



HAL
open science

Reoptimization Nearly Solves Weakly Coupled Markov Decision Processes

Nicolas Gast, Bruno Gaujal, Chen Yan

► **To cite this version:**

Nicolas Gast, Bruno Gaujal, Chen Yan. Reoptimization Nearly Solves Weakly Coupled Markov Decision Processes. 2024. hal-04570177

HAL Id: hal-04570177

<https://inria.hal.science/hal-04570177v1>

Preprint submitted on 6 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Reoptimization Nearly Solves Weakly Coupled Markov Decision Processes

Nicolas GAST

Univ. Grenoble Alpes, Inria, nicolas.gast@inria.fr

Bruno GAUJAL

Univ. Grenoble Alpes, Inria, bruno.gaujal@inria.fr

Chen YAN

Univ. Grenoble Alpes, Inria, chen.yan@inria.fr

We propose a new policy, called the *LP-update policy*, to solve finite horizon weakly-coupled Markov decision processes. The latter can be seen as multi-constraint multi-action bandits, and generalize the classical restless bandit problems. Our solution is based on re-solving periodically a relaxed version of the original problem, that can be cast as a linear program (LP). When the problem is made of N statistically identical sub-components, we show that the LP-update policy becomes asymptotically optimal at rate $\mathcal{O}(T^2/\sqrt{N})$. This rate can be improved to $\mathcal{O}(T/\sqrt{N})$ if the problem satisfies some ergodicity property and to $\mathcal{O}(1/N)$ if the problem is *non-degenerate*. The definition of non-degeneracy extends the same notion for restless bandits. By using this property, we also improve the computational efficiency of the LP-update policy. We illustrate the performance of our policy on randomly generated examples, as well as a generalized applicant screening problem, and show that it outperforms existing heuristics.

Key words: Markov processes, stochastic optimization, linear programming, large scale optimization, restless bandit

1. INTRODUCTION

Markov Decision Processes (MDPs) have proven tremendously useful as models of stochastic sequential planning problems. Dynamic programming is the principal method to address these problems under uncertainty. While MDPs are in theory applicable to many problems, applying classic dynamic programming algorithms to realistic problems is computationally difficult. This has triggered much research into techniques to deal with large state and action spaces. One such technique is decomposition, for which the very large global MDP is decomposed into N loosely dependent sub-processes, each of a size exponentially smaller compared to the global problem. Each local problem can then be solved independently at relatively low cost. If these local solutions can be pieced together effectively, and used to guide the search for a global solution that performs well, then dramatic improvements in the overall solution time can be obtained.

In this paper we study *weakly-coupled* MDPs that fall into this situation. The model originates from sequential stochastic resource allocation problems: A number of tasks must be addressed and an action consists in assigning various resource at every decision epoch to each of these tasks. Weakly-coupled MDPs have been used to model scheduling, queueing, supply chain as well as health care problems [16, 17, 24, 30, 33]. In such a model, the tasks are *additive utility independent*: the utility of achieving any collection of tasks is the sum of rewards associated with each task. Moreover, for a given allocation of resource, each task can be viewed as an independent sub-process whose rewards and transitions are independent of the others. The tasks are linked only through the global resource constraints at each decision epoch. This explains the terminology "weakly-coupled".

An interesting sub-class of weakly-coupled MDPs, for which many results have been obtained, is the class of multi-armed restless bandits. In such a problem, the bandit (the global MDP) consists of N arms (the sub-MDPs). Two actions are possible for each arm: The passive action and the

active action. The decision maker must respect the constraint of activating at most αN arms at each decision epoch. In the seminal paper [37], Whittle proposes a relaxation of the problem that decomposes the N -armed problem into N independent sub-problems, one for each arm. Piecing these subproblems together to obtain a feasible solution leads to the famous Whittle index policy, that is known to be extremely efficient. Under several additional technical assumptions, the Whittle index policy is proven to be asymptotically optimal in Weber and Weiss [36], in the sense that the gap between the performance of the Whittle index policy and the optimal policy converges to zero when N goes to infinity. In the recent years, many asymptotic optimality results have been obtained for the multi-armed bandit model, under either finite [5, 14, 41, 42] or infinite horizons [35, 13, 42]. Some of them assume that arms can have several actions [38, 39].

Contributions We consider a general model of homogenous weakly-coupled MDPs, and we design a policy that becomes optimal as the number of components grows. We provide upper bounds on the sub-optimality gap of the LP-update policy, which is a bound on the rates at which the policy becomes optimal. More precisely, our paper makes two main contributions:

1. We generalize the LP-update policy, that first appeared in Gast et al. [14] for two-action restless bandits, to the broader class of weakly-coupled MDPs. We call this policy LP-update because it re-solves periodically a new LP to update its future control decisions. We show that the sub-optimality gap of the LP-update policy is at most $\mathcal{O}(T^2/\sqrt{N})$, where N is the number of sub-components of the weakly-coupled MDP and T is the time horizon. Under the condition that the ergodic coefficient of each homogenous sub-component is strictly smaller than one, we show that the sub-optimality gap can be refined to $\mathcal{O}(T/\sqrt{N})$. These rates improve upon the bound $\mathcal{O}(T^3/\sqrt{N})$ in [6], under the more general case of heterogeneous sub-components.
2. We introduce the notion of a *non-degenerate* weakly-coupled MDP using a rank condition inspired from the linear independence constraint qualification in optimization (Bazaraa et al. [3]), and show that for such a problem, the sub-optimality gap of the LP-update policy is $\mathcal{O}(1/N)$. If the problem satisfies an additional perfect rounding condition, then the rate can be further improved to $e^{-\Omega(N)}$. The non-degeneracy is a key notion to achieve faster convergence rates. There already exists definitions of non-degeneracy for particular sub-classes of weakly-coupled MDPs: The ones presented in [14, 41, 6] focus on particular sub-cases, and are equivalent to our definition when restricted to those. The one given in Zhang [40] for general weakly-coupled MDPs is more restrictive than ours, in the sense that the class of non-degenerate problems according to Zhang [40] is a strict subset compared to using our definition. This difference is important in practice: on randomly generated examples, about half of our non-degenerate examples are not captured by the definition of [40].

In addition, we show that all the convergence rates claimed previously are tight, in the sense that for each case, we are able to exhibit an example whose sub-optimality gap does not decrease faster than our upper bound. This shows that we have identified the essential features of the problem for achieving a fast convergence rate. We also present a few numerical results that illustrate the benefit of using our LP-update policy compared with the occupation measure policy presented in [38, 39]. We illustrate this on randomly generated examples, as well as on a case study that is a generalization of the applicant screening problem studied in Brown and Smith [5] and Gast et al. [14].

Outline The rest of the paper is organized as follows. We discuss related work in Section 2. We introduce the weakly-coupled MDPs model in Section 3. The LP-update policy for weakly-coupled MDPs with full updates is given in Section 4. The improved version of the policy with selective updates is given in Section 5, together with the definition of non-degeneracy. The case study on the applicant screening problem is given in Section 6. We conclude the paper and discuss some future research directions in Section 7. The proofs of the technical results are collected in Section A. Additional discussions are collected in the other appendices.

2. Related work

There are several branches of works in the existing literature that are related to our results in this paper. The first branch is concentrated on the study of weakly-coupled MDPs under finite horizon, e.g. Adelman and Mersereau [1], Astaraky and Patrick [2], Dolgov and Durfee [12], Gocgun and Ghate [15], Gocgun and Ghate [16], Meuleau et al. [24], Patrick et al. [30], as well as the two PhD thesis Hawkins [17] and Salemi Parizi [33]. In these works the authors consider the situation when the sub-MDPs are not statistically identical. The authors of these papers use a Lagrange decomposition technique to solve approximately the above problem, which works as follows: By relaxing the constraints with state independent multipliers, the N -dimensional optimization problem decouples into N one-dimensional subproblems, each involves the corresponding sub-MDP alone; the key is in choosing the value of multipliers and in transforming the optimal controls of the decoupled problems into a feasible policy to the original problem (as has been done for the simpler multi-armed bandit model). The idea of decomposition also appears in other domains, for instance in network congestion control as in Palomar and Chiang [28]. More broadly, in a series of papers [9, 8, 27, 26] as well as in the textbook Carpentier et al. [10], the authors study a very general large-scale stochastic optimization problem with many distributed sub-components, using a guiding principle called "decomposition-coordination".

Our results are also close to the branch of researches on asymptotic optimality results on multi-armed bandits under finite horizon as in Brown and Smith [5], Zhang and Frazier [41], Gast et al. [14], Zayas-Cabán et al. [39] and Xiong et al. [38]. The finite horizon two-action single-constraint multi-armed bandit is the subject of the first three papers, the more general multi-action single-constraint multi-armed bandit is the subject of the two latter papers. Except in Gast et al. [14], all the policies considered in these references are one-pass policies that involve solving a linear program only once at the very beginning. The non-degenerate property on two-action single-constraint multi-armed bandit is proposed independently in both Zhang and Frazier [41] and Gast et al. [14], together with the $\mathcal{O}(1/N)$ rate on non-degenerate models proven therein.

There is another related branch of works that consider the problem under infinite horizon. In Whittle [37], Weber and Weiss [36], Verloop [35], Gast et al. [13] and Zhang and Frazier [42] the two-action bandit has been studied. It is shown in Weber and Weiss [36] that, for indexable model, Whittle index policy is asymptotically optimal under an additional global attractor assumption. Later in Verloop [35] the indexability assumption has been removed. In Gast et al. [13] the optimal convergence rate of these asymptotic results are proven. In Zhang and Frazier [42], by using the discount criterion, instead of the time-average criterion considered in other works, the authors construct an asymptotically optimal policy that does not require the indexability nor the global attractor assumption. This asymptotic result relies on a truncation over a finite horizon T , that depends on the discount factor. When studying the problem under finite horizon as we do here, none of the assumptions of indexability nor the global attractor property are needed for the asymptotic optimality results. On the other hand, the policies become time-dependent and the focus shifts to the transient behavior. The notion of index is generalized to the multi-action single-constraint case in Hodge and Glazebrook [18] and Niño-Mora [25]. It is not clear if such a notion can be generalized to the case of multiple constraints, since the index of a sub-process can only be defined with respect to a single constraint, and it is not obvious to incorporate multiple constraints into one real-valued meaningful index. The infinite-horizon two-action bandit problem is also explored in the recent study by [20], wherein the authors present an algorithm named *Follow-the-Virtual-Advice* (FTVA). Assuming a mild condition termed synchronization by the authors, the suboptimality gap of FTVA is shown to be $r_{\max} \tau^{\text{sync}} / \sqrt{N}$, where τ^{sync} represents a coupling time. This synchronization condition bears resemblance to our ergodicity condition $\lambda^\tau > 0$ in Theorem 2. Despite our analysis focusing on the finite horizon with multiple constraints, as opposed to their single constraint infinite horizon framework, direct comparison of the two performance bounds is challenging. On one side, our bound

incurs an additional factor of \sqrt{d} ; on the other, their bound is influenced by a "passive" coupling time of independent trajectories, while ours relies on an "active" mixing time typically much shorter than the coupling time for large d .

Our proposed heuristic for weakly-coupled MDPs, that we call the LP-update policy, generalizes the one for two-action single-constraint bandits proposed in Gast et al. [14], with one difference: in the current paper, we improve its computational complexity by re-solving a LP only when necessary. We also obtain a faster rate of convergence than the $\mathcal{O}(1/\sqrt{N})$ proven in [14]. The LP-update policy combines ideas from the certainty equivalent control in dynamic programming (Chapter 6 of Bertsekas [4]), with the model predictive control in optimal control theory (Rakovic and Levine [32]). The same technique has been applied in other stochastic optimization problems. For example, the discrete-review policy for multi-class queueing network in Maglaras [22], or the re-solving heuristic for the network revenue management in Jasin and Kumar [21]. In this paper we use a similar idea as the ones used in these papers that consists of four ingredients: (a) find an appropriate deterministic approximation to the stochastic problem (see Section 4.1); (b) use the optimal solution to the deterministic problem to derive an implementable solution for the original problem (see Section 4.2 and Section D); (c) decide the moments for re-optimization (see Section 5.2); (d) analyse the performance of the resulting policy (see Theorems 1, 2, 3 and 4).

Concurrently with the development of this study, an independent research project led by Brown and Zhang [6] explores a related issue. Their work primarily delves into heterogeneous weakly-coupled MDPs, contrasting with our investigation of homogeneous sub-MDPs. The convergence rate identified by Brown and Zhang [6] for their "reoptimized feasible fluid policy" in a broader context is $\mathcal{O}(T^3/\sqrt{N})$, where T is the finite horizon and N is the number of sub-components. In comparison, our approach yields a convergence rate of $\mathcal{O}(T^2/\sqrt{N})$, and potentially $\mathcal{O}(T/\sqrt{N})$ when the ergodic coefficient of the sub-components is less than one. This improved convergence rate in our model relies on the premise of homogeneity across all arms. Furthermore, our bounds exhibit simple (essentially linear) dependence on additional problem parameters, such as the cardinality of the state and action space, as articulated in Theorem 1. This is distinct from bounds of the theorems of Section 4 of [6] that have more complex dependence on the problems parameters.

3. MODEL DESCRIPTION

3.1. Weakly-coupled MDPs

We consider a finite-horizon discrete-time weakly-coupled MDP composed of N statistically identical sub-MDPs (or *components*, or *arms* in the bandit literature), indexed by $n \in \{1 \dots N\}$. The finite state space of each sub-MDP is the set $\mathcal{S} := \{1, 2, \dots, d\}$, and its finite action space¹ is $\mathcal{A} := \{0, 1, \dots, A\}$. The state space of the weakly-coupled MDP is therefore \mathcal{S}^N and the action space is a subset of \mathcal{A}^N . There are J types of resources, and the decision maker is allowed to use up to $N \cdot b_j > 0$ resource of type j at each decision epoch. We assume that taking the action a_n for component n being in state s_n uses $D_j(s_n, a_n) \geq 0$ of resource j , and that the action 0 consumes no resource: $D_j(s_n, 0) = 0$ for all $s_n \in \mathcal{S}$. Hence, the set of *feasible* actions in state $\mathbf{s} = (s_1, s_2, \dots, s_N)$ is the set of $\mathbf{a} \in \mathcal{A}^N$ such that for all $j \in \{1 \dots J\}$: $\sum_{n=1}^N D_j(s_n, a_n) \leq N \cdot b_j$.

The sub-processes are weakly-coupled, in the sense that the N sub-MDPs are only linked through the J resource constraints, i.e. for a given feasible action \mathbf{a} , the system transitions from a state \mathbf{s} to state $\mathbf{s}' = (s'_1, s'_2, \dots, s'_N)$ with probability

$$p(\mathbf{s}' | \mathbf{s}, \mathbf{a}) = \prod_{n=1}^N p(s'_n | s_n, a_n) = \prod_{n=1}^N P_{s_n, s'_n}^{a_n}, \quad (1)$$

¹ For notational simplicity, we assume that the sub-action space of each sub-MDP n is independent of its state S_n . This is without loss of generality as one may forbid some actions by giving them an extremely high cost.

where for each action a , the matrix \mathbf{P}^a is a probability transition matrix of dimension $d \times d$.

Upon choosing an action \mathbf{a} in state \mathbf{s} , the decision maker receives an immediate reward $\sum_{n=1}^N r_{s_n}^{a_n}$, where $r_{s_n}^{a_n} \in \mathbb{R}$ depends on the state s_n and action a_n , and the rewards among components are additive. We denote by $r_{\max} := \max_{s,a} |r^{s,a}| < \infty$. It is finite as the state and action spaces are finite.

The aforementioned model makes a number of assumptions, that are classical in the literature (e.g. [38, 39]):

- We assume that all terms in $D(s, a)$ and \mathbf{b} are non-negative numbers, and that $D(s, 0) = \mathbf{0}$. This is a natural assumption under the resource allocation context in which $a = 0$ corresponds to a *passive* action that consumes no resource. The later also implies that our resource constraint problem has at least a feasible solution by always choosing the passive action. .
- We assume that each sub-MDP is statistically identical. This is needed under our scaling of the arm population N . However, this assumption can be relaxed, in the sense that we can incorporate the case where there is a finite number of types of sub-MDPs, by making direct sum of the state spaces of each type into a single larger state space.
- In our formulation of the problem, the rewards, constraints and transition probabilities do not dependent on time. This choice is only to lighten the notations: As we consider finite-horizon problem, all the results apply to the case of time-dependent parameters—it would suffice to add a dependence on t on all parameters.

Finally, we remark that this model includes the classical restless multi-armed bandit that corresponds to the case $\mathcal{A} = \{0, 1\}$ with $J = 1$ constraint and $D_1(s, a) = a$. This simpler optimization problem is already PSPACE-hard (see Papadimitriou and Tsitsiklis [29]). In the rest of the paper, we focus on developing approximate solutions whose performance are provably close to optimal.

3.2. Symmetry simplification and population representation

Since we assume that the N sub-components are statistically identical, the problem can be reformulated by keeping track of the number (population) of sub-MDPs in each of the d states, as well as the number of sub-MDPs for each action taken in that state. For the purposes of our subsequent analyses, it proves beneficial to normalize each quantity by dividing it by N .

Hence, we denote by $\mathbf{X}^{(N)}(t) = (X_s^{(N)}(t))_{s \in \mathcal{S}}$ the vector of proportions of sub-MDPs that are in each state $s \in \mathcal{S}$ at decision epoch t . We denote by $\Delta^{(N)}$ the possible values for $\mathbf{X}^{(N)}$, which is the set of vectors $\mathbf{x} \in \mathbb{R}^d$ such that $x_s \geq 0$, $\sum_{s \in \mathcal{S}} x_s = 1$ and $N \cdot x_s$ is an integer for all $s \in \mathcal{S}$. It is a finite collection of points in the simplex Δ^d . Although this representation reduces the state space size of the MDP considerably from d^N to $\binom{N+d-1}{d-1}$, it is still intractable when N and d are large.

Let $\mathcal{U} = \mathcal{S} \times \mathcal{A}$ be the set of state-action pairs. Upon observing $\mathbf{X}^{(N)}(t)$, the action taken by the decision maker is represented by $\mathbf{y} \in \mathbb{R}^u$ where $y_{s,a}$ is the proportion of sub-MDPs that are in state s and for which action a is taken. For a given $\mathbf{x} \in \Delta^{(N)}$, we denote by $\mathcal{Y}^{(N)}(\mathbf{x})$ the set of feasible actions in state \mathbf{x} . It is the set of possible $\mathbf{y} \in \mathcal{Y}(\mathbf{x})$ such that for all s, a , $N \cdot y_{s,a}$ is an integer, where $\mathcal{Y}(\mathbf{x})$ is:

$$\mathcal{Y}(\mathbf{x}) := \left\{ \mathbf{y} \geq 0 \text{ such that } \forall s \in \mathcal{S} : \sum_{a \in \mathcal{A}} y_{s,a} = x_s; D\mathbf{y} \leq \mathbf{b} \right\}. \quad (2)$$

In the above definition, the first set of equality constraints is due to the mass conservation, and the second represents the budget constraints. They are written in matrix form by viewing $D_j(s, a)$ as a matrix whose lines are indexed by $j \in \{1 \dots J\}$ and whose columns are indexed by (s, a) . Similarly, the vector \mathbf{b} is a column vector indexed by $j \in \{1 \dots J\}$.

3.3. Optimal control formulation

The decision maker's goal is to maximize the total expected reward over a finite horizon T , given an initial state vector $\mathbf{x} \in \Delta^{(N)}$ of the system. We denote by $V_{\text{opt}}^{(N)}(\mathbf{x}, T)$ the maximal expected gain per arm that can be obtained by the decision maker.

This optimization problem can be written as follows:

$$V_{\text{opt}}^{(N)}(\mathbf{x}, T) = \max_{\mathbf{Y}} \mathbb{E} \left[\sum_{t=0}^{T-1} \sum_{(s,a) \in \mathcal{U}} r_s^a Y_{s,a}^{(N)}(t) \right] \quad (3a)$$

$$\text{s.t. } X_s^{(N)}(0) = x_s \quad \forall s, \quad (3b)$$

$$\mathbf{X}^{(N)}(t+1) \text{ follows the Markov transitions (1) given } \mathbf{Y}^{(N)}(t), \quad (3c)$$

$$\mathbf{Y}^{(N)}(t) \in \mathcal{Y}^{(N)}(\mathbf{X}^{(N)}(t)) \quad \forall t \quad (3d)$$

Notation and Terminology Convention

A super-script N indicates that the problem has population N , with N being a positive integer number. This implies that variables with super-script (N) (for example $X_s^{(N)}$) are integral multiples of $1/N$ (i.e. $N \cdot X_s^{(N)} \in \mathbb{N}$). When we talk about asymptotic results, this variable N is understood as a scaling parameter that tends to infinity. By "asymptotic optimality", implicitly we are always comparing with the LP relaxation bound, which is an upper bound on the value of the optimal policy. We do so since the later is in general not easily computable. A bold letter (e.g. \mathbf{y} , \mathbf{x}) denotes a vector whereas a normal letter (e.g. $y_{s,a}(t)$, $x_s(t)$) denotes a scalar. Unless otherwise specified, vectors are column vectors. We use super-script $*$ to denote optimal quantities (e.g. \mathbf{x}^* , y^*). The function $\mathbf{1}_E$ is a random variable that equals 1 if the event E occurs and 0 otherwise. The vector \mathbf{r}^\top denotes the transpose of the vector \mathbf{r} . For two real valued functions f and g , we write $f(x) = \mathcal{O}(g(x))$ if there exists $C > 0$ and x_0 real such that $|f(x)| \leq Cg(x)$ for all $x \geq x_0$, and we write $f(x) = \Omega(g(x))$ if $\limsup_{x \rightarrow \infty} |f(x)/g(x)| > 0$.

4. THE LP-UPDATE POLICY FOR WEAKLY COUPLED MDPs

In this section, we first introduce the linear program after relaxing the resource constraints in Section 4.1, which is the starting point for all policies we consider afterwards. We then define the LP-update policy with full updates, and present the $\mathcal{O}(1/\sqrt{N})$ performance guarantee that holds in all cases in Section 4.2.

4.1. The relaxed problem as a linear program

The main difficulty of the optimization problem (3) is that the constraint $D\mathbf{Y}^{(N)}(t) \leq \mathbf{b}$ couples all sub-MDPs. To overcome this difficulty, a by-now classical approach [5, 14, 38, 39, 41] is to relax these constraints and consider a problem where they have to be satisfied only in expectation: $D\mathbb{E}[\mathbf{Y}^{(N)}(t)] \leq \mathbf{b}$. This leads us to write a relaxed optimization problem in terms of the variables $\mathbf{y}(t) = \mathbb{E}[\mathbf{Y}^{(N)}(t)]$. Indeed, equation (1) implies that the expectation of $\mathbf{X}^{(N)}(t+1)$ given $\mathbf{Y}^{(N)}(t)$ can be rewritten as a linear map ϕ as follows:

$$\mathbb{E}[X_s^{(N)}(t+1) | \mathbf{Y}^{(N)}(t) = \mathbf{y}] = (\phi(\mathbf{y}))_s := \sum_{(s',a) \in \mathcal{U}} y_{s',a} P_{s',s}^a. \quad (4)$$

This shows that the relaxed optimization problem is the following linear program with decision variables $\mathbf{y}(t) = \mathbb{E}[\mathbf{Y}^{(N)}(t)]$:

$$V_{\text{rel}}(\mathbf{x}, T) = \max_{\mathbf{y} \geq 0} \sum_{t=0}^{T-1} \sum_{(s,a) \in \mathcal{U}} r_s^a y_{s,a}(t) \quad (5a)$$

$$\text{s.t.} \quad \sum_{a \in \mathcal{A}} y_{s,a}(0) = x_s \quad \forall s, \quad (5b)$$

$$\sum_{a \in \mathcal{A}} y_{s,a}(t+1) = (\phi(\mathbf{y}(t)))_s \quad \forall s, t, \quad (5c)$$

$$D\mathbf{y}(t) \leq \mathbf{b} \quad \forall t, \quad (5d)$$

In the above formulation, the constraint (5b) corresponds to the condition on the initial state. The constraint (5c) corresponds to the time-evolution (1), plus the fact that $x_s(t+1) = \sum_a y_{s,a}(t+1)$. Lastly, the constraint (5d) corresponds to $D\mathbb{E}[\mathbf{Y}^{(N)}(t)] \leq \mathbf{b}$, which is the relaxed version of the one in (2). Equivalently it can be written as $\mathbf{y}(t) \in \mathcal{Y}(\mathbf{x}(t))$, where $x_s(t) = \sum_{a \in \mathcal{A}} y_{s,a}(t)$. For later reference, we call the linear program (5) being parameterized by (\mathbf{x}, T) , where \mathbf{x} is the initial state vector and T is the horizon.

By the assumptions that $D(s, 0) = \mathbf{0}$ and $D(s, a), \mathbf{b} \geq \mathbf{0}$, the linear program (5) is feasible (e.g. it suffices to always choose the passive action $a = 0$). In the following, we denote by \mathbf{y}^* one of its optimal solution, and by \mathbf{x}^* the sequence of vectors $\mathbf{x}^*(t)$ such that $x_s^*(t) = \sum_{a \in \mathcal{A}} y_{s,a}^*(t)$. It is the optimal state vector on the relaxed problem at time-step t .

4.2. The LP-update policy with full updates

To construct a policy for the system of size N , the LP solution suggests to use $\mathbf{Y}^{(N)}(t) = \mathbf{y}^*(t)$ at decision epoch t . Yet, this is in general not possible because of random fluctuations: Indeed, it is likely that $\mathbf{X}^{(N)}(t) \neq \mathbf{x}^*(t)$, which implies that in general, $\mathbf{y}^*(t)$ is not feasible for $\mathbf{X}^{(N)}(t)$, that is

$$\text{In general: } \mathbf{y}^*(t) \notin \mathcal{Y}^{(N)}(\mathbf{X}^{(N)}(t)). \quad (6)$$

The classical way to solve this problem in the literature is to construct a sequence of decision rules $\pi_t: \Delta^d \rightarrow \Delta^{d(A+1)}$ such that $\pi_t(\mathbf{x}) \in \mathcal{Y}^{(N)}(\mathbf{x})$ and $\pi_t(\mathbf{x}^*(t)) = \mathbf{y}^*(t)$. This is what has been used to build the randomized activation control policy in Zayas-Cabán et al. [39], the fluid-priority policies in Zhang and Frazier [41], the LP-index policy in Gast et al. [14], the occupancy-measured-reward index policy in Xiong et al. [38], to name a few. In particular, it is shown in [14] that any such policy is $\mathcal{O}(1/\sqrt{N})$ -optimal if all the decision rules $\pi_t(\cdot)$ are locally Lipschitz-continuous functions.

In this paper, we adopt another approach, that we call the *LP-update* policy, which is a generalization of the LP-update policy of Gast et al. [14] introduced for two-action bandits. The policy of [14] is as follows: at each decision epoch, we solve a new LP starting from $\mathbf{X}^{(N)}(t)$ with horizon $T - t$. This guarantees that the newly computed $\mathbf{Y}(t)$, that corresponds to the solution $\mathbf{y}^*(0)$ of the LP, is in $\mathcal{Y}(\mathbf{X}^{(N)}(t))$, by the mass conservation (5b). However, this control is not necessarily feasible for the system of size N , since $N \cdot Y_{s,a}(t)$ is not necessarily an integer. A simple² way to obtain a feasible solution is to use a rounding procedure:

$$Y_{s,a}^{(N)}(t) = \begin{cases} N^{-1} \lfloor N \cdot Y_{s,a}(t) \rfloor & \text{if } a \neq 0. \\ X_s^{(N)}(t) - \sum_{a>0} Y_{s,a}^{(N)}(t) & \text{if } a = 0. \end{cases} \quad (7)$$

which satisfy the resource constraints because the resource consumptions of action 0 are all 0.

This leads to our first LP-update algorithm, that is summarized in Algorithm 1.

4.3. Performance guarantees: general case

The forthcoming Theorem 1 demonstrates that the LP-update policy, as outlined in Algorithm 1, achieves an $\mathcal{O}(T^2/\sqrt{N})$ -optimal performance. This optimality further improves to $\mathcal{O}(T/\sqrt{N})$ when the probability transition matrices meet a certain ergodic condition. The proof of this theorem is anchored on two pivotal elements. Firstly, it leverages the characteristics of the LP-update algorithm

²We will discuss more sophisticated rounding procedures in Section D.

Algorithm 1: The LP-update policy for weakly-coupled MDPs (full updates).

Input: Initial configuration vector $\mathbf{X}^{(N)}(0)$ over time span $[0, T]$.

for $t = 0, 1, 2, \dots, T - 1$ **do**

Solve LP (5) with parameters $(\mathbf{X}^{(N)}(t), T - t)$. Output is \mathbf{y}^* over $[0, T - t]$;
 Set $Y_{s,a}^{(N)}(t) := N^{-1} [N y_{s,a}^*(0)]$ for $a \neq 0$ and $Y_{s,0}^{(N)}(t) := X_s^{(N)}(t) - \sum_{a \neq 0} Y_{s,a}^{(N)}(t)$;
 Use actions $Y_{s,a}^{(N)}(t)$ over all sub-MDPs to advance to the next time-step;

end

to establish a connection between the sub-optimality gap and the Lipschitz coefficient of the value function $\mathbf{x} \mapsto V_{\text{rel}}(\mathbf{x}, t)$. Secondly, it derives an explicit bound on this Lipschitz coefficient.

It is important to note that the first phase of this proof is inapplicable to LP-based policies that do not involve resolving the LP at every step. This distinction is crucial, as it is the resolution of the LP at each step in our approach that enables us to achieve a polynomial dependence on the time horizon T for the optimality gap.

THEOREM 1. Denote by $V_{\text{LP-update}}^{(N)}(\mathbf{x}, T)$ the value of the expected performance obtained by applying Algorithm 1 and by $V_{\text{rel}}(\mathbf{x}, T)$ the value of the linear program (5).

(i) For any weakly-coupled MDP with statistically identical arms and all N , we have:

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \leq \frac{r_{\max} \sqrt{d}}{4} \frac{T^2}{\sqrt{N}} + \frac{d|\mathcal{A}|Tr_{\max}}{N}.$$

(ii) For any weakly-coupled MDP with statistically identical arms such that³ $\gamma = \min_{i \neq j, a} \sum_k \min(P_{ik}^0, P_{jk}^a) > 0$, we have for all N :

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \leq \frac{r_{\max} \sqrt{d}}{2\gamma} \frac{T}{\sqrt{N}} + \frac{d|\mathcal{A}|Tr_{\max}}{N}.$$

(iii) There exists a weakly-coupled MDP with statistically identical arms and a constant $C' > 0$ such that for all N :

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \geq \frac{C'}{\sqrt{N}}.$$

Proof. The proof of the lower bound (iii) is done in Section A.4. Below, we prove Items (i) and (ii).

We start our analysis with Algorithm 1 that executes an update at every time step. Specifically, at each decision epoch t , the LP-update algorithm selects a vector $\mathbf{Y}(t)$ deemed optimal for the LP defined in (5), using the parameters $(\mathbf{X}^{(N)}(t), T - t)$. Following this, a decision vector $\mathbf{Y}^{(N)}(t)$ is chosen from the set $\mathcal{Y}^{(N)}(\mathbf{X}^{(N)}(t))$. This vector $\mathbf{Y}^{(N)}(t)$ is designed to be close to $\mathbf{Y}(t)$, with a deviation bounded by $\|\mathbf{Y}^{(N)}(t) - \mathbf{Y}(t)\|_1 \leq c/N$, as stipulated in our construction in (7). Upon implementing this decision, the controller receives a reward equal to $\mathbf{r}^\top \mathbf{Y}^{(N)}(t)$. Subsequently, the system progresses to the next state $\mathbf{X}(t + 1)$. This process translates into

$$V_{\text{LP-update}}^{(N)}(\mathbf{X}^{(N)}(t), T - t) = \mathbb{E} \left[\mathbf{r}^\top \mathbf{Y}^{(N)}(t) + V_{\text{LP-update}}^{(N)}(\mathbf{X}^{(N)}(t + 1), T - t - 1) \right]. \quad (8)$$

On the other hand, Bellman's principle of optimality says

$$V_{\text{rel}}(\mathbf{X}^{(N)}(t), T - t) = \mathbf{r}^\top \mathbf{Y}(t) + V_{\text{rel}}(\phi(\mathbf{Y}(t)), T - t - 1). \quad (9)$$

³The value of γ is very similar to what is call the ergodic coefficient of a MDP in [31]. The ergodic coefficient contains an additional minimization on action and does not impose $i \neq j$. It is equal to $\min_{i,j,a,b} \sum_k \min(P_{ik}^b, P_{jk}^a)$ and is smaller or equal to γ .

Denote by $Z(t) := \mathbb{E} \left[V_{\text{LP-update}}^{(N)}(\mathbf{X}^{(N)}(t), T-t) - V_{\text{rel}}(\mathbf{X}^{(N)}(t), T-t) \right]$ for $0 \leq t \leq T+1$, where we interpret $Z(T) = Z(T+1) = 0$. Combining (8) and (9) shows that $Z(t) - Z(t+1)$ is equal to

$$\underbrace{\mathbf{r}^\top \mathbb{E} [\mathbf{Y}^{(N)}(t) - \mathbf{Y}(t)]}_{\text{Term A}} + \underbrace{\mathbb{E} [V_{\text{rel}}(\mathbf{X}^{(N)}(t+1), T-t-1) - V_{\text{rel}}(\phi(\mathbf{Y}(t)), T-t-1)]}_{\text{Term B}}. \quad (10)$$

By construction, **Term A** is smaller than $r_{\max} d |\mathcal{A}| / N$. Moreover, we shall show in Lemma 2 that the function $\mathbf{x} \mapsto V_{\text{rel}}(\mathbf{x}, t)$ is L_t -Lipschitz-continuous, with $L_t \leq \min(t/2, 1/(2\gamma))$. Hence, the absolute value of **Term B** is such that

$$|(\text{Term B})| \leq \min \left(\frac{T-t-1}{2}, \frac{1}{2\gamma} \right) \mathbb{E} [\|\mathbf{X}^{(N)}(t+1) - \phi(\mathbf{Y}(t))\|_1].$$

By Lemma 1, the inequality $\mathbb{E} [\|\mathbf{X}^{(N)}(t+1) - \phi(\mathbf{Y}(t))\|_1] \leq \sqrt{d}/\sqrt{N}$ holds. The result then follows by summing over t from $t=0$ to $t=T$ for the differences $Z(t) - Z(t+1)$, depending on whether $\gamma > 0$ or not. Q.E.D.

Theorem 1(ii) shows that the sub-optimality gap of the LP-update policy is at most linear in T when $\gamma > 0$. The quantity γ corresponds to the minimum probability that two components starting in two different states couple after one time-step, where one component takes the action 0 and the other takes an arbitrary action. In fact, it is possible to generalize this notion by considering the probability that two components couple after more than one time step. More precisely, for any $\tau \in \{1, 2, \dots\}$, we define λ^τ as:

$$\lambda^\tau := \min_{i \neq j, a_1 \dots a_\tau} \sum_k \min(\underbrace{(P^0 \dots P^0)}_{\tau \text{ times}})_{ik}, (P^{a_1} \dots P^{a_\tau})_{jk} \quad (11)$$

Note that a sufficient condition for the existence of a τ such that $\lambda^\tau > 0$ is that the Markov chain corresponding to P^0 is irreducible and aperiodic. Indeed, by classical Markov chain theory, ergodicity and aperiodicity of P^0 implies that there exists a time $\tau \geq 1$ such the matrix $(P^0)^\tau$ has all of its coefficient strictly positive. This implies that $\lambda^\tau > 0$ for this value of τ . We obtain the following theorem, that is a generalization of Theorem 1(ii).

THEOREM 2. For any weakly-coupled MDP with statistically identical arms such that there exists an integer $\tau \geq 1$ such that $\lambda^\tau > 0$, we have:

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \leq \frac{\tau r_{\max} \sqrt{d}}{2\lambda^\tau} \frac{T}{\sqrt{N}} + \frac{d |\mathcal{A}| r_{\max} T}{N}.$$

Proof. The proof of this result is identical to the one of Theorem 1(ii), but we replace the use of Lemma 2 by the one of Lemma 4 (these lemma are given in Section A.2).

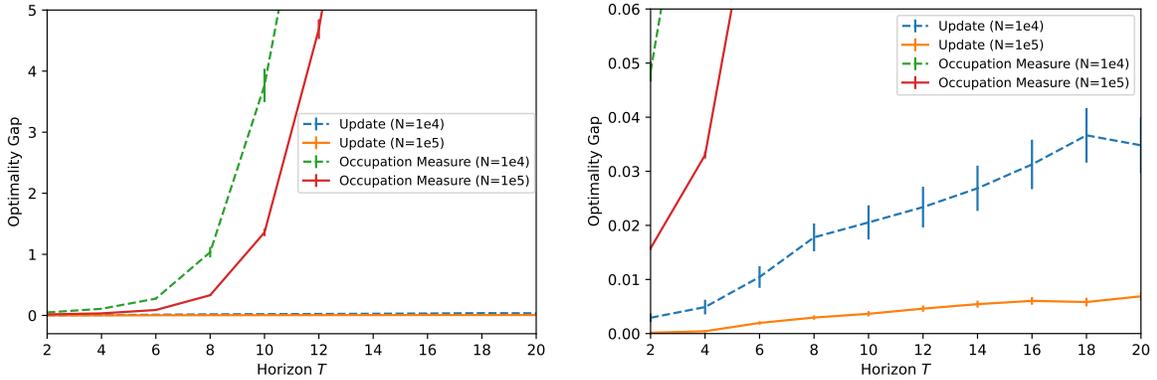
Theorem 1(i) shows that our algorithm offers a performance guarantee of $\mathcal{O}(1/\sqrt{N})$. As established in point (iii) of the Theorem, this rate is optimal for general models and cannot be further enhanced within that context. However, in Section 5, we will introduce the concept of *non-degenerate* problems. This concept builds upon and extends the ideas presented by Gast et al. [14], Zhang and Frazier [41]. For these specific non-degenerate problems, we demonstrate that the sub-optimality gap can be improved to $\mathcal{O}(1/N)$. Moreover, we will explore how leveraging the non-degenerate property of problems not only refines our theoretical understanding but also yields computational advantages.

4.4. Resolving LP brings performance benefits

Algorithm 1 solves a new LP at each decision epoch. This can be computationally intensive, making its time complexity high compared with the methods proposed by [5, 14, 38, 39, 41]. In this subsection, we show that this time cost is compensated by its performance gain. For that, we focus on a numerical

comparison between our LP-update policy and the *occupation measure* policy, introduced in [39, 38]. This algorithm solves an LP only once and then bases all subsequent decisions on this initial solution. The occupation measure policy has an asymptotic sub-optimality gap in $\mathcal{O}(1/\sqrt{N})$, as established in [39].

However, a closer look at the proof of the occupation measure’s optimality shows that its hidden constant in the $\mathcal{O}(\cdot)$ notation grows exponentially with the time horizon T , due to estimations similar to those in Gronwall’s lemma. In contrast, the sub-optimality gap of our LP-update algorithm has a more modest growth rate of $\mathcal{O}(T^2/\sqrt{N})$ or $\mathcal{O}(T/\sqrt{N})$.



(a) LP-update vs occupation measure

(b) Zoom of the left panel.

Figure 1 Sub-optimality gap as a function of the time-horizon T . The right panel is a zoom-in of the left panel.

In Figure 1, we illustrate the sub-optimality gap of both policies as a function of the time horizon T for a model generated at random. This model assumes a large number of identical arms, with N being either 10^4 or 10^5 . Our observations indicate that the sub-optimality gap for the occupation measure policy appears to increase exponentially with T . In contrast, the sub-optimality gap of the LP-update algorithm displays a linear growth relative to T . This finding underscores that solving a new LP at each time step can offer significant performance advantages over an algorithm that resolves the LP only once at the initial step.

In the upcoming section, we will explore how the computational demands associated with repeatedly solving a new LP can be substantially diminished, if the problem satisfies the non-degenerate condition.

5. NON-DEGENERATE PROBLEMS AND IMPROVED CONVERGENCE RATE

In this section, we define what we call a non-degenerate problem in Section 5.1, and show how it allows one to design an improved LP-update policy with selective updates that is more efficient in Section 5.2. As we will see, when a problem is non-degenerate, the solutions to the LP starting from initial conditions $\mathbf{x} \approx \mathbf{x}^*(t)$ can be expressed as linear functions of \mathbf{x} . We will show that this can be used to improve both the computational efficiency of the algorithm and the rate at which the algorithm becomes asymptotically optimal: we prove that the new LP-update policy has an $\mathcal{O}(1/N)$ performance guarantee for non-degenerate problems. We discuss questions related to rounding in Section D, which is an important step when applying the policy. We provide an alternative formulation of non-degeneracy in Section C, and compare it with various existing notions in the literature. The proofs of the technical results are postponed to Section A.

5.1. The non-degenerate property

To ease the notation, we denote by $u := d(A + 1)$. We start by remarking that the linear program (5) can be decomposed recursively: By Bellman's principle of optimality, for each $0 \leq t \leq T - 1$, we have

$$V_{\text{rel}}(\mathbf{x}, T - t) = \max_{\mathbf{y} \in \mathbb{R}^u} \mathbf{r}^\top \mathbf{y} + V_{\text{rel}}(\phi(\mathbf{y}), T - t - 1) \quad (12a)$$

$$\text{s.t.} \quad \mathbf{y} \geq \mathbf{0}, \quad (12b)$$

$$D\mathbf{y} \leq \mathbf{b}, \quad (12c)$$

$$E\mathbf{y} = \mathbf{x}. \quad (12d)$$

where E is the matrix encoding the equality constraints $\sum_{a \in \mathcal{A}} y_{s,a} = x_s$, so that it has d rows and u columns, for which the s -th row is such that from the $(s - 1)(A + 1) + 1$ -th to the $s(A + 1)$ -th coordinates are 1, the rest are 0.

Let \mathbf{y}^* be an optimal solution to the linear program (5), with the corresponding configuration vectors \mathbf{x}^* . Write $\mathbf{y}^*(t)$ (resp. $\mathbf{x}^*(t)$) the part of \mathbf{y} (resp. \mathbf{x}^*) at time-step t . Define $\mathcal{J}^*(t)$ as the set of indices for which the budget constraint (12c) is an equality: $(D\mathbf{y}^*(t))_j = b_j$ for all $j \in \mathcal{J}^*(t)$ and $(D\mathbf{y}^*(t))_j < b_j$ for $j \in \mathcal{J} \setminus \mathcal{J}^*(t)$. We now consider the following optimization problem:

$$F_{\mathbf{y}^*}(\mathbf{x}, T - t) = \max_{\mathbf{y} \in \mathbb{R}^u} \mathbf{r}^\top \mathbf{y} + V_{\text{rel}}(\phi(\mathbf{y}), T - t - 1) \quad (13a)$$

$$\text{s.t.} \quad \mathbf{y} \geq \mathbf{0}, \quad (13b)$$

$$(D\mathbf{y})_j < b_j \quad \forall j \notin \mathcal{J}^*(t), \quad (13c)$$

$$(D\mathbf{y})_j = b_j \quad \forall j \in \mathcal{J}^*(t), \quad (13d)$$

$$E\mathbf{y} = \mathbf{x} \quad (13e)$$

As (13) is more restrictive than (12), we have $F_{\mathbf{y}^*}(\mathbf{x}, T - 1) \leq V_{\text{rel}}(\mathbf{x}, T - t)$. Moreover, by definition, when $\mathbf{x} = \mathbf{x}^*(t)$, we have $F_{\mathbf{y}^*}(\mathbf{x}^*(t), T - 1) = V_{\text{rel}}(\mathbf{x}^*(t), T - t)$. In what follows, we look for a condition which guarantees that this equality is preserved in a neighbourhood of $\mathbf{x}^*(t)$, for all $1 \leq t \leq T - 1$. In other words, the sets of saturated and non-saturated constraints of an optimal LP solution to (5) are unchanged upon a small perturbation to the state vectors $\mathbf{x}^*(t)$.

To do so, let $\mathcal{U}^*(t) \subset \mathcal{U}$ be the set of indices (s, a) for which $y_{s,a}^*(t) = 0$, and $\mathcal{S}^*(t)$ be the set of states for which $x_s^*(t) > 0$. We collect all the equality constraints from (13b), (13d) and (13e) below:

$$y_{s,a} = 0 \quad \forall (s, a) \in \mathcal{U}^*(t) \quad (14)$$

$$\sum_{s,a} D_j(s, a) y_{s,a} = b_j \quad \forall j \in \mathcal{J}^*(t) \quad (15)$$

$$\sum_a E(s, a) y_{s,a} = x_s(t) \quad \forall s \in \mathcal{S}^*(t) \quad (16)$$

The above equalities (14)–(16) can be represented in compact form using a matrix $C^*(t)$ that has $|\mathcal{J}^*(t)| + |\mathcal{S}^*(t)| + |\mathcal{U}^*(t)|$ rows and $u = |\mathcal{U}|$ columns:

$$C^*(t)\mathbf{y} = [\mathbf{0}_{|\mathcal{U}^*(t)}; \mathbf{b}_{|\mathcal{J}^*(t)}; \mathbf{x}_{|\mathcal{S}^*(t)}]^\top, \quad (17)$$

where the notations $\mathbf{b}_{|\mathcal{J}^*(t)}$ and $\mathbf{x}_{|\mathcal{S}^*(t)}$ indicate that the vectors are restricted to the coordinates $\mathcal{J}^*(t)$ or $\mathcal{S}^*(t)$. We wish to establish an inverse to the above linear relation (17), so that if $\mathbf{x}^*(t)$ is slightly perturbed as $\tilde{\mathbf{x}}$ (where the perturbation comes from the stochastic noise), the inverse equation gives a $\tilde{\mathbf{y}}$ that remains optimal.

We are prepared to introduce the concept of non-degeneracy, a pivotal term in our discussion.

DEFINITION 1 (NON-DEGENERACY). An LP (5) is *non-degenerate* if there exists a solution \mathbf{y}^* to the LP, for which at *all* time-step $t \in \{1 \dots T - 1\}$, the matrix $C^*(t)$ in (17) has rank $|\mathcal{J}^*(t)| + |\mathcal{S}^*(t)| + |\mathcal{U}^*(t)|$.

Note that in the above definition we exclude $t = 0$, since the perturbation comes from stochastic noise, which only appears from $t = 1$ onward.

Recall that our model is a generalization of the restless bandit model studied in [14, 41, 6], for which a notion of non-degeneracy is already introduced in the aforementioned papers. We show in Section C that our definition coincides with their definition when we restrict to the model studied in [14, 41, 6]. This shows that our new definition is indeed an extension of theirs to the broader class of multi-action multi-constraint bandit problems.

To the best of our knowledge, the most comparable notion of non-degeneracy at the same level of generality as ours is found in Zhang [40]. However, as discussed in Section C, their definition is more restrictive: basically, it asks that a single constraint is saturated. Therefore, our definition encompasses a broader spectrum of problems as non-degenerate compared to theirs. This distinction is visually represented in Figure 2, where we quantified the proportion of non-degenerate problems according to both definitions. For this, we used randomly generated transition matrices (both dense and tri-diagonal) and reward vectors, plotting the percentage of non-degenerate examples against the number of states d . These plots reveal that about half of the problems deemed non-degenerate by our definition are classified as degenerate by the definition in [40].

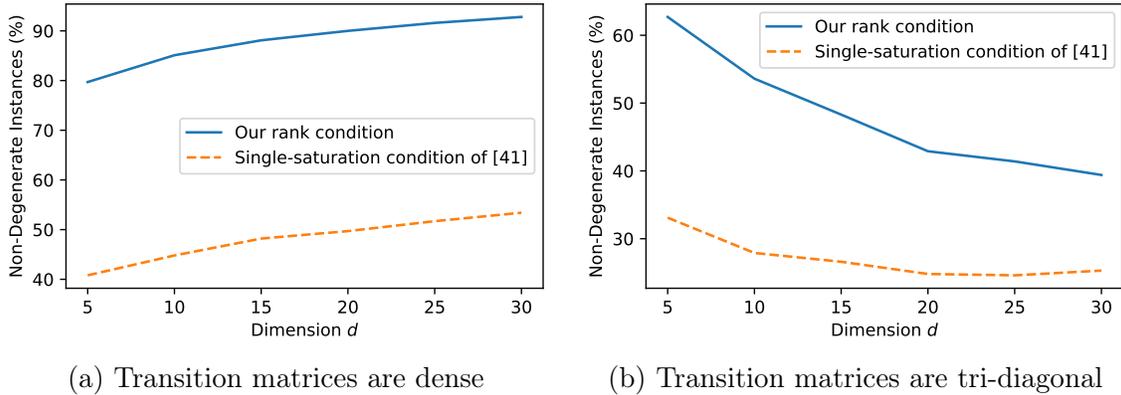


Figure 2 Percentage of non-degenerate examples as a function of the sub-component's state space dimension d .

Our concept of non-degeneracy, defined through a rank condition, bears a significant resemblance to the Linear Independence Constraint Qualification found in general nonlinear optimization problems, as discussed in works such as Bazarra et al. [3]. The latter stipulates that the gradients of active constraints in an optimization problem must be linearly independent. We adopt this principle in a unique way, by breaking down the linear program outlined in (5) into a T -step optimization problem, following Bellman's principle of optimality. In our approach, the notion of non-degeneracy implies that the saturated constraints in the optimal solution of the LP should be distributed "evenly" across the T steps. As we shall see, this property is key to ensuring that we have sufficient rank to create an inverse linear mapping around the original optimal solution at *any* given time-step t , while accounting for the stochastic uncertainties in the system states $\mathbf{x}(t)$.

5.2. The improved LP-update policy

For a given $\mathbf{x} \in \Delta^d$ and $\varepsilon > 0$, we define the neighbourhood of \mathbf{x} of size ε as $\mathcal{B}(\mathbf{x}, \varepsilon) := \{\mathbf{x}' \in \Delta^d \mid |x'_s - x_s| < \varepsilon \text{ and } x'_s = 0 \text{ for all } s \text{ such that } x_s = 0\}$. As we show below, being non-degenerate implies that the optimal solution $\mathbf{y}(\mathbf{x})$ to the LP (12) can be chosen as a *locally linear* function of \mathbf{x} in a small neighbourhood of \mathbf{x}^* .

PROPOSITION 1. Assume that the problem is non-degenerate as in Definition 1. Then, for all time-step $t \geq 1$, the matrix $C^*(t)$ has a right inverse $C^+(t)$. Moreover, there exists $\varepsilon > 0$ such that:

- The function $\mathbf{x} \mapsto V_{\text{rel}}(\mathbf{x}, T - t)$ is linear on $\mathcal{B}(\mathbf{x}^*(t), \varepsilon)$.
- Choosing

$$\mathbf{y}(\mathbf{x}) := \mathbf{y}^*(t) + C^+(t) \begin{bmatrix} \mathbf{0}|_{\mathcal{U}^*(t)} \\ \mathbf{0}|_{\mathcal{J}^*(t)} \\ (\mathbf{x} - \mathbf{x}^*(t))|_{\mathcal{S}^*(t)} \end{bmatrix} \quad (18)$$

is an optimal solution of (12) for all $\mathbf{x} \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon)$.

Proof. By standard linear algebra argument, a matrix of dimension $d_1 \times d_2$ with rank d_1 has a right inverse. This implies that there exists a matrix $C^+(t)$ such that $C^*(t)C^+(t)$ is the identity matrix. Hence if $\mathbf{y}(\mathbf{x})$ is defined by (18), then

$$C^*(t)\mathbf{y}(\mathbf{x}) = C^*(t)\mathbf{y}^*(t) + C^*(t)C^+(t) \begin{bmatrix} \mathbf{0}|_{\mathcal{U}^*(t)} \\ \mathbf{0}|_{\mathcal{J}^*(t)} \\ (\mathbf{x} - \mathbf{x}^*(t))|_{\mathcal{S}^*(t)} \end{bmatrix} = \begin{bmatrix} \mathbf{0}|_{\mathcal{U}^*(t)} \\ \mathbf{b}|_{\mathcal{J}^*(t)} \\ \mathbf{x}|_{\mathcal{S}^*(t)} \end{bmatrix}$$

In particular, $\mathbf{y}(\mathbf{x})$ satisfies (14)–(16). This shows that there exists $\varepsilon > 0$ such that $\mathbf{y}(\mathbf{x}) \in \mathcal{Y}(\mathbf{x})$ (i.e. satisfy all constraints of (12)) for all $\mathbf{x} \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon)$, because the constraints that are not covered by (14)–(16) are either satisfied by *strict* inequalities for $\mathbf{x}^*(t)$, or correspond to $x_s^*(t) = 0$. It is worth noting that the selection of ε is explicitly based on the *slack* of the non-saturated constraints.

We now prove by a backward induction on t , that for all t , there exists $\varepsilon_t > 0$ such that the function $\mathbf{x} \mapsto V_{\text{rel}}(\mathbf{x}, T - t)$ is linear on $\mathcal{B}(\mathbf{x}^*(t), \varepsilon_t)$. This is clearly true for $t = T$ for which $V_{\text{rel}}(\mathbf{x}, 0) = 0$ for all \mathbf{x} .

Assume now that it holds for some $t + 1 \leq T$, and denote by $g_t(\mathbf{x}, \mathbf{y})$ the reward provided by the control \mathbf{y} with initial state vector \mathbf{x} and horizon $T - t$. As shown before, for \mathbf{x} close enough to $\mathbf{x}^*(t)$, the control $\mathbf{y}(\mathbf{x})$ is feasible for \mathbf{x} . Moreover, the induction hypothesis implies that $V_{\text{rel}}(\phi(\mathbf{y}(\mathbf{x})), T - t - 1)$ is locally linear in \mathbf{x} for all \mathbf{x} close enough to $\mathbf{x}^*(t)$. This shows that $\mathbf{x} \mapsto g_t(\mathbf{x}, \mathbf{y}(\mathbf{x}))$ is locally linear on $\mathcal{B}(\mathbf{x}^*(t), \varepsilon_t)$. We argue that this implies that $\mathbf{y}(\mathbf{x})$ is the optimal control for all $\mathbf{x} \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon_t)$. Indeed:

- As $V_{\text{rel}}(\mathbf{x}, T - t)$ is the solution of a linear program where \mathbf{x} can be seen as a linear constraint, the function $\mathbf{x} \mapsto V_{\text{rel}}(\mathbf{x}, T - t)$ is concave in \mathbf{x} . Moreover, by construction $\mathbf{y}(\mathbf{x}^*(t))$ provides the optimal solution for $\mathbf{x}^*(t)$.
- Denote by $\mathbf{x}' := 2\mathbf{x}^*(t) - \mathbf{x}$. By construction $\mathbf{x}' \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon_t)$. By concavity of the $V_{\text{rel}}(\cdot, T - t)$ function on \mathbf{x} , the possible sub-optimality of $\mathbf{y}(\mathbf{x})$, and together with the local linearity of $g_t(\mathbf{x}, \mathbf{y}(\mathbf{x}))$, we deduce that:

$$\begin{aligned} V_{\text{rel}}(\mathbf{x}^*, T - t) &\geq (V_{\text{rel}}(\mathbf{x}, T - t) + V_{\text{rel}}(\mathbf{x}', T - t))/2 \\ &\geq (g_t(\mathbf{x}', \mathbf{y}(\mathbf{x}')) + g_t(\mathbf{x}, \mathbf{y}(\mathbf{x}))/2 \\ &= g_t(\mathbf{x}^*, \mathbf{y}(\mathbf{x}^*)) = V_{\text{rel}}(\mathbf{x}^*, T - t). \end{aligned}$$

This shows that the inequalities must be equalities, which implies that $\mathbf{y}(\mathbf{x})$ defined as in Equation (18) is optimal on $\mathcal{B}(\mathbf{x}^*(t), \varepsilon_t)$. Q.E.D.

The definition of a non-degenerate problem suggests that the original LP-update Algorithm 1 can be implemented by only recomputing an update (i.e. re-solving) when necessary:

1. When at a time t the rank condition on $C^*(t)$ is not satisfied, then one cannot compute the right inverse $C^+(t)$.
2. When at a time t , the rank condition is satisfied but the stochastic trajectory has deviated too much from the optimal deterministic trajectory, so that the suggested decision vector $\mathbf{y}(\mathbf{X}^{(N)}(t))$ in (18) no longer gives a feasible decision vector, i.e. is not in $\mathcal{Y}(\mathbf{X}^{(N)}(t))$.

When one of these two situations occurs, the algorithm solves a new LP with initial configuration vector $\mathbf{X}^{(N)}(t)$ over the horizon $[t, T]$. This leads to the improved LP update algorithm that is described in Algorithm 2.

Algorithm 2: The LP-update policy for weakly-coupled MDPs (selective updates).

Input: Initial configuration vector $\mathbf{X}^{(N)}(0)$.

for $t = 0, 1, 2, \dots, T - 1$ **do**

if $t \geq 1$ **and** $C^*(t)$ **has rank** $|\mathcal{J}^*(t)| + |\mathcal{S}^*(t)| + |\mathcal{U}^*(t)|$ **and** $\mathbf{y}(\mathbf{X}^{(N)}(t)) \in \mathcal{Y}(\mathbf{X}^{(N)}(t))$ **then**

 Set $\mathbf{y}(t) := \mathbf{y}(\mathbf{X}^{(N)}(t))$ as in (18);

else

 Solve LP (5) with parameters $(\mathbf{X}^{(N)}(t), T - t)$ and call the output $\mathbf{y}^*(\tau)$ for $\tau \in \{t \dots T\}$;

 Set $\mathbf{y}(t) := \mathbf{y}^*(t)$;

end

 Set $Y_{s,a}^{(N)}(t) := N^{-1} \lfloor Ny_{s,a}(t) \rfloor$ for $a \neq 0$ and $Y_{s,0}^{(N)}(t) := X_s^{(N)}(t) - \sum_{a \neq 0} Y_{s,a}^{(N)}(t)$;

 Use actions $Y_{s,a}^{(N)}(t)$ over all sub-MDPs to advance to the next time-step;

end

It is worth noting that $\mathbf{y}(t) \in \mathcal{Y}(\mathbf{X}^{(N)}(t))$ is not a sufficient condition for optimality (this only holds if $\mathbf{X}^{(N)}(t)$ is in the neighborhood of $x^*(t)$). Consequently, Algorithm 2 may not have the same output as Algorithm 1, even in cases where no update is performed in Algorithm 2. Nonetheless, the square-root convergence rate established in Theorem 1 is also applicable to Algorithm 2, as can be demonstrated using analogous reasoning.

Additionally, it is important to highlight that Algorithm 2 requires fewer updates than Algorithm 1, enhancing its computational efficiency, see Section 6.4 for a demonstration. In our forthcoming proof, we will show that the discrepancy in the outcomes of the two algorithms is exponentially small for non-degenerate problems. Therefore, in the subsequent theorem, we use $V_{\text{LP-update}}^{(N)}(\mathbf{x}, T)$ to refer interchangeably to the value derived from either the LP-update policy in Algorithm 1 or Algorithm 2.

5.3. Performance guarantees

The next theorem shows that when a problem is non-degenerate, the sub-optimality of the LP-update is of order $\mathcal{O}(1/N)$, which is an order of magnitude faster than the $\mathcal{O}(1/\sqrt{N})$ guaranty of Theorem 1. One drawback, though, is that the hidden dependence on T in the $\mathcal{O}(1/N)$ can be exponential in T . This exponential dependence on T comes from the proof of the theorem which can be decomposed in two steps: first, we show that $\mathbf{X}^{(N)}(t)$ is close to $\mathbf{x}^*(t)$ that when N is large, and second we show that V_{rel} is local linear close to $\mathbf{x}^*(t)$. The first step leads to an exponential dependence in T in the worst case. This dependence disappears in the proof of Theorem 3 because we only use that V_{rel} is Lipschitz-continuous but not that $\mathbf{X}^{(N)}(t)$ is close to $\mathbf{x}^*(t)$. That being said, in practice, the LP-update becomes asymptotically much faster for non-degenerate examples than for degenerate examples (see Section 5.4).

THEOREM 3. Denote by $V_{\text{LP-update}}^{(N)}(\mathbf{x}, T)$ the value of the LP-update policy defined in Algorithm 2 (or Algorithm 1), and by $V_{\text{rel}}(\mathbf{x}, T)$ the value of the linear program (5).

(i) For any non-degenerate weakly-coupled MDP with statistically identical arms, there exists constant $C > 0$ such that for all N :

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \leq \frac{C}{N}.$$

(ii) There exists a non-degenerate weakly-coupled MDP with statistically identical arms and a constant $C' > 0$ such that for all N :

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \geq \frac{C'}{N}.$$

Proof. As for Theorem 1, we prove the upper bound (i) and postpone the proof of the lower bound (ii) to Section A.4.

Let \mathbf{y}^* be the optimal solution of the LP solved at time 0 and let $x_s^*(t) = \sum_a y_{s,a}^*(t)$. In what follows, we show that with very high probability, $\mathbf{X}^{(N)}(t)$ remains close to $\mathbf{x}^*(t)$ so that no update is necessary, and the sequence of $\mathbf{Y}(t)$ computed by Algorithm 2 will always be an optimal solution to the LP problems. To this end, we shall bound the probability of the event when there is some t for which $\mathbf{X}^{(N)}(t)$ deviates far away from $\mathbf{x}^*(t)$ from above.

By Proposition 1, there exists $\varepsilon > 0$ such that $V_{\text{rel}}(\mathbf{x}, T-t)$ is linear in $\mathcal{B}(\mathbf{x}^*(t), \varepsilon)$ and such that the control $\mathbf{y}(\mathbf{x})$ defined in Proposition 1 and used in Algorithm 2 is optimal when $\mathbf{X}^{(N)}(t) \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon)$. Call \mathcal{E} the event $\mathbf{X}^{(N)}(t) \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon)$. By Lemma 5, there exists $C_1, C_2 > 0$ such that the event \mathcal{E} holds with probability at least $C_1 \cdot e^{-C_2 N}$. Hence, when \mathcal{E} is true, Algorithm 2 behaves as Algorithm 1. This shows that (8) holds also for Algorithm 2, up to an $\mathcal{O}(e^{-C_2 N})$ term due to the (exponentially small) probability that \mathcal{E} does not hold. Similarly (10) also holds for Algorithm 2 up to an $\mathcal{O}(e^{-C_2 N})$ term. As $V_{\text{rel}}(\mathbf{x}, T-t)$ is locally linear when \mathcal{E} holds, (Term B) of (10) is smaller than $C_1 \cdot e^{-C_2 N}$.

This shows that the sub-optimality gap is $\sum_{t=0}^T \mathbf{r}^\top \mathbb{E} [\mathbf{Y}^{(N)}(t) - \mathbf{Y}(t)] + \mathcal{O}(e^{-C_2 N})$. Moreover, by our explicit construction of $\mathbf{Y}^{(N)}(t)$ as in (7), it satisfies $\|\mathbf{Y}^{(N)}(t) - \mathbf{Y}(t)\|_1 \leq d|\mathcal{A}|/N$, which gives Theorem 3. Q.E.D.

It is noteworthy that the term $\mathcal{O}(1/N)$ in our analysis arises totally from the rounding procedure detailed in (7). This aspect is further explored in Section D, where we introduce the concept of a *perfect rounding condition*. Under this ideal condition, the rounding error is effectively eliminated, thereby enhancing the convergence rate to $\mathcal{O}(e^{-C_2 N})$. Additionally, this section presents an alternative rounding strategy, which has demonstrated superior performance in practical applications.

5.4. Sub-optimality gap for degenerate and non-degenerate problems

To demonstrate the distinct behaviors between degenerate and non-degenerate cases, we conducted an experiment with two randomly chosen examples – one degenerate and one non-degenerate. The results, showcased in Figure 3, compare the sub-optimality gap of the LP-update policy, as well as the occupation measure policy as a function of the number of arms, N . This comparison unequivocally shows that our LP-update policy consistently outperforms the others.

Additionally, the experiment highlights notable differences in the rate at which these policies approach asymptotic optimality. Consistent with our theorem, the sub-optimality gap for the LP-update policy is $\mathcal{O}(1/\sqrt{N})$ in degenerate cases and improves to $\mathcal{O}(1/N)$ in non-degenerate cases. In contrast, the occupation measure policy suffers a sub-optimality gap of $\mathcal{O}(1/\sqrt{N})$ in the two scenarios.

We remark that the zigzag pattern in the performance curves of the LP-update policy is caused by cutoff rounding errors, as outlined in Equation (7). This pattern is mitigated when a more sophisticated rounding approach is employed, as showcased in Section D.1.

6. CASE STUDY: GENERALIZED APPLICANT SCREENING PROBLEM

The applicant screening problem is proposed in Brown and Smith [5], and has been subsequently studied in Gast et al. [14]. The problem studied in these papers can be modeled as a two-action single-constraint restless bandit problem. We study in this section a generalization of the problem allowing multiple actions while also having multiple constraints, so that the model fits naturally into our weakly-coupled MDP framework.

6.1. Problem description

Consider a group of N applicants applying for a job. The decision maker’s goal is to hire the best possible βN applicants. The applicant n has an unknown quality level $p_n \in [0, 1]$. At each decision epoch t , the decision maker chooses for each candidate either to interview this candidate with one or

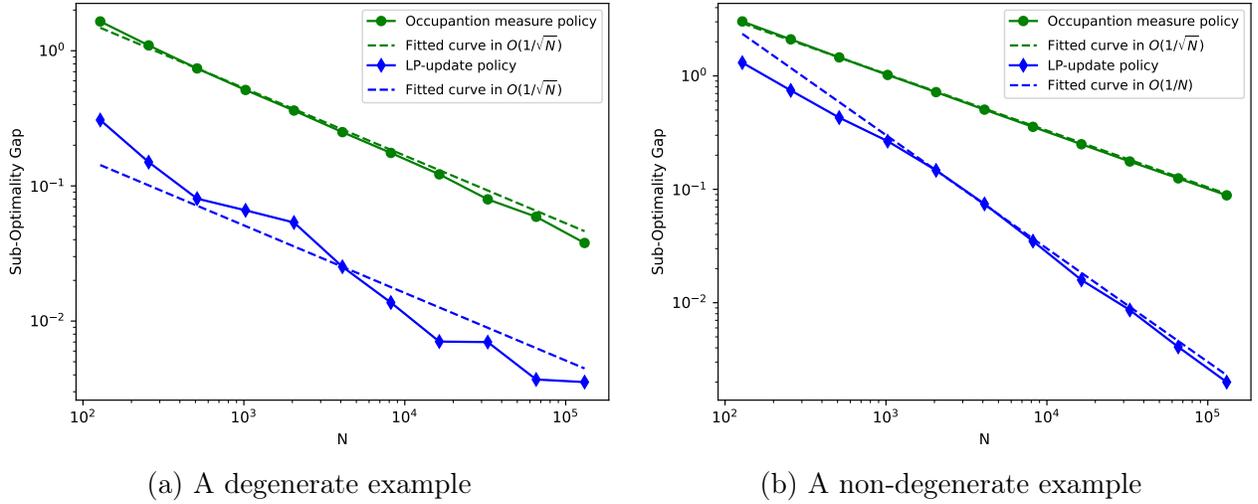


Figure 3 Sub-optimality gap in degenerate and non-degenerate examples: for non-degenerate examples, the sub-optimality gap of our LP-update policy decreases at a faster rate $O(1/N)$ than the occupation measure policy of [39].

two questions, or chooses not to interview this candidate, giving the action set $A = \{0, 1, 2\}$. For each interviewed applicant (i.e. $a \in \{1, 2\}$), a signal $q_n(t) \sim \text{binomial}(a, p_n)$ is returned, indicating how many among the a questions have been solved correctly by the applicant. All variables $q_n(t)$ are supposed to be independent.

Choosing an action a on an applicant consumes $D(a)$ units of resource (time, space, organization cost, etc.), for which we choose to be

$$D(s, a) = D(a) = \begin{cases} 0, & \text{if } a = 0; \\ 1, & \text{if } a = 1; \\ 1.5, & \text{if } a = 2. \end{cases} \quad (19)$$

The value "1.5" on action $a = 2$ is interpreted as asking a single applicant consecutively two questions consumes less resource than asking separately two applicants each with one question. At each decision epoch a total amount of αN resource is available. There is a total number of T interviewing rounds, and in the final $(T + 1)$ -th round, the decision maker admits βN applicants, based on the results of the interviewing rounds.

We assume that the applicants belong to two different groups, each having a population of $N/2$. And we will consider two scenarios:

- No fairness: in this scenario, there is a single budget constraint of αN for the whole population.
- Fair selection: in addition to the above constraint, the decision maker cannot spend more than γN budget on each of the single group, where $\gamma < \alpha < 2\gamma$.

Notice that the above applicant screening problem generalizes the problem studied in Brown and Smith [5] and Gast et al. [14] in two ways: we allow for more than one question whereas the cited paper consider $A = \{0, 1\}$; and we consider a fairness constraint, whereas the cited paper only have a single resource constraint that the number of interviewed candidate should not be larger than αN .

6.2. Modeling as weakly-coupled MDP

To cast the problem into a weakly-coupled MDP as described in Section 3, we consider a Bayesian model in which the quality level p of an applicant from each group is generated from some beta distribution, and the decision maker's estimation on each applicant's p is updated using Bayes' rule.

The state s of an applicant is hence a 2-tuple (a, b) , indicating the current estimation of her quality level and is distributed according to the beta distribution $beta(a, b)$.

For the first T interviewing rounds, the action set is $A = \{0, 1, 2\}$. Upon taking action 0 on an applicant, the estimation is unchanged, so the matrix \mathbf{P}^0 is the identity matrix. Upon taking action 1, the state is updated according to

$$(a, b) \xrightarrow{\text{action 1}} \begin{cases} (a + 1, b), & \text{with probability } a/(a + b); \\ (a, b + 1), & \text{with probability } b/(a + b). \end{cases}$$

This gives the matrix \mathbf{P}^1 . Likewise, for action 2 the state is updated as

$$(a, b) \xrightarrow{\text{action 2}} \begin{cases} (a + 2, b), & \text{with probability } \frac{a(1+a)}{(a+b)(1+a+b)}; \\ (a + 1, b + 1), & \text{with probability } \frac{2ab}{(a+b)(1+a+b)}; \\ (a, b + 2), & \text{with probability } \frac{b(1+b)}{(a+b)(1+a+b)}. \end{cases}$$

This gives the matrix \mathbf{P}^2 . The function $D(s, a)$ is given by (19) and is independent of the state s . The rewards are all zero during the first T rounds.

For the final $(T + 1)$ -th admitting round, the action set is $\{\text{admit}, \text{not admit}\}$. The reward on an admitted applicant in state (a, b) is $a/(a + b)$; the reward on a non admitted applicant is zero.

6.3. Discussion on simulation results

For our simulations, we choose $\beta = 0.1$, and we consider two scenarios. The first one is such that $\alpha = 0.15$, $\gamma = 0.1$. This is the scenario where the resource is "scarce". The second scenario where the resource is "abundant" is such that $\alpha = 0.3$ and $\gamma = 0.2$ (resource being doubled). In each scenario we shall compare the effect of adding or removing the fairness constraints. Without the fairness constraints, the decision maker can distribute the total αN units of resource freely among the two groups at each interviewing round.

We assume that the decision maker's prior estimations on the two groups of applicants are respectively $beta(1, 1)$ and $beta(2, 2)$, so that they have the same mean but the second group has a lower variance. Throughout the horizon is fixed to $T = 10$. To make the problem more realistic, and in the same time to reduce the total number of possible states of the MDP, we require in addition that no more than 10 questions can be asked on a single applicant during the 10 interviewing rounds.

The simulations are done for N ranging from 20 to $10240 = 20 \times 2^{10}$. For each value of N , we generate 1600 instances of N applicants according to the beta distributions we described previously in each scenario,. We apply the LP-update policy and occupation measure policy on each instance, with or without the fairness constraints. In each simulation the performance is evaluated based on the average quality level of the final admitted βN applicants. The results are reported in Figure 4.

Note that as guaranteed by Theorem 1 and Theorem 1 of Xiong et al. [38], both LP-update policy and occupation measure policy converge to the LP-relaxed bounds, as N goes to infinity. This is what we observe in Figure 4: all policies converge to the relaxed upper bound as N goes to infinity. The situation is however different for smaller values of N : In all cases, the LP-update policy outperforms the occupation measure policy. The smaller the value of N , the more apparent is the advantage of the LP-update policy.

We shall now discuss the effect of adding fairness constraints. We observe that in the first scarce resource scenario, the LP-relaxed performance bound becomes smaller when fairness is imposed. This should be expected since adding more constraints on a linear program can only decrease its optimal value. However, in the second scenario, when the resource is doubled, adding or removing the fairness constraints result in the same upper-bound value, and the performances of LP-update policy are identical in the two situations. Nevertheless, these fairness constraints still play a role in

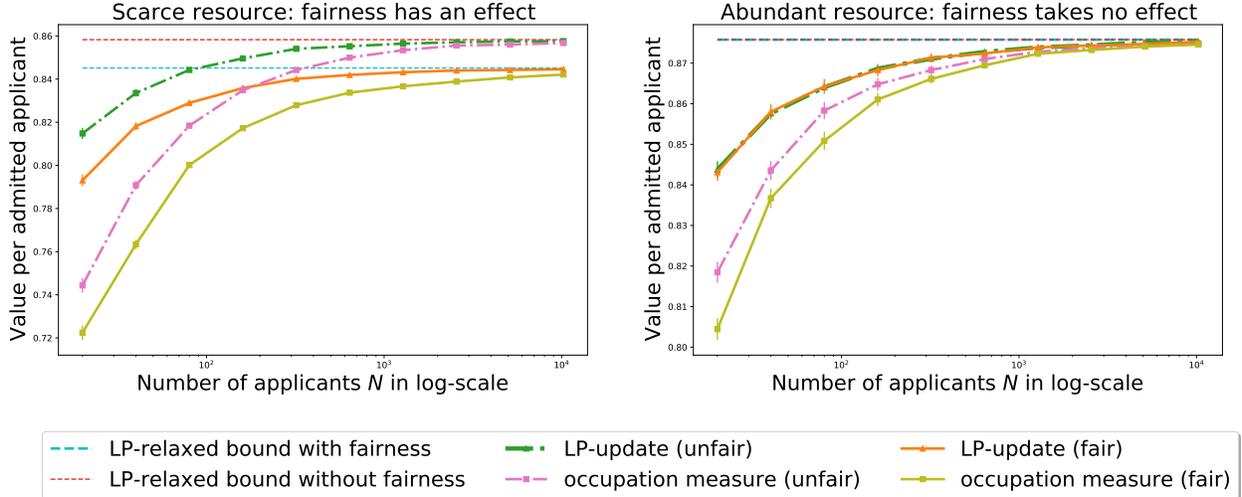


Figure 4 Performance on generalized applicant screening problems when the resource is scarce (left panel) or abundant (right panel), with or without fairness constraints. Overall the LP-update policy outperforms the occupation measure policy in all situations, and the advantage is more apparent for small values of N . Fairness has a negative impact in the scarce resource scenario, whereas in the abundant case, it has no effect asymptotically, but can still influence the performance of the occupation measure policy.

both policies: for the LP-update policy we observe that the fairness constraints can be saturated when solving new linear programs by applying updates, but it turns out that this does not influence its performance at all. For the occupation measure policy, however, sampling an action is more restrictive if there are more constraints, and as a result this influences its performance, even though asymptotically they converge to the same limit.

6.4. Computation time analysis

Although the LP-update policy outperforms the occupation measure policy under all situations, it comes with a price of using more computations. This extra computation time is due to the fact that the LP-update policy periodically solves a new LP problem. We record the execution time for different values of N in Table 1, averaging over 100 runs on one problem instance under each circumstance, together with the 95% confidence interval. Here by execution time we mean the time needed to apply a policy on an instance of N applicants problem until the final phase of admission round. The program is written in Python using NumPy for data structure and PuLP for solving LP's. Note that our code is not particularly optimized and is tested on an ordinary personal laptop.

| | | Number of applicants | $N = 20$ | $N = 100$ | $N = 1000$ |
|---------------------|----------------------------------|----------------------|---------------------------------------|---------------------------------------|--------------------------------------|
| With fairness | Occupation measure | | 2.70 ± 0.01 | 2.72 ± 0.01 | 2.83 ± 0.01 |
| | LP-update (Number of updates) | | 12.69 ± 0.39 (≈ 6.4) | 10.44 ± 0.41 (≈ 5.2) | 7.06 ± 0.19 (≈ 3.9) |
| Without fairness | Occupation measure | | 2.62 ± 0.01 | 2.64 ± 0.01 | 2.75 ± 0.01 |
| | LP-update (Number of updates) | | 8.95 ± 0.31 (≈ 4.5) | 7.93 ± 0.23 (≈ 3.6) | 6.75 ± 0.16 (≈ 2.8) |

Table 1 Computation time (in seconds) of the LP-update policy and the occupation measure policy for different values of N , as well as the number of times that the LP-update policy solves a new linear program (the first initial one not included).

We notice that the variance of the computation time is much larger for the LP-update policy, since the number of time-steps when it solves a new LP for an update varies on each run. Overall this

variance is decreasing as N becomes large. We also remark that the LP-update policy takes less time for $N = 1000$ compared to $N = 20$. This is because for larger N the stochastic trajectory is closer to the deterministic one, consequently it is less likely to perform updates caused by a non-feasible action obtained from (18) (situation 2 described before Algorithm 2). Indeed, we observe from Table 1 that the number of updates is decreasing with N .

7. Conclusion, Discussions and Future Works

In this paper, we study finite horizon weakly-coupled Markov decision processes with symmetric components. We propose a re-solving heuristic that is asymptotically optimal under a population scaling of the model. Furthermore, we identify a non-degenerate condition on the model that can improve both the convergence rate and efficiency of the algorithm.

We conclude the main part of the paper by making the following remarks and pointing out several future research directions:

1. In our numerical experiments, we observe that applying re-solving in the LP-update policy improves the performance, as compared to the one-pass policies that do not. However, under other similar problem settings, it is not always better to re-solve. For instance, in the revenue management problem (e.g. [11, 21, 7]), there exists a counter-example in Cooper [11] where using re-solving gives strictly worse performance in expectation than not using it. A theoretical analysis on the effect of re-solving is provided in Secomandi [34] for this model, giving sufficient conditions so that re-solving is never detrimental. It is unclear whether some similar sufficient conditions exist in our model. Nevertheless, if the problem is non-degenerate, then as can be seen from the proof of Theorems 3 and 4, it is indifferent whether we re-solve or not in the large N limit. Yet, a complete understanding of the re-solving effect in our case is still open.
2. There exists various possibilities to extend the weakly-coupled MDPs model considered in this paper. These include:
 - We could have additional resource constraints that span multiple time-steps, instead of covering only a single time-step.
 - The population of the sub-MDPs may not be fixed by N , instead there are systems in which sub-component arrive and depart from the system. A typical example is the case of queueing networks where the sub-components are tasks [35].
 - The re-solving nature of the LP-update policy requires that the problem be formulated in finite horizon, or to discounted infinite horizon by truncating the time as in [6, 42]. One may wish to consider the same general weakly-coupled MDPs under infinite horizon. To the best of our knowledge, an asymptotically optimal policy under this generality is still an open question.
 - The reward may not be additive over the sub-components (but still additive across time), since it is common to have a utility function that is concave and non-linear. Linearity plays a crucial role in our methodology, especially in the non-degenerate condition and subsequently the faster convergence rate. It is worthwhile to investigate how far we can go by relaxing this linear condition.

References

- [1] Adelman D, Mersereau AJ (2008) Relaxations of weakly coupled stochastic dynamic programs. *Oper. Res.* 56(3):712–727, ISSN 0030-364X.
- [2] Astaraky D, Patrick J (2015) A simulation based approximate dynamic programming approach to multi-class, multi-resource surgical scheduling. *European Journal of Operational Research* 245(1):309–319, ISSN 0377-2217.
- [3] Bazaraa MS, Sherali HD, Shetty CM (2013) *Nonlinear programming: theory and algorithms* (John Wiley & Sons).
- [4] Bertsekas D (2012) *Dynamic programming and optimal control: Volume I*, volume 1 (Athena scientific).

-
- [5] Brown DB, Smith JE (2020) Index policies and performance bounds for dynamic selection problems. *Manag. Sci.* 66:3029–3050.
 - [6] Brown DB, Zhang J (2023) Fluid policies, reoptimization, and performance guarantees in dynamic resource allocation. *Operations Research* .
 - [7] Bumpensanti P, Wang H (2020) A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science* 66(7):2993–3009.
 - [8] Carpentier P, Chancelier JP, De Lara M, Pacaud F (2020) Mixed spatial and temporal decompositions for large-scale multistage stochastic optimization problems. *Journal of Optimization Theory and Applications* 186:985–1005.
 - [9] Carpentier P, Chancelier JP, Leclère V, Pacaud F (2018) Stochastic decomposition applied to large-scale hydro valleys management. *European Journal of Operational Research* 270(3):1086–1098.
 - [10] Carpentier P, Cohen G, et al. (2017) *Décomposition-coordination en optimisation déterministe et stochastique*, volume 81 (Springer).
 - [11] Cooper WL (2002) Asymptotic behavior of an allocation policy for revenue management. *Operations Research* 50(4):720–727.
 - [12] Dolgov DA, Durfee EH (2004) Optimal resource allocation and policy formulation in loosely-coupled Markov decision processes. *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 04)*, 315–324 (Whistler, BC).
 - [13] Gast N, Gaujal B, Yan C (2023) Exponential asymptotic optimality of whittle index policy. *Queueing Systems* 1–44.
 - [14] Gast N, Gaujal B, Yan C (2023) Linear program-based policies for restless bandits: Necessary and sufficient conditions for (exponentially fast) asymptotic optimality. *Mathematics of Operations Research* .
 - [15] Gocgun Y, Ghaté A (2010) A lagrangian approach to dynamic resource allocation. *Proceedings of the 2010 Winter Simulation Conference*, 3330–3340.
 - [16] Gocgun Y, Ghaté A (2012) Lagrangian relaxation and constraint generation for allocation and advanced scheduling. *Computers and Operations Research* 39(10):2323–2336, ISSN 0305-0548.
 - [17] Hawkins JT (2003) *A Lagrangian decomposition approach to weakly coupled dynamic optimization problems and its applications*. Ph.D. thesis, Massachusetts Institute of Technology.
 - [18] Hodge DJ, Glazebrook KD (2015) On the asymptotic optimality of greedy index heuristics for multi-action restless bandits. *Advances in Applied Probability* 47(3):652–667.
 - [19] Hoffman AJ, Kruskal JB (2010) Integral boundary points of convex polyhedra. Jünger M, Liebling TM, Naddef D, Nemhauser GL, Pulleyblank WR, Reinelt G, Rinaldi G, Wolsey LA, eds., *50 Years of Integer Programming 1958-2008 - From the Early Years to the State-of-the-Art*, 49–76 (Springer).
 - [20] Hong Y, Xie Q, Chen Y, Wang W (2023) Restless bandits with average reward: Breaking the uniform global attractor assumption. *Advances in Neural Information Processing Systems*.
 - [21] Jasin S, Kumar S (2012) A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research* 37(2):313–345.
 - [22] Maglaras C (2000) Discrete-review policies for scheduling stochastic networks: Trajectory tracking and fluid-scale asymptotic optimality. *The Annals of Applied Probability* 10(3):897–929.
 - [23] Matoušek J, Vondrák J (2001) The probabilistic method. *Lecture Notes, Department of Applied Mathematics, Charles University, Prague* .
 - [24] Meuleau N, Hauskrecht M, Kim KE, Peshkin L, Kaelbling LP, Dean TL, Boutilier C (1998) Solving very large weakly coupled markov decision processes. *AAAI/IAAI*, 165–172.
 - [25] Niño-Mora J (2022) Multi-gear bandits, partial conservation laws, and indexability. *Mathematics* 10(14), ISSN 2227-7390.
 - [26] Pacaud F, Carpentier P, Chancelier JP, De Lara M (2022) Optimization of a domestic microgrid equipped with solar panel and battery: Model predictive control and stochastic dual dynamic programming approaches. *Energy Systems* 1–25.

- [27] Pacaud F, De Lara M, Chancelier JP, Carpentier P (2021) Distributed multistage optimization of large-scale microgrids under stochasticity. *IEEE Transactions on Power Systems* 37(1):204–211.
- [28] Palomar DP, Chiang M (2006) A tutorial on decomposition methods for network utility maximization. *IEEE Journal on Selected Areas in Communications* 24(8):1439–1451.
- [29] Papadimitriou CH, Tsitsiklis JN (1999) The complexity of optimal queuing network control. *Math. Oper. Res* 293–305.
- [30] Patrick J, Puterman ML, Queyranne M (2008) Dynamic multipriority patient scheduling for a diagnostic resource. *Operations research* 56(6):1507–1525.
- [31] Puterman ML (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (New York, NY, USA: John Wiley & Sons, Inc.), 1st edition.
- [32] Rakovic SV, Levine WS (2018) *Handbook of model predictive control* (Springer).
- [33] Salemi Parizi M (2018) *Approximate dynamic programming for weakly coupled Markov decision processes with perfect and imperfect information*. Ph.D. thesis.
- [34] Secomandi N (2008) An analysis of the control-algorithm re-solving issue in inventory and revenue management. *Manufacturing & Service Operations Management* 10(3):468–483.
- [35] Verloop M (2016) Asymptotically optimal priority policies for indexable and nonindexable restless bandits. *Annals of Applied Probability* 26(4):1947–1995.
- [36] Weber RR, Weiss G (1990) On an index policy for restless bandits. *Journal of Applied Probability* 27(3):637–648, ISSN 00219002.
- [37] Whittle P (1988) Restless bandits: activity allocation in a changing world. *Journal of Applied Probability* 25A:287–298.
- [38] Xiong G, Li J, Singh R (2021) Reinforcement learning for finite-horizon restless multi-armed multi-action bandits. *arXiv preprint arXiv:2109.09855* .
- [39] Zayas-Cabán G, Jasin S, Wang G (2017) An asymptotically optimal heuristic for general non-stationary finite-horizon restless multi-armed multi-action bandits. *Ross: Technology & Operations (Topic)* .
- [40] Zhang X (2022) *Near-Optimality for Multi-action Multi-resource Restless Bandits with Many Arms*. Ph.D. thesis.
- [41] Zhang X, Frazier PI (2021) Restless bandits with many arms: Beating the central limit theorem. *arXiv preprint arXiv:2107.11911* .
- [42] Zhang X, Frazier PI (2022) Near-optimality for infinite-horizon restless bandits with many arms. *arXiv preprint arXiv:2203.15853* .

Appendix A: PROOF OF THE TECHNICAL RESULTS

This section is dedicated to the proofs of some technical lemmas that are used to prove Theorem 1, Theorem 3 and Theorem 4. We also provide counter-examples for the lower bounds of the three theorems, as well as an example to illustrate the equivalence of the two notions of non-degeneracy.

A.1. One-step transition and concentration arguments

Recall that the linear function ϕ maps a decision vector \mathbf{y} to a configuration vector $\phi(\mathbf{y}) = (\phi_1(\mathbf{y}), \dots, \phi_d(\mathbf{y})) \in \Delta^d$ whose s -th component is

$$\phi_s(\mathbf{y}) = \sum_{s', a} y_{s', a} P_{s', s}^a. \quad (20)$$

This is the deterministic behavior of the Markov transition at time-step t . We claim that:

LEMMA 1. Define the random vector $\mathbf{E}^{(N)}(t) := \mathbf{X}^{(N)}(t+1) - \phi(\mathbf{Y}^{(N)}(t))$. We have

$$\mathbb{E} [\mathbf{E}^{(N)}(t) \mid \mathbf{Y}^{(N)}(t)] = \mathbf{0}, \quad (21)$$

$$\mathbb{E} [\|\mathbf{E}^{(N)}(t)\|_1 \mid \mathbf{Y}^{(N)}(t)] \leq \frac{\sqrt{d}}{\sqrt{N}}, \quad (22)$$

$$\mathbb{P} (\|\mathbf{E}^{(N)}(t)\|_1 \geq \epsilon \mid \mathbf{Y}^{(N)}(t)) \leq 2d \cdot e^{-2N\epsilon^2/d^2}. \quad (23)$$

Proof of Lemma 1 For simplicity of notation, let us denote by $\mathbf{y} := \mathbf{Y}^{(N)}(t)$. There are $Ny_{s,a}$ arms in state s and whose action is a , and each of these arms makes a transition to state s' with probability $P_{s,s'}^a$. This shows that $X^{(N)}(t+1)$ can be written as a sum of independent random variables as follows:

$$X_{s'}^{(N)}(t+1) = \frac{1}{N} \sum_{s,a} \sum_{i=1}^{Ny_{s,a}} \mathbf{1}_{\{\sum_{s''=1}^{s'-1} P_{s,s''}^a \leq U_{s,a,i} < \sum_{s''=1}^{s'} P_{s,s''}^a\}},$$

where $U_{s,a,i}$ are i.i.d uniform random variables in $[0, 1]$. Taking expectation then gives $\mathbb{E} \left[X_{s'}^{(N)}(t+1) \mid \mathbf{Y}^{(N)}(t) \right] = \phi_{s'}(\mathbf{Y}^{(N)}(t))$, which gives (21). It also implies that

$$\begin{aligned} \mathbb{E} \left[|E_{s'}^{(N)}(t+1)|^2 \mid \mathbf{Y}^{(N)}(t) = \mathbf{y} \right] &= \text{var} \left[X_{s'}^{(N)}(t+1) \mid \mathbf{Y}^{(N)}(t) = \mathbf{y} \right] \\ &= \frac{1}{N^2} \sum_{s,a} Ny_{s,a} P_{s,s'}^a (1 - P_{s,s'}^a) \\ &\leq \frac{\sum_{s,a} y_{s,a} P_{s,s'}^a}{N}. \end{aligned}$$

This shows that

$$\mathbb{E} \left[\|\mathbf{E}^{(N)}(t+1)\|_1 \mid \mathbf{Y}^{(N)}(t) = \mathbf{y} \right] \leq \sqrt{d} \frac{\sqrt{\sum_{s'} \sum_{s,a} y_{s,a} P_{s,s'}^a}}{\sqrt{N}} = \frac{\sqrt{d}}{\sqrt{N}},$$

where the first inequality comes from Cauchy-Schwartz, and this gives (22).

Equation (23) is a direct consequence of Hoeffding's inequality. Indeed, one has

$$\mathbb{P} \left(|E_s^{(N)}(t)| \geq \varepsilon/d \mid \mathbf{Y}^{(N)}(t) \right) \leq 2e^{-N\varepsilon^2/d^2}.$$

By using the union bound, this implies that

$$\mathbb{P} \left(\|\mathbf{E}^{(N)}(t)\|_1 \geq \varepsilon \mid \mathbf{Y}^{(N)}(t) \right) \leq d \cdot \mathbb{P} \left(|E_s^{(N)}(t)| \geq \varepsilon/d \mid \mathbf{Y}^{(N)}(t) \right) \leq 2d \cdot e^{-N\varepsilon^2/d^2}.$$

Q.E.D.

A.2. Lipschitz-constant of $V_{\text{rel}}(\cdot, t)$

Recall that γ is equal to

$$\gamma = \min_{i \neq j, a} \sum_k \min(P_{ik}^0, P_{jk}^a). \quad (24)$$

LEMMA 2. Recall that $r_{\max} = \max_{s,a} |r_s^a|$. Then, for all $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta$, the value function satisfies:

$$|V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t)| \leq (1 + (1 - \gamma) + \dots + (1 - \gamma)^{t-1}) r_{\max} \|\mathbf{x} - \tilde{\mathbf{x}}\|_1.$$

In particular, this result implies that:

- The Lipschitz constant of $\mathbf{x} \mapsto V_{\text{rel}}(\mathbf{x}, t)$ is upper bounded by tr_{\max} ;
- When $\gamma > 0$, the Lipschitz constant of $\mathbf{x} \mapsto V_{\text{rel}}(\mathbf{x}, t)$ is bounded by r_{\max}/γ (for all t).

Proof of Lemma 2 The main idea is to work with vectors of N components $\mathbf{s} \in \mathcal{S}$ instead of the population state \mathbf{x} . The use of vectors of components will allow the usage of a sort of coupling between trajectories. To simplify our notation, we will need a way to go from vectors of N components to population vectors. For a population vector $\mathbf{s} \in \mathcal{S}^N$, we denote by $\chi(\mathbf{s})$ the vector such that for $i \in \mathcal{S}$:

$$\chi_i(\mathbf{s}) = \frac{1}{N} \sum_{n=1}^N \mathbf{1}_{\{s_n=i\}}.$$

Recall that Δ is the simplex of dimension $d = |\mathcal{S}|$ and that $\Delta^{(N)}$ is the set $\mathbf{x} \in \Delta$, such that Nm_s is an integer for all $s \in \{1 \dots d\}$. For all $\mathbf{s} \in \mathcal{S}^N$, $\chi(\mathbf{s}) \in \Delta^{(N)}$. Respectively, for all $\mathbf{x} \in \Delta^{(N)}$, there exists $\mathbf{s} \in \mathcal{S}^N$ such that $\chi(\mathbf{s}) = \mathbf{x}$ (unique up to a permutation of its components).

Our proof of the result uses the following lemma, which shows that the property holds for some particular values of \mathbf{x} and $\tilde{\mathbf{x}}$:

LEMMA 3. The result of Lemma 2 holds for all $\mathbf{x}, \tilde{\mathbf{x}}$ for which there exists a $N > 0$ such that $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta^{(N)}$ and $\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 \leq 2/N$.

Lemma 3 implies that the result of Lemma 2 holds for all N and for all $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta^{(N)}$. Indeed, take any $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta^{(N)}$, there exists an integer $k \in \mathbb{N}$ such that $\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 = 2k/N$. We can then construct a sequence of $k+1$ state vectors $\mathbf{x} = \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k, \mathbf{x}_{k+1} = \tilde{\mathbf{x}}$, such that $\|\mathbf{x}_i - \mathbf{x}_{i+1}\|_1 = 2/N$ for all $1 \leq i \leq k$. Hence

$$\begin{aligned} |V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t)| &\leq \sum_{i=1}^k |V_{\text{rel}}(\mathbf{x}_i, t) - V_{\text{rel}}(\mathbf{x}_{i+1}, t)| \\ &\leq (1 + (1 - \gamma) + \dots + (1 - \gamma)^{t-1}) r_{\max} \sum_{i=1}^k \|\mathbf{x}_i - \mathbf{x}_{i+1}\|_1 \\ &= (1 + (1 - \gamma) + \dots + (1 - \gamma)^{t-1}) r_{\max} \|\mathbf{x} - \tilde{\mathbf{x}}\|_1 \end{aligned}$$

By continuity of the value function and since $\cup_{N>0} \Delta^{(N)}$ is dense in Δ , we can conclude that the result holds for all $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta$. Q.E.D.

Let us now return to the proof the Lemma 3.

Proof of Lemma 3 We prove the result by induction on t . For $t = 0$, the result is clearly true as $V_t(\mathbf{x}) = 0$ for all $\mathbf{x} \in \Delta$.

We now assume that it holds for $t-1 \geq 0$ and look at what happens for t . Let $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta^{(N)}$ be two population vectors such that $\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 \leq 2/N$ and $\mathbf{x} \neq \tilde{\mathbf{x}}$. This implies that there exist two states $i \neq j \in \{1 \dots d\}$ such that $x_i = \tilde{x}_i + 1/N$ and $x_j = \tilde{x}_j - 1/N$ (all others being equal), so that actually $\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 = 2/N$. We denote by $\mathbf{s}, \tilde{\mathbf{s}} \in \mathcal{S}^N$ two vectors such that $s_1 = i, s'_1 = j, s_n = s'_n$ for $n \geq 2$, and $\chi(\mathbf{s}) = \mathbf{x}, \chi(\tilde{\mathbf{s}}) = \tilde{\mathbf{x}}$.

For a given vector of states $\mathbf{s} \in \mathcal{S}^N$ and a corresponding vector of actions $\mathbf{a} = a_1 \dots a_N$, with each a_n representing a probability measure on \mathcal{A} as part of the randomized policy for arm n in the relaxed problem, we define $Q(\mathbf{s}, \mathbf{a})$ as the Q -value obtained by executing the action vector \mathbf{a} for the initial time step and subsequently adhering to the optimal policy:

$$Q(\mathbf{s}, \mathbf{a}) := \frac{1}{N} \sum_{n=1}^N r_{s_n}^{a_n} + \sum_{\mathbf{s}' \in \mathcal{S}^N} V_{\text{rel}}(\chi(\mathbf{s}'), t-1) \prod_{n=1}^N P_{s_n, s'_n}^{a_n}. \quad (25)$$

In the formulation above, $r_{s_n}^{a_n}$ and $P_{s_n, s'_n}^{a_n}$ should be interpreted as the *expected* rewards and transition probabilities under the randomized action a_n across \mathcal{A} , that is $r_{s_n}^{a_n} = \sum_a r_{s_n}^a \mathbb{P}(a_n = a)$ and $P_{s_n, s'_n}^{a_n} = \sum_a P_{s_n, s'_n}^a \mathbb{P}(a_n = a)$. Consider \mathbf{a}^* as the optimal vector of actions for the state \mathbf{s} , such that $V_{\text{rel}}(\mathbf{x}) = Q(\mathbf{s}, \mathbf{a}^*)$. Let \mathbf{a}' denote a vector of actions where $a'_1 = 0$ with probability one, and for all $n \in \{2 \dots N\}$, $a'_n = a_n^*$. The action vector \mathbf{a}' represents an *admissible* set of actions for $\tilde{\mathbf{s}}$, given that action 0 does not utilize any resources. However, it may not be optimal. Thus, it follows that $V_{\text{rel}}(\tilde{\mathbf{x}}) \geq Q(\tilde{\mathbf{s}}, \mathbf{a}')$, leading to

$$\begin{aligned} V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t) &\leq Q(\mathbf{s}, \mathbf{a}^*) - Q(\tilde{\mathbf{s}}, \mathbf{a}') \\ &= \frac{1}{N} (r_{s_1}^{a_1^*} - r_{\tilde{s}_1}^0) + \sum_{\mathbf{s}' \in \mathcal{S}^N} V_{\text{rel}}(\chi(\mathbf{s}'), t-1) \left(\prod_{n=1}^N P_{s_n, s'_n}^{a_n^*} - \prod_{n=1}^N P_{\tilde{s}_n, s'_n}^{a'_n} \right) \\ &= \frac{1}{N} (r_{s_1}^{a_1^*} - r_{\tilde{s}_1}^0) + \sum_{\mathbf{s}' \in \mathcal{S}^N} \prod_{n=2}^N P_{s_n, s'_n}^{a_n^*} V_{\text{rel}}(\chi(\mathbf{s}'), t-1) (P_{s_1, s'_1}^{a_1^*} - P_{\tilde{s}_1, s'_1}^0) \\ &= \frac{1}{N} (r_{s_1}^{a_1^*} - r_{\tilde{s}_1}^0) + \sum_{s'_2, \dots, s'_N \in \mathcal{S}^{N-1}} \prod_{n=2}^N P_{s_n, s'_n}^{a_n^*} \sum_{k \in \mathcal{S}} V_{\text{rel}}(\chi(k, s'_2 \dots s'_N), t-1) (P_{s_1, k}^{a_1^*} - P_{\tilde{s}_1, k}^0), \quad (26) \end{aligned}$$

where we have used Equation (25) for the first equality and the fact that \mathbf{s} and $\tilde{\mathbf{s}}$, and \mathbf{a} and \mathbf{a}' are equal for all components except the first one. The last equality is a rewriting of the above equation where the special role of the first component has been emphasized.

The first term of Equation (26) is smaller than $2r_{\max}/N$. For the second term, denote by $\gamma_{s_1 \tilde{s}_1, k} := \min(P_{s_1, k}^{a_1^*}, P_{\tilde{s}_1, k}^0)$. For any $s'_2, \dots, s'_N \in \mathcal{S}^{N-1}$, we have

$$\begin{aligned} &\sum_{k \in \mathcal{S}} V_{\text{rel}}(\chi(k, s'_2 \dots s'_N), t-1) (P_{s_1, k}^{a_1^*} - P_{\tilde{s}_1, k}^0) \\ &= \sum_{k \in \mathcal{S}} V_{\text{rel}}(\chi(k, s'_2 \dots s'_N), t-1) (P_{s_1, k}^{a_1^*} - \gamma_{s_1 \tilde{s}_1, k}) - \sum_{k \in \mathcal{S}} V_{\text{rel}}(\chi(k, s'_2 \dots s'_N), t-1) (P_{\tilde{s}_1, k}^0 - \gamma_{s_1 \tilde{s}_1, k}) \end{aligned}$$

$$\begin{aligned}
&\leq \left(1 - \sum_{k \in \mathcal{S}} \gamma_{s_1 \bar{s}_1, k}\right) \left(\max_{\bar{k}} V_{\text{rel}}(\chi(\bar{k}, s'_2 \dots s'_N), t-1) - \min_{\underline{k}} V_{\text{rel}}(\chi(\underline{k}, s'_2 \dots s'_N), t-1) \right) \\
&\leq (1-\gamma) \left(\max_{\bar{k}} V_{\text{rel}}(\chi(\bar{k}, s'_2 \dots s'_N), t-1) - \min_{\underline{k}} V_{\text{rel}}(\chi(\underline{k}, s'_2 \dots s'_N), t-1) \right)
\end{aligned} \tag{27}$$

where the last step follows from our definition of γ .

Since $\chi(\bar{k}, s'_2 \dots s'_N)$ and $\chi(\underline{k}, s'_2 \dots s'_N)$ are at \mathcal{L}^1 -distance $2/N$ when $\bar{k} \neq \underline{k}$, and 0 otherwise, we can apply our induction hypothesis to Equation (27) and deduce from Equation (26) that

$$V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t) \leq \frac{2r_{\max}}{N} (1 + (1-\gamma)(1 + (1-\gamma) + \dots + (1-\gamma)^{t-2}))$$

By exchanging the roles of \mathbf{x} and $\tilde{\mathbf{x}}$ and using that $\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 = 2/N$, we finally obtain

$$\begin{aligned}
|V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t)| &\leq \frac{2r_{\max}}{N} (1 + (1-\gamma)(1 + (1-\gamma) + \dots + (1-\gamma)^{t-2})) \\
&= (1 + (1-\gamma) + \dots + (1-\gamma)^{t-1}) r_{\max} \|\mathbf{x} - \tilde{\mathbf{x}}\|_1,
\end{aligned}$$

which completes the inductions step. Q.E.D.

We remark that to obtain a Lipschitz constant of $V_{\text{rel}}(\cdot, t)$ that does not depend of t , Lemma 2 has imposed the condition that $\gamma > 0$. The constant γ corresponds to the minimum probability that two arms starting in different states couple at the next time-step, when the action 0 is applied to the first arm and any action is applied to the second arm. This condition can be weakened substantially. In particular, a less demanding sufficient condition is that there exists an integer $\tau \geq 1$ such that $\lambda^\tau > 0$, where

$$\lambda^\tau := \min_{i \neq j, a_1 \dots a_\tau} \sum_k \min(\underbrace{(P^0 \dots P^0)}_{\tau \text{ times}})_{ik}, (P^{a_1} \dots P^{a_\tau})_{jk}.$$

The quantity λ^τ is the minimum probability that two arms starting in two different states couple after τ time-steps, where one arm applies the constant action 0, while the other arm chooses arbitrary actions. Similar to Lemma 2, we can prove the following result:

LEMMA 4. For all $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta$, and all τ such that $\lambda^\tau > 0$, the value function satisfies:

$$|V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t)| \leq \frac{\tau r_{\max}}{\lambda^\tau} \|\mathbf{x} - \tilde{\mathbf{x}}\|_1.$$

Sketch of proof. The proof of this results follows the same line as the proof of Lemma 2 and we only highlight the main steps:

- As in the proof Lemma 2, the first step is to restrict our attention to points that are at distance $2/N$ in $\Delta^{(N)}$, which by using the representation \mathbf{s} implies that only one of the components differs between \mathbf{s} and $\tilde{\mathbf{s}}$.
- The main step of the proof of Lemma 2 is to write $V_{\text{rel}}(\mathbf{x}, t)$ as a function of $V_{\text{rel}}(\mathbf{x}, t-1)$ as in Equation (25), and uses Equations (26) and (27) to show that if $\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 \leq 2/N$, then

$$|V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t)| \leq \frac{2r_{\max}}{N} + (1-\gamma) \max_{\mathbf{x}', \tilde{\mathbf{x}}' \text{ s.t. } \|\mathbf{x}' - \tilde{\mathbf{x}}'\|_1 \leq 2/N} |V_{\text{rel}}(\mathbf{x}', t-1) - V_{\text{rel}}(\tilde{\mathbf{x}}', t-1)|.$$

To prove Lemma 4, we can use a similar reasoning to show that if $\|\mathbf{x} - \tilde{\mathbf{x}}\|_1 \leq 2/N$, then

$$|V_{\text{rel}}(\mathbf{x}, t) - V_{\text{rel}}(\tilde{\mathbf{x}}, t)| \leq \frac{2\tau r_{\max}}{N} + (1-\lambda^\tau) \max_{\mathbf{x}', \tilde{\mathbf{x}}' \text{ s.t. } \|\mathbf{x}' - \tilde{\mathbf{x}}'\|_1 \leq 2/N} |V_{\text{rel}}(\mathbf{x}', t-\tau) - V_{\text{rel}}(\tilde{\mathbf{x}}', t-\tau)|.$$

The result then follows from the same reasoning. Q.E.D.

A.3. Non-degenerate problem and concentration on a trajectory

The previous lemma can be extended to show that the stochastic system $\mathbf{X}^{(N)}(t)$ concentrates on a neighbourhood of $\mathbf{x}^*(t)$.

LEMMA 5. Assume that the problem is non-degenerate, let \mathbf{y}^* be the optimal solution to the LP computed at time-step 0 and $\mathbf{X}^{(N)}(t)$ be the sequence of configuration vectors obtained when applying Algorithm 2. Then for all $\varepsilon > 0$, there exists $C_1, C_2 > 0$ such that for all N :

$$\mathbb{P}(\mathbf{X}^{(N)}(t) \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon)) \geq 1 - C_1 \cdot e^{-C_2 N}. \tag{28}$$

Moreover, if the LP has a unique solution starting from any initial point, then (28) also holds if $\mathbf{X}^{(N)}(t)$ is the output of Algorithm 1.

Proof. We first consider what happens when Algorithm 2 is used. We proceed by induction on t . This is clearly true for $t=0$. Assume that this holds for some $t \geq 0$. By Proposition 1, there exists ε_t such that the control $\mathbf{y}(\mathbf{x})$ defined in Proposition 1 is optimal for all $\mathbf{x} \in \mathcal{B}(\mathbf{x}^*(t), \varepsilon_t)$. The induction hypothesis and the continuity of $\mathbf{y}(\mathbf{x})$ therefore imply that for all $\epsilon > 0$, there exists $C_1, C_2 > 0$ such that $\|\mathbf{Y}^{(N)}(t) - \mathbf{y}^*(t)\|_1 \leq \epsilon$ with probability at least $1 - C_1 \cdot e^{-C_2 N}$, where $\mathbf{Y}^{(N)}(t)$ is the control used for $\mathbf{X}^{(N)}(t)$. Write

$$\|\mathbf{X}^{(N)}(t+1) - \mathbf{x}^*(t+1)\|_1 \leq \|\mathbf{X}^{(N)}(t+1) - \phi(\mathbf{Y}^{(N)}(t))\|_1 + \|\phi(\mathbf{Y}^{(N)}(t)) - \mathbf{x}^*(t+1)\|_1.$$

Hence, by using Lemma 1 and the union bound, for all $\epsilon > 0$, there exists $C'_1, C'_2 > 0$ such that $\|\mathbf{X}^{(N)}(t+1) - \mathbf{x}^*(t+1)\|_1 \leq \epsilon$ with probability at least $1 - C'_1 \cdot e^{C'_2 N}$.

To show (28), the only remaining point is to see that if s' is a state such that $x_{s'}^*(t+1) = 0$, then so is $X_{s'}^{(N)}(t+1)$. By definition of the deterministic evolution $\mathbf{x}^*(t+1) = \phi(\mathbf{y}^*(t))$ (see e.g. Equation (20)), we have:

$$x_{s'}^*(t+1) = \sum_{a,s} y_{s,a}^*(t) P_{s,s'}^a.$$

Hence, $x_{s'}^*(t+1)$ equals 0 if for all s, a , either $y_{s,a}^*(t) = 0$ or $P_{s,s'}^a = 0$. By construction of $\mathbf{Y}^{(N)}(t)$ from $\mathbf{y}^*(t)$, $y_{s,a}^*(t) = 0$ implies that $Y_{s,a}^{(N)}(t) = 0$, consequently $X_{s'}^{(N)}(t+1) = 0$.

To study what happens when Algorithm 1 is used, we remark that if the solution to the LP is unique, and $\mathbf{X}^{(N)}(t)$ is close enough to $\mathbf{x}^*(t)$, then the new LP solution computed by Algorithm 1 is the same $\mathbf{y}(\mathbf{X}^{(N)}(t))$ used by Algorithm 2. Hence, the proof for Algorithm 2 also applies in this case. Q.E.D.

A.4. Proof of the lower bounds

Consider the following two-action restless bandit with two states $\mathcal{S} = \{1, 2\}$, with parameters $\mathbf{x}(0) = (0.5, 0.5)$, $T = 2$, $A = \{0, 1\}$, $\mathbf{r}^0 = [0, 0]$, $\mathbf{r}^1 = [1, 0]$, and where the transition matrices are

$$\mathbf{P}^0 = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}, \mathbf{P}^1 = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}.$$

We consider that $D(s, 0) = 0$, $D(s, 1) = 1$ for any state s , and we distinguish the resource constraints $b = 0.3$ and $b = 0.5$.

For any resource constraint, the solution to the LP is to choose action 1 for as many arms as possible. This gives a reward $2b$. If $b = 0.3$, the problem is non-degenerate whereas if $b = 0.5$ the problem is degenerate.

For the stochastic system with N components, this gives a reward $\frac{1}{N} \lfloor Nb \rfloor$ at time-step 0 and a reward $\min(b, \frac{1}{N} \lfloor NX_1^{(N)}(1) \rfloor)$ at time-step 1. Since $X_1^{(N)}(1)$ follows a binomial distribution of parameter $(N, 0.5)$, the total reward of the LP-update policy is equal to

$$\frac{1}{N} \lfloor Nb \rfloor + b + \mathbb{E} \left[\min\left(\frac{1}{N} \lfloor NX_1^{(N)}(1) \rfloor - b, 0\right) \right].$$

By the central limit theorem, $\lim_{N \rightarrow \infty} \sqrt{N} \mathbb{E} \left[\min(X_1^{(N)}(1) - 0.5, 0) \right] = \sqrt{2}\pi > 0$. This provides a counter-example for the lower bound of Theorem 1.

If $b = 0.3$ and N is not a multiple of 10, the problem does not admit a perfect rounding, and we have $2b - \left(\frac{1}{N} \lfloor Nb \rfloor + b + \mathbb{E} \left[\min\left(\frac{1}{N} \lfloor NX_1^{(N)}(1) \rfloor - b, 0\right) \right]\right) \geq b - \frac{1}{N} \lfloor Nb \rfloor \geq 0.1/N$. This provides a counter-example for the lower-bound of Theorem 3.

If $b = 0.3$ and N is a multiple of 10, then the problem admits a perfect rounding. In this case, classical anti-concentration arguments (Matoušek and Vondrák [23]) show that

$$\mathbb{P} \left(X_1^{(N)}(1) \leq 0.2 \right) \geq \frac{1}{15} \exp(-16N(0.6 - 0.2)^2).$$

This shows that

$$\mathbb{E} \left[\min\left(\frac{1}{N} \lfloor NX_1^{(N)}(1) \rfloor - 0.3, 0\right) \right] \geq -0.1 \mathbb{P} \left(X_1^{(N)}(1) \leq 0.2 \right) = -\frac{1}{150} \exp(-2.56N).$$

It provides a counter-example for the lower-bound of Theorem 4. Q.E.D.

Algorithm 3: Occupation measure policy for weakly-coupled MDPs.

Input: Time horizon T and initial configuration vector $\mathbf{x}(0)$.

Solve the linear program (5) with time horizon T and initial configuration vector $\mathbf{x}(0)$, obtain an optimal solution \mathbf{y}^* and the corresponding \mathbf{x}^* ;

for $t = 0, 1, 2, \dots, T - 1$ **do**

 Compute from the LP solution the occupation measure

$$\mu_{s,a}^*(t) := \begin{cases} \frac{y_{s,a}^*(t)}{x_s^*(t)}, & \text{if } x_s^*(t) > 0 \\ \mathbf{1}_{\{a=0\}}, & \text{otherwise.} \end{cases} \quad (29)$$

 Observe the current states of the N sub-MDP's $\mathbf{s} = (s_1, s_2, \dots, s_N)$. Initialize $\mathbf{B} := N \cdot \mathbf{b}$, and actions on the N sub-MDP's as $A(\mathbf{s}) = (0, 0, \dots, 0)$;

for $n = 1, 2, \dots, N$ **do**

 Sample an action a_n according to the probability vector $(\mu_{s_n,a}^*(t))_{a \in A(s_n)}$;

if $\mathbf{B} - D(s_n, a_n) \geq \mathbf{0}$ (*component-wise*) **then**

$A(s_n) := a_n$;

$\mathbf{B} := \mathbf{B} - D(s_n, a_n)$;

end

end

 Apply the actions $A(\mathbf{s})$ to the N sub-MDP's ;

end

Appendix B: THE OCCUPATION MEASURE POLICY

To provide a benchmark for evaluation of the performance of the LP-update policy, we describe in this section a policy that we call the *occupation measure policy*. This policy is a straightforward generalization of the randomized activation control policy in Zayas-Cabán et al. [39] and the occupancy-measured-reward index policy in Xiong et al. [38] to the multi-constraint multi-action case. The occupation measure policy is proven to be asymptotically optimal for the multi-action single-constraint multi-armed bandits in Xiong et al. [38]. Adapting their proof to the multi-constraint case is direct.

The occupation measure policy is a one-pass policy that solves the linear program only once at the very beginning, and constructs the probability vectors $\mu_{s,a}^*(t)$ defined in (29) from the solution. At each decision epoch t , the budget left is initialized as $\mathbf{B} := N \cdot \mathbf{b}$, the action on each arm n is initialized as the passive 0. It then samples a new action a_n from the distribution $(\mu_{s_n,a}^*(t))_{a \in A(s_n)}$ on each arm n . If choosing action a_n instead of action 0 on arm n does not violate any of the budget constraints, then we apply action a_n on arm n , and we decrease the budget left \mathbf{B} ; otherwise we keep action 0 on arm n and continue to sample an action on the next arm. The detailed implementation is given in Algorithm 3.

Appendix C: COMPARISON WITH EXISTING NOTIONS OF NON-DEGENERACY

The non-degenerate property as defined in Definition 1 uses a rank condition from the optimal LP solution \mathbf{y}^* . In this section, we provide a detailed comparison of our notion to the existing notions presented [41, 14, 6, 40]. To facilitate the comparison, we reformulate our definition.

Let \mathbf{y}^* be an optimal solution to the LP (at time 0) and let us fix a time t . For each state $s \in \mathcal{S}$, we define by $A_s^+ := |\{a \in \mathcal{A} \mid y_{s,a}^*(t) > 0\}|$ the number of different actions that are chosen in state s at time t . Similar to Brown and Zhang [6], we define:

$$\begin{aligned} \mathcal{S}^+ &:= \{s \in \mathcal{S} \mid A_s^+ = 1\}; \\ \mathcal{S}^0 &:= \{s \in \mathcal{S} \mid A_s^+ > 1\}; \\ \mathcal{S}^- &:= \{s \in \mathcal{S} \mid A_s^+ = 0\}. \end{aligned}$$

We see that \mathcal{S}^+ , \mathcal{S}^0 , \mathcal{S}^- form a partition of \mathcal{S} . Their interpretations are that (for the LP), a state in \mathcal{S}^+ has a unique optimal action, a state in \mathcal{S}^0 has at least two equally optimal actions, and a state in \mathcal{S}^- is not reached at time t .

As $\mathcal{U}^*(t)$ is the set of indices (s, a) for which $y_{s,a}^*(t) = 0$, we have:

$$\begin{aligned} |\mathcal{U}^*| &= \sum_{s \in \mathcal{S}} |\mathcal{A}| - A_s^+ \\ &= \sum_{s \in \mathcal{S}^-} A_s + \sum_{s \in \mathcal{S}^+} A_s - |\mathcal{S}^+| + \sum_{s \in \mathcal{S}^0} |\mathcal{A}| - \sum_{s \in \mathcal{S}^0} A_s^+ \\ &= |\mathcal{U}| - |\mathcal{S}^+| - \sum_{s \in \mathcal{S}^0} A_s^+, \end{aligned}$$

while $|\mathcal{S}^*| = |\mathcal{S}^0| + |\mathcal{S}^+|$.

As $|\mathcal{J}^*|$ is the number of saturated constraints, we have

$$|\mathcal{J}^*| + |\mathcal{S}^*| + |\mathcal{U}^*| = |\mathcal{U}| + |\mathcal{J}^*| + |\mathcal{S}^0| - \sum_{s \in \mathcal{S}^0} A_s^+.$$

The matrix C^* has $|\mathcal{J}^*| + |\mathcal{S}^*| + |\mathcal{U}^*|$ rows and $|\mathcal{U}|$ columns. To simplify the discussion, in this section, we assume that:

$$\text{All rows of } C^*(t) \text{ are linearly independent.} \quad (30)$$

So that the rank of $C^*(t)$ boils down to verifying an inequality among dimension. Indeed, under this assumption (30), in order for C^* to have rank $|\mathcal{J}^*| + |\mathcal{S}^*| + |\mathcal{U}^*|$ as in Definition 1, it is necessary and sufficient that

$$|\mathcal{J}^*| + |\mathcal{S}^0| \leq \sum_{s \in \mathcal{S}^0} A_s^+. \quad (31)$$

We are now ready to compare the formulation (31) with various existing notions of non-degeneracy in the literature:

1. In Zhang and Frazier [41] and Gast et al. [14] that study two-action bandits with a single *equality* resource constraint $b = \alpha$. Their notion of non-degeneracy using our notation is to say that $|\mathcal{S}^0| \geq 1$. Since in their model the assumption (30) always holds, and $|\mathcal{J}^*|$ is always 1. In order for (31) to hold, it is necessary and sufficient that $|\mathcal{S}^0| \geq 1$, which shows that the two definitions of non-degeneracy are indeed equivalent.
2. In Brown and Zhang [6], the authors define another notion of non-degeneracy in Definition 5.2 for multi-action single-constraint bandits, which by using the notations in this section can be stated as follows (assuming (30)): the problem is non-degenerate if at least one of the conditions hold: (i) the single resource constraint is not saturated; (ii) $|\mathcal{S}^0| \geq 1$. This is again equivalent to (31). Note that additional care is taken in Definition 5.2 of Brown and Zhang [6], in order to account for the non-satisfaction of (30).
3. In Zhang [40], another definition of non-degeneracy for multi-action multi-constraint bandits is given in Section 3.3. According to this definition, a multi-action multi-constraint bandit is non-degenerate if $|\mathcal{J}^*| \leq 1$ and $|\mathcal{S}^0| \geq 1$. Clearly this implies (31), but under this general model, there are other possible problems that satisfy (31) while not covered by the non-degenerate definition of Zhang [40].

Appendix D: IMPROVING THE ROUNDING METHOD

When the LP is solved with a given initial condition $\mathbf{X}^{(N)}$, it yields a first-step control $\mathbf{Y} := \mathbf{y}(0) \in \mathcal{Y}(\mathbf{X}^{(N)})$. However, this solution may not be directly transferable to the original problem because the quantities $N \cdot Y_{s,a}$ are not necessarily integers. Throughout our paper, we address this issue by employing a simple truncation method to obtain a vector $\mathbf{Y}^{(N)}$ that is defined as follows:

$$\mathbf{Y}_{s,a}^{(N)} = \begin{cases} N^{-1} \lfloor NY_{s,a} \rfloor & \text{if } a \neq 0 \\ X_s^{(N)} - \sum_{a \neq 0} \tilde{Y}_{s,a} & \text{if } a = 0. \end{cases} \quad (32)$$

This approach results in an admissible solution $\mathbf{Y}^{(N)} \in \mathcal{Y}^{(N)}(\mathbf{X}^{(N)})$, as all elements in $D(s, a)$ and \mathbf{b} are non-negative, and notably, $D(s, 0) = 0$. However, this truncation solution introduces an $\mathcal{O}(1/N)$ sub-optimality gap.

In this section, we delve into potential strategies to enhance performance and reduce this sub-optimality gap. By examining alternative methods to handle the non-integer nature of the solution, we aim to improve the applicability and efficiency of our approach, ensuring closer alignment with the optimal solution of the original problem.

D.1. Improvement by using MILP

As previously mentioned, the rounding process we employ in (32) addresses the issue that the first step $\mathbf{y}(0)$ from the LP might not yield integer values for $N \cdot y_{s,a}(0)$. To circumvent this, another more refined approach is to explicitly incorporate integer constraints into our LP, leading to the following mixed-integer optimization problem (MILP):

$$\max_{\mathbf{y} \geq 0} \sum_{t=0}^{T-1} \sum_{(s,a) \in \mathcal{U}} r_s^a y_{s,a}(t) \quad (33a)$$

$$\text{s.t.} \quad \sum_{a \in \mathcal{A}} y_{s,a}(0) = x_s \quad \forall s, \quad (33b)$$

$$\sum_{a \in \mathcal{A}} y_{s,a}(t+1) = (\phi(\mathbf{y}(t)))_s \quad \forall s, t, \quad (33c)$$

$$D\mathbf{y}(t) \leq \mathbf{b} \quad \forall t, \quad (33d)$$

$$N \cdot y_{s,a}(0) \in \mathbb{N} \quad \forall s, a, \quad (33e)$$

It is important to note that since the solution primarily utilizes control at the first time-step, we apply the integrity constraint solely at $t=0$.

We introduce $V_{\text{MILP-update}}^{(N)}(\mathbf{x}, T)$ to denote the expected performance of the policy obtained by implementing Algorithm 1, wherein the solution of the mixed-integer linear program (33) is utilized instead of (5). While this modified problem is more constrained than the original LP, it still represents a relaxation of the initial problem. Therefore, for any admissible policy $\pi \in \text{LP-update, MILP-update}$, the following holds:

$$V_{\pi}^{(N)}(\mathbf{x}, T) \leq V_{\text{rel}}(\mathbf{x}, T)$$

In Figure 5, we present a performance comparison between the LP-update and MILP-update policies. As anticipated, the MILP-update demonstrates superior performance compared to the original LP-update, particularly for smaller values of N .

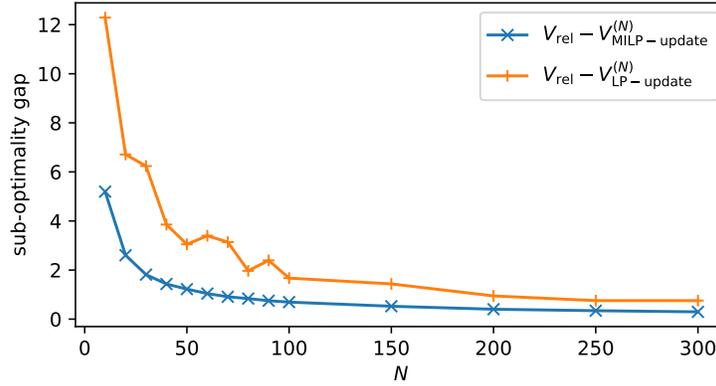


Figure 5 Sub-optimality gap of the LP-update policy compared to the MILP-update policy.

From a computational perspective, addressing (33) is more demanding in terms of resources compared to solving the original LP problem. Our experimental findings revealed that solving the MILP required no more than twice the time needed for the original LP for this specific model having 10 states. Notably, the duration for solving the MILP appeared to be unaffected by the size of N , suggesting that increases in N do not proportionally complicate the MILP solution process. This consistency in computational time, regardless of the scale of N , offers an interesting topic for further investigation, especially to determine if this trend is consistent across various scenarios and applications. It is also important to observe that the zigzag effect in the LP-update policy, resulting from the cutoff rounding, is absent in the MILP-update policy. This suggests that the MILP-update policy might be a more suitable and natural choice for implementation.

D.2. Perfect rounding

The constructed $\mathbf{Y}^{(N)}$, even using MILP, might still be far from \mathbf{Y} for small values of N . To obtain a $\mathbf{Y}^{(N)}$ closer to \mathbf{Y} , the approach developed in Gast et al. [14] for two-action bandits is to use the so-called *randomized rounding*, that consists of using a random vector $\mathbf{Y}^{(N)}$ such that $\mathbb{E}[\mathbf{Y}^{(N)}] = \mathbf{Y}$ and $\mathbf{Y}^{(N)} \in \mathcal{Y}^{(N)}(\mathbf{X}^{(N)})$ almost surely. If such a random vector exists, we say that the problem admits a *perfect rounding*.

DEFINITION 2 (PERFECT ROUNDING). An LP problem admits a perfect rounding for an integer N , if for all $\mathbf{x} \in \Delta^{(N)}$ and for all $\mathbf{y} \in \mathcal{Y}(\mathbf{x})$, there exists a random vector $\mathbf{Y}^{(N)}$ such that $\mathbf{Y}^{(N)} \in \mathcal{Y}^{(N)}(\mathbf{x})$ almost surely and $\mathbb{E}[\mathbf{Y}^{(N)}] = \mathbf{y}$.

The next theorem shows that, if there exists a perfect rounding, we obtain the much faster exponential convergence rate by avoiding the rounding errors. The result is summarized in the following convergence theorem:

THEOREM 4. Denote by $V_{\text{LP-update}}^{(N)}(\mathbf{x}, T)$ the value of the LP-update policy defined in Algorithm 1 or 2, and by $V_{\text{rel}}(\mathbf{x}, T)$ the value of the linear program (5). Then

- If the weakly-coupled MDP with statistically identical arms is non-degenerate, then there exist constants $C_1, C_2 > 0$ such that for any N for which the problem admits a perfect rounding:

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \leq C_1 \cdot e^{-C_2 N}.$$

- There exists a non-degenerate weakly-coupled MDP with statistically identical arms that admits a perfect rounding for an infinite number of N , and two constants C'_1, C'_2 such that for all such N :

$$\left| V_{\text{LP-update}}^{(N)}(\mathbf{x}, T) - V_{\text{rel}}(\mathbf{x}, T) \right| \geq C'_1 \cdot e^{-C'_2 N}.$$

Proof. The proof of the first item is identical to the proof of Theorem 3 until the last equation, for which if we have a perfect rounding, then $\mathbb{E}[\mathbf{Y}^{(N)}(t) - \mathbf{Y}(t)] = 0$. The proof of the lower bound in the second item is provided in Section A.4. Q.E.D.

Comparison with existing literature The authors of [14] study the single constraint case where $\mathcal{A} = \{0, 1\}$, $D(s, a) = a$ and $b = \alpha$. They show that a perfect rounding exists for all N such that αN is an integer. Another example is when the constraints (12b)-(12c)-(12d) form a totally unimodular matrix (see for instance Hoffman and Kruskal [19]) for all t : in that case, the solution to the LP satisfies $\mathbf{y} \in \mathcal{Y}^{(N)}(\mathbf{X}^{(N)})$ automatically and no rounding is needed. There are, however, many situations for which perfect rounding is impossible. In such cases, a method to improve the simple truncation (32) is to find $\mathbf{Y}^{(N)} \in \mathcal{Y}^{(N)}(\mathbf{X}^{(N)})$ that minimizes the distance $\|\mathbf{Y}^{(N)} - \mathbf{Y}\|_1$. This can be computed by solving an integer linear program.

The researchers in [39, 38] implement a randomized approach to derive an admissible $\mathbf{Y}^{(N)}$ from \mathbf{y}^* . This method is detailed in our description of their algorithm in Algorithm 3. Initially, their algorithm sets all arms to action 0. Subsequently, it iteratively processes each arm. For an arm in state s , the algorithm probabilistically selects a new action a based on the ratio $y_{s,a}^*(t)/x_s^*(t)$. This action a is then assigned to the arm, provided it does not breach any budget constraints. The central limit theorem underpins the rationale for this approach, ensuring that $\mathbf{Y}^{(N)}$ approximates $\mathbf{y}^* + \mathcal{O}(1/\sqrt{N})$. However, it does not extend to achieving a tighter approximation of $\mathbf{y}^* + \mathcal{O}(1/N)$. Consequently, even in the context of non-degenerate problems, the randomized process employed by these algorithms results in a sub-optimality gap of $\mathcal{O}(1/\sqrt{N})$ rather than the refined $\mathcal{O}(1/N)$.