–1961, 2017.
- Brown, D. B. and Smith, J. E. Index policies and performance bounds for dynamic selection problems. *Management Science*, 66:3029–3050, 2020.
- Brown, D. B. and Zhang, J. Fluid policies, reoptimization, and performance guarantees in dynamic resource allocation. *Operations Research*, 2023.
- del Río, A. F., Leong, M. B., Saraiva, P., Nazarov, I., Rastogi, A., Hassan, M., Tang, D., and África Perriñez. Adaptive user journeys in pharma e-commerce with reinforcement learning: Insights from swiperx, 2024a.
- del Río, A. F., Leong, M. B., Saraiva, P., Nazarov, I., Rastogi, A., Hassan, M., Tang, D., and África Perriñez. Adaptive behavioral ai: Reinforcement learning to enhance pharmacy services, 2024b.
- Feng, Q., Ma, T., and Zhu, R. Satisficing exploration in bandit optimization. *arXiv preprint arXiv:2406.06802*, 2024.
- Gast, N., Gaujal, B., and Yan, C. Linear program-based policies for restless bandits: Necessary and sufficient conditions for (exponentially fast) asymptotic optimality. *Mathematics of Operations Research*, 49(4):2468–2491, 2024.
- Hajiabolhassan, H. and Ortner, R. Online regret bounds for satisficing in mdps. In *Sixteenth European Workshop on Reinforcement Learning*, 2023.
- Hong, Y., Xie, Q., Chen, Y., and Wang, W. Achieving exponential asymptotic optimality in average-reward restless bandits without global attractor assumption, 2024.
- Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055 – 1080, 2021.
- Hsu, Y.-P. Age of information: Whittle index for scheduling stochastic arrivals. pp. 2634–2638. IEEE Press, 2018.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.
- Jung, Y. H. and Tewari, A. Regret bounds for thompson sampling in episodic restless bandit problems. *Advances in Neural Information Processing Systems*, 32, 2019.
- Kadota, I., Uysal-Biyikoglu, E., Singh, R., and Modiano, E. Minimizing the age of information in broadcast wireless networks. In *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 844–851, 2016.
- Kakarapalli, P., Kayande, D., and Meshram, R. Faster q-learning algorithms for restless bandits. *arXiv preprint arXiv:2409.05908*, 2024.
- 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*, pp. 871–881, 2021.
- Killian, J. A., Biswas, A., Xu, L., Verma, S., Nair, V., Taneja, A., Hegde, A., Madhiwalla, N., Diaz, P. R., Johnson-Yu, S., et al. Robust planning over restless groups: engagement interventions for a large-scale maternal telehealth program. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 14295–14303, 2023.
- Liang, B., Xu, L., Taneja, A., Tambe, M., and Janison, L. A bayesian approach to online learning for contextual restless bandits with applications to public health. *arXiv preprint arXiv:2402.04933*, 2024.
- Mate, A., Madaan, L., Taneja, A., Madhiwalla, N., Verma, S., Singh, G., Hegde, A., Varakantham, P., and Tambe, M. Field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health, 2021.

Michel, T., Hajiabolhassan, H., and Ortner, R. Regret bounds for satisficing in multi-armed bandit problems. *Transactions on Machine Learning Research*, 2022.

Ortner, R., Ryabko, D., Auer, P., and Munos, R. Regret bounds for restless markov bandits. In *Proceedings of the 23rd International Conference on Algorithmic Learning Theory, ALT'12*, pp. 214–228, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 9783642341052.

Papadimitriou, C. H. and Tsitsiklis, J. N. The complexity of optimal queuing network control. *Mathematics of Operations Research*, 24(2):293–305, 1999. ISSN 0364765X, 15265471.

Prieto-Cerdeira, R., Perez-Fontan, F., Burzigotti, P., Bolea-Alamañac, A., and Sanchez-Lago, I. Versatile two-state land mobile satellite channel model with first application to dvb-sh analysis. *International Journal of Satellite Communications and Networking*, 28(5-6):291–315, 2010.

Simchi-Levi, D. and Wang, C. Multi-armed bandit experimental design: Online decision-making and adaptive inference. *Management Science*, 71(6):4828–4846, 2025. doi: 10.1287/mnsc.2023.00492.

Simchi-Levi, D. and Xu, Y. Bypassing the monster: A faster and simpler optimal algorithm for contextual bandits under realizability. *Mathematics of Operations Research*, 47:1904–1931, 2022.

Tamatsukuri, A. and Takahashi, T. Guaranteed satisficing and finite regret: Analysis of a cognitive satisficing value function. *Biosystems*, 180:46–53, 2019.

Verloop, I. M. Asymptotically optimal priority policies for indexable and nonindexable restless bandits. 2016.

Villar, S. S., Bowden, J., and Wason, J. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 30:199, 2015.

Wang, K., Verma, S., Mate, A., Shah, S., Taneja, A., Madhiwalla, N., Hegde, A., and Tambe, M. Scalable decision-focused learning in restless multi-armed bandits with application to maternal and child health. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 12138–12146, 2023a.

Wang, K., Xu, L., Taneja, A., and Tambe, M. Optimistic whittle index policy: Online learning for restless bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 10131–10139, 2023b.

Wang, S., Huang, L., and Lui, J. Restless-ucb, an efficient and low-complexity algorithm for online restless bandits. *Advances in Neural Information Processing Systems*, 33:11878–11889, 2020.

Weber, R. R. and Weiss, G. On an index policy for restless bandits. *Journal of Applied Probability*, 27(3):637–648, 1990.

Whittle, P. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 25(A):287–298, 1988.

Xiong, G. and Li, J. Finite-time analysis of whittle index based q-learning for restless multi-armed bandits with neural network function approximation. *Advances in Neural Information Processing Systems*, 36:29048–29073, 2023.

Xiong, G., Li, J., and Singh, R. Reinforcement learning augmented asymptotically optimal index policy for finite-horizon restless bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 8726–8734, 2022a.

Xiong, G., Wang, S., and Li, J. Learning infinite-horizon average-reward restless multi-action bandits via index awareness. *Advances in Neural Information Processing Systems*, 35:17911–17925, 2022b.

Xiong, G., Wang, S., Li, J., and Singh, R. Whittle index-based q-learning for wireless edge caching with linear function approximation. *IEEE/ACM Transactions on Networking*, 2024.

Zhang, X. and Frazier, P. I. Restless bandits with many arms: Beating the central limit theorem. *arXiv preprint arXiv:2107.11911*, 2021.

Zhang, Z. and Wang, Z. Online experimental design with estimation-regret trade-off under network interference. *arXiv preprint arXiv:2412.03727*, 2024.

Zhong, Z., Cheung, W. C., and Tan, V. Y. Achieving the pareto frontier of regret minimization and best arm identification in multi-armed bandits. *arXiv preprint arXiv:2110.08627*, 2021.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]

- (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
- (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Appendix

A REAL-WORLD RESTLESS BANDIT EXAMPLES

Table A.1: Examples of good and bad states and corresponding budgeted interventions in various domains

Domain	Good state	Bad state	Budgeted Intervention Example
General Usage	Frequent logins, feature exploration, meaningful use	Rare/brief logins, session abandonment	Reach out, offer help, gather feedback
E-commerce	High purchase frequency, browsing multiple categories	Low/no purchases, high cart abandonment	Provide a personalized discount or free shipping
E-learning	Completes modules, passes quizzes, revisits content	Seldom opens modules, fails quizzes	Progress-based reward (finish by deadline for free certificate)
Loyalty	Completes surveys, redeems rewards, advances loyalty	Not collecting points, ignoring surveys, no rewards	Special bonus for feedback (earn points by reviewing new features)
Patient management	Updates records, completes follow-ups, tracks adherence	Rarely logs updates, misses assessments or follow-ups	Points or bonus for timely follow-ups

We will further use an example to show case how the state, action, reward and budget is defined in practice. SwipeRX is an online platform in Indonesia where pharmacists typically place orders on Mondays. The platform aims to increase engagement and revenue by providing targeted coupons to pharmacists. Since pharmacists operate in different local markets with varying demand patterns and competition from other platforms, we model them as heterogeneous agents, each represented by an MDP.

- **State:** The pharmacist’s engagement level with the platform
- **Action:** Binary decision of whether to provide a coupon to a given pharmacist
- **Action-dependent State Transitions:** Based on whether an action is taken or not, the agent’s engagement evolves accordingly based on environmental factors and platform interactions
- **Reward:** Engagement metrics (e.g., time spent on platform, order frequency)
- **Budget:** The platform has a budget on the number of coupons available at each day.

B WHITTLE’S INDEX

Whittle’s index is computed for each agent. At each time t , given \mathbf{s}_t , the top B agents with the highest indices are selected. Whittle’s index policy has two variants, each aimed at maximizing a different objective in an *infinite* horizon setting: one for maximizing average reward and the other for maximizing discounted reward.

Average Reward: In this case, Whittle’s index policy is designed to approximately find the policy π_d that maximizes the average reward,

$$\max_{\pi_d} \liminf_{T \rightarrow \infty} \frac{1}{TM} \mathbb{E} \left[\sum_{m=1}^M \sum_{t=0}^T R^m(s_t^m) \mid \mathbf{s}_0 \right]$$

$$\text{s.t. } \sum_{m=1}^M a_t^m = B, \forall t \in [T]. \quad (\text{B.5})$$

We will denote J^* to be the optimal value of Problem B.5. Instead of directly solving Problem B.5, Whittle deals with the Lagrangian relaxation of Problem B.5, which is stated in Problem B.6

$$\max_{\pi_d} \liminf_{T \rightarrow \infty} \frac{1}{TM} \mathbb{E} \left[\sum_{m=1}^M \sum_{t=0}^T R^m(s_t^m) + \lambda (1 - a_{t+1}^m) \mid \mathbf{s}_0 \right]. \quad (\text{B.6})$$

The key observation to derive Whittle's index policy is that Problem B.6 can be solved by addressing the following dynamic programming separately for each agent

$$\begin{aligned} V_\lambda^m(s) &= \max_{a \in \{0,1\}} Q_\lambda^m(s, a) \\ Q_\lambda^m(s, a) + \beta_m^* &= R^m(s) - \lambda \mathbb{I}\{a = 1\} \\ &\quad + \sum_{s' \in \mathcal{S}} P_a^m(s, s') V^m(s'), \end{aligned}$$

where β_m^* is the maximized average reward for arm m without any constraint, i.e.

$$\beta_m^* := \max \liminf_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=0}^T R^m(s_t^m) - \lambda \mathbb{I}\{a = 1\} \mid s_0^m \right].$$

The Whittle's index λ_s^m is defined as the multiplier λ_s^m when it is equally favorable between choosing action 1 and action 0, i.e. $Q_{\lambda_s^m}^m(s, 1) = Q_{\lambda_s^m}^m(s, 0)$. The Whittle's index policy would then give action to B agents with the highest Whittle's index at each time step t . Denote this as π^{Whittle} and the average reward of Whittle's index policy as J^{Whittle} . The performance of Whittle's index policy is characterized in Theorem B.1.

Theorem B.1 (Theorem 2, [Weber & Weiss 1990](#)). *Suppose all arms are homogeneous and assumptions therein, i.e., $P_a^m = P_a^{m'}$, $R^m = R^{m'}$, $\forall m \neq m'$, then the Whittle's index policy is asymptotically optimal,*

$$\lim_{M \rightarrow \infty} J^{\text{Whittle}} = \lim_{M \rightarrow \infty} J^*.$$

Discounted Reward: In this case, Whittle's index policy is designed to approximately find the policy π_d that maximizes the discounted reward for a fixed discount factor β ,

$$\begin{aligned} \max_{\pi_d} \sum_{t=0}^{\infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=1}^M \beta^t R^m(s_t^m) \mid \mathbf{s}_0 \right] \\ \text{s.t. } \sum_{m=1}^M a_t^m = B, \forall t \in [T]. \end{aligned} \quad (\text{B.7})$$

Similarly, instead of directly solving Problem B.7, Whittle deals with the Lagrangian relaxation of Problem B.7, which is stated in Problem B.8

$$\max_{\pi_d} \sum_{t=0}^{\infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=1}^M \beta^t (R^m(s_t^m) + \lambda (1 - a_{t+1}^m)) \mid \mathbf{s}_0 \right]. \quad (\text{B.8})$$

Again, the same key observation is that Problem B.8 can be solved by addressing the following dynamic programming separately for each agent

$$\begin{aligned} V_\lambda^m(s) &= \max_{a \in \{0,1\}} Q_\lambda^m(s, a) \\ Q_\lambda^m(s, a) &= R^m(s) - \lambda \mathbb{I}\{a = 1\} \\ &\quad + \beta \sum_{s' \in \mathcal{S}} P_a^m(s, s') V^m(s'). \end{aligned} \quad (\text{B.9})$$

The definition of Whittle's index will be the same in the average reward case.

B.1 Generalized Regret in Online Restless Bandits

Given K episodes each with a horizon length of T and the set of online Whittle’s index policy at each episode $k \in [K]$, $\omega := \{\omega_k\}_{k \in [K]}$, the regret $\mathcal{R}_{\text{Restless}}^{\pi^{\text{Whittle}}}(\omega)$:

$$\mathcal{R}_f(\omega) := \frac{1}{K} \sum_{k=1}^K \left(\mathbb{E}_{\pi^{\text{Whittle}}} \left[\sum_{m=1}^M \sum_{t=0}^T R^m(s_t^m) \mid \mathbf{s}_0 \right] - \mathbb{E}_{\omega_k} \left[\sum_{m=1}^M \sum_{t=0}^T R^m(s_t^m) \mid \mathbf{s}_0 \right] \right). \tag{B.10}$$

We note that our problem setting corresponds to the scenario where $K = 1$.

Existing regret bounds on online Whittle’s index policy are only meaningful when the number of episodes is sufficiently large. To illustrate this, Wang et al. (2023b) shows that their algorithm has a regret bound of the order $\mathcal{O}(T\sqrt{K})$. If we only have 1 episode, this will be equivalent to saying that the regret bound is $\mathcal{O}(T)$. The hardness lies in estimating the MDP: Auer et al. 2008 show that the lower bound of the regret is $\Omega(T\sqrt{K})$. This lower bound indicates that when directly dealing with an MDP, it is impossible to get a regret bound that is sublinear in T .

C COUNTER EXAMPLE FOR HETEROGENEOUS AGENTS

In this section, we will construct an instance involving heterogeneous agents where the greedy policy is not optimal. Say that each agent is associated with two state: State 0 and State 1, and the agent can take either Action 1 (intervention) or Action 0 (no intervention). Assume for now that there are 4 agents with different transition probabilities. The first two agents fall under Type 1: their transition probability under Action 1 is shown in Figure C.3 and their transition probability under Action 0 is shown in Figure C.4. The other two agents fall under Type 2: their transition probability under Action 1 is shown in Figure C.5 and their transition probability under Action 0 is shown in Figure C.6.




Figure C.3: Dynamic of Agent Type 1 with Intervention




Figure C.4: Dynamic of Agent Type 1 without Intervention

Suppose the budget equals one, meaning that we can only give one agent intervention at each time. We observe that under our problem setup Type 2 agents always have a higher incremental reward $\phi^2(1, 1) := \frac{3}{4} \cdot 1, \phi^2(0, 1) := \frac{3}{4} \cdot 1, \phi^2(s, 0) = 0, I^2(1) = I^2(0) = \frac{3}{4}$, opposed to $\phi^1(0, 1) = \frac{2}{3} \cdot 1, \phi^1(1, 1) = \frac{1}{2} \cdot 1, \phi^1(0, 0) = 0, \phi^1(1, 0) = \frac{1}{3}, I(1) = \frac{2}{3}, I(0) = \frac{1}{6}$. However, this does not imply that the optimal policy is to always take Action 1 for Type 2 agents. If all resources are allocated to Type 2 agents, the expected reward is $3/4$ per time step summed across all agents. Alternatively, if resources are devoted to Type 1 agents, alternating actions between the two Type 1 agents, the expected reward is $2/3 + 2/3 \times 1/3 = 8/9$ per Type 1 agent over two time periods (where an action is taken in one period but not in the other). To illustrate this point, we ran the following experiment: we set the horizon to




Figure C.5: Dynamic of Agent Type 2 with Intervention




Figure C.6: Dynamic of Agent Type 2 without Intervention

be 20 and repeated the simulation 300 times. We compare the performance of the following two policies in terms of cumulative rewards: Policy 1) choosing one of the Type 2 agents for intervention at each time, and Policy 2) alternating intervention between the two Type 1 agents. The observed cumulative for Policies 1 and 2 are 15.02 and 16.78, respectively. This shows that the greedy policy is not optimal under all heterogeneous agents cases.

D COUNTER EXAMPLE FOR THREE-STATE, HOMOGENEOUS AGENTS

In this section, we will construct an instance involving homogeneous agents with three states where the greedy policy is not optimal. Say that each agent is associated with three states: State 0, State 1 and State 2, and the agent can take either Action 1 (intervention) or Action 0 (no intervention). Assume for now that there are two agents with the same transition probabilities. Their transition probability under Action 1 is shown in Figure D.7 and their transition probability under Action 0 is shown in Figure D.8.




Figure D.7: Dynamic with Intervention




Figure D.8: Dynamic without Intervention

Suppose that the budget equals one, we observe that State 2 and State 1 have the same incremental reward $I(2) = \frac{3}{7} - \frac{5}{28} = \frac{1}{4} = \frac{3}{2} - \frac{5}{4} = I(1)$. However, we will show that when $T = 3$, greedy policy is not optimal.

Specifically, we will show that when $t = 0$ and $\mathbf{s}_0 = (1, 2)$, the optimal policy would be giving State 2 Action 1 instead of State 1 Action 1. This will show that greedy policy is not optimal as greedy policy could give State 1 Action 1 by definition. Suppose we are giving State 1 Action 1, then we have

$$\begin{aligned} V((1, 2)) &= 3 + \frac{5}{56} V_1^*((1, 2)) + \frac{5}{56} V_1^*((1, 1)) + \frac{23}{28} V_1^*((1, 0)) \\ &= 3 + \frac{5}{56} \frac{47 + 3 * 28}{28} + \frac{5}{56} \frac{19}{4} + \frac{23}{28} \frac{5}{2}. \end{aligned}$$

Similarly, if we are giving State 2 Action 1, then we have

$$\begin{aligned} V((1, 2)) &= 3 + \frac{1}{28} V_1^*((2, 2)) + \frac{1}{7} V_1^*((1, 2)) + \frac{15}{28} V_1^*((1, 0)) + \frac{5}{28} V_1^*((2, 0)) + \frac{3}{28} V_1^*((1, 1)) \\ &= 3 + \frac{1}{28} \left(\frac{17}{28} + 4 \right) + \frac{1}{7} \left(3 + \frac{47}{28} \right) + \frac{15}{28} \left(\frac{5}{2} \right) + \frac{5}{28} \left(\frac{3}{7} + 2 \right) + \frac{3}{28} \left(\frac{11}{4} + 2 \right). \end{aligned}$$

By calculation, we will have the results. This shows that the greedy policy is not optimal under three states setting.

E TECHNICAL DETAILS IN SECTION 3

First, recall that at each time t of an offline restless bandit problem, the learner observes the set of states of M agents, denoted by $\mathbf{s}_t := (s_t^1, s_t^2, \dots, s_t^M)$, and receives the corresponding reward $\mathbf{r}_t := (r_t^1, \dots, r_t^M)$. Based on the observed states, the learner determines the action $\mathbf{a}_{t+1} := (a_{t+1}^1, a_{t+1}^2, \dots, a_{t+1}^M)$, according to a deterministic policy $\pi_{t+1}^d(\mathbf{s}) : \mathcal{S}^M \rightarrow \{0, 1\}^M$. The learner is provided with a budget constraint, allowing at most B arms to be pulled at each time step, i.e., $\sum_{m \in [M]} a_{t+1}^m = B$. After pulling up to B arms, the learner observes the next state \mathbf{s}_{t+1} , generated according to $s_{t+1}^m \sim P_{a_{t+1}^m}^m(s_t^m, s_{t+1}^m)$, and receives a feedback \mathbf{r}_{t+1} .

To ease notation, in Section E we will use $q_a^m(s, s')$ to denote $P_a^m(s, s')$. Before presenting the proofs, we reformulate the online restless bandit problem as a single large MDP, where the state and action are defined as tuples comprising the composite states and actions of all agents. This reformulation is valid because the agents/arms in the restless bandit problem are independent. This MDP can be described as $(\mathcal{S}^M, \mathcal{F}, q, \mathbf{R})$, where \mathcal{F} is defined as

$$\mathcal{F} := \left\{ (a^1, a^2, \dots, a^M) : (a^1, a^2, \dots, a^M) \in \mathcal{A}^M, \sum_{m=1}^M a^m \leq B \right\}.$$

Let $q_{\mathbf{s}\mathbf{s}'}^{\mathbf{a}} := \mathbb{P}(\mathbf{s}' | \mathbf{s}, \mathbf{a}) = P_{\mathbf{a}}(\mathbf{s}, \mathbf{s}') = \prod_{m=1}^M q_{s s'}^a(m)$, and let $\mathbf{R}(\mathbf{s}) = \sum_{m=1}^M r^m(s^m)$. Then for each policy π and each time-step h , we can define the value function $V_h^\pi(\mathbf{s})$ as

$$\mathbb{E}_\pi \left[\sum_{t=h}^T \mathbf{R}(\mathbf{s}_t) \mid \mathbf{s}_h = \mathbf{s} \right]. \quad (\text{E.11})$$

With this definition, Problem (1) is equivalent to

$$\max_{\pi} V_0^\pi(\mathbf{s}_0).$$

Again, we use $\pi^* := \{\pi_t^*\}_{t \in [T]}$ to denote the maximizer, and let V_h^* be the value function under policy π^* . We will use this notation in Appendix E.1, and F.

E.1 Proof of Theorem 3.1

Theorem 3.1. *Assuming all agents are homogeneous, i.e., $P^m = P^{m'}, R^m = R^{m'}, \forall m, m' \in [M]$ and $\mathcal{S} = \{0, 1\}$, $\pi^{\text{greedy}} = \pi^*$ for any $0 \leq B \leq M < \infty$.*

Proof. WLOG, assume that $I(1) \leq I(0)$. For simplicity, we first look at the simple case where $M = 2, B = 1$. The generalization to multi-agent case is straightforward and will be deferred to the last part of the proof. First

of all, we will derive some useful equations that are useful in the proof.

$$\begin{aligned}
 & \sum_{s_{h+1}^1, s_{h+1}^2} q_{1s_{h+1}^1}^0 q_{0s_{h+1}^2}^1 V_{h+1}^* ((s_{h+1}^1, s_{h+1}^2)) - \sum_{s_{h+1}^1, s_{h+1}^2} q_{1s_{h+1}^1}^1 q_{0s_{h+1}^2}^0 V_{h+1}^* ((s_{h+1}^1, s_{h+1}^2)) \\
 &= q_{11}^0 q_{01}^1 V_{h+1}^* ((1, 1)) + q_{11}^0 q_{00}^1 V_{h+1}^* ((1, 0)) + q_{10}^0 q_{01}^1 V_{h+1}^* ((1, 0)) + q_{10}^0 q_{00}^1 V_{h+1}^* ((0, 0)) \\
 &\quad - q_{11}^1 q_{01}^0 V_{h+1}^* ((1, 1)) - q_{11}^1 q_{00}^0 V_{h+1}^* ((1, 0)) - q_{10}^1 q_{01}^0 V_{h+1}^* ((1, 0)) - q_{10}^1 q_{00}^0 V_{h+1}^* ((0, 0)) \\
 &= q_{11}^0 q_{01}^1 V_{h+1}^* ((1, 1)) - q_{11}^0 q_{01}^1 V_{h+1}^* ((1, 0)) + q_{11}^0 V_{h+1}^* ((1, 0)) \\
 &\quad + q_{01}^1 V_{h+1}^* ((1, 0)) - q_{11}^0 q_{01}^1 V_{h+1}^* ((1, 0)) \\
 &\quad + V_{h+1}^* ((0, 0)) + q_{11}^0 q_{01}^1 V_{h+1}^* ((0, 0)) - (q_{11}^0 + q_{01}^1) V_{h+1}^* ((0, 0)) \\
 &\quad - q_{11}^1 q_{01}^0 V_{h+1}^* ((1, 1)) + q_{11}^1 q_{01}^0 V_{h+1}^* ((1, 0)) - q_{11}^1 V_{h+1}^* ((1, 0)) \\
 &\quad - q_{01}^0 V_{h+1}^* ((1, 0)) + q_{11}^1 q_{01}^0 V_{h+1}^* ((1, 0)) \\
 &\quad - V_{h+1}^* ((0, 0)) - q_{11}^1 q_{01}^0 V_{h+1}^* ((0, 0)) + (q_{11}^1 + q_{01}^0) V_{h+1}^* ((0, 0)) \\
 &= (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &\quad + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))). \tag{E.12}
 \end{aligned}$$

Next, we will simplify the difference between the value function of different states, assuming that the optimal policy is the greedy policy. For simplicity, we will only give the detailed derivation of the equation for $V_h^* ((1, 1)) - V_h^* ((1, 0))$. The equation for $V_h^* ((1, 0)) - V_h^* ((0, 0))$ holds for similar reasons.

$$\begin{aligned}
 V_h^* ((1, 1)) - V_h^* ((1, 0)) &= 2R(1) + \sum_{s_{h+1}^1, s_{h+1}^2} q_{1s_{h+1}^1}^1 q(s_{h+1}^2 | 1, 0) V_{h+1}^* ((s_{h+1}^1, s_{h+1}^2)) \\
 &\quad - R(1) - R(0) - \sum_{s_{h+1}^1, s_{h+1}^2} q_{1s_{h+1}^1}^0 q_{0s_{h+1}^2}^1 V_{h+1}^* ((s_{h+1}^1, s_{h+1}^2)) \\
 &= R(1) - R(0) + q_{11}^1 q_{11}^0 V_{h+1}^* ((1, 1)) + q_{11}^1 q_{10}^0 V_{h+1}^* ((1, 0)) \\
 &\quad + q_{10}^1 q_{11}^0 V_{h+1}^* ((1, 0)) + q_{10}^1 q_{10}^0 V_{h+1}^* ((0, 0)) \\
 &\quad - q_{11}^0 q_{01}^1 V_{h+1}^* ((1, 1)) - q_{11}^0 q_{00}^1 V_{h+1}^* ((1, 0)) \\
 &\quad - q_{10}^0 q_{01}^1 V_{h+1}^* ((1, 0)) - q_{10}^0 q_{00}^1 V_{h+1}^* ((0, 0)) \\
 &= R(1) - R(0) + q_{11}^1 q_{11}^0 V_{h+1}^* ((1, 1)) + q_{11}^1 V_{h+1}^* ((1, 0)) - q_{11}^1 q_{11}^0 V_{h+1}^* ((1, 0)) \\
 &\quad + q_{11}^0 V_{h+1}^* ((1, 0)) - q_{11}^1 q_{11}^0 V_{h+1}^* ((1, 0)) \\
 &\quad + V_{h+1}^* ((0, 0)) + q_{11}^0 q_{11}^1 V_{h+1}^* ((0, 0)) - (q_{11}^1 + q_{11}^0) V_{h+1}^* ((0, 0)) \\
 &\quad - q_{11}^0 q_{01}^1 V_{h+1}^* ((1, 1)) - q_{11}^0 V_{h+1}^* ((1, 0)) + q_{01}^1 q_{11}^0 V_{h+1}^* ((1, 0)) \\
 &\quad - q_{01}^1 V_{h+1}^* ((1, 0)) + q_{11}^0 q_{01}^1 V_{h+1}^* ((1, 0)) \\
 &\quad - V_{h+1}^* ((0, 0)) - q_{11}^0 q_{01}^1 V_{h+1}^* ((0, 0)) + (q_{11}^0 + q_{01}^1) V_{h+1}^* ((0, 0)) \\
 &= R(1) - R(0) + (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &\quad + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))). \tag{E.13}
 \end{aligned}$$

Similarly, we can have

$$\begin{aligned}
 V_h^* ((1, 0)) - V_h^* ((0, 0)) &= R(1) - R(0) + (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &\quad + q_{01}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))). \tag{E.14}
 \end{aligned}$$

We will derive the final useful equation for the incremental reward.

$$\begin{aligned}
 I(1) &= \phi(1, 1) - \phi(1, 0) \\
 &= q_{11}^1 R(1) + q_{10}^1 R(0) - q_{11}^0 R(1) - q_{10}^0 R(0) \tag{E.15}
 \end{aligned}$$

$$= (q_{11}^1 - q_{11}^0) (R(1) - R(0)). \tag{E.16}$$

Similarly, we also have

$$I(0) = (q_{01}^1 - q_{01}^0) (R(1) - R(0)). \quad (\text{E.17})$$

In the following, we will consider four cases.

Case 1: $q_{11}^0 \geq q_{01}^0, q_{11}^1 \leq q_{01}^1, q_{11}^1 \geq q_{11}^0$. We will use backward induction to prove the following statements:

$$\forall t \leq T, 0 \leq V_t^* ((1, 0)) - V_t^* ((0, 0)) \leq \frac{R(1) - R(0)}{q_{10}^0} \quad (\text{E.18})$$

$$0 \leq V_t^* ((1, 1)) - V_t^* ((1, 0)) \leq R(1) - R(0) \quad (\text{E.19})$$

$$\pi_{t+1}^{\text{greedy}} = \pi_{t+1}^* \quad (\text{E.20})$$

The statements hold trivially when $t = T$. Suppose the statements hold when $t \geq h + 1$, when $t = h$, we will first prove Eq. (E.20) for $t = h$. Recall that the policy satisfying the following Bellman optimality equation is the optimal policy:

$$V_h^* ((s_h^1, s_h^2)) = R(s_h^1) + R(s_h^2) + \max_{a_{h+1}^1, a_{h+1}^2} \sum_{s_{h+1}^1, s_{h+1}^2} q_{s_h^1 s_{h+1}^1}^{a_{h+1}^1} q_{s_h^2 s_{h+1}^2}^{a_{h+1}^2} V_{h+1}^* (s_{h+1}^1, s_{h+1}^2). \quad (\text{E.21})$$

By the problem formulation, one of the $a_{h+1}^i = 1$. Therefore, it's easy to see that when $s_h^1 = s_h^2$, randomly giving action will satisfy this equation. We will focus on the case where $s_h^1 \neq s_h^2$. In that case, to show Eq. (E.20), we only need to show that

$$\sum_{s_{h+1}^1, s_{h+1}^2} q_{1s_{h+1}^1}^0 q_{0s_{h+1}^2}^1 V_{h+1}^* (s_{h+1}^1, s_{h+1}^2) \geq \sum_{s_{h+1}^1, s_{h+1}^2} q_{1s_{h+1}^1}^1 q_{0s_{h+1}^2}^0 V_{h+1}^* (s_{h+1}^1, s_{h+1}^2). \quad (\text{E.22})$$

By Eq. (E.12), this is equivalent to showing

$$\begin{aligned} & (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\ & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \geq 0. \end{aligned}$$

Then we have

$$\begin{aligned} & (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\ & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\ & \geq (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\ & - (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\ & = \frac{I(0) - I(1)}{R(1) - R(0)} (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) - \frac{(q_{11}^1 I(0) - q_{01}^1 I(1))}{R(1) - R(0)} (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\ & = \frac{(q_{10}^1 I(0) - q_{00}^1 I(1))}{R(1) - R(0)} (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\ & \geq 0, \end{aligned}$$

where the first inequality holds because of the induction hypothesis and $q_{11}^0 \geq q_{01}^0, q_{01}^1 \geq q_{11}^1$, the first equality holds because of Eq. (E.17) and Eq. (E.16) and the last inequality holds because $q_{10}^1 \geq q_{00}^1, I(0) \geq I(1)$ and the assumption that Eq. (E.18) holds for $t = h + 1$.

Next, we first show the LHS of Eq. (E.18) for $t = h$ under the assumption that Eq. (E.20) holds for $t = h$ and Eq. (E.18) holds for $t = h + 1$.

$$\begin{aligned} V_h^* ((1, 0)) - V_h^* ((0, 0)) & = R(1) - R(0) + (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\ & \quad + q_{01}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\ & \geq R(1) - R(0) + (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \end{aligned}$$

$$\begin{aligned}
 & -q_{01}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 & \geq R(1) - R(0) + q_{00}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 & \geq 0,
 \end{aligned}$$

where the first equality holds because of Eq. (E.14) and the last inequality holds since $q_{11}^0 \geq q_{01}^0$. Under the same condition, for the RHS of Eq. (E.18), we have

$$\begin{aligned}
 q_{10}^0 (V_h^* ((1, 0)) - V_h^* ((0, 0))) &= q_{10}^0 (R(1) - R(0)) \\
 &+ q_{10}^0 q_{00}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &+ q_{10}^0 q_{01}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0))) \\
 &\leq q_{10}^0 (R(1) - R(0)) \\
 &+ q_{00}^1 (q_{11}^0 - q_{01}^0) (R(1) - R(0)) \\
 &+ q_{10}^0 q_{01}^1 (q_{11}^0 - q_{01}^0) (R(1) - R(0)) \\
 &\leq q_{10}^0 (R(1) - R(0)) \\
 &+ q_{00}^1 q_{11}^0 (R(1) - R(0)) + q_{01}^1 q_{11}^0 (R(1) - R(0)) \\
 &= R(1) - R(0).
 \end{aligned}$$

Then, we only need to show Eq. (E.19) for $t = h$. Assuming that Eq. (E.20) holds for $t = h$ and Eq. (E.19) holds for $t = h + 1$, for the RHS, we have

$$\begin{aligned}
 V_h^* ((1, 1)) - V_h^* ((1, 0)) &= R(1) - R(0) + (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &+ q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\
 &= R(1) - R(0) + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0))) \\
 &+ q_{10}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &\leq R(1) - R(0),
 \end{aligned}$$

where the last inequality holds since $q_{11}^1 \leq q_{01}^1$. Under the same condition, for the LHS, we have

$$\begin{aligned}
 V_h^* ((1, 1)) - V_h^* ((1, 0)) &= R(1) - R(0) + (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &+ q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\
 &= R(1) - R(0) + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0))) \\
 &+ q_{10}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 &\geq R(1) - R(0) + q_{11}^0 (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \\
 &+ (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \\
 &= (q_{00}^1 + q_{11}^1 - q_{11}^0 q_{01}^1 + q_{11}^0 q_{11}^1) (R(1) - R(0)) \\
 &\geq 0,
 \end{aligned}$$

where the last inequality holds since $q_{11}^1 \geq q_{11}^0$.

Case 2: $q_{11}^0 \leq q_{01}^0, q_{11}^1 \leq q_{01}^1$.

First we consider the sub-case where $\boxed{q_{11}^1 q_{01}^0 \leq q_{01}^1 q_{11}^0}$. We will use backward induction to prove the following statements:

$$\forall t \leq T, 0 \leq V_t^* ((1, 0)) - V_t^* ((0, 0)) \leq R(1) - R(0) \quad (\text{E.23})$$

$$0 \leq V_t^* ((1, 1)) - V_t^* ((1, 0)) \leq R(1) - R(0) \quad (\text{E.24})$$

$$\pi_{t+1}^{\text{greedy}} = \pi_{t+1}^*. \quad (\text{E.25})$$

When $t = T$, the statements hold trivially. Suppose the statements hold for $t \geq h + 1$, when $t = h$, we will first prove Eq. (E.25) under the assumption that Eq. (E.23) and Eq. (E.24) holds for $t = h + 1$. Similarly we only

need to show Eq. (E.22).

$$\begin{aligned}
& (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
& + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\
& \geq \frac{(q_{10}^1 I(0) - q_{00}^1 I(1))}{R(1) - R(0)} (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
& \geq 0,
\end{aligned}$$

where the first inequality holds because the assumption we make and the last inequality holds because $q_{10}^1 \geq q_{00}^1$, $I(0) \geq I(1)$ and the assumption that Eq. (E.23) holds.

Next we want to prove the RHS of Eq. (E.24) under the assumption that Eq. (E.25) holds for $t = h$ and that Eq. (E.23) and Eq. (E.24) hold for $t = h + 1$.

$$\begin{aligned}
V_h^* ((1, 1)) - V_h^* ((1, 0)) &= R(1) - R(0) + (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
&+ q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\
&= R(1) - R(0) + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0))) \\
&+ q_{10}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
&\leq R(1) - R(0).
\end{aligned}$$

Under the same condition, for the LHS of Eq. (E.24), we have

$$\begin{aligned}
V_h^* ((1, 1)) - V_h^* ((1, 0)) &= R(1) - R(0) + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0))) \\
&+ q_{10}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
&\geq R(1) - R(0) + q_{10}^0 (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \\
&+ q_{11}^0 (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \\
&= (q_{11}^1 + q_{00}^1)(R(1) - R(0)) \\
&\geq 0.
\end{aligned}$$

Similarly, we can show Eq. (E.23) for $t = h$ under the same condition.

Secondly, we consider the sub-case where $\boxed{q_{11}^1 q_{01}^0 \geq q_{01}^1 q_{11}^0}$. We want to use backward induction to prove the following statements:

$$\forall t \leq T, 0 \leq V_t^* ((1, 0)) - V_t^* ((0, 0)) \leq R(1) - R(0) \quad (\text{E.26})$$

$$0 \leq V_t^* ((1, 1)) - V_t^* ((1, 0)) \leq R(1) - R(0) \quad (\text{E.27})$$

$$V_t^* ((1, 1)) - V_t^* ((1, 0)) \leq V_t^* ((1, 0)) - V_t^* ((0, 0)) \quad (\text{E.28})$$

$$\pi_{t+1}^{\text{greedy}} = \pi_{t+1}^* \quad (\text{E.29})$$

When $t = T$, the statements hold trivially. Suppose this holds for $t \geq h + 1$, when $t = h$, we will first show Eq. (E.29) for $t = h$ under the assumption that Eq. (E.26), Eq. (E.27) and Eq. (E.28) hold for $t = h + 1$. Similarly, this is equivalent to show Eq. (E.22).

$$\begin{aligned}
& (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
& + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\
& = \frac{I(0) - I(1)}{R(1) - R(0)} (V_{h+1}^* (1, 0) - V_{h+1}^* (0, 0)) \\
& + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\
& \geq 0,
\end{aligned}$$

where the first equality holds because of Eq. (E.17) and Eq. (E.16) and the first inequality holds since $q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 \leq 0$ and the assumption that Eq. (E.28) holds for $t = h + 1$.

Using the same logic as the previous case, we can prove that

$$\begin{aligned} 0 &\leq V_h^*((1, 0)) - V_h^*((0, 0)) \leq R(1) - R(0) \\ 0 &\leq V_h^*((1, 1)) - V_h^*((1, 0)) \leq R(1) - R(0). \end{aligned}$$

Therefore, the only thing we need to show is the Eq. (E.28) for $t = h$ under the assumption that Eq. (E.29) holds for $t = h$ and that Eq. (E.28) holds for $t = h + 1$. We have

$$\begin{aligned} &V_h^*((1, 1)) - 2V_h^*((1, 0)) + V_h^*((0, 0)) \\ &= q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 1)) - 2V_{h+1}^*((1, 0)) + V_{h+1}^*((0, 0))) \\ &+ (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))) \\ &- q_{01}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^*((1, 1)) - 2V_{h+1}^*((1, 0)) + V_{h+1}^*((0, 0))) \\ &- (q_{11}^0 - q_{01}^0) (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))) \\ &= q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 1)) - 2V_{h+1}^*((1, 0)) + V_{h+1}^*((0, 0))) \\ &- q_{01}^1 (q_{11}^0 - q_{01}^0) (V_{h+1}^*((1, 1)) - 2V_{h+1}^*((1, 0)) + V_{h+1}^*((0, 0))) \\ &+ \frac{I(1) - I(0)}{R(1) - R(0)} (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))). \end{aligned}$$

Since $I(1) \leq I(0)$, $V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0)) \geq 0$, we only need to prove that

$$q_{11}^0 (q_{11}^1 - q_{01}^1) \geq q_{01}^1 (q_{11}^0 - q_{01}^0).$$

Since we have

$$q_{11}^1 q_{01}^0 \geq q_{01}^1 q_{11}^0,$$

we have

$$\begin{aligned} q_{11}^1 q_{01}^0 - q_{11}^0 q_{01}^1 &\geq q_{01}^1 q_{11}^0 - q_{11}^1 q_{01}^1 \\ q_{11}^1 (q_{01}^0 - q_{11}^0) &\geq q_{11}^0 (q_{01}^1 - q_{11}^1) \\ q_{01}^1 (q_{01}^0 - q_{11}^0) &\geq q_{11}^1 (q_{01}^0 - q_{11}^0) \geq q_{11}^0 (q_{01}^1 - q_{11}^1) \\ q_{01}^1 (q_{11}^0 - q_{01}^0) &\leq q_{11}^0 (q_{11}^1 - q_{01}^1). \end{aligned}$$

This will give us the desired result.

Case 3: $q_{11}^1 \geq q_{01}^1, q_{11}^0 \geq q_{01}^0$.

We will use backward induction to prove the following statements:

$$\forall t \leq T, \frac{R(1) - R(0)}{q_{10}^1 + q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-t+1}\right) \leq V_t^*((1, 0)) - V_t^*((0, 0)) \quad (\text{E.30})$$

$$\frac{R(1) - R(0)}{q_{10}^0 + q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-t+1}\right) \geq V_t^*((1, 0)) - V_t^*((0, 0)) \quad (\text{E.31})$$

$$\frac{R(1) - R(0)}{q_{10}^1 + q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-t+1}\right) \leq V_t^*((1, 0)) - V_t^*((0, 0)) \quad (\text{E.32})$$

$$\frac{R(1) - R(0)}{q_{10}^0 + q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-t+1}\right) \geq V_t^*((1, 0)) - V_t^*((0, 0)) \quad (\text{E.33})$$

$$\pi_t^{\text{greedy}} = \pi_t^*. \quad (\text{E.34})$$

When $t = T$, this holds trivially, suppose this holds for $t \geq h + 1$. When $t = h$, we will first show that Eq. (E.34) holds for $t = h$ under the assumption that Eq. (E.30), Eq. (E.31), Eq. (E.32) and Eq. (E.33) hold for $t = h + 1$. This again is equivalent to show that

$$\begin{aligned} &(q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))) \\ &+ (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^*((1, 1)) - 2V_{h+1}^*((1, 0)) + V_{h+1}^*((0, 0))) \geq 0. \end{aligned}$$

We have

$$\begin{aligned}
 & (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - 2V_{h+1}^* ((1, 0)) + V_{h+1}^* ((0, 0))) \\
 & = (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0 - q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0))).
 \end{aligned}$$

If $q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0 - q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 \geq 0$, then using our assumption that, $V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0)) \geq 0$, $V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0)) \geq 0$, will give us the desired result.

If $q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0 - q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 \leq 0$, with our hypothesis, we have

$$\begin{aligned}
 & (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0 - q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 0)) - V_{h+1}^* ((0, 0))) \\
 & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (V_{h+1}^* ((1, 1)) - V_{h+1}^* ((1, 0))) \\
 & \geq (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0 - q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \\
 & \times \frac{R(1) - R(0)}{q_{10}^0 + q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-h}\right) \\
 & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \frac{R(1) - R(0)}{q_{10}^1 + q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-h}\right). \tag{E.35}
 \end{aligned}$$

Denote $\frac{R(1)-R(0)}{q_{10}^0+q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-h}\right)$ as d_{upper} and let k be defined as

$\frac{R(1)-R(0)}{q_{10}^1+q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-h}\right) / d_{\text{upper}}$. Then to show that Eq. (E.35) is larger than 0, we only need to show that

$$q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0 - (k-1) (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \geq 0,$$

which is equivalent to show $k \geq \frac{q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 - q_{11}^0 - q_{01}^1 + q_{11}^1 + q_{01}^0}{q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1}$.

First, we will show

$$\frac{q_{10}^0 + q_{01}^0}{q_{10}^1 + q_{01}^1} \geq \frac{q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 - q_{11}^0 - q_{01}^1 + q_{11}^1 + q_{01}^0}{q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1}.$$

To do this, we only need to show

$$\begin{aligned}
 & (q_{10}^0 + q_{01}^0) (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \\
 & - (q_{10}^1 + q_{01}^1) (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \\
 & - (q_{10}^1 + q_{01}^1) (q_{11}^1 + q_{01}^0 - q_{11}^0 - q_{01}^1) \geq 0.
 \end{aligned}$$

We have

$$\begin{aligned}
 & (q_{10}^0 + q_{01}^0) (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \\
 & - (q_{10}^1 + q_{01}^1) (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \\
 & - (q_{10}^1 + q_{01}^1) (q_{11}^1 + q_{01}^0 - q_{11}^0 - q_{01}^1) \\
 & = (q_{10}^0 + q_{01}^0 - q_{10}^1 - q_{01}^1) (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) \\
 & - (q_{10}^1 + q_{01}^1) (q_{11}^1 + q_{01}^0 - q_{11}^0 - q_{01}^1) \\
 & = (q_{11}^1 + q_{01}^0 - q_{11}^0 - q_{01}^1) (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 - q_{10}^1 - q_{01}^1).
 \end{aligned}$$

Since $q_{11}^1 + q_{01}^0 - q_{11}^0 - q_{01}^1 = \frac{I(1)-I(0)}{R(1)-R(0)} \leq 0$ and $q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 \leq 0$, we have the desired result.

Second, we want to show that $1 - (q_{11}^1 - q_{01}^1)^{T-h} \geq 1 - (q_{11}^0 - q_{01}^0)^{T-h}$. To do this, we only need to show that $(q_{11}^1 - q_{01}^1)^{T-h} \leq (q_{11}^0 - q_{01}^0)^{T-h}$, which comes directly from the fact that $0 \leq q_{11}^1 - q_{01}^1 \leq q_{11}^0 - q_{01}^0$.

Then we have

$$k = \frac{q_{10}^0 + q_{01}^0}{q_{10}^1 + q_{01}^1} \frac{1 - (q_{11}^1 - q_{01}^1)^{T-h}}{1 - (q_{11}^0 - q_{01}^0)^{T-h}}$$

$$\begin{aligned}
 &\geq \frac{q_{10}^0 + q_{01}^0}{q_{10}^1 + q_{01}^1} \\
 &\geq \frac{q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1 - q_{11}^0 - q_{01}^1 + q_{11}^1 + q_{01}^0}{q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1}.
 \end{aligned}$$

This will give us the desired result.

Next, we will first prove Eq. (E.30) for $t = h$ under the assumption that Eq. (E.30), Eq. (E.31), Eq. (E.32) and Eq. (E.33) hold for $t = h + 1$ and Eq. (E.29) holds for $t = h$. We have

$$\begin{aligned}
 V_h^*((1, 1)) - V_h^*((1, 0)) &= R(1) - R(0) + (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))) \\
 &\quad + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 1)) - 2V_{h+1}^*((1, 0)) + V_{h+1}^*((0, 0))) \\
 &= R(1) - R(0) + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 1)) - V_{h+1}^*((1, 0))) \\
 &\quad + q_{10}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))) \\
 &\geq R(1) - R(0) \\
 &\quad + (q_{11}^1 - q_{01}^1) \frac{R(1) - R(0)}{q_{10}^1 + q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-h}\right) \\
 &= \frac{R(1) - R(0)}{q_{10}^1 + q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-h+1}\right).
 \end{aligned}$$

Under the same condition, for Eq. (E.31) when $t = h$, we have

$$\begin{aligned}
 V_h^*((1, 1)) - V_h^*((1, 0)) &= R(1) - R(0) + (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))) \\
 &\quad + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 1)) - 2V_{h+1}^*((1, 0)) + V_{h+1}^*((0, 0))) \\
 &= R(1) - R(0) + q_{11}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 1)) - V_{h+1}^*((1, 0))) \\
 &\quad + q_{10}^0 (q_{11}^1 - q_{01}^1) (V_{h+1}^*((1, 0)) - V_{h+1}^*((0, 0))) \\
 &\leq R(1) - R(0) \\
 &\quad + (q_{11}^0 - q_{01}^0) \frac{R(1) - R(0)}{q_{10}^0 + q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-h}\right) \\
 &= \frac{R(1) - R(0)}{q_{10}^0 + q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-h+1}\right),
 \end{aligned}$$

where the inequality holds because we have $q_{11}^1 - q_{01}^1 \leq q_{11}^0 - q_{01}^0$. Using the same logic will give us the Eq. (E.33) and Eq. (E.32) for $t = h$.

Case 4: Suppose $q_{11}^1 \leq q_{01}^1, q_{11}^0 \geq q_{01}^0, q_{11}^1 \geq q_{11}^0$. We will directly show that $V_t^{\pi^{\text{greedy}}}$ satisfies the optimal Bellman equation. With the same logic, we only need to prove that Eq. (E.22) holds for each $t \leq T$. We will first assume that the following holds and will prove this later:

$$V_t^{\pi^{\text{greedy}}}((1, 1)) + V_t^{\pi^{\text{greedy}}}((0, 0)) - 2V_t^{\pi^{\text{greedy}}}((1, 0)) \leq 0, \quad \forall t \leq T. \quad (\text{E.36})$$

Then we have

$$\begin{aligned}
 &(q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) - V_{t+1}^{\pi^{\text{greedy}}}((0, 0))\right) \\
 &\quad + (q_{01}^1 q_{11}^0 - q_{01}^0 q_{11}^1) \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 1)) + V_{t+1}^{\pi^{\text{greedy}}}((0, 0)) - 2V_{t+1}^{\pi^{\text{greedy}}}((1, 0))\right) \\
 &= (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) - V_{t+1}^{\pi^{\text{greedy}}}((0, 0))\right) \\
 &\quad + (q_{11}^1 (q_{11}^0 - q_{01}^0) - q_{11}^0 (q_{11}^1 - q_{01}^1)) \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 1)) + V_{t+1}^{\pi^{\text{greedy}}}((0, 0)) - 2V_{t+1}^{\pi^{\text{greedy}}}((1, 0))\right) \\
 &\geq (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) - V_{t+1}^{\pi^{\text{greedy}}}((0, 0))\right) \\
 &\quad + (q_{01}^1 (q_{11}^0 - q_{01}^0) - q_{11}^0 (q_{11}^1 - q_{01}^1)) \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 1)) + V_{t+1}^{\pi^{\text{greedy}}}((0, 0)) - 2V_{t+1}^{\pi^{\text{greedy}}}((1, 0))\right)
 \end{aligned}$$

$$\begin{aligned}
 &= -V_t^{\pi^{\text{greedy}}}((1, 1)) - V_t^{\pi^{\text{greedy}}}((0, 0)) + 2V_t^{\pi^{\text{greedy}}}((1, 0)) \\
 &\geq 0.
 \end{aligned}$$

Next, we will prove Eq. (E.36). To ease notation, we will let $a := 1 - q_{01}^1, b := 1 - q_{11}^0, c = q_{11}^0 - q_{01}^0, d = q_{11}^1 - q_{01}^1$. By Eq. (E.13) and (E.14), we have the following iterative equations

$$\begin{aligned}
 V_t^{\pi^{\text{greedy}}}((1, 0)) - V_t^{\pi^{\text{greedy}}}((0, 0)) &= R(1) - R(0) + ac \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) - V_{t+1}^{\pi^{\text{greedy}}}((0, 0)) \right) \\
 &\quad + (1 - a)c \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 1)) - V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) \right) \\
 V_t^{\pi^{\text{greedy}}}((1, 1)) - V_t^{\pi^{\text{greedy}}}((1, 0)) &= R(1) - R(0) + bd \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) - V_{t+1}^{\pi^{\text{greedy}}}((0, 0)) \right) \\
 &\quad + (1 - b)d \left(V_{t+1}^{\pi^{\text{greedy}}}((1, 1)) - V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) \right).
 \end{aligned}$$

Let $\mathbf{V}_t = \begin{pmatrix} V_t^{\pi^{\text{greedy}}}((1, 0)) - V_t^{\pi^{\text{greedy}}}((0, 0)) \\ V_t^{\pi^{\text{greedy}}}((1, 1)) - V_t^{\pi^{\text{greedy}}}((1, 0)) \end{pmatrix}, A = \begin{pmatrix} ac & (1 - a)c \\ bd & (1 - b)d \end{pmatrix}$. We have

$$\mathbf{V}_t = (R(1) - R(0))\mathbf{1} + A\mathbf{V}_{t+1}.$$

First we consider the subcase where $I - A$ is not invertible, which is A has eigenvalue 1. This is equivalent to $\boxed{q_{11}^0 = 1, q_{01}^0 = q_{01}^1 = q_{11}^1 = 0}$. Under this case, we have $V_t^{\pi^{\text{greedy}}}((1, 1)) - V_t^{\pi^{\text{greedy}}}((1, 0)) = R(1) - R(0)$. We could also easily use induction to show that

$$V_t^{\pi^{\text{greedy}}}((1, 0)) - V_t^{\pi^{\text{greedy}}}((0, 0)) = R(1) - R(0) + V_{t+1}^{\pi^{\text{greedy}}}((1, 0)) - V_{t+1}^{\pi^{\text{greedy}}}((0, 0)) \geq R(1) - R(0).$$

This will give us the desired result.

Next, we assume $I - A$ is invertible, then we have

$$\begin{aligned}
 \mathbf{V}_t &= (R(1) - R(0))(I - A)^{-1}\mathbf{1} + (R(1) - R(0))A^{T-t}(\mathbf{1} - (I - A)^{-1}\mathbf{1}) \\
 &= (R(1) - R(0))(I - A)^{-1}\mathbf{1} - (R(1) - R(0))A^{T-t+1}(I - A)^{-1}\mathbf{1} \\
 &= \sum_{k=0}^{T-t} (R(1) - R(0))A^k\mathbf{1}.
 \end{aligned}$$

It is easy to see that A has two eigenvalues:

$$\lambda_1 = \frac{\eta + \sqrt{\eta^2 - 4\delta}}{2}, \lambda_2 = \frac{\eta - \sqrt{\eta^2 - 4\delta}}{2},$$

where $\eta = ac + (1 - b)d, \delta = (a - b)cd$.

Then, we first consider the case $\boxed{\lambda_i \in \mathbb{R}}$. We have

$$\begin{aligned}
 -V_t^{\pi^{\text{greedy}}}((1, 1)) - V_t^{\pi^{\text{greedy}}}((0, 0)) + 2V_t^{\pi^{\text{greedy}}}((1, 0)) &= (1, -1)\mathbf{V}_t\mathbf{1} \\
 &= (R(1) - R(0))(c - d) \sum_{k=0}^{T-t} \frac{\lambda_1^k - \lambda_2^k}{\lambda_1 - \lambda_2}.
 \end{aligned}$$

Since under Case 4, $c \geq d$, we only need to prove that $\sum \frac{\lambda_1^k - \lambda_2^k}{\lambda_1 - \lambda_2} \geq 0$.

When $\eta \geq 0, |\lambda_1| \geq |\lambda_2|$. Therefore, for all $k, \lambda_1^k \geq \lambda_2^k$, which will give us the desired result.

When $\eta \leq 0, \delta \geq 0$, we have $-1 \leq \eta \leq \lambda_2 \leq \lambda_1 \leq 0$. First, we will prove $\lambda_1 + \lambda_1^2 \geq \lambda_2 + \lambda_2^2$:

$$\begin{aligned}
 \lambda_1 + \lambda_1^2 - \lambda_2 - \lambda_2^2 &= (\lambda_1 - \lambda_2)(\lambda_1 + \lambda_2 + 1) \\
 &= \delta(1 + \eta)
 \end{aligned}$$

$$\geq 0.$$

Therefore, we have $\frac{\lambda_1}{\lambda_2} \leq \frac{1+\lambda_2}{1+\lambda_1}$. Then we can have for each odd k

$$\frac{\lambda_1^k}{\lambda_2^k} \leq \frac{\lambda_1}{\lambda_2} \leq \frac{1+\lambda_2}{1+\lambda_1}.$$

This means for each odd k , $\lambda_1^k + \lambda_1^{k+1} \geq \lambda_2^k + \lambda_2^{k+1}$. Then we have for all $T-t$, $\sum_{k=0}^{T-t} \lambda_1^k - \lambda_2^k \geq 0$. This will give us the desired result.

When $\eta \leq 0, \delta \leq 0$, we can numerically show that $\lambda_2 \geq -1$. Specifically, we use gradient descent with a sufficient small step size (10^{-5}) and multiple initialization by choosing the grid points to calculate the solutions of the KKT conditions. The results shows that the minimum is larger than -1 . Then we have $\sum_{k=0}^{T-t} \lambda_2^k \leq 1$. We also have that $\sum_{k=0}^{T-t} \lambda_1^k \geq 1$ since $\lambda_1 \geq 0$. This will give us the desired result.

Finally, we consider the subcase where $\boxed{\lambda_i \in \mathbb{C}}$. We note that this will only happen when $\delta \geq 0$ which means $q_{11}^0 \leq q_{01}^1$. We will first rewrite the eigenvalues using its spherical form, i.e. $\lambda_i = \rho \exp(\pm i\theta)$, where

$$\rho = \sqrt{(a-b)cd}, \cos(\theta) = \frac{ac + (1-b)d}{2\sqrt{(a-b)cd}}.$$

Then we have

$$\begin{aligned} & -V_t^{\pi^{\text{greedy}}}((1,1)) - V_t^{\pi^{\text{greedy}}}((0,0)) + 2V_t^{\pi^{\text{greedy}}}((1,0)) = (1,-1)\mathbf{V}_t\mathbf{1} \\ & = (R(1) - R(0))(c-d) \sum_{k=0}^{T-t} \frac{\lambda_1^k - \lambda_2^k}{\lambda_1 - \lambda_2} \\ & = (R(1) - R(0)) \frac{c-d}{1-2\rho\cos(\theta) + \rho^2} \left(1 - \rho^{T-t} \frac{\sin((T-t+1)\theta) - \rho\sin((T-t)\theta)}{\sin(\theta)} \right). \end{aligned}$$

Since $c-d \geq 0, 1-2\rho\cos(\theta) + \rho^2 = |1 - \rho \exp(i\theta)|^2 \geq 0$, we only need to show that

$$1 - \rho^{T-t} \frac{\sin((T-t+1)\theta) - \rho\sin((T-t)\theta)}{\sin(\theta)} \geq 0.$$

Since $0 \leq \rho \leq 1$, we only need to show the following:

$$\begin{aligned} \rho^{T-t} \frac{\sin((T-t+1)\theta) - \sin((T-t)\theta)}{\sin(\theta)} &\leq 1 \\ \rho^{T-t} \frac{\sin((T-t+1)\theta)}{\sin(\theta)} &\leq 1. \end{aligned}$$

We will use induction to prove the following:

$$\begin{aligned} \rho^k \frac{\sin((k+1)\theta) - \sin((k)\theta)}{\sin(\theta)} &\leq 1 \\ \rho^k \frac{\sin((k+1)\theta)}{\sin(\theta)} &\leq 1. \end{aligned}$$

When $k=0$, the statements trivially hold. Suppose it holds for $k=t$, when $k=t+1$, we have

$$\begin{aligned} \rho^{t+1} \frac{\sin((t+2)\theta) - \sin((t+1)\theta)}{\sin(\theta)} &= \rho^{t+1} \frac{\cos((t+1)\theta + \theta/2)}{\cos(\theta/2)} \\ &= \rho^{t+1} \frac{\cos(t\theta + \theta/2) \cos(\theta) - \sin(t\theta + \theta/2) \sin(\theta)}{\cos(\theta/2)} \\ &\leq \rho \cos(\theta) + 2 \sin(\theta/2) \rho. \end{aligned}$$

Similarly, we have

$$\begin{aligned}
 \rho^{t+1} \frac{\sin((t+2)\theta)}{\sin(\theta)} &= \rho^{t+1} \frac{\sin((t+1)\theta) \cos(\theta) + \sin(\theta) \cos((t+1)\theta)}{\sin(\theta)} \\
 &\leq \rho \cos(\theta) + \rho \\
 &= \frac{1}{2} \left(ac + (1-b)d + 2\sqrt{(q_{01}^1 - q_{11}^1)(q_{01}^1 - q_{11}^0)(q_{11}^0 - q_{01}^0)} \right) \\
 &\leq \frac{1}{2} \left((1 - q_{01}^1)q_{11}^0 + q_{11}^0(q_{11}^1 - q_{01}^1) + 2(q_{01}^1 - q_{11}^1) \right) \\
 &\leq \frac{1}{2} \left((1 - q_{01}^1)q_{11}^0 + (2 - q_{11}^0)q_{01}^1 \right) \\
 &\leq 1.
 \end{aligned}$$

Now we are going to prove the theorem for arbitrary number of agents and budget. We will again prove the result under three cases.

Case 1: Suppose $q_{11}^1 \leq q_{01}^1, q_{11}^0 \geq q_{01}^0, q_{11}^1 \geq q_{11}^0$.

We will first define three sets: $\mathcal{S}_1^{M-2}, \mathcal{S}_2^{M-2}, \mathcal{S}_3^{M-2}$.

$$\begin{aligned}
 \mathcal{S}_1^{M-2} &:= \{\mathbf{s} : \mathbf{s} \in \mathcal{S}^{M-2}, n_{\mathbf{s}} \geq B\} \\
 \mathcal{S}_2^{M-2} &:= \{\mathbf{s} : \mathbf{s} \in \mathcal{S}^{M-2}, n_{\mathbf{s}} = B - 1\} \\
 \mathcal{S}_3^{M-2} &:= \{\mathbf{s} : \mathbf{s} \in \mathcal{S}^{M-2}, n_{\mathbf{s}} \leq B - 2\},
 \end{aligned}$$

where $n_{\mathbf{s}}$ denotes the number of state 0 of \mathbf{s} .

We will use backward induction to prove that

$$\forall t \leq T, \forall \mathbf{s} \in \mathcal{S}_3^{M-2}, 0 \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \leq R(1) - R(0) \quad (\text{E.37})$$

$$0 \leq V_t^*((1, 1, \mathbf{s})) - V_t^*((1, 0, \mathbf{s})) \leq R(1) - R(0) \quad (\text{E.38})$$

$$\forall \mathbf{s} \in \mathcal{S}_2^{M-2}, 0 \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \leq \frac{R(1) - R(0)}{q_{10}^0} \quad (\text{E.39})$$

$$0 \leq V_t^*((1, 1, \mathbf{s})) - V_t^*((1, 0, \mathbf{s})) \leq R(1) - R(0) \quad (\text{E.40})$$

$$\forall \mathbf{s} \in \mathcal{S}_1^{M-2}, 0 \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \leq \frac{R(1) - R(0)}{q_{10}^0} \quad (\text{E.41})$$

$$0 \leq V_t^*((1, 1, \mathbf{s})) - V_t^*((1, 0, \mathbf{s})) \leq \frac{R(1) - R(0)}{q_{10}^0} \quad (\text{E.42})$$

$$\pi_{t+1}^{\text{greedy}} = \pi_{t+1}^* \quad (\text{E.43})$$

This holds trivially when $t = T$. Suppose $t \geq h + 1$, this holds, we will first prove Eq. (E.43) for $t = h$ under the assumption that Eq. (E.37), Eq. (E.38), Eq. (E.39), Eq. (E.40), Eq. (E.41) and Eq. (E.42) hold for $t = h + 1$. Similarly we have the following Bellman optimality equation.

$$\begin{aligned}
 V_h^*((1, 0, \mathbf{s}_h)) &= R(1) + R(0) + \sum_{\mathbf{s} \in \mathbf{s}_h} R(\mathbf{s}) \\
 &+ \max_{a_{h+1}^1, a_{h+1}^2, \dots, a_{h+1}^M} \sum_{s_{h+1}^1, s_{h+1}^2} q_{1, s_{h+1}^1}^{a_{h+1}^1} q_{0, s_{h+1}^2}^{a_{h+1}^2} \sum_{\mathbf{s}_{h+1}} q_{\mathbf{s}_h, \mathbf{s}_{h+1}}^{a_{h+1}} V_{h+1}^*((s_{h+1}^1, s_{h+1}^2, \mathbf{s}_{h+1})).
 \end{aligned}$$

We want to show that fix any policy of choosing $a_{h+1}^3, \dots, a_{h+1}^M$, giving action 1 to agent who is in state 0 will give a higher value, which is

$$\forall a_{h+1}, \sum_{s_{h+1}^1, s_{h+1}^2} q_{1, s_{h+1}^1}^0 q_{0, s_{h+1}^2}^1 \sum_{\mathbf{s}_{h+1}} q_{\mathbf{s}_h, \mathbf{s}_{h+1}}^{a_{h+1}} V_{h+1}^*((s_{h+1}^1, s_{h+1}^2, \mathbf{s}_{h+1}))$$

$$\geq \sum_{s_{h+1}^1, s_{h+1}^2} q_{1s_{h+1}^1}^1 q_{0s_{h+1}^2}^0 \sum_{\mathbf{s}_{h+1}} q_{\mathbf{s}_h, \mathbf{s}_{h+1}}^{\mathbf{a}_{h+1}} V_{h+1}^* ((s_{h+1}^1, s_{h+1}^2, \mathbf{s}_{h+1})). \quad (\text{E.44})$$

Suppose Eq. (E.44) holds, starting from any policy of allocating action that is not greedy, we can find an agent who is in state 1 and gets action 1 while another agent who is in state 0 gets action 0. By using the result, switching the action between them will give us a better policy. Suppose there are less than B agents who are in state 0, then we can keep switching until all agents who are in state 0 are given action 1, this is the greedy policy. Suppose there are more than B agents who are in state 0, then we can keep switching until B agents who are in state 0 get action 1, which is the greedy policy.

Now, we need to prove Eq. (E.44). Eq. (E.44) is equivalent to

$$\begin{aligned} & (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (F_{\mathbf{a}_{h+1}, h+1}((1, 1)) - F_{\mathbf{a}_{h+1}, h+1}((1, 0))) \\ & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (F_{\mathbf{a}_{h+1}, h+1}((1, 1)) - 2F_{\mathbf{a}_{h+1}, h+1}((1, 0)) + F_{\mathbf{a}_{h+1}, h+1}((0, 0))) \geq 0, \end{aligned} \quad (\text{E.45})$$

where $F_{\mathbf{a}_{h+1}, h+1}((s_{h+1}^1, s_{h+1}^2)) = \sum_{\mathbf{s}_{h+1}} q_{\mathbf{s}_h, \mathbf{s}_{h+1}}^{\mathbf{a}_{h+1}} V_{h+1}^* ((s_{h+1}^1, s_{h+1}^2, \mathbf{s}_{h+1}))$. The reason why this holds is similar to what we proved in binary agent case. We know from the assumption that Eq. (E.37)-(E.42) hold and the definition of F that

$$\begin{aligned} 0 & \leq F_{\mathbf{a}_{h+1}, h+1}((1, 1)) - F_{\mathbf{a}_{h+1}, h+1}((1, 0)) \leq \frac{R(1) - R(0)}{q_{10}^0} \\ 0 & \leq F_{\mathbf{a}_{h+1}, h+1}((1, 0)) - F_{\mathbf{a}_{h+1}, h+1}((0, 0)) \leq \frac{R(1) - R(0)}{q_{10}^0}. \end{aligned}$$

Then, we have

$$\begin{aligned} & (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (F_{\mathbf{a}_{h+1}, h+1}((1, 0)) - F_{\mathbf{a}_{h+1}, h+1}((0, 0))) \\ & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (F_{\mathbf{a}_{h+1}, h+1}((1, 1)) - 2F_{\mathbf{a}_{h+1}, h+1}((1, 0)) + F_{\mathbf{a}_{h+1}, h+1}((0, 0))) \\ & \geq (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (F_{\mathbf{a}_{h+1}, h+1}((1, 0)) - F_{\mathbf{a}_{h+1}, h+1}((0, 0))) \\ & - (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (F_{\mathbf{a}_{h+1}, h+1}((1, 0)) - F_{\mathbf{a}_{h+1}, h+1}((0, 0))) \\ & = \frac{I(0) - I(1)}{R(1) - R(0)} (F_{\mathbf{a}_{h+1}, h+1}((1, 0)) - F_{\mathbf{a}_{h+1}, h+1}((0, 0))) - \frac{(q_{11}^1 I(0) - q_{01}^1 I(1))}{R(1) - R(0)} \\ & (F_{\mathbf{a}_{h+1}, h+1}((1, 0)) - F_{\mathbf{a}_{h+1}, h+1}((0, 0))) \\ & = \frac{(q_{10}^1 I(0) - q_{00}^1 I(1))}{R(1) - R(0)} (F_{\mathbf{a}_{h+1}, h+1}((1, 0)) - F_{\mathbf{a}_{h+1}, h+1}((0, 0))) \\ & \geq 0. \end{aligned}$$

This shows Eq. (E.44), which, by our argument, shows Eq. (E.43) for $t = h$. Next, we will first define the specific greedy policy instead of randomly choosing agents who have higher incremental reward at time h for different \mathbf{s} . When $\mathbf{s} \in \mathcal{S}_1^{M-2}$, we will only give action 1 to agent 3, 4, \dots , M ; when $\mathbf{s} \in \mathcal{S}_2^{M-2}$, we will give one action 1 to agent 1, 2; when $\mathbf{s} \in \mathcal{S}_3^{M-2}$, we will give two action 1 to agent 1, 2.

For $\mathbf{s} \in \mathcal{S}_1^{M-2}$, we will show Eq. (E.41) holds for $t = h$ under the assumption that Eq.(E.37)-(E.42) hold for $t = h + 1$ and Eq. (E.43) holds for $t = h$. We have

$$\begin{aligned} V_h^* ((1, 0, \mathbf{s})) - V_h^* ((0, 0, \mathbf{s})) & = R(1) - R(0) + (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\ & + q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 1)) - 2F_{h+1}^* ((1, 0)) + F_{h+1}^* ((0, 0))) \\ & = R(1) - R(0) + q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 1)) - F_{h+1}^* ((1, 0))) \\ & + q_{00}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\ & \geq 0. \end{aligned}$$

Next, we will show the RHS of Eq. (E.41) for $t = h$ under the same condition.

$$q_{10}^0 (V_h^* ((1, 0, \mathbf{s})) - V_h^* ((0, 0, \mathbf{s}))) = q_{10}^0 (R(1) - R(0)) + q_{10}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0)))$$

$$\begin{aligned}
 & + q_{10}^0 q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 1)) - 2F_{h+1}^* ((1, 0)) + F_{h+1}^* ((0, 0))) \\
 & = q_{10}^0 (R(1) - R(0)) + q_{10}^0 q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 1)) - F_{h+1}^* ((1, 0))) \\
 & + q_{10}^0 q_{00}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\
 & \leq q_{10}^0 (R(1) - R(0)) + q_{01}^0 q_{11}^0 (R(1) - R(0)) + q_{00}^0 q_{11}^0 (R(1) - R(0)) \\
 & = R(1) - R(0).
 \end{aligned}$$

Meanwhile, when $t = h$, again under the same condition, the LHS of Eq. (E.42) satisfies

$$\begin{aligned}
 V_h^* ((1, 1, \mathbf{s})) - V_h^* ((1, 0, \mathbf{s})) & = R(1) - R(0) + (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\
 & + q_{11}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 1)) - 2F_{h+1}^* ((1, 0)) + F_{h+1}^* ((0, 0))) \\
 & = R(1) - R(0) + q_{11}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 1)) - F_{h+1}^* ((1, 0))) \\
 & + q_{10}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\
 & \geq 0.
 \end{aligned}$$

With the same reasoning, we can show the RHS of Eq. (E.42) under the same condition for $t = h$.

For $\mathbf{s} \in \mathcal{S}_3^{M-2}$, we have

$$\begin{aligned}
 V_h^* ((1, 0, \mathbf{s})) - V_h^* ((0, 0, \mathbf{s})) & = R(1) - R(0) + (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\
 & + q_{01}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 1)) - 2F_{h+1}^* ((1, 0)) + F_{h+1}^* ((0, 0))) \\
 & = R(1) - R(0) + q_{01}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 1)) - F_{h+1}^* ((1, 0))) \\
 & + q_{00}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\
 & = \frac{q_{10}^0}{q_{10}^0} (R(1) - R(0) + q_{01}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 1)) - F_{h+1}^* ((1, 0)))) \\
 & + \frac{q_{10}^0}{q_{10}^0} (q_{00}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0)))) \\
 & \geq \frac{q_{10}^0}{q_{10}^0} \left(R(1) - R(0) + \frac{q_{01}^1}{q_{10}^0} (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \right) \\
 & + \frac{q_{00}^1}{q_{10}^0} (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \\
 & = \frac{1}{q_{10}^0} (q_{10}^0 + q_{00}^1 q_{11}^1 - q_{00}^1 q_{01}^1) (R(1) - R(0)) \\
 & + \frac{1}{q_{10}^0} (q_{01}^1 (q_{11}^1 - q_{01}^1)) (R(1) - R(0)) \\
 & \geq \frac{1}{q_{10}^0} (1 - q_{01}^1 + q_{11}^1 - q_{11}^0) (R(1) - R(0)) \\
 & \geq 0.
 \end{aligned}$$

Meanwhile, we also have

$$\begin{aligned}
 V_h^* ((1, 1, \mathbf{s})) - V_h^* ((1, 0, \mathbf{s})) & = R(1) - R(0) + (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\
 & + q_{11}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 1)) - 2F_{h+1}^* ((1, 0)) + F_{h+1}^* ((0, 0))) \\
 & = R(1) - R(0) + q_{11}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 1)) - F_{h+1}^* ((1, 0))) \\
 & + q_{10}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0))) \\
 & = \frac{q_{10}^0}{q_{10}^0} (R(1) - R(0) + q_{11}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 1)) - F_{h+1}^* ((1, 0)))) \\
 & + \frac{q_{10}^0}{q_{10}^0} (q_{10}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^* ((1, 0)) - F_{h+1}^* ((0, 0)))) \\
 & \geq \frac{q_{10}^0}{q_{10}^0} \left(R(1) - R(0) + \frac{q_{11}^1}{q_{10}^0} (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \right)
 \end{aligned}$$

$$\begin{aligned}
 & + \frac{q_{10}^1}{q_{10}^0} (q_{11}^1 - q_{01}^1) (R(1) - R(0)) \\
 & = \frac{1}{q_{10}^0} (q_{10}^0 + q_{10}^1 q_{11}^1 - q_{10}^1 q_{01}^1) (R(1) - R(0)) \\
 & + \frac{1}{q_{10}^0} (q_{11}^1 (q_{11}^1 - q_{01}^1)) (R(1) - R(0)) \\
 & \geq \frac{1}{q_{10}^0} (1 - q_{01}^1 + q_{11}^1 - q_{11}^0) (R(1) - R(0)) \\
 & \geq 0.
 \end{aligned}$$

For $\mathbf{s} \in \mathcal{S}_2^{M-2}$, we can use the same reasoning used in the two agent case to get the result.

Case 2: $q_{11}^0 \leq q_{01}^0, q_{11}^1 \leq q_{01}^1$.

First we consider the sub-case where $\boxed{q_{11}^1 q_{01}^0 \leq q_{01}^1 q_{11}^0}$. We will use backward induction to prove the following statements:

$$\forall t \leq T, 0 \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \leq R(1) - R(0) \quad (\text{E.46})$$

$$0 \leq V_t^*((1, 1, \mathbf{s})) - V_t^*((1, 0, \mathbf{s})) \leq R(1) - R(0) \quad (\text{E.47})$$

$$\pi_{t+1}^{\text{greedy}} = \pi_{t+1}^* \quad (\text{E.48})$$

They hold trivially when $t = T$. Suppose they hold when $t \geq h + 1$, when $t = h$, we will first prove Eq. (E.48) for $t = h$ under the assumption that Eq. (E.46) and Eq. (E.24) hold for $t = h + 1$. Using the same reasoning, we only need to show Eq. (E.45). Again, from the assumptions and the definition of F , we know that

$$\begin{aligned}
 0 & \leq F_{\mathbf{a}_{h+1, h+1}}((1, 1)) - F_{\mathbf{a}_{h+1, h+1}}((1, 0)) \leq R(1) - R(0) \\
 0 & \leq F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0)) \leq R(1) - R(0).
 \end{aligned}$$

Therefore, we have

$$\begin{aligned}
 & (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0))) \\
 & + (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (F_{\mathbf{a}_{h+1, h+1}}((1, 1)) - 2F_{\mathbf{a}_{h+1, h+1}}((1, 0)) + F_{\mathbf{a}_{h+1, h+1}}((0, 0))) \\
 & \geq (q_{11}^0 + q_{01}^1 - q_{11}^1 - q_{01}^0) (F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0))) \\
 & - (q_{11}^0 q_{01}^1 - q_{01}^0 q_{11}^1) (F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0))) \\
 & = \frac{I(0) - I(1)}{R(1) - R(0)} (F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0))) - \frac{(q_{11}^1 I(0) - q_{01}^1 I(1))}{R(1) - R(0)} \\
 & (F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0))) \\
 & = \frac{(q_{10}^1 I(0) - q_{00}^1 I(1))}{R(1) - R(0)} (F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0))) \\
 & \geq 0.
 \end{aligned}$$

This will give us the desired result. Assuming Eq. (E.48) holds for $t = h$, we will select the greedy policy for \mathcal{S}_i^{M-2} as we did in case 1. Again, the only case that is different from the 2-agent case is when $\mathbf{s} \in \mathcal{S}_1^{M-2}, \mathcal{S}_3^{M-2}$. Therefore, we will only show Eq. (E.46) and Eq. (E.47) for those two cases when $t = h$ under the assumption that Eq. (E.46) and Eq. (E.47) hold for $t = h + 1$ and Eq. (E.48) holds for $t = h$. We will first prove the RHS of Eq. (E.46). When $\mathbf{s} \in \mathcal{S}_1^{M-2}$, we have

$$\begin{aligned}
 V_h^*((1, 0, \mathbf{s})) - V_h^*((0, 0, \mathbf{s})) & = R(1) - R(0) + (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0))) \\
 & + q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 1)) - 2F_{h+1}^*((1, 0)) + F_{h+1}^*((0, 0))) \\
 & = R(1) - R(0) + q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 1)) - F_{h+1}^*((1, 0))) \\
 & + q_{00}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0)))
 \end{aligned}$$

$$\leq R(1) - R(0).$$

Similarly, when $\mathbf{s} \in \mathcal{S}_3^{M-2}$, we have

$$\begin{aligned} V_h^*((1, 0, \mathbf{s})) - V_h^*((0, 0, \mathbf{s})) &= R(1) - R(0) + (q_{11}^1 - q_{01}^1) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0))) \\ &\quad + q_{01}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^*((1, 1)) - 2F_{h+1}^*((1, 0)) + F_{h+1}^*((0, 0))) \\ &= R(1) - R(0) + q_{01}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^*((1, 1)) - F_{h+1}^*((1, 0))) \\ &\quad + q_{00}^1 (q_{11}^1 - q_{01}^1) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0))) \\ &\leq R(1) - R(0). \end{aligned}$$

Now, we are going to prove the LHS of Eq. (E.46). When $\mathbf{s} \in \mathcal{S}_1^{M-2}$, we have

$$\begin{aligned} V_h^*((1, 0, \mathbf{s})) - V_h^*((0, 0, \mathbf{s})) &= R(1) - R(0) + (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0))) \\ &\quad + q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 1)) - 2F_{h+1}^*((1, 0)) + F_{h+1}^*((0, 0))) \\ &= R(1) - R(0) + q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 1)) - F_{h+1}^*((1, 0))) \\ &\quad + q_{00}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0))) \\ &\geq R(1) - R(0) + q_{01}^0 (q_{11}^0 - q_{01}^0) (R(1) - R(0)) + q_{00}^0 (q_{11}^0 - q_{01}^0) (R(1) - R(0)) \\ &= (1 + q_{11}^0 + q_{00}^0)(R(1) - R(0)) \\ &\geq 0. \end{aligned}$$

Using the same calculation, we can show the case for $\mathbf{s} \in \mathcal{S}_3^{M-2}$ and Eq. (E.47).

Secondly, we consider the sub-case where $\boxed{q_{11}^1 q_{01}^0 \geq q_{01}^1 q_{11}^0}$. We want to use backward induction to prove the following statements:

$$\forall t \leq T, 0 \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \leq R(1) - R(0) \quad (\text{E.49})$$

$$0 \leq V_t^*((1, 1, \mathbf{s})) - V_t^*((1, 0, \mathbf{s})) \leq R(1) - R(0) \quad (\text{E.50})$$

$$V_t^*((1, 1, \mathbf{s})) - V_t^*((1, 0, \mathbf{s})) \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \quad (\text{E.51})$$

$$\pi_{t+1}^{\text{greedy}} = \pi_{t+1}^* \quad (\text{E.52})$$

When $t = T$, the statements hold trivially. Suppose this holds for $t \geq h + 1$, when $t = h$, we will first show Eq. (E.52) under the assumption that Eq. (E.49)-(E.51) hold for $t = h + 1$. Using the same reasoning, we only need to show Eq. (E.45). Again, from the above assumptions and the definition of F , we know that

$$0 \leq F_{\mathbf{a}_{h+1, h+1}}((1, 1)) - F_{\mathbf{a}_{h+1, h+1}}((1, 0)) \leq R(1) - R(0)$$

$$0 \leq F_{\mathbf{a}_{h+1, h+1}}((1, 0)) - F_{\mathbf{a}_{h+1, h+1}}((0, 0)) \leq R(1) - R(0).$$

Then combined with the fact that $q_{11}^1 q_{01}^0 \geq q_{01}^1 q_{11}^0$, it is easy to see that Eq. (E.45) holds.

Using the same logic used in the first sub-case, we can prove Eq. (E.49),(E.50) for $t = h$. The only thing we need to show is Eq. (E.51) for $t = h$. Again, when $\mathbf{s} \in \mathcal{S}_2^{M-2}$, this is the same as the two agent case. Therefore, we will only prove the other two cases under the assumption that Eq. (E.49)-(E.51) hold for $t = h + 1$ and Eq. (E.52) holds for $t = h$. We will prove the result for $\mathbf{s} \in \mathcal{S}_1^{M-2}$, the other cases will hold for similar reasons. We have

$$\begin{aligned} &V_h^*((1, 1, \mathbf{s})) - 2V_h^*((1, 0, \mathbf{s})) + V_h^*((0, 0, \mathbf{s})) \\ &= (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0))) \\ &\quad + q_{11}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 1)) - 2F_{h+1}^*((1, 0)) + F_{h+1}^*((0, 0))) \\ &\quad - (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 0)) - F_{h+1}^*((0, 0))) \\ &\quad - q_{01}^0 (q_{11}^0 - q_{01}^0) (F_{h+1}^*((1, 1)) - 2F_{h+1}^*((1, 0)) + F_{h+1}^*((0, 0))) \\ &= (q_{11}^0 - q_{01}^0)^2 (F_{h+1}^*((1, 1)) - 2F_{h+1}^*((1, 0)) + F_{h+1}^*((0, 0))) \\ &\leq 0, \end{aligned}$$

where the last inequality holds because of the assumption that Eq. (E.51) holds for $t = h + 1$.

Case 3: $q_{11}^1 \geq q_{01}^1, q_{11}^0 \geq q_{01}^0$. Similar to the two agent case, we can prove the following statements using backward induction:

$$\forall t \leq T, \frac{R(1) - R(0)}{q_{10}^1 + q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-t+1}\right) \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \quad (\text{E.53})$$

$$\frac{R(1) - R(0)}{q_{10}^0 + q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-t+1}\right) \geq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \quad (\text{E.54})$$

$$\frac{R(1) - R(0)}{q_{10}^1 + q_{01}^1} \left(1 - (q_{11}^1 - q_{01}^1)^{T-t+1}\right) \leq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \quad (\text{E.55})$$

$$\frac{R(1) - R(0)}{q_{10}^0 + q_{01}^0} \left(1 - (q_{11}^0 - q_{01}^0)^{T-t+1}\right) \geq V_t^*((1, 0, \mathbf{s})) - V_t^*((0, 0, \mathbf{s})) \quad (\text{E.56})$$

$$\pi_{t+1}^{\text{greedy}} = \pi_{t+1}^* \quad (\text{E.57})$$

The proof will be exactly using the idea used in 2 agent case and the calculation for multiple agents case discussed above, thus will be omitted.

Case 4: Suppose $q_{11}^1 \leq q_{01}^1, q_{11}^0 \geq q_{01}^0, q_{11}^1 \leq q_{11}^0$. We can use the same way used in 2 agent case to directly calculate $2V_t^{\pi^{\text{greedy}}}((1, 0, \mathbf{s})) - V_t^{\pi^{\text{greedy}}}((0, 0, \mathbf{s})) - V_t^{\pi^{\text{greedy}}}((1, 1, \mathbf{s}))$ and have the result. \square

F TECHNICAL DETAILS IN SECTION 4

We will continue to use the notation defined in Appendix E.

F.1 Proof of Lemma 4.3

Lemma 4.3 (Existence of Sufficient Good Agent-State Pairs). *Under Assumptions 4.2 and 4.1, for any joint state $\mathbf{s} \in \mathcal{S}^M$, the number of good agent-state pairs observed in \mathbf{s} is at least B . That is, $|\{m \in [M] : (m, \mathbf{s}^m) \in \mathcal{G}\}| \geq B$.*

Proof. By Bellman equation, we have

$$V_0^{\pi^b}(\mathbf{s}_0) = \mathbb{E} \left[\sum_{m=1}^M c_0^m \right] + \mathbb{E}_{\pi_1^b} \left[V_1^{\pi^b}(\mathbf{s}_1) \right].$$

Since $V_0^{\pi^b}(\mathbf{s}_0) = 0, V_1^{\pi^b}(\mathbf{s}_1) \geq 0, P_{\pi^b}(\mathbf{s}_0, \mathbf{s}_1) > 0$, we have $V_1^{\pi^b}(\mathbf{s}_1) = 0$ for all $\mathbf{s}_1 \in \mathcal{S}^M$. Since again by Bellman equation, we have $V_1^{\pi^b}(\mathbf{s}_1) \geq \sum_{m=1}^M c_1^m (\pi_2^b(\mathbf{s}_1^m))$, we have this sum should be 0. This will indicate that there are enough good arms for state \mathbf{s}_1 . This will give us the desired result. \square

Before proving Theorem 4.4, we first introduce Lemma F.1 which decomposes the cumulative regret to per-timestep regret.

Lemma F.1. *For any $t \leq T$, let per-timestep regret $Regret_t(\pi, \mathcal{H}_t)$ for the policy π be defined as:*

$$\begin{aligned} Regret_t(\pi, \mathcal{H}_t) = & \mathbb{E}_{\pi_{t+1}} \left[\sum_{m=1}^M c_t^m(a_{t+1}^m) \mid \mathcal{H}_t \right] + \mathbb{E}_{\pi_{t+1}} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=t+1}^T c_k^m(a_{t+1}^m) \mid \mathbf{r}_{t+1}, \mathbf{s}_{t+1}, \mathbf{a}_{t+1}, \mathcal{H}_t \right] \right] \\ & - \mathbb{E}_{\pi_{t+1}^b} \left[\sum_{m=1}^M c_t^m(a_{t+1}^m) \mid \mathcal{H}_t \right] - \mathbb{E}_{\pi_{t+1}^b} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=t+1}^T c_k^m(a_{t+1}^m) \mid \mathbf{r}_{t+1}, \mathbf{s}_{t+1}, \mathbf{a}_{t+1}, \mathcal{H}_t \right] \right]. \end{aligned}$$

Then the cumulative regret $\mathcal{R}(\pi)$ can be decomposed as following:

$$\mathcal{R}(\pi) = \mathbb{E} \left[\sum_{t=0}^T Regret_t(\pi, \mathcal{H}_t) \right].$$

F.2 Proof of Lemma F.1

Proof. We will prove the result iteratively from $t = 1$ to T . All we need is calculation. We have

$$\begin{aligned}
 \mathcal{R}(\pi) &= \mathbb{E}_\pi \left[\sum_{m=1}^M \sum_{t=0}^T c_t^m(a_{t+1}^m) \mid \mathbf{s}_0 \right] - \mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{t=0}^T c_t^m(a_{t+1}^m) \mid \mathbf{s}_0 \right] \\
 &= \mathbb{E}_{\pi_1} \left[\sum_{m=1}^M c_0^m(a_1^m) \mid \mathcal{H}_0 \right] + \mathbb{E}_{\pi_1} \left[\mathbb{E}_\pi \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathbf{r}_1, \mathbf{s}_1, \mathbf{a}_1, \mathcal{H}_0 \right] \right] \\
 &\quad - \mathbb{E}_{\pi_1} \left[\sum_{m=1}^M c_0^m(a_1^m) \mid \mathcal{H}_0 \right] - \mathbb{E}_{\pi_1} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathbf{r}_1, \mathbf{s}_1, \mathbf{a}_1, \mathcal{H}_0 \right] \right] \\
 &\quad + \mathbb{E}_{\pi_1} \left[\sum_{m=1}^M c_0^m(a_1^m) \mid \mathcal{H}_0 \right] + \mathbb{E}_{\pi_1} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathbf{r}_1, \mathbf{s}_1, \mathbf{a}_1, \mathcal{H}_0 \right] \right] \\
 &\quad - \mathbb{E}_{\pi_1^b} \left[\sum_{m=1}^M c_0^m(a_1^m) \mid \mathcal{H}_0 \right] - \mathbb{E}_{\pi_1^b} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathbf{r}_1, \mathbf{s}_1, \mathbf{a}_1, \mathcal{H}_0 \right] \right] \\
 &= \mathbb{E} \left[\mathbb{E}_\pi \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathcal{H}_1 \right] - \mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathcal{H}_1 \right] \right] \\
 &\quad + \text{Regret}(\pi, \mathcal{H}_0).
 \end{aligned}$$

We can continue this trick to $\mathbb{E}_\pi \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathcal{H}_1 \right] - \mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=1}^T c_k^m(a_{k+1}^m) \mid \mathcal{H}_1 \right]$, this will give us the desired result. \square

F.3 Proof of Theorem 4.4

Theorem 4.4 (Regret Upperbound). *Assume r_t^m is 1-subgaussian random variable for all $t \in [T]$. Let $D := \max \left\{ \frac{8M}{B}, \frac{8M}{M-B}, 2 \right\}$, $p = 1, \eta = \Delta$. Under Assump.s 4.2 and 4.1, the regret, $\mathcal{R}(\pi)$, defined by Def. 2.2 with the loss function l^m given by Eq. (3), satisfies:*

$$\begin{aligned}
 \mathcal{R}(\pi) &= \min \left\{ \mathcal{O} \left(\frac{|\mathcal{S}|(M|\mathcal{S}| - |\mathcal{G}|)MD}{\Delta} \log T \right), \right. \\
 &\quad \left. \mathcal{O} \left(\sqrt{MDB(|\mathcal{S}|(M|\mathcal{S}| - |\mathcal{G}|))T \log T} \right) \right\}. \tag{4}
 \end{aligned}$$

Before we prove the theorem, we first provide a technical lemma showing the concentration property of the empirical mean estimator $\hat{I}_t^m(s)$, which is a direct result from Howard et al. 2021.

Lemma F.2. *For any $\alpha > 0$, $0 \leq t \leq T$, we have*

$$\mathbb{P} \left(\left| \hat{I}_t^m(s) - I^m(s) \right| \geq 1.7 \sqrt{\frac{\log \log(2N_{s,m}(t)) + 0.72 \log(10.4/\alpha)}{N_{s,m}(t)}} \right) \leq \alpha.$$

Proof. The proof will proceed by first showing that the cumulative regret can be decomposed into per-timestep regret $\text{Regret}_t(\pi, \mathcal{H}_t)$. Then we will use the standard technique to bound the number of times suboptimal agents are given action 1.

By Lemma F.1, we only need to bound $\text{Regret}_t(\pi, \mathcal{H}_t)$. We have

$$\text{Regret}_t(\pi, \mathcal{H}_t) = \mathbb{E}_{\pi_{t+1}} \left[\sum_{m=1}^M c_t^m(a_{t+1}^m) \mid \mathcal{H}_t \right] + \mathbb{E}_{\pi_{t+1}} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=t+1}^T c_k^m(a_{k+1}^m) \mid \mathbf{r}_{t+1}, \mathbf{s}_{t+1}, \mathbf{a}_{t+1}, \mathcal{H}_t \right] \right]$$

$$\begin{aligned}
 & - \mathbb{E}_{\pi_{t+1}^b} \left[\sum_{m=1}^M c_t^m(a_{t+1}^m) \mid \mathcal{H}_t \right] - \mathbb{E}_{\pi_{t+1}^b} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=t+1}^T c_k^m(a_{t+1}^m) \mid \mathbf{r}_{t+1}, \mathbf{s}_{t+1}, \mathbf{a}_{t+1}, \mathcal{H}_t \right] \right] \\
 & = \mathbb{E}_{\pi_{t+1}} \left[\sum_{m=1}^M c_t^m(a_{t+1}^m) \mid \mathcal{H}_t \right] + \mathbb{E}_{\pi_{t+1}} \left[\mathbb{E}_{\pi^b} \left[\sum_{m=1}^M \sum_{k=t+1}^T c_k^m(a_{t+1}^m) \mid \mathbf{r}_{t+1}, \mathbf{s}_{t+1}, \mathbf{a}_{t+1}, \mathcal{H}_t \right] \right] \\
 & \leq \mathbb{E} \left[\sum_{I^m(s) \leq \gamma} \Delta_s^m \mathbb{I} \{a_{t+1}^m = 1, s_t^m = 1\} \right] + \mathbb{E}_{\pi_{t+1}} \left[V_{t+1}^{\pi^b}(\mathbf{s}_{t+1}) \mid \mathbf{s}_{t+1}, \mathbf{a}_{t+1}, \mathcal{H}_t \right] \\
 & = \mathbb{E} \left[\sum_{I^m(s) \leq \gamma} \Delta_s^m \mathbb{I} \{a_{t+1}^m = 1, s_t^m = 1\} \right],
 \end{aligned}$$

where the second and the third equality holds by Lemma 4.3. Then we have the cumulative regret is upper bounded by

$$\sum_{I^m(s) \leq \gamma} \Delta_s^m \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s\} \right].$$

We have

$$\mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s\} \right] = \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s, \text{random}\} \right] \tag{F.58}$$

$$+ \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s, \text{not random}\} \right], \tag{F.59}$$

where random stands for doing random exploration. We will bound the two terms separately. For Eq. (F.58), we have

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s, \text{random}\} \right] \\
 & = \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \left\{ a_{t+1}^m = 1, s_t^m = s, \text{random}, \exists (m', s'), s_t^{m'} = s', N_{s', m'}^{(t)} \leq N_{s_t^{m'}, m''}^{(t)}, \forall m'' \right\} \right] \\
 & \leq \sum_{m'} \sum_{s'} \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \left\{ a_{t+1}^m = 1, s_t^m = s, \text{random}, s_t^{m'} = s', N_{s', m'}^{(t)} \leq N_{s_t^{m'}, m''}^{(t)}, \forall m'' \right\} \right] \\
 & \leq \sum_{m'} \sum_{s'} \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \left\{ \text{random}, s_t^{m'} = s', N_{s', m'}^{(t)} \leq N_{s_t^{m'}, m''}^{(t)}, \forall m'' \right\} \right] \\
 & \leq \sum_{m'} \sum_{s'} \mathbb{E} \left[\sum_{k=0}^{\infty} \mathbb{I} \left\{ \text{random}, N_{s', m'}^{(\tau_k^{m'}(s'))} \leq N_{s_t^{m'}, m''}^{(\tau_k^{m'}(s'))}, \forall m'', \tau_k^{m'}(s') \leq T \right\} \right] \\
 & \leq \sum_{m'} \sum_{s'} \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ \text{random}, N_{s', m'}^{(\tau_k^{m'}(s'))} \leq N_{s_t^{m'}, m''}^{(\tau_k^{m'}(s'))}, \forall m'' \right\} \right] \\
 & = \sum_{m'} \sum_{s'} \sum_{k=0}^T \mathbb{P} \left(\text{random} \mid N_{s', m'}^{(\tau_k^{m'}(s'))} \leq N_{s_t^{m'}, m''}^{(\tau_k^{m'}(s'))}, \forall m'' \right) \mathbb{P} \left(N_{s', m'}^{(\tau_k^{m'}(s'))} \leq N_{s_t^{m'}, m''}^{(\tau_k^{m'}(s'))}, \forall m'' \right) \\
 & \leq \sum_{m'} \sum_{s'} \sum_{k=0}^T \mathbb{P} \left(\text{random} \mid N_{s', m'}^{(\tau_k^{m'}(s'))} \leq N_{s_t^{m'}, m''}^{(\tau_k^{m'}(s'))}, \forall m'' \right) \\
 & \leq \sum_{m'} \sum_{s'} \sum_{k=0}^T \frac{D}{k \Delta^2}
 \end{aligned}$$

$$\leq SMD \frac{\log(T)}{\Delta^2}, \quad (\text{F.60})$$

where $\tau_k^m(s) := \inf\{t : N_{s,m}^{(t)} = k\}$.

For Eq. (F.59), we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s, \text{not random}\} \right] \\ &= \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s, \text{not random}, \exists(m', s') \in \mathcal{G}, \hat{I}_t^m(s) \geq \hat{I}_t^{m'}(s'), s_t^{m'} = s'\} \right] \\ &\leq \sum_{(m', s') \in \mathcal{G}} \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s, \hat{I}_t^m(s) \geq \hat{I}_t^{m'}(s'), s_t^{m'} = s'\} \right] \end{aligned}$$

For each $(m', s') \in \mathcal{G}$, we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{a_{t+1}^m = 1, s_t^m = s, \hat{I}_t^m(s) \geq \hat{I}_t^{m'}(s'), s_t^{m'} = s'\} \right] \\ &\leq \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \left\{ a_{t+1}^m = 1, s_t^m = s, \hat{I}_t^m(s) \geq \gamma + \frac{\Delta}{2}, s_t^{m'} = s' \right\} \right] \end{aligned} \quad (\text{F.61})$$

$$+ \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \left\{ a_{t+1}^m = 1, s_t^m = s, \hat{I}_t^{m'}(s') \leq \gamma + \frac{\Delta}{2}, s_t^{m'} = s' \right\} \right]. \quad (\text{F.62})$$

Again, for Eq. (F.61), we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \left\{ a_{t+1}^m = 1, s_t^m = s, \hat{I}_t^m(s) \geq \gamma + \frac{\Delta}{2}, s_t^{m'} = s' \right\} \right] \\ &\leq \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \left\{ s_t^m = s, \hat{I}_t^m(s) \geq \gamma + \frac{\Delta}{2} \right\} \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{\infty} \mathbb{I} \left\{ s_{\tau_k^m(s)}^m = s, \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2}, \tau_k^m(s) \leq T \right\} \right] \\ &\leq \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2} \right\} \right] \\ &= \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \in [k_1 t_{\text{random}}^k, k_2 t_{\text{random}}^k], \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2} \right\} \right] \\ &+ \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \notin [k_1 t_{\text{random}}^k, k_2 t_{\text{random}}^k], \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2} \right\} \right] \\ &\leq \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \in [k_1 t_{\text{random}}^k, k_2 t_{\text{random}}^k], \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2}, t_{\text{random}}^k \geq k_3^k \right\} \right] \\ &+ \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \notin [k_1 t_{\text{random}}^k, k_2 t_{\text{random}}^k] \right\} \right] + \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \{t_{\text{random}}^k \leq k_3^k\} \right], \end{aligned}$$

where $N_{s,m}^{(\tau_k^m(s))}(r) := \sum_{l=1}^k \mathbb{I} \{a_{\tau_l^m(s)}^m = 0, s_{\tau_l^m(s)}^m = s, \text{random}\}$ denotes the total number of times that action 0 is given to agent m when m is in state s and the algorithm is doing random exploration, $t_{\text{random}}^k := \sum_{l=1}^m \sum_{l=1}^k \mathbb{I} \{\text{random}\}$ denotes the total number of times when agent m is in state s and the algorithm is doing random explorations. We let $k_1 = \max \left\{ \frac{M-B}{2M}, \frac{2M-3B}{2M} \right\}$, $k_2 = \min \left\{ \frac{3(M-B)}{2M}, \frac{2M-B}{2(M-B)} \right\}$, $k_3^k = \frac{C \log(k)}{2 \min\{\Delta^2, 1/2\}}$

For the first term, we have

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \in [k_1 t_{\text{random}}^k, k_2 t_{\text{random}}^k], \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2}, t_{\text{random}}^k \geq k_3^k \right\} \right] \\
 & \leq \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \in [k_1 t_{\text{random}}^k, k_2 t_{\text{random}}^k], N_{s,m,1}^{(\tau_k^m(s))} \geq (1 - k_2) t_{\text{random}}^k, \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2}, t_{\text{random}}^k \geq k_3^k \right\} \right] \\
 & \leq \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \geq k_1 k_3^k, N_{s,m,1}^{(\tau_k^m(s))} \geq (1 - k_2) k_3^k, \hat{I}_{\tau_k^m(s)}^m(s) \geq \gamma + \frac{\Delta}{2} \right\} \right] \\
 & \leq \mathbb{E} \left[\sum_{k=0}^T \mathbb{I} \left\{ N_{s,m}^{(\tau_k^m(s))}(r) \geq k_1 k_3^k, N_{s,m,1}^{(\tau_k^m(s))} \geq (1 - k_2) k_3^k, \left| \hat{I}_{\tau_k^m(s)}^m(s) - I^m(s) \right| \geq \frac{\Delta}{2} \right\} \right] \\
 & \leq \sum_{k=0}^T \frac{1}{k^2} \\
 & \leq 2,
 \end{aligned}$$

where the second last inequality holds due to Lemma F.2. Similarly, we can prove the last two terms are upper bounded by 2 using Lemma F.2.

Using the same way, we can upper bound Eq. (F.62) by 6. Therefore, Eq. (F.59) can be upper bounded as

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{ a_{t+1}^m = 1, s_t^m = s, \text{not random} \} \right] \\
 & = \sum_{(m',s') \in \mathcal{G}} \mathbb{E} \left[\sum_{t=0}^T \mathbb{I} \{ a_{t+1}^m = 1, s_t^m = s, \hat{I}_t^m(s) \geq \hat{I}_t^{m'}(s'), s_t^{m'} = s' \} \right] \\
 & \leq 12|\mathcal{G}|.
 \end{aligned} \tag{F.63}$$

Further assume that $k|\Delta| \geq \Delta_s^m$, for all $I^m(s) \leq \gamma$. Then combining Eq. (F.60) and Eq. (F.63) will give us the desired result.

We also have

$$\begin{aligned}
 \mathcal{R}(\pi) & \leq kB\Delta T + \frac{k|\mathcal{S}|MD(M|\mathcal{S}| - |\mathcal{G}|)}{\Delta} \log T \\
 & \leq k\sqrt{BMD|\mathcal{S}|(M|\mathcal{S}| - |\mathcal{G}|)T \log T}.
 \end{aligned}$$

□

G IMPLEMENTATION DETAIL

All the experiments are conducted on MacBook Pro with Apple M3 Pro chip and 18 GB memory.

G.1 Regret Curve

To numerically verify Theorem 4.4, we set $M = 2$, $B = 1$, and $\mathcal{S} = \{0, 1\}$. We set the expected reward of the state as $R^m(s) = s$. We randomly generated 10 instances and chose three γ values satisfying Assumption 4.2, labeling each γ by the corresponding gap Δ between the threshold and the smallest incremental reward. Specifically, we will randomly generate a number between $[0, 1]$ independently and assign this number to be $P_a^m(s, 1), \forall a, m, s \in \{0, 1\}$.

Without loss of generality, assume that $I^1(1) \geq I^1(0) \geq I^2(1) \geq I^2(0)$. We choose γ to be between $I^1(0)$ and $I^2(0)$. Specifically, let $\theta \in (0, 1)$ be a tunable parameter. We set $\gamma := \theta I^2(1) + (1 - \theta)I^2(0)$. To make sure that the gap between γ and $I^2(1)$ is not too small, we choose those instances where $I^2(1) - I^2(0) \geq 0.15$. To generate the three different γ values in Figure G.9, we selected three different choices of θ : 0.1, 0.5, 0.9, yielding the observed Δ values.




Figure G.9: Regret Curve as horizon T grows.

For each γ , the regret was averaged over 100 repetitions per instance (Figure G.9). Figure G.9 demonstrates sublinear regret for all choices of γ , with regret stabilizing at a constant as the horizon length increases, consistent with Theorem 4.4. Further, as Δ decreases, regret increases as it becomes harder to distinguish whether the incremental reward is above the threshold or not. This matches the results in Theorem 4.4.

G.2 Cumulative Reward

We first describe the benchmarks we chose in detail:

- Oracle Whittle: A variant of Whittle’s index policy with access to the ground truth transitions and rewards, optimized for maximizing discounted reward (as described in Appendix B).
- Oracle Greedy: The greedy policy we defined (Algorithm 1) with knowledge of the ground truth incremental reward.
- OMR: A modified LP-based policy for maximizing finite-horizon reward.
- WIQL (Biswas et al., 2021): A Q-learning based approach to learn the value function of each arm at each state by interacting with the RMAB instance. We use the same learning rate (α) as in Biswas et al. (2021).
- UCWhittle-penalty (Wang et al., 2023b): It constructs the confidence interval of the unknown transition probability and then chooses one to calculate Whittle’s index policy. Specifically, it constructs a confidence interval and chooses the transition probability that will maximize $V_\lambda^m(s)$ (Eq. (B.9)) for some adaptively updated λ .
- R(MA)²B: It uses the explore-then-commit framework. It first randomly gives action and constructs the confidence interval of the unknown transition. After the exploration period, it uses the estimated transition to calculate the OMR policy.

In all of the settings, the discount factor β was set to 0.9. We note that though $\beta = 1$ was chosen in Biswas et al. (2021), $\beta = 0.9$ has better empirical performance, which aligns with the choice in Wang et al. (2023b).

Then, we will discuss more on the implementation of the algorithms. First of all for the UCWhittle families introduced in Wang et al. (2023b), the algorithms reinitialize the initial state after certain time-steps. We do not include this episodic structure as this contradicts our goal of maximizing the finite-horizon cumulative reward. We keep the rest structures of their algorithms as the same.

We additionally note that the code provided by Wang et al. (2023b) contains errors and could not be used to reproduce the results shown in Wang et al. (2023b). A mistake was found in Line 73 of uc_whittle.py that they provided. Specifically, they use the current state of an agent as the index to get the current action of this agent. As a result, we build our simulation environment from scratch.

Furthermore, we note that the Whittle’s index could be negative. Therefore, we make modifications to the binary search routine of getting Whittle’s index. Specifically, we let the initial lower bound be -10 to ensure a more accurate calculation of Whittle’s index.

We also note that in the original $R(MA)^2B$ algorithm, it assumes that they can reinitialize at any state for arbitrary number of times. This makes their exploration criteria, which will only stop when sufficient number of samples are collected for each agent, state, action pair, possible. We think this assumption is not impractical and do not use this structure. Therefore, we also change the algorithm so that the exploration period will end after the same number of time-steps instead of number of samples.




Figure G.10: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 50 instances with 13 repetitions each. Top: noiseless reward; bottom: noisy reward. Ten states.

Randomly Generated Instances With 10 States In this paragraph, we increase the size of the state space to 10 states to randomly generate 50 instances to test the performance of ϵ -GT comparing to the benchmarks. Similarly, we also test two settings: 1) a noiseless case where the reward is deterministic, and 2) a noisy case where the rewards at state $s \in \mathcal{S}$ follow a normal distribution with mean $R^m(s)$ and variance 0.1 . We would like to note that since UCWhittlePv and RRMABB behave like random policy for 2 states, we do not include that in the 10 states setting.

As shown in Figure G.10, ϵ -GT with threshold $\gamma = -8.9$ consistently outperform WIQL in both noiseless and noisy setting. Further, we would like to emphasize that the gap between ϵ -GT and WIQL is larger than the 2 states setting, demonstrating that our algorithm works well in larger states setting and adapting to the real-life setting better.

G.3 Ablation Study

In this section, we will test the performance of Alg. 1 with different choices of η and γ .

Choice of η : We will choose four choices of $\eta \in \{0.1, 0.5, 1, 1.5\}$. We will fix the γ and p to align with the choice of γ used for experiments in Section 5. As shown in Figures G.11-G.13 Alg. 1 is stable over different choices of η for all of the synthetic binary state, ten state instances and real instance we have tested.

Choice of γ In this paragraph, we will compare ϵ -GT with different threshold. We will fix other hyperparameters to align with choices used in Section 5. As shown in Figure G.14-Figure G.16, for all instances, as γ increases, the convergence rate of Alg. 1 is slower. This translates to a lower average reward. Also, we observe that among all instances, a low threshold that satisfies Assump. 4.2 will achieve the best performance.




Figure G.11: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 50 binary state instances with 13 repetitions each. Top: noiseless reward; bottom: noisy reward.




Figure G.12: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 50 10 state instances with 13 repetitions each. Top: noiseless reward; bottom: noisy reward.




Figure G.13: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 100 repetitions of real dataset. Top: noiseless reward; bottom: noisy reward.




Figure G.14: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 50 binary state instances with 13 repetitions each. Top: noiseless reward; bottom: noisy reward.




Figure G.15: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 50 ten state instances with 13 repetitions each. Top: noiseless reward; bottom: noisy reward.




Figure G.16: Average (over time) cumulative reward for budgets $B = 5$ (left), 10 (middle), 20 (right), averaged over 100 repetitions of real dataset. Top: noiseless reward; bottom: noisy reward.