
Deterministic Policies for Constrained Reinforcement Learning in Polynomial Time

Jeremy McMahan
University of Wisconsin-Madison
jcmahan@wisc.edu

Abstract

We present a novel algorithm that efficiently computes near-optimal deterministic policies for constrained reinforcement learning (CRL) problems. Our approach combines three key ideas: (1) value-demand augmentation, (2) action-space approximate dynamic programming, and (3) time-space rounding. Our algorithm constitutes a fully polynomial-time approximation scheme (FPTAS) for any time-space recursive (TSR) cost criteria. A TSR criteria requires the cost of a policy to be computable recursively over both time and (state) space, which includes classical expectation, almost sure, and anytime constraints. Our work answers three open questions spanning two long-standing lines of research: polynomial-time approximability is possible for 1) anytime-constrained policies, 2) almost-sure-constrained policies, and 3) deterministic expectation-constrained policies.

1 Introduction

Constrained Reinforcement Learning (CRL) traditionally produces stochastic, expectation-constrained policies that can behave undesirably - imagine a self-driving car that randomly changes lanes or runs out of fuel. However, artificial decision-making systems must be predictable, trustworthy, and robust. One approach to ensuring these qualities is to focus on deterministic policies, which are inherently predictable, easily implemented [19], reliable for autonomous vehicles [30, 23], and effective for multi-agent coordination [38]. Similarly, almost sure and anytime constraints [36] provide inherent trustworthiness and robustness, essential for applications in medicine [15, 37, 32], disaster relief [18, 50, 45], and resource management [35, 34, 40, 7]. Despite the advantages of deterministic policies and stricter constraints, even the computation of approximate solutions has remained an open challenge since NP-hardness was proven nearly 25 years ago [19]. Our work addresses this challenge by studying the computational complexity of computing deterministic policies for general constraint criteria.

Consider a constrained Markov Decision Process (cMDP) denoted by M . Let C represent an arbitrary cost criterion and B be the available budget. We focus on the set of deterministic policies denoted by Π^D . Our objective is to compute: $\max_{\pi \in \Pi^D} V_M^\pi$ s.t. $C_M^\pi \leq B$, where V_M^π is the value and C_M^π is the cost of π in M . This objective generalizes the example of a self-driving car calculating the fastest fixed route without running out of fuel. Our main question is the following:

Can near-optimal deterministic policies for constrained reinforcement learning problems be computed in polynomial time?

Although optimal stochastic policies for expectation-constrained problems are efficiently computable [3], the situation drastically changes when we require deterministic policies and general constraints. Computing optimal deterministic policies is NP-hard for most popular constraints, including expectation [19], chance [51], almost sure, and anytime constraints [36]. This complexity remains even if we relax our goal to finding just one feasible policy, provided that we are dealing

with a single chance constraint [51], or at least two of the other mentioned constraints [36]. Beyond these computational challenges, traditional solution methods, such as backward induction [41, 3], fail to apply due to the cyclic dependencies among subproblems: the value of any decision may depend on the costs of both preceding and concurrent decisions, preventing a solution from being computed in a single backward pass.

Past work. Past approaches fail to simultaneously achieve computational efficiency, feasibility, and optimality. Optimal and feasible algorithms, albeit inefficient, utilize Mixed-Integer Linear Programs [17] and Dual-guided heuristic forward searches [29] for expectation-constraints, and cost-augmented MDPs for almost sure [11] and anytime constraints [36]. Conversely, optimal and efficient, though infeasible, algorithms are known for expectation [43], almost sure, and anytime constraints [36]. A fully polynomial-time approximation scheme (FPTAS) [49] is known for expectation constraints, but it requires strong assumptions such as a constant horizon [31]. Balancing computational efficiency, feasibility, and optimality remains the bottleneck to efficient approximation.

Our contributions. We present an FPTAS for computing deterministic policies under any time-space recursive (TSR) constraint criteria. A TSR criteria requires the cost of a policy to be computable recursively in both time and (state) space, which captures expectation, almost sure, and anytime constraints. Since non-TSR criteria, such as chance constraints [51], are provably inapproximable, TSR seems pivotal for efficient computation. Overall, our general framework answers three open complexity questions spanning two longstanding lines of work: we prove polynomial-time approximability for 1) anytime-constrained policies, 2) almost-sure-constrained policies, and 3) deterministic expectation-constrained policies, which have been open for nearly 25 years [19].

Our approach breaks down into three main ideas: (1) value-demand augmentation, (2) action-space approximate dynamic programming, and (3) time-space rounding. We augment the states with value demands and the actions with future value demands to break cyclic subproblem dependencies, enabling dynamic programming methods. Importantly, we use values because they can be rounded without compromising feasibility [36] and can capture constraints that are not predictable from cumulative costs. However, this results in an exponential action space that makes solving the Bellman operator as hard as the knapsack problem. By exploiting the space-recursive nature of the criterion, we can efficiently approximate the Bellman operator with dynamic programming. Finally, rounding value demands result in approximation errors over both time and space, but carefully controlling these errors ensures provable guarantees.

1.1 Related work

Approximate packing. Many stochastic packing problems, which generalize the knapsack problem, are captured by our problem. Dean et al. [16], Frieze and Clarke [21] derived optimal approximation ratio algorithms for stochastic packing and integer packing with multiple constraints, respectively. Yang et al. [52], Bhalgat et al. [6] designed efficient approximation algorithms for variations of the stochastic knapsack problem. Then, Halman et al. [27] derived an FPTAS for a general class of stochastic dynamic programs, which was then further improved in [26, 1]. These methods require a single-dimensional state space that captures the constraint. In contrast, our problems have an innate state space in addition to the constraint. Our work forms a similar general dynamic programming framework for the more complex MDP setting.

Constrained RL. It is known that stochastic expectation-constrained policies are polynomial-time computable via linear programming [3], and many planning and learning algorithms exist for them [39, 46, 8, 28]. Some learning algorithms can even avoid violation during the learning process under certain assumptions [48, 4]. Furthermore, Brantley et al. [10] developed no-regret algorithms for cMDPs and extended their algorithms to the setting with a constraint on the cost accumulated over all episodes, which is called a knapsack constraint [10, 13].

Safe RL. The safe RL community [22, 25] has mainly focused on no-violation learning for stochastic expectation-constrained policies [14, 9, 2, 12, 5] and solving chance constraints [47, 53], which capture the probability of entering unsafe states. Performing learning while avoiding dangerous states [53] is a special case of expectation constraints that has also been studied [42, 44] under

non-trivial assumptions. In addition, instantaneous constraints, which require the expected cost to be within budget at all times, have also been studied [33, 20, 24].

2 Cost criteria

In this section, we formalize our problem setting. We also define our conditions for cost criteria.

Constrained Markov decision processes. A (tabular, finite-horizon) *Constrained Markov Decision Process* (cMDP) is a tuple $M = (\mathcal{S}, \mathcal{A}, P, r, c, H)$, where (i) \mathcal{S} is the finite set of *states*, (ii) \mathcal{A} is the finite set of *actions*, (iii) $P_h(s, a) \in \Delta(\mathcal{S})$ is the *transition* distribution, (iv) $r_h(s, a) \in \mathbb{R}$ is the *reward*, (v) $c_h(s, a) \in \mathbb{R}$ is the *cost*, and (vi) $H \in \mathbb{N}$ is the finite *time horizon*. We let $S := |\mathcal{S}|$, $A := |\mathcal{A}|$, $[H] := \{1, \dots, H\}$, and \mathcal{M} denote the set of all cMDPs. We also let $r_{max} \stackrel{\text{def}}{=} \max_{h,s,a} |r_h(s, a)|$ denote the maximum magnitude reward, $r_{min} \stackrel{\text{def}}{=} \min_{h,s,a} r_h(s, a)$ denote the true minimum reward, and $p_{min} \stackrel{\text{def}}{=} \min_{h,s,a,s'} P_h(s' | s, a)$ denote the minimum transition probability. Since \mathcal{S} is a finite set, we often assume $S = [S]$ WLOG. Lastly, for any predicate p , we use the Iverson bracket notation $[p]$ to denote 1 if p is true and 0 otherwise, and we let χ_p denote the characteristic function which evaluates to 0 if p is true and ∞ otherwise.

Interaction protocol. The agent interacts with M using a policy $\pi = (\pi_h)_{h=1}^H$. In the fullest generality, $\pi_h : \mathcal{H}_h \rightarrow \Delta(\mathcal{A})$ is a mapping from the observed history at time h to a distribution of actions. In contrast, a deterministic policy takes the form $\pi_h : \mathcal{H}_h \rightarrow \mathcal{A}$. We let Π denote the set of all possible policies and Π^D denote the set of all deterministic policies. The agent starts at the initial state $s_0 \in \mathcal{S}$ with observed history $\tau_1 = (s_0)$. For any $h \in [H]$, the agent chooses an action $a_h \sim \pi_h(\tau_h)$. Then, the agent receives immediate reward $r_h(s_h, a_h)$ and cost $c_h(s_h, a_h)$. Lastly, M transitions to state $s_{h+1} \sim P_h(s_h, a_h)$ and the agent updates the history to $\tau_{h+1} = (\tau_h, a_h, s_{h+1})$. This process is repeated for H steps; the interaction ends once s_{H+1} is reached.

Objective. For any cost criterion $C : \mathcal{M} \times \Pi \rightarrow \mathbb{R}$ and budget $B \in \mathbb{R}$, the agent's goal is to compute a solution to the following optimization problem:

$$\max_{\pi \in \Pi} \mathbb{E}_M^\pi \left[\sum_{h=1}^H r_h(s_h, a_h) \right] \quad \text{s.t.} \quad \begin{cases} C_M^\pi \leq B \\ \pi \text{ deterministic} \end{cases} \quad (\text{CON})$$

Here, \mathbb{P}_M^π denotes the probability law over histories induced from the interaction of π with M , and \mathbb{E}_M^π denotes the expectation defined by this law. We let $V_M^\pi \stackrel{\text{def}}{=} \mathbb{E}_M^\pi \left[\sum_{t=1}^H r_t(s_t, a_t) \right]$ denote the value of a policy π , and V_M^* denote the optimal solution value to (CON).

Cost criteria. We consider a broad family of cost criteria that satisfy a strengthening of the standard policy evaluation equations [41]. This strengthening requires not only the cost of a policy to be computable recursively in the time horizon, but at each time the cost should also break down recursively in (state) space.

Definition 1 (TSR). We call a cost criterion C *time-recursive* (TR) if for any cMDP M and policy $\pi \in \Pi^D$, π 's cost decomposes recursively into $C_M^\pi = C_1^\pi(s_0)$. Here, $C_{H+1}^\pi(\cdot) = \mathbf{0}$ and for any $h \in [H]$ and $\tau_h \in \mathcal{H}_h$,

$$C_h^\pi(\tau_h) = c_h(s, a) + f \left((P_h(s' | s, a), C_{h+1}^\pi(\tau_h, a, s'))_{s' \in P_h(s, a)} \right), \quad (\text{TR})$$

where $s = s_h(\tau_h)$, $a = \pi_h(\tau_h)$, and f is a non-decreasing function¹ computable in $O(S)$ time. For technical reasons, we also require that $f(x) = \infty$ whenever $\infty \in x$.

We further say C is *time-space-recursive* (TSR) if the f term above is equal to $g_h^{\tau_h, a}(1)$. Here, $g_h^{\tau_h, a}(S+1) = 0$ and for any $t \leq S$,

$$g_h^{\tau_h, a}(t) = \alpha \left(\beta \left(P_h(t | s, a), C_{h+1}^\pi(\tau_h, a, t) \right), g_h^{\tau_h, a}(t+1) \right), \quad (\text{SR})$$

¹When we say a multivariate function is non-decreasing, we mean it is non-decreasing with respect to the partial ordering induced by component-wise ordering.

where α is a non-decreasing function, and both α, β are computable in $O(1)$ time. We also assume that $\alpha(\cdot, \infty) = \infty$, and β satisfies $\alpha(\beta(0, \cdot), x) = x$ to match f 's condition.

Since the TR condition is a slight generalization of traditional policy evaluation, it is easy to see that we can solve for minimum-cost policies using backward induction.

Proposition 1 (TR Intuition). *If C is TR, then C satisfies the usual optimality equations. Furthermore, $\arg \min_{\pi \in \Pi^D} C_M^\pi$ can be computed using backward induction in $O(HS^2A)$ time.*

Although the TR condition is straightforward, the TSR condition is more strict. We will see the utility of the TSR condition in [Section 4](#) when computing Bellman updates. For now, we point out that the TSR condition is not too restrictive: it is satisfied by many popular criteria studied in the literature.

Proposition 2 (TSR examples). *The following classical constraints can be modeled by a TSR cost constraint.*

1. (Expectation Constraints) are captured by $C_M^\pi \stackrel{\text{def}}{=} \mathbb{E}_M^\pi \left[\sum_{h=1}^H c_h(s_h, a_h) \right] \leq B$. We see C is TSR by defining $\alpha(x, y) \stackrel{\text{def}}{=} x + y$ and $\beta(x, y) \stackrel{\text{def}}{=} xy$.
2. (Almost Sure Constraints) are captured by $C_M^\pi \stackrel{\text{def}}{=} \max_{\tau \in \mathcal{H}_{H+1}, \mathbb{P}_M^\pi[\tau] > 0} \sum_{h=1}^H c_h(s_h, a_h) \leq B$. We see C is TSR by defining $\alpha(x, y) \stackrel{\text{def}}{=} \max(x, y)$ and $\beta(x, y) \stackrel{\text{def}}{=} [x > 0]y$.
3. (Anytime Constraints) are captured by $C_M^\pi \stackrel{\text{def}}{=} \max_{t \in [H]} \max_{\tau \in \mathcal{H}_{H+1}, \mathbb{P}_M^\pi[\tau] > 0} \sum_{h=1}^t c_h(s_h, a_h) \leq B$. We see C is TSR by defining $\alpha(x, y) \stackrel{\text{def}}{=} \max(0, \max(x, y))$ and $\beta(x, y) \stackrel{\text{def}}{=} [x > 0]y$.

Remark 1 (Extensions). Our methods can also handle stochastic costs and infinite discounting. We defer the details to [Appendix F](#). Moreover, we can handle multiple constraints using vector-valued criteria so long as the comparison operator is a total ordering of the vector space.

Remark 2 (Inapproximability). Our methods cannot handle chance constraints or more than one of our example constraints. However, this is not a limitation of our framework as the problem becomes provably inapproximable under said constraints [[51](#), [36](#)].

3 Covering algorithm

In this section, we propose an algorithm to solve (CON). Our approach relies on converting the original problem into an equivalent covering problem that can be solved using an unconstrained MDP. This covering MDP is derived using the key idea of value augmentation.

Packing and covering. We can view (CON) as a *packing program*, which wishes to maximize V_M^π subject to $C_M^\pi \leq B$. However, we could also tackle the problem by reversing the objective: attempt to minimize C_M^π subject to $V_M^\pi \geq V_M^*$. If (CON) is feasible, then any optimal solution π to this *covering program* satisfies $V_M^\pi \geq V_M^*$ and $C_M^\pi \leq B$. Thus, we can solve the original packing program by solving the covering program.

Proposition 3 (Packing-Covering Reduction). *Suppose that $C_M^* \stackrel{\text{def}}{=} \min_{\pi \in \Pi^D} C_M^\pi$ s.t. $V_M^\pi \geq V_M^*$. Then, $C_M^* \leq B \iff V_M^* > -\infty$. Furthermore, if $V_M^* > -\infty$, then,*

$$\arg \min_{\substack{\pi \in \Pi^D \\ V_M^\pi \geq V_M^*}} C_M^\pi \subseteq \arg \max_{\substack{\pi \in \Pi^D \\ C_M^\pi \leq B}} V_M^\pi. \quad (\text{PC})$$

Thus, any solution to the covering program is a solution to the packing program.

We focus on the covering program for several reasons. To optimize the value recursively, we would need to predict the final cost resulting from intermediate decisions to ensure feasibility. Generally, such predictions would require strict assumptions on the cost criteria. By treating the value as the constraint instead, we only need to assume the cost can be optimized efficiently. Moreover, values are well understood in RL and are more amenable to approximation [[36](#)]. Thus, the covering program allows us to capture many criteria, ensure feasibility, and compute accurate value approximations.

Algorithm 1 Reduction to RL

Input: (M, C, B)

- 1: $\bar{M}, \bar{C} \leftarrow \text{Definition 2}(M, C)$
 - 2: $\pi, \bar{C}^* \leftarrow \text{SOLVE}(\bar{M}, \bar{C})$
 - 3: **if** $\bar{C}_1^*(s_0, v) > B$ for all $v \in \mathcal{V}$ **then**
 - 4: **return** “Infeasible”
 - 5: **else**
 - 6: **return** π
-

Value augmentation. We can solve the covering program by solving a cost-minimizing MDP \bar{M} . The key idea is to augment the state space with value demands, (s, v) . Then, the agent can recursively reason how to minimize its cost while meeting the current value demand. If the agent starts at (s_0, V_M^*) , then an optimal policy for \bar{M} should be a solution to the covering program.

The key invariant we desire is that any feasible policy π for \bar{M} should satisfy $\bar{V}_h^\pi(s, v) \geq v$. To ensure this invariance, we recall the policy evaluation equations [41]. If $\pi_h(s) = a$, then,

$$\bar{V}_h^\pi(s, v) = r_h(s, a) + \sum_{s'} P_h(s' | s, a) \bar{V}_{h+1}^\pi(s', v_{s'}). \quad (\text{PE})$$

For the value invariant to be satisfied, it suffices for the agent to choose an action a and commit to future value demands $v_{s'}$ satisfying,

$$r_h(s, a) + \sum_{s'} P_h(s' | s, a) v_{s'} \geq v. \quad (\text{DEM})$$

We can view choosing future value demands as part of the agent’s augmented actions. Then, at any augmented state (s, v) , the agent’s augmented action space includes all $(a, \mathbf{v}) \in \mathcal{A} \times \mathbb{R}^S$ satisfying (DEM). When M transitions to $s' \sim P_h(s, a)$, the agent’s new augmented state should consist of the environment’s new state in addition to its chosen demand for that state, $(s', v_{s'})$. Putting these pieces together yields the definition of the cover MDP, **Definition 2**.

Definition 2 (Cover MDP). The *cover MDP* $\bar{M} \stackrel{\text{def}}{=} (\bar{\mathcal{S}}, \bar{\mathcal{A}}, \bar{P}, \bar{c}, H)$ where,

1. $\bar{\mathcal{S}} \stackrel{\text{def}}{=} \mathcal{S} \times \mathcal{V}$ where $\mathcal{V} \stackrel{\text{def}}{=} \{v \mid \exists \pi \in \Pi^D, h \in [H + 1], \tau_h \in \mathcal{H}_h, V_h^\pi(\tau_h) = v\}$
2. $\bar{\mathcal{A}}_h(s, v) \stackrel{\text{def}}{=} \{(a, \mathbf{v}) \in \mathcal{A} \times \mathcal{V}^S \mid r_h(s, a) + \sum_{s'} P_h(s' | s, a) v_{s'} \geq v\}$.
3. $\bar{P}_h((s', v') \mid (s, v), (a, \mathbf{v})) \stackrel{\text{def}}{=} P_h(s' | s, a) [v' = v_{s'}]$.
4. $\bar{c}_h((s, v), (a, \mathbf{v})) \stackrel{\text{def}}{=} c_h(s, a)$.

The objective for \bar{M} is to minimize the cost function $\bar{C} \stackrel{\text{def}}{=} C_{\bar{M}}$ with modified base case $\bar{C}_{H+1}^\pi(s, v) \stackrel{\text{def}}{=} \chi_{\{v \leq 0\}}$.

Covering algorithm. Importantly, the action space definition ensures the value constraint is satisfied. Meanwhile, the minimum cost objective ensures optimal cost. So long as our cost is TR, \bar{M} can be solved using fast RL methods instead of the brute force computation required for general covering programs. These properties ensure our method, **Algorithm 1**, is correct.

Theorem 1 (Reduction). *If SOLVE is any finite-time MDP solver, then **Algorithm 1** correctly solves (CON) in finite time for any TR cost criterion.*

Remark 3 (Execution). Given a value-augmented policy π output from **Algorithm 1**, the agent can execute π using **Algorithm 2**. To compute V_M^* as the starting value, it suffices for the agent to compute,

$$V_M^* = \max \{v \in \mathcal{V} \mid \bar{C}_1^*(s_0, v) \leq B\}. \quad (1)$$

This computation can be easily done given $\bar{C}_1^*(s_0, \cdot)$ in $O(|\mathcal{V}|)$ time.

Algorithm 2 Augmented interaction

Input: π

- 1: $\bar{s}_1 = (s_0, V_M^*)$
 - 2: **for** $h \leftarrow 1$ to H **do**
 - 3: $(a, \mathbf{v}) \leftarrow \pi_h(\bar{s}_h)$
 - 4: $r_h = r_h(s, a)$ and $s_{h+1} \sim P_h(s_h, a)$
 - 5: $\bar{s}_{h+1} = (s_{h+1}, v_{s_{h+1}})$
-

4 Fast Bellman updates

In this section, we present an algorithm to solve \bar{M} from [Definition 2](#) efficiently. Although the Bellman updates can be as hard to solve as the knapsack problem, we use ideas from knapsack approximation algorithms to create an efficient method. Our approach exploits (SR) through approximate dynamic programming on the action space.

Even if \mathcal{V} were small, solving \bar{M} would still be challenging due to the exponentially large action space. Even a single Bellman update requires the solution of a constrained optimization problem:

$$\begin{aligned} \bar{C}_h^*(s, v) = \min_{a, \mathbf{v}} c_h(s, a) + f \left((P_h(s' | s, a), \bar{C}_{h+1}^*(s', v_{s'}))_{s' \in P_h(s, a)} \right) \\ \text{s.t. } r_h(s, a) + \sum_{s'} P_h(s' | s, a) v_{s'} \geq v. \end{aligned} \quad (\text{BU})$$

Above, we used the fact that $(s', v') \in \bar{P}_h((s, v), (a, \mathbf{v}))$ iff $s' \in P_h(s, a)$ and $v' = v_{s'}$ to simplify f 's input. Observe that even when each $v_{s'}$ only takes on two possible values, $\{0, w_{s'}\}$, the optimization above can capture the minimization version of the knapsack problem, implying that it is NP-hard to compute.

Recursive approach. Fortunately, we can use the connection to the Knapsack problem positively to efficiently approximate the Bellman update. For any fixed $(s, v) \in \bar{\mathcal{S}}$ and $a \in \mathcal{A}$, we focus on the inner constrained minimization over \mathbf{v} :

$$\min_{\substack{\mathbf{v} \in \mathcal{V}^S, \\ r_h(s, a) + \sum_{s'} P_h(s' | s, a) v_{s'} \geq v}} f \left((P_h(s' | s, a), \bar{C}_{h+1}^*(s', v_{s'}))_{s' \in P_h(s, a)} \right) \quad (2)$$

We use (SR) to transform this minimization over \mathbf{v} into a sequential decision-making problem that decides each $v_{s'}$. As above, we can use the definition of \bar{P} to simplify $g_h^{(s, v), (a, \mathbf{v})}(t, v')$ into a function of t alone:

$$g_h^{(s, v), (a, \mathbf{v})}(t) = \alpha \left(\beta (P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t)), g_h^{(s, v), (a, \mathbf{v})}(t+1) \right). \quad (3)$$

Since v only constrains the valid (a, \mathbf{v}) pairs, we can discard v and use the simplified notation $g_{h, \mathbf{v}}^{s, a}(t)$ instead of $g_h^{(s, v), (a, \mathbf{v})}(t)$. It is then clear that we can recursively optimize the value of v_t by focusing on $g_{h, \mathbf{v}}^{s, a}(t)$.

To recursively encode the value constraint, we can record the partial value $u = r_h(s, a) + \sum_{s'=1}^{t-1} P_h(s' | s, a) v_{s'}$ that we have accumulated so far. Then, we can check if our choices for \mathbf{v} satisfied the constraint with the inequality $u \geq v$. The formal recursion is defined in [Definition 3](#).

Definition 3. For any $h \in [H]$, $s \in \mathcal{S}$, $v \in \mathcal{V}$, and $u \in \mathbb{R}$, we define, $g_{h, v}^{s, a}(S+1, u) = \chi_{\{u \geq v\}}$ and for $t \leq S$,

$$g_{h, v}^{s, a}(t, u) = \min_{v_t \in \mathcal{V}} \alpha \left(\beta (P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t)), g_{h, v}^{s, a}(t+1, u + P_h(t | s, a) v_t) \right). \quad (\text{DP})$$

Recursive rounding. This approach can still be slow due to the exponential number of partial values u induced. Similarly to the knapsack problem, the key is to round each input u to ensure fewer subproblems. Unlike the knapsack problem, however, we do not have an easily computable lower bound on the optimal value. Thus, we turn to a more aggressive recursive rounding. Since rounding may cause originally feasible values to violate the demand constraint, we also relax the demand constraint to $u \geq \kappa(v)$ for some lower bound function κ .

Algorithm 3 Approx Bellman Update

Input: $(h, s, v, \bar{C}_{h+1}^*)$

- 1: **for** $a \in \mathcal{A}$ **do**
- 2: $\hat{g}_{h,v}^{s,a}(S+1, u) \leftarrow \chi_{\{u \geq v\}} \forall u \in \hat{\mathcal{U}}_h^{s,a}(S+1)$
- 3: **for** $t \leftarrow S$ **down to** 1 **do**
- 4: **for** $u \in \hat{\mathcal{U}}_h^{s,a}(t)$ **do**
- 5: $v_{t,a}, \hat{g}_{h,v}^{s,a}(t, u) \leftarrow (\text{ADP})$
- 6: $a^*, \hat{C}_h^*(s, v) \leftarrow \min_{a \in \mathcal{A}} c_h(s, a) + \hat{g}_{h,v}^{s,a}(1, r_h(s, a))$
- 7: **return** $(a^*, v_{1,a^*}, \dots, v_{S,a^*})$ and $\hat{C}_h^*(s, v)$

Algorithm 4 Approx Solve

Input: (\bar{M}, \bar{C})

- 1: $\hat{C}_{H+1}^*(s, v) \leftarrow \chi_{\{v \leq 0\}}$ for all $(s, v) \in \bar{\mathcal{S}}$
- 2: **for** $h \leftarrow H$ **down to** 1 **do**
- 3: **for** $(s, v) \in \bar{\mathcal{S}}$ **do**
- 4: $\hat{a}, \hat{C}_h^*(s, v) \leftarrow \text{Algorithm 3}(h, s, v, \hat{C}_{h+1}^*)$
- 5: $\pi_h(s, v) \leftarrow \hat{a}$
- 6: **return** π and \hat{C}^*

Definition 4. Fix a rounding function $\lfloor \cdot \rfloor_{\mathcal{G}}$ and a lower bound function κ . For any $h \in [H]$, $s \in \mathcal{S}$, $v \in \mathcal{V}$, and $u \in \mathbb{R}$, we define, $\hat{g}_{h,v}^{s,a}(S+1, u) = \chi_{\{u \geq v\}}$ and for $t \leq S$,

$$\hat{g}_{h,v}^{s,a}(t, u) \stackrel{\text{def}}{=} \min_{v_t \in \mathcal{V}} \alpha \left(\beta (P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t)), \hat{g}_{h,v}^{s,a}(t+1, \lfloor u + P_h(t | s, a)v_t \rfloor_{\mathcal{G}}) \right). \quad (\text{ADP})$$

Fortunately, the approximate version behaves similarly to the original. The main difference is the constraint now ensures the rounded sums are at least the value lower bound. This is formalized in [Lemma 1](#).

Lemma 1. For any $t \in [S+1]$ and $u \in \mathbb{R}$, we have that,

$$\begin{aligned} \hat{g}_{h,v}^{s,a}(t, u) &= \min_{v \in \mathcal{V}^{S-t+1}} g_{h,v}^{s,a}(t) \\ \text{s.t. } \hat{\sigma}_{h,v}^{s,a}(t, u) &\geq \kappa(v), \end{aligned} \quad (4)$$

where $\hat{\sigma}_{h,v}^{s,a}(t, u) \stackrel{\text{def}}{=} \lfloor [u + P_h(t | s, a)v_t]_{\mathcal{G}} + \dots + P_h(S | s, a)v_S \rfloor_{\mathcal{G}}$.

To turn this recursion into a usable dynamic programming algorithm, we must also pre-compute the inputs to any sub-computation. Unlike in standard RL, this computation must be done with a forward recursion. The details for the approximate Bellman update are given in [Definition 5](#).

Definition 5 (Approx Bellman). For any $h \in [H]$, $s \in \mathcal{S}$, and $a \in \mathcal{A}$, we define $\hat{\mathcal{U}}_h^{s,a}(1) \stackrel{\text{def}}{=} \{r_h(s, a)\}$ and for any $t \in [S]$,

$$\hat{\mathcal{U}}_h^{s,a}(t+1) \stackrel{\text{def}}{=} \bigcup_{v_t \in \mathcal{V}} \bigcup_{u \in \hat{\mathcal{U}}_h^{s,a}(t)} \{ \lfloor u + P_h(t | s, a)v_t \rfloor_{\mathcal{G}} \}. \quad (5)$$

Then, an approximation to the Bellman update can be computed using [Algorithm 3](#).²

Proposition 4. *Algorithm 4* runs in $O(HS^2A|\mathcal{V}|^2\hat{U})$ time, where $\hat{U} \stackrel{\text{def}}{=} \max_{h,s,a} |\hat{\mathcal{U}}_h^{s,a}|$. When $\lfloor \cdot \rfloor_{\mathcal{G}}$ and κ are the identity function, *Algorithm 4* outputs an optimal solution to \bar{M} .

Remark 4 (Speedups). The runtime of our methods can be quadratically improved by rounding the differences instead of the sums. We defer the details to [Appendix F](#).

²We use the notation $x, o \leftarrow \min_x z(x)$ to say that x is the minimizer and o the value of the optimization.

5 Approximation algorithms

In this section, we present our approximation algorithms for solving (CON). We carefully round the value demands over both time and space to induce an approximate MDP. Solving this approximate MDP with [Algorithm 4](#) yields our FPTAS.

Although we can avoid exponential-time Bellman updates, the running time of the approximate Bellman update will still be slow if $|\mathcal{V}|$ is large. To reduce the complexity, we instead use a smaller set of approximate values by rounding elements of $|\mathcal{V}|$. By rounding down, we effectively relax the value-demand constraint. More aggressive rounding not only leads to smaller augmented state spaces but also to smaller cost policies. The trade-off is aggressive rounding leads to weaker guarantees on the computed policy's value. Thus, it is critical to carefully design the rounding and lower bound functions to balance this trade-off.

Value approximation. Given a rounding down function $\lfloor \cdot \rfloor_{\mathcal{G}}$, we would ideally use the rounded set $\{\lfloor v \rfloor_{\mathcal{G}} \mid v \in \mathcal{V}\}$ to form our approximate state space. To avoid having to compute \mathcal{V} explicitly, we instead use the rounded superset $\{\lfloor v \rfloor_{\mathcal{G}} \mid v \in [v_{min}, v_{max}]\}$, where v_{min} and v_{max} are bounds on the extremal values that we specify later. To ensure we can use [Algorithm 4](#) to find solutions efficiently, we must also relax the augmented action space to only include vectors that lead to feasible subproblems for (ADP). From [Lemma 1](#), we know this is exactly the set of (a, \hat{v}) for which $\hat{\sigma}_{h, \hat{v}}^{s, a}(1, r_h(s, a)) \geq \kappa(v)$. Combining these ideas yields the new approximate MDP, defined in [Definition 6](#).

Definition 6 (Approximate MDP). Given a rounding function $\lfloor \cdot \rfloor_{\mathcal{G}}$ and lower bound function κ , the approximate MDP $\hat{M} \stackrel{\text{def}}{=} (\hat{\mathcal{S}}, \hat{\mathcal{A}}, \hat{P}, \hat{c}, H)$ where,

1. $\hat{\mathcal{S}} \stackrel{\text{def}}{=} \mathcal{S} \times \hat{\mathcal{V}}$ where $\hat{\mathcal{V}} \stackrel{\text{def}}{=} \{\lfloor v \rfloor_{\mathcal{G}} \mid v \in [v_{min}, v_{max}]\}$.
2. $\hat{\mathcal{A}}_h(s, \hat{v}) \stackrel{\text{def}}{=} \{(a, \hat{v}) \in \mathcal{A} \times \hat{\mathcal{V}}^S \mid \hat{\sigma}_{h, \hat{v}}^{s, a}(1, r_h(s, a)) \geq \kappa(\hat{v})\}$.
3. $\hat{P}_h((s', \hat{v}') \mid (s, \hat{v}), (a, \hat{v})) \stackrel{\text{def}}{=} P_h(s' \mid s, a)[\hat{v}' = \hat{v}_{s'}]$.
4. $\hat{c}_h((s, \hat{v}), (a, \hat{v})) \stackrel{\text{def}}{=} c_h(s, a)$.

The objective for \hat{M} is to minimize the cost function $\hat{C} \stackrel{\text{def}}{=} C_{\hat{M}}$ with modified base case $\hat{C}_{H+1}^\pi(s, \hat{v}) \stackrel{\text{def}}{=} \chi_{\{\hat{v} \leq 0\}}$.

We can show that rounding down in [Definition 6](#) achieves our goal of producing smaller cost policies. This ensures feasibility is even easier to achieve. We formalize this observation in [Lemma 2](#).

Lemma 2 (Optimistic Costs). *For our later choices of $\lfloor \cdot \rfloor_{\mathcal{G}}$ and κ , the following holds: for any $h \in [H + 1]$ and $(s, v) \in \hat{\mathcal{S}}$, we have $\hat{C}_h^*(s, \lfloor v \rfloor_{\mathcal{G}}) \leq \bar{C}_h^*(s, v)$.*

Thus, [Algorithm 5](#) always outputs a policy with better than optimal cost when the instance is feasible, $V_M^* > -\infty$. If the instance is infeasible, all policies have cost larger than B by definition and so [Algorithm 5](#) correctly indicates the instance is infeasible. The remaining question is whether [Algorithm 5](#) outputs policies having near-optimal value.

Time-Space errors. To assess the optimality gap of [Algorithm 5](#) policies, we must first explore the error accumulated by our rounding approach. Rounding each value naturally accumulates approximation error over time. Rounding the partial values while running [Algorithm 3](#) accumulates additional error over (state) space. Thus, solving \hat{M} using [Algorithm 4](#) accumulates error over both time and space, unlike other approximate methods in RL. As a result, our rounding and threshold functions will generally depend on both H and S .

Arithmetic rounding. Our first approach is to round each value down to its closest element in a δ -cover. This guarantees that $v - \delta \leq \lfloor v \rfloor_{\mathcal{G}} \leq v$. Thus, $\lfloor v \rfloor_{\mathcal{G}}$ is an underestimate that is not too far from the true value. By setting δ to be inversely proportional to SH , we control the errors over time and space. The lower bound must also be a function of S since it controls the error over space.

Algorithm 5 Approximation Scheme

Input: (M, C, B)
1: **Hyperparameters:** $\lfloor \cdot \rfloor_{\mathcal{G}}$ and κ
2: $\hat{M}, \hat{C} \leftarrow$ **Definition 6** $(M, C, \lfloor \cdot \rfloor_{\mathcal{G}}, \kappa)$
3: $\pi, \hat{C}^* \leftarrow$ **Algorithm 4** (\hat{M}, \hat{C})
4: **if** $\hat{C}_1^*(s_0, \hat{v}) > B$ for all $\hat{v} \in \hat{\mathcal{V}}$ **then**
5: **return** “Infeasible”
6: **else**
7: **return** π

Definition 7 (Additive Approx). Fix $\epsilon > 0$. We define,

$$\lfloor v \rfloor_{\mathcal{G}} \stackrel{\text{def}}{=} \left\lfloor \frac{v}{\delta} \right\rfloor \delta \text{ and } \kappa(v) \stackrel{\text{def}}{=} v - \delta(S + 1), \quad (6)$$

where $\delta \stackrel{\text{def}}{=} \frac{\epsilon}{H(S+1)+1}$, $v_{\min} \stackrel{\text{def}}{=} -Hr_{\max}$, and $v_{\max} \stackrel{\text{def}}{=} Hr_{\max}$.

Theorem 2 (Additive FPTAS). For any $\epsilon > 0$, **Algorithm 5** using **Definition 7** given any cMDP M and TSR criteria C either correctly outputs the instance is infeasible, or produces a policy π satisfying $\hat{V}^{\pi} \geq V_M^* - \epsilon$ in $O(H^7 S^5 Ar_{\max}^3 / \epsilon^3)$ time. Thus, it is an additive-FPTAS for the class of cMDPs with polynomial-bounded r_{\max} and TSR criteria.

Geometric rounding. Since the arithmetic approach can be slow when r_{\max} is large, we can instead round values down to their closest power of $1/(1 - \delta)$. This guarantees the number of approximate values needed is upper bounded by a function of $\log(r_{\max})$, which is polynomial in the input size. We choose a geometric scheme satisfying $v(1 - \delta) \leq \lfloor v \rfloor_{\mathcal{G}} \leq v$ so that the rounded value is an underestimate and a relative approximation to the true value. To ensure this property, we must now require that all rewards are non-negative.

Definition 8 (Relative Approx). Fix $\epsilon > 0$. We define,

$$\lfloor v \rfloor_{\mathcal{G}} \stackrel{\text{def}}{=} v_{\min} \left(\frac{1}{1 - \delta} \right)^{\left\lfloor \log_{\frac{1}{1-\delta}} \frac{v}{v_{\min}} \right\rfloor} \text{ and } \kappa(v) \stackrel{\text{def}}{=} v(1 - \delta)^{S+1}, \quad (7)$$

where $\delta \stackrel{\text{def}}{=} \frac{\epsilon}{H(S+1)+1}$, $v_{\min} = p_{\min}^H r_{\min}$, and $v_{\max} = Hr_{\max}$.

Theorem 3 (Relative FPTAS). For $\epsilon > 0$, **Algorithm 5** using **Definition 8** given any cMDP M and TSR criteria C either correctly outputs the instance is infeasible, or produces a policy π satisfying $\hat{V}^{\pi} \geq V_M^*(1 - \epsilon)$ in $O(H^7 S^5 A \log(r_{\max}/r_{\min} p_{\min})^3 / \epsilon^3)$ time. Thus, it is a relative-FPTAS for the class of cMDPs with non-negative rewards and TSR criteria.

Remark 5 (Assumption Necessity). We also note the mild reward assumptions we made to guarantee efficiency are unavoidable. Without reward bounds, (CON) captures the knapsack problem which does not admit additive approximations. Similarly, without non-negativity, relative approximations for maximization problems are generally not computable.

6 Conclusions

In this paper, we studied the computational complexity of computing deterministic policies for CRL problems. Our main contribution was the design of an FPTAS, **Algorithm 5**, that solves (CON) for any cMPD and TSR criteria under mild reward assumptions. In particular, our method is an additive-FPTAS if the cMDP’s rewards are polynomially bounded, and is a relative-FPTAS if the cMDP’s rewards are non-negative. We note these assumptions are necessary for efficient approximation, so our algorithm achieves the best approximation guarantees possible under worst-case analysis. Moreover, our algorithmic approach, which uses approximate dynamic programming over time and the state space, highlights the importance of the TSR condition in making (CON) tractable. Our work finally resolves the long-standing open questions of polynomial-time approximability for 1) anytime-constrained policies, 2) almost-sure-constrained policies, and 3) deterministic expectation-constrained policies.

Future work. Several interesting questions remain unanswered. First, it remains unresolved whether an FPTAS exists asymptotically faster than ours. Second, whether our TSR condition is necessary for efficient computation or whether a more general condition could be derived is unclear. Lastly, it is open whether there exist algorithms that can feasibly handle multiple constraints from [Proposition 2](#). Although computing feasible policies for multiple constraints is NP-hard, special cases may be approximable efficiently. Moreover, an average-case or smoothed-case analysis could circumvent this worst-case hardness.

Acknowledgments and Disclosure of Funding

This work was supported in part by NSF grant 2023239.

References

- [1] T. Alon and N. Halman. Automatic generation of fptases for stochastic monotone dynamic programs made easier. *SIAM Journal on Discrete Mathematics*, 35(4):2679–2722, 2021. doi: 10.1137/19M1308633. URL <https://doi.org/10.1137/19M1308633>.
- [2] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu. Safe reinforcement learning via shielding. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), Apr. 2018. doi: 10.1609/aaai.v32i1.11797. URL <https://ojs.aaai.org/index.php/AAAI/article/view/11797>.
- [3] E. Altman. *Constrained Markov Decision Processes*. Chapman and Hall/CRC, 1999. doi: 10.1201/9781315140223.
- [4] Q. Bai, A. Singh Bedi, and V. Aggarwal. Achieving zero constraint violation for constrained reinforcement learning via conservative natural policy gradient primal-dual algorithm. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6):6737–6744, 6 2023. doi: 10.1609/aaai.v37i6.25826. URL <https://ojs.aaai.org/index.php/AAAI/article/view/25826>.
- [5] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause. Safe model-based reinforcement learning with stability guarantees. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/766ebcd59621e305170616ba3d3dac32-Paper.pdf.
- [6] A. Bhalgat, A. Goel, and S. Khanna. Improved approximation results for stochastic knapsack problems. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’11, page 1647–1665, USA, 2011. Society for Industrial and Applied Mathematics.
- [7] A. Bhatia, P. Varakantham, and A. Kumar. Resource constrained deep reinforcement learning. *Proceedings of the International Conference on Automated Planning and Scheduling*, 29(1): 610–620, 5 2021. doi: 10.1609/icaps.v29i1.3528. URL <https://ojs.aaai.org/index.php/ICAPS/article/view/3528>.
- [8] V. Borkar. An actor-critic algorithm for constrained markov decision processes. *Systems & Control Letters*, 54(3):207–213, 2005. ISSN 0167-6911. doi: <https://doi.org/10.1016/j.sysconle.2004.08.007>. URL <https://www.sciencedirect.com/science/article/pii/S0167691104001276>.
- [9] D. M. Bossens and N. Bishop. Explicit explore, exploit, or escape (e4): Near-optimal safety-constrained reinforcement learning in polynomial time. *Mach. Learn.*, 112(3):817–858, 6 2022. ISSN 0885-6125. doi: 10.1007/s10994-022-06201-z. URL <https://doi.org/10.1007/s10994-022-06201-z>.
- [10] K. Brantley, M. Dudík, T. Lykouris, S. Miryoosefi, M. Simchowitz, A. Slivkins, and W. Sun. Constrained episodic reinforcement learning in concave-convex and knapsack settings. In *NeurIPS*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/bc6d753857fe3dd4275dff707dedf329-Abstract.html>.

- [11] A. Castellano, H. Min, E. Mallada, and J. A. Bazerque. Reinforcement learning with almost sure constraints. In R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, and M. Kochenderfer, editors, *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, volume 168 of *Proceedings of Machine Learning Research*, pages 559–570. PMLR, 6 2022. URL <https://proceedings.mlr.press/v168/castellano22a.html>.
- [12] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):3387–3395, Jul. 2019. doi: 10.1609/aaai.v33i01.33013387. URL <https://ojs.aaai.org/index.php/AAAI/article/view/4213>.
- [13] W. C. Cheung. Regret minimization for reinforcement learning with vectorial feedback and complex objectives. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL <https://proceedings.neurips.cc/paper/2019/file/a02ffd91ece5e7efeb46db8f10a74059-Paper.pdf>.
- [14] Y. Chow, O. Nachum, E. Duenez-Guzman, and M. Ghavamzadeh. A lyapunov-based approach to safe reinforcement learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/4fe5149039b52765bde64beb9f674940-Paper.pdf.
- [15] A. Coronato, M. Naeem, G. De Pietro, and G. Paragliola. Reinforcement learning for intelligent healthcare applications: A survey. *Artificial Intelligence in Medicine*, 109:101964, 2020. ISSN 0933-3657. doi: <https://doi.org/10.1016/j.artmed.2020.101964>. URL <https://www.sciencedirect.com/science/article/pii/S093336572031229X>.
- [16] B. C. Dean, M. X. Goemans, and J. Vondrák. Adaptivity and approximation for stochastic packing problems. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05*, page 395–404, USA, 2005. Society for Industrial and Applied Mathematics. ISBN 0898715857.
- [17] D. A. Dolgov and E. H. Durfee. Stationary deterministic policies for constrained mdps with multiple rewards, costs, and discount factors. In *IJCAI*, volume 19, pages 1326–1331, 2005.
- [18] C. Fan, C. Zhang, A. Yahja, and A. Mostafavi. Disaster city digital twin: A vision for integrating artificial and human intelligence for disaster management. *International Journal of Information Management*, 56:102049, 2021. ISSN 0268-4012. doi: <https://doi.org/10.1016/j.ijinfomgt.2019.102049>. URL <https://www.sciencedirect.com/science/article/pii/S0268401219302956>.
- [19] E. A. Feinberg. Constrained discounted markov decision processes and hamiltonian cycles. *Mathematics of Operations Research*, 25(1):130–140, 2000. doi: 10.1287/moor.25.1.130.15210. URL <https://doi.org/10.1287/moor.25.1.130.15210>.
- [20] J. F. Fisac, N. F. Lugovoy, V. Rubies-Royo, S. Ghosh, and C. J. Tomlin. Bridging hamilton-jacobi safety analysis and reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, page 8550–8556. IEEE Press, 2019. doi: 10.1109/ICRA.2019.8794107. URL <https://doi.org/10.1109/ICRA.2019.8794107>.
- [21] A. Frieze and M. Clarke. Approximation algorithms for the m-dimensional 0–1 knapsack problem: Worst-case and probabilistic analyses. *European Journal of Operational Research*, 15(1):100–109, 1984. ISSN 0377-2217. doi: [https://doi.org/10.1016/0377-2217\(84\)90053-5](https://doi.org/10.1016/0377-2217(84)90053-5). URL <https://www.sciencedirect.com/science/article/pii/0377221784900535>.
- [22] J. García, Fern, and o Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(42):1437–1480, 2015. URL <http://jmlr.org/papers/v16/garcia15a.html>.
- [23] F. Geißer, G. Povéda, F. Trevizan, M. Bondouy, F. Teichteil-Königsbuch, and S. Thiébaux. Optimal and heuristic approaches for constrained flight planning under weather uncertainty. *Proceedings of the International Conference on Automated Planning and Scheduling*, 30(1): 384–393, Jun. 2020. doi: 10.1609/icaps.v30i1.6684. URL <https://ojs.aaai.org/index.php/ICAPS/article/view/6684>.

- [24] S. Gros, M. Zanon, and A. Bemporad. Safe reinforcement learning via projection on a safe set: How to achieve optimality? *IFAC-PapersOnLine*, 53(2):8076–8081, 2020. ISSN 2405-8963. doi: <https://doi.org/10.1016/j.ifacol.2020.12.2276>. URL <https://www.sciencedirect.com/science/article/pii/S2405896320329360>. 21st IFAC World Congress.
- [25] S. Gu, L. Yang, Y. Du, G. Chen, F. Walter, J. Wang, and A. Knoll. A review of safe reinforcement learning: Methods, theory and applications, 2024. URL <https://arxiv.org/abs/2205.10330>.
- [26] N. Halman and G. Nannicini. Toward breaking the curse of dimensionality: An fptas for stochastic dynamic programs with multidimensional actions and scalar states. *SIAM Journal on Optimization*, 29(2):1131–1163, 2019. doi: 10.1137/18M1208423. URL <https://doi.org/10.1137/18M1208423>.
- [27] N. Halman, D. Klabjan, C.-L. Li, J. Orlin, and D. Simchi-Levi. Fully polynomial time approximation schemes for stochastic dynamic programs. *SIAM Journal on Discrete Mathematics*, 28(4):1725–1796, 2014. doi: 10.1137/130925153. URL <https://doi.org/10.1137/130925153>.
- [28] A. HasanzadeZonuzi, A. Bura, D. Kalathil, and S. Shakkottai. Learning with safety constraints: Sample complexity of reinforcement learning for constrained mdps. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(9):7667–7674, 5 2021. doi: 10.1609/aaai.v35i9.16937. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16937>.
- [29] S. Hong and B. C. Williams. An anytime algorithm for constrained stochastic shortest path problems with deterministic policies. *Artificial Intelligence*, 316:103846, 2023. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2022.103846>. URL <https://www.sciencedirect.com/science/article/pii/S0004370222001862>.
- [30] S. Hong, S. U. Lee, X. Huang, M. Khonji, R. Alyassi, and B. C. Williams. An anytime algorithm for chance constrained stochastic shortest path problems and its application to aircraft routing. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 475–481, 2021. doi: 10.1109/ICRA48506.2021.9561229.
- [31] M. Khonji, A. Jasour, and B. Williams. Approximability of constant-horizon constrained pomdp. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5583–5590. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/775. URL <https://doi.org/10.24963/ijcai.2019/775>.
- [32] P. Kolesar. A markovian model for hospital admission scheduling. *Management Science*, 16(6):B384–B396, 1970. ISSN 00251909, 15265501. URL <http://www.jstor.org/stable/2628725>.
- [33] J. Li, D. Fridovich-Keil, S. Sojoudi, and C. J. Tomlin. Augmented lagrangian method for instantaneously constrained reinforcement learning problems. In *2021 60th IEEE Conference on Decision and Control (CDC)*, page 2982–2989. IEEE Press, 2021. doi: 10.1109/CDC45484.2021.9683088. URL <https://doi.org/10.1109/CDC45484.2021.9683088>.
- [34] R. Li, Z. Zhao, Q. Sun, C.-L. I, C. Yang, X. Chen, M. Zhao, and H. Zhang. Deep reinforcement learning for resource management in network slicing. *IEEE Access*, 6:74429–74441, 2018. doi: 10.1109/ACCESS.2018.2881964.
- [35] H. Mao, M. Alizadeh, I. Menache, and S. Kandula. Resource management with deep reinforcement learning. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks, HotNets ’16*, page 50–56, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450346610. doi: 10.1145/3005745.3005750. URL <https://doi.org/10.1145/3005745.3005750>.
- [36] J. McMahan and X. Zhu. Anytime-constrained reinforcement learning. In S. Dasgupta, S. Mandt, and Y. Li, editors, *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 4321–4329. PMLR, 02–04 May 2024. URL <https://proceedings.mlr.press/v238/mcmahan24a.html>.

- [37] G. Paragliola, A. Coronato, M. Naeem, and G. De Pietro. A reinforcement learning-based approach for the risk management of e-health environments: A case study. In *2018 14th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*, pages 711–716, 2018. doi: 10.1109/SITIS.2018.00114.
- [38] P. Paruchuri, M. Tambe, F. Ordonez, and S. Kraus. Towards a formalization of teamwork with resource constraints. In *Third International Joint Conference on Autonomous Agents and Multiagent Systems*, United States, 2004. IEEE Computer Society. Place of conference:USA.
- [39] S. Paternain, L. Chamon, M. Calvo-Fullana, and A. Ribeiro. Constrained reinforcement learning has zero duality gap. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/c1aeb6517a1c7f33514f7ff69047e74e-Paper.pdf.
- [40] H. Peng and X. Shen. Multi-agent reinforcement learning based resource management in mec- and uav-assisted vehicular networks. *IEEE Journal on Selected Areas in Communications*, 39(1):131–141, 2021. doi: 10.1109/JSAC.2020.3036962.
- [41] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., USA, 1st edition, 1994. ISBN 0471619779.
- [42] M. Roderick, V. Nagarajan, and Z. Kolter. Provably safe pac-mdp exploration using analogies. In A. Banerjee and K. Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 1216–1224. PMLR, 4 2021. URL <https://proceedings.mlr.press/v130/roderick21a.html>.
- [43] M. A. Taleghan and T. G. Dietterich. Efficient exploration for constrained mdps. In *2018 AAAI Spring Symposium Series*, 2018.
- [44] G. Thomas, Y. Luo, and T. Ma. Safe reinforcement learning by imagining the near future. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 13859–13869. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/73b277c11266681122132d024f53a75b-Paper.pdf.
- [45] Y. L. Tsai, A. Phatak, P. K. Kitanidis, and C. B. Field. Deep Reinforcement Learning for Disaster Response: Navigating the Dynamic Emergency Vehicle and Rescue Team Dispatch during a Flood. In *AGU Fall Meeting Abstracts*, volume 2019, pages NH33B–14, Dec. 2019.
- [46] S. Vaswani, L. Yang, and C. Szepesvari. Near-optimal sample complexity bounds for constrained mdps. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 3110–3122. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/14a5ebc9cd2e507cd811df78c15bf5d7-Paper-Conference.pdf.
- [47] Y. Wang, S. S. Zhan, R. Jiao, Z. Wang, W. Jin, Z. Yang, Z. Wang, C. Huang, and Q. Zhu. Enforcing hard constraints with soft barriers: Safe reinforcement learning in unknown stochastic environments. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 36593–36604. PMLR, 7 2023. URL <https://proceedings.mlr.press/v202/wang23as.html>.
- [48] H. Wei, X. Liu, and L. Ying. Triple-q: A model-free algorithm for constrained reinforcement learning with sublinear regret and zero constraint violation. In G. Camps-Valls, F. J. R. Ruiz, and I. Valera, editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 3274–3307. PMLR, 3 2022. URL <https://proceedings.mlr.press/v151/wei22a.html>.
- [49] D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, USA, 1st edition, 2011. ISBN 0521195276.

- [50] C. Wu, B. Ju, Y. Wu, X. Lin, N. Xiong, G. Xu, H. Li, and X. Liang. Uav autonomous target search based on deep reinforcement learning in complex disaster scene. *IEEE Access*, 7: 117227–117245, 2019. doi: 10.1109/ACCESS.2019.2933002.
- [51] H. Xu and S. Mannor. Probabilistic goal markov decision processes. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Three*, IJCAI'11, page 2046–2052. AAAI Press, 2011. ISBN 9781577355151.
- [52] S. Yang, S. Khuller, S. Choudhary, S. Mitra, and K. Mahadik. Correlated Stochastic Knapsack with a Submodular Objective. In S. Chechik, G. Navarro, E. Rotenberg, and G. Herman, editors, *30th Annual European Symposium on Algorithms (ESA 2022)*, volume 244 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 91:1–91:14, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-247-1. doi: 10.4230/LIPIcs.ESA.2022.91. URL <https://drops.dagstuhl.de/opus/volltexte/2022/17029>.
- [53] W. Zhao, T. He, R. Chen, T. Wei, and C. Liu. State-wise safe reinforcement learning: a survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, IJCAI '23, 2023. ISBN 978-1-956792-03-4. doi: 10.24963/ijcai.2023/763. URL <https://doi.org/10.24963/ijcai.2023/763>.

A Appendix / supplemental material

B Proofs for Section 2

B.1 Proof of Proposition 1

The proof follows from the standard proof of backward induction [41]. The main ideas for the proof can also be seen in the proof of Lemma 4 and Lemma 5.

B.2 Proof of Proposition 2

Proof.

- (Expectation Constraints) We claim that C_M^π captures expectation constraints. This is immediate as an expectation constraint takes the form $\mathbb{E}_M^\pi \left[\sum_{h=1}^H c_h(s_h, a_h) \right] \leq B$ and by definition $C_M^\pi = \mathbb{E}_M^\pi \left[\sum_{h=1}^H c_h(s_h, a_h) \right]$. Moreover, the standard policy evaluation equations for deterministic policies immediately imply,

$$C_h^\pi(\tau_h) = c_h(s, a) + \sum_{s'} P_h(s' | s, a) C_{h+1}^\pi(\tau_h, a, s'). \quad (\text{EC})$$

Thus, (TR) holds. It is also easy to see that $\sum_{s'} P_h(s' | s, a) C_{h+1}^\pi(\tau_h, a, s')$ can be computed recursively state-wise by,

$$P_h(1 | s, a) C_{h+1}^\pi(\tau_h, a, 1) + \sum_{s'=2}^S P_h(s' | s, a) C_{h+1}^\pi(\tau_h, a, s'), \quad (8)$$

and so (SR) holds. The infinity conditions and non-decreasing requirements are also easy to verify.

- (Almost Sure Constraints) We claim that C_M^π captures almost sure constraints. This is because that for tabular MDPs, $\mathbb{P}_M^\pi[\sum_{h=1}^H c_h(s_h, a_h) \leq B] = 1$ if and only if for all $\tau \in \mathcal{H}_{H+1}$ with $\mathbb{P}_M^\pi[\tau] > 0$ it holds that $\sum_{h=1}^H c_h(s_h, a_h) \leq B$ if and only if $C_M^\pi = \max_{\tau \in \mathcal{H}_{H+1}: \mathbb{P}_M^\pi[\tau] > 0} \sum_{h=1}^H c_h(s_h, a_h) \leq B$.

Let $c(\tau) = \sum_{h=1}^H c_h(s_h, a_h)$ denote the cost of a full history $\tau \in \mathcal{H}_{H+1}$ and let $c_{h:t}(\tau) = \sum_{k=h}^t c_k(s_k, a_k)$ denote the partial cost of τ from time h to time t . Our choice of α and β imply that,

$$C_h^\pi(\tau_h) = c_h(s, a) + \max_{s' \in P_h(s, a)} C_{h+1}^\pi(\tau_h, a, s'). \quad (\text{ASC})$$

To show that C_M^π satisfies (TR), we prove for all $h \in [H + 1]$ and all $\tau_h \in \mathcal{H}_h$ that

$$C_h(\tau_h) = \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h] > 0}} c_{h:H}(\tau). \quad (9)$$

Then, we see that $C_1^\pi(s_0) = \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | s_0] > 0}} c_{1:H}(\tau) = \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau] > 0}} \sum_{h=1}^H c_h(s_h, a_h) = C_M^\pi$. Thus, C_M^π satisfies (TR). Furthermore, it is clear that $\max_{s' \in P_h(s, a)} C_{h+1}^\pi(\tau_h, a, s')$ can be computed state-recursively by,

$$\max(C_{h+1}^\pi(\tau_h, a, 1)[P_h(1 | s, a) > 0], \max_{s'=2}^S C_{h+1}^\pi(\tau_h, a, s')[P_h(s' | s, a) > 0]), \quad (10)$$

and so C_M^π satisfies (SR). The infinity conditions and non-decreasing requirements are also easy to verify.

We proceed by induction on h .

- (Base Case) For the base case, we consider $h = H + 1$. Observe that for any history τ , we have $c_{H+1:H}(\tau) = 0$ since it is an empty sum. Then, by definition of C_M^π , we see that $C_{H+1}^\pi(\tau_{H+1}) = 0 = \max_\tau 0 = \max_\tau c_{H+1:H}(\tau)$.

- (Inductive Step) For the inductive step, we consider $h \leq H$. Let $s = s_h(\tau_h)$ and $a = \pi_h(\tau_h)$. For any $\tau \in \mathcal{H}_{H+1}$ for which $\mathbb{P}_M^\pi[\tau | \tau_h] > 0$, we can decompose its cost by $c_{h:H}(\tau) = c_h(s, a) + c_{h+1:H}(\tau)$. Since a is fixed, we can remove $c_h(s, a)$ from the optimization to get,

$$\max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h] > 0}} c_{h:H}(\tau) = c_h(s, a) + \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h] > 0}} c_{h+1:H}(\tau).$$

Next, we observe by the Markov property that $\mathbb{P}_M^\pi[\tau | \tau_h] = \sum_{s'} \mathbb{P}_M^\pi[\tau | \tau_h, a, s'] P_h(s' | s, a)$. Thus, $\mathbb{P}_M^\pi[\tau | \tau_h] > 0$ if and only if there exists some $s' \in P_h(s, a)$ satisfying $\mathbb{P}_M^\pi[\tau | \tau_h, a, s'] > 0$. This implies that,

$$\max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h] > 0}} c_{h+1:H}(\tau) = \max_{s' \in P_h(s, a)} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h, a, s'] > 0}} c_{h+1:H}(\tau).$$

By applying the induction hypothesis, we see that,

$$\begin{aligned} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h] > 0}} c_{h:H}(\tau) &= c_h(s, a) + \max_{s' \in P_h(s, a)} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h, a, s'] > 0}} c_{h+1:H}(\tau) \\ &= c_h(s, a) + \max_{s' \in P_h(s, a)} C_{h+1}^\pi(\tau_h, a, s') \\ &= C_h(\tau_h). \end{aligned}$$

The second line used the induction hypothesis and the third line used the definition of C_M^π .

3. (Anytime Constraints) We claim that C_M^π captures anytime constraints. This is because that for tabular MDPs, $\mathbb{P}_M^\pi[\forall t \in [H], \sum_{h=1}^t c_h(s_h, a_h) \leq B] = 1$ if and only if for all $t \in [H]$ and $\tau \in \mathcal{H}_{H+1}$ with $\mathbb{P}_M^\pi[\tau] > 0$ it holds that $\sum_{h=1}^t c_h(s_h, a_h) \leq B$ if and only if $C_M^\pi = \max_{t \in [H]} \max_{\tau \in \mathcal{H}_{H+1}: \mathbb{P}_M^\pi[\tau] > 0} \sum_{h=1}^t c_h(s_h, a_h) \leq B$.

Our choice of α and β imply that,

$$C_h^\pi(\tau_h) = c_h(s, a) + \max \left(0, \max_{s' \in P_h(s, a)} C_{h+1}^\pi(\tau_h, a, s') \right). \quad (\text{AC})$$

To show that C_M^π satisfies (TR), we show that for all $h \in [H+1]$ and all $\tau_h \in \mathcal{H}_h$ that

$$C_h(\tau_h) = \max_{t \geq h} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | \tau_h] > 0}} c_{h:t}(\tau). \quad (11)$$

Then, we see that $C_1^\pi(s_0) = \max_{t \in [H]} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau | s_0] > 0}} c_{1:t}(\tau) = \max_{t \in [H]} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau] > 0}}$

$\sum_{h=1}^t c_h(s_h, a_h) = C_M^\pi$. Thus, C_M^π satisfies (TR). Furthermore, it is clear that $\max(0, \max_{s' \in P_h(s, a)} C_{h+1}^\pi(\tau_h, a, s'))$ can be computed state-recursively by,

$$\begin{aligned} &\max \left(\max(0, C_{h+1}^\pi(\tau_h, a, 1)[P_h(1 | s, a) > 0]), \right. \\ &\quad \left. \max(0, \max_{s'=2}^S C_{h+1}^\pi(\tau_h, a, s')[P_h(s' | s, a) > 0]) \right), \end{aligned} \quad (12)$$

and so C_M^π satisfies (SR). The infinity conditions and non-decreasing requirements are also easy to verify.

We proceed by induction on h .

- (Base Case) For the base case, we consider $h = H+1$. Observe that for any history τ and t , we have $c_{H+1:t}(\tau) = 0$ since it is an empty sum. Then, by definition of C_M^π , we see that $C_{H+1}^\pi(\tau_{H+1}) = 0 = \max_t \max_\tau 0 = \max_t \max_\tau c_{H+1:t}(\tau)$ ³.

³Technically, there is no $t \in [H]$ satisfying $t \geq H+1$. We instead interpret the $t \geq h$ condition in the max as over all integers and define the immediate costs to be 0 for all future times to simplify the base case.

- (Inductive Step) For the inductive step, we consider $h \leq H$. Let $s = s_h(\tau_h)$ and $a = \pi_h(\tau_h)$. By separately considering the case where $t = h$ and $t \geq h + 1$ in the $\max_{t \geq h}$, we see that $\max_{t \geq h} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau|\tau_h] > 0}} c_{h:t}(\tau)$

$$\begin{aligned}
&= \max \left(\max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau|\tau_h] > 0}} c_{h:h}(\tau), \max_{t \geq h+1} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau|\tau_h] > 0}} c_{h:t}(\tau) \right) \\
&= \max \left(c_h(s, a), c_h(s, a) + \max_{t \geq h+1} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau|\tau_h] > 0}} c_{h+1:t}(\tau) \right) \\
&= c_h(s, a) + \max \left(0, \max_{t \geq h+1} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau|\tau_h] > 0}} c_{h+1:t}(\tau) \right) \\
&= c_h(s, a) + \max \left(0, \max_{t \geq h+1} \max_{s' \in P_h(s, a)} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau|\tau_h, a, s'] > 0}} c_{h+1:t}(\tau) \right) \\
&= c_h(s, a) + \max \left(0, \max_{s' \in P_h(s, a)} \max_{t \geq h+1} \max_{\substack{\tau \in \mathcal{H}_{H+1}: \\ \mathbb{P}_M^\pi[\tau|\tau_h, a, s'] > 0}} c_{h+1:t}(\tau) \right) \\
&= c_h(s, a) + \max \left(0, \max_{s' \in P_h(s, a)} C_{h+1}^\pi(\tau_h, a, s') \right) \\
&= C_h(\tau_h).
\end{aligned}$$

The second line used the fact that $c_{h:h}(\tau) = c_h(s, a)$ and the recursive definition of $c_{h:t}(\tau)$. The fourth line used the result proven for the almost sure case above. The sixth line used the induction hypothesis. The last line used the definition of C_M^π .

□

C Proofs for Section 3

C.1 Helpful Technical Lemmas

Here, we use a different, inductive definition for \mathcal{V} then in the main text. However, the following lemma shows they are equivalent.

Definition 9 (Value Space). For any $s \in \mathcal{S}$, we define $\mathcal{V}_{H+1}(s) \stackrel{\text{def}}{=} \{0\}$, and for any $h \in [H]$,

$$\mathcal{V}_h(s) \stackrel{\text{def}}{=} \bigcup_a \bigcup_{v \in \times_{s'} \mathcal{V}_{h+1}(s')} \left\{ r_h(s, a) + \sum_{s'} P_h(s' | s, a) v_{s'} \right\}. \quad (13)$$

We define $\mathcal{V} \stackrel{\text{def}}{=} \bigcup_{h,s} \mathcal{V}_h(s)$.

Lemma 3 (Value Intution). For all $s \in \mathcal{S}$ and $h \in [H + 1]$,

$$\mathcal{V}_h(s) = \{v \in \mathbb{R} \mid \exists \pi \in \Pi^D, \tau_h \in \mathcal{H}_h, (s = s_h(\tau_h) \wedge V_h^\pi(\tau_h) = v)\}, \quad (14)$$

and $|\mathcal{V}_h(s)| \leq A^{\sum_{t=h}^H S^{H-t}}$. Thus, \mathcal{V} can be computed in finite time using backward induction.

Lemma 4 (Cost). For any $h \in [H + 1]$, $\tau_h \in \mathcal{H}_h$, and $v \in \mathcal{V}$, if $s = s_h(\tau_h)$, then,

$$\begin{aligned}
\bar{C}_h^*(s, v) &\leq \min_{\pi \in \Pi^D} C_h^\pi(\tau_h) \\
&\text{s.t. } V_h^\pi(\tau_h) \geq v.
\end{aligned} \quad (15)$$

Lemma 5 (Value). *Suppose that $\pi \in \Pi^D$. For all $h \in [H + 1]$ and $(s, v) \in \bar{\mathcal{S}}$, if $\bar{C}_h^\pi(s, v) < \infty$, then $\bar{V}_h^\pi(s, v) \geq v$.*

Remark 6 (Technical Subtlety). Technically, $V_h^\pi(\tau_h)$ is only well defined if $\mathbb{P}_M^\pi[\tau_h] > 0$ and all of our arguments technically should assume this is the case. However, it is standard in MDP theory to define the policy evaluation equations on non-reachable trajectories using the standard recursion to simplify proofs, as we have done here. Formally, this is equivalent to assuming the process starts initially at τ_h instead of just conditioning on reaching τ_h , or defining the values to correspond to policy evaluation equations directly. This is consistent with the usual definition when $\mathbb{P}_M^\pi[\tau_h] > 0$ but gives it a defined value also when $\mathbb{P}_M^\pi[\tau_h] = 0$. In either case, this detail only means our recursive definition of \mathcal{V} is a superset rather than exactly the set of all values as we defined in the main text. This does not effect the final results since unreachable trajectories do not effect π 's overall value in the MDP anyway, and only effects the interpretations of some intermediate variables.

C.2 Proof of Proposition 3

Proof. By definition of V_M^* and C_M^* ,

$$\begin{aligned} V_M^* > -\infty &\iff \exists \pi \in \Pi^D, C_M^\pi \leq B \wedge V_M^\pi \geq V_M^* \\ &\iff C_M^* \leq B. \end{aligned}$$

For the second claim, we observe that if $V_M^* > -\infty$ then by the above argument any optimal deterministic policy π for COVER satisfies $C_M^\pi = C_M^* \leq B$ and $V_M^\pi \geq V_M^*$. Thus, $COVER \subseteq PACK$. \square

C.3 Proof of Lemma 3

Proof. We proceed by induction on h . Let $s \in \mathcal{S}$ be arbitrary.

Base Case. For the base case, we consider $h = H + 1$. In this case, we know that for any $\pi \in \Pi^D$ and any $\tau \in \mathcal{H}_{H+1}$, $V_{H+1}^\pi(\tau_{H+1}) = 0 \in \{0\} = \mathcal{V}_{H+1}(s)$ by definition. Furthermore, $|\mathcal{V}_{H+1}(s)| = 1 = A^0 = A^{\sum_{t=H+1}^H S^t}$.

Inductive Step. For the inductive step, we consider $h \leq H$. In this case, we know that for any $\pi \in \Pi^D$ and any $\tau_h \in \mathcal{H}_h$, if $s = s_h(\tau_h)$ and $a = \pi_h(\tau_h)$, then the policy evaluation equations imply,

$$V_h^\pi(\tau_h) = r_h(s, a) + \sum_{s'} P_h(s' | s, a) V_{h+1}^\pi(\tau_h, a, s').$$

We know by the induction hypothesis that $V_{h+1}^\pi(\tau_h, a, s') \in \mathcal{V}_{h+1}(s')$. Thus, by (13), $V_h^\pi(\tau_h) \in \mathcal{V}_h(s)$. Lastly, we see by (13) and the induction hypothesis that,

$$|\mathcal{V}_h(s)| \leq A \prod_{s'} |\mathcal{V}_{h+1}(s')| \leq A \prod_{s'} A^{\sum_{t=h+1}^H S^{H-t}} = A^{1+S \sum_{t=h+1}^H S^{H-t}} = A^{\sum_{t=h}^H S^{H-t}}.$$

This completes the proof. \square

C.4 Proof of Lemma 4

Proof. We proceed by induction on h . Let $\tau_h \in \mathcal{H}_h$ and $v \in \mathcal{V}$ be arbitrary and suppose that $s = s_h(\tau_h)$. We let $C_h^*(\tau_h, v)$ denote the minimum for the RHS of (15).

Base Case. For the base case, we consider $h = H + 1$. Observe that for any $\pi \in \Pi^D$, $V_{H+1}^\pi(\tau_{H+1}) = 0$ by definition. Thus, there exists a $\pi \in \Pi^D$ satisfying $V_{H+1}^\pi(\tau_{H+1}) \geq v$ if and only if $v \leq 0$. We also know by definition that any such policy π satisfies $C_{H+1}^\pi(\tau_{H+1}) = 0$ and if no such policy exists $C_{H+1}^*(\tau_{H+1}, v) = \infty$ by convention. Therefore, we see that $C_{H+1}^*(\tau_{H+1}, v) = \chi_{\{v \leq 0\}}$. Then, by definition of the base case for \bar{C} , it follows that,

$$\bar{C}_{H+1}^*(s, v) = \chi_{\{v \leq 0\}} = C_{H+1}^*(\tau_{H+1}, v).$$

Inductive Step. For the inductive step, we consider $h \leq H$. If $C_h^*(\tau_h, v) = \infty$, then trivially $\bar{C}_h^*(s, v) \leq C_h^*(\tau_h, v)$. Instead, suppose that $C_h^*(\tau_h, v) < \infty$. Then, there must exist a feasible $\pi \in \Pi^D$ satisfying $V_h^\pi(\tau_h) \geq v$. Let $a^* = \pi_h(\tau_h)$. By the policy evaluation equations, we know that,

$$V_h^\pi(\tau_h) = r_h(s, a^*) + \sum_{s'} P_h(s' | s, a^*) V_{h+1}^\pi(\tau_h, a^*, s').$$

For each $s' \in \mathcal{S}$, define $v_{s'}^* \stackrel{\text{def}}{=} V_{h+1}^\pi(\tau_h, a^*, s')$ and observe that $v_{s'}^* \in \mathcal{V}_{h+1}(s') \subseteq \mathcal{V}$ by **Lemma 3**. Thus, we see that $(a^*, \mathbf{v}^*) \in \mathcal{A} \times \mathcal{V}^{\mathcal{S}}$ and $r_h(s, a) + \sum_{s'} P_h(s' | s, a) v_{s'} \geq v$, which implies $(a^*, \mathbf{v}^*) \in \bar{\mathcal{A}}_h(s, v)$.

Since π satisfies $V_{h+1}^\pi(\tau_h, a^*, s') \geq v_{s'}^*$, it is clear that $C_{h+1}^*(s', v_{s'}^*) \leq C_{h+1}^\pi(\tau_h, a^*, s')$. Thus, the induction hypothesis implies that $\bar{C}_{h+1}^*(s', v_{s'}^*) \leq C_{h+1}^*(s', v_{s'}^*) \leq C_{h+1}^\pi(\tau_h, a^*, s')$. The optimality equations for \bar{M} then imply that,

$$\begin{aligned} \bar{C}_h^*(s, v) &= \min_{(a, \mathbf{v}) \in \bar{\mathcal{A}}_h(s, v)} c_h(s, a) + f\left(\left(P_h(s' | s, a), \bar{C}_{h+1}^*(s', v_{s'}^*)\right)_{s' \in P_h(s, a)}\right) \\ &\leq c_h(s, a^*) + f\left(\left(P_h(s' | s, a^*), \bar{C}_{h+1}^*(s', v_{s'}^*)\right)_{s' \in P_h(s, a^*)}\right) \\ &\leq c_h(s, a^*) + f\left(\left(P_h(s' | s, a), C_{h+1}^\pi(\tau_h, a^*, s')\right)_{s' \in P_h(s, a^*)}\right) \\ &= C_h^\pi(\tau_h). \end{aligned}$$

The first inequality used the fact that $(a^*, \mathbf{v}^*) \in \bar{\mathcal{A}}_h(s, v)$. The second inequality relied on f being non-decreasing and the induction hypothesis. The final equality used **(TR)**.

Since π was an arbitrary feasible policy for the optimization defining $C_h^*(\tau_h, v)$, we see that $\bar{C}_h^*(s, v) \leq C_h^*(\tau_h, v)$. This completes the proof. \square

C.5 Proof of **Lemma 5**

Proof. We proceed by induction on h . Let $(s, v) \in \bar{\mathcal{S}}$ be arbitrary.

Base Case. For the base case, we consider $h = H+1$. By definition and assumption, $\bar{C}_{H+1}^\pi(s, v) = \chi_{\{v \leq 0\}} < \infty$. Thus, it must be the case that $v \leq 0$ and so by definition $\bar{V}_{H+1}^\pi(s, v) = 0 \geq v$.

Inductive Step. For the inductive step, we consider $h \leq H$. We decompose $\pi_h(s, v) = (a, \mathbf{v})$ where we know $(a, \mathbf{v}) \in \bar{\mathcal{A}}_h(s, v)$ since π has finite cost⁴. Moreover, it must be the case that for any $s' \in \mathcal{S}$ with $P_h(s' | s, a) > 0$ that $\bar{C}_{h+1}^\pi(s', v_{s'}) < \infty$ otherwise the property that f outputs ∞ when inputted an ∞ would imply a contradiction:

$$\begin{aligned} \bar{C}_h^\pi(s, v) &= c_h(s, a) + f\left(\left(P_h(s' | s, a), \bar{C}_{h+1}^\pi(s', v_{s'})\right)_{s' \in P_h(s, a)}\right) \\ &= c_h(s, a) + f(\dots, \infty, \dots) \\ &= \infty. \end{aligned}$$

Thus, the induction hypothesis implies that $\bar{V}_{h+1}^\pi(s', v_{s'}) \geq v_{s'}$ for any such $s' \in \mathcal{S}$. By the policy evaluation equations, we see that,

$$\begin{aligned} \bar{V}_h^\pi(s, v) &= r_h(s, a) + \sum_{s'} P_h(s' | s, a) \bar{V}_{h+1}^\pi(s', v_{s'}) \\ &\geq r_h(s, a) + \sum_{s'} P_h(s' | s, a) v_{s'} \\ &\geq v. \end{aligned}$$

The third line uses the definition of $\bar{\mathcal{A}}_h(s, v)$. This completes the proof. \square

⁴By convention, we assume $\min \emptyset = \infty$

C.6 Proof of Theorem 1

Proof. If $\bar{C}_1^*(s_0, v) > B$ for all $v \in \mathcal{V}$, then $C_M^* > B$ since otherwise we would have $\bar{C}_1^*(s_0, v) \leq C_1^*(s_0, v) = C_M^* \leq B$ by Lemma 4. Thus, if Algorithm 1 outputs “infeasible” it is correct.

On the other hand, suppose that there exists some $v \in \mathcal{V}$ for which $\bar{C}_1^*(s_0, v) \leq B$. By standard MDP theory, we know that since $\pi \in \Pi^D$ is a solution to \bar{M} , it must satisfy the optimality equations. In particular, $\bar{C}_1^\pi(s_0, v) = \bar{C}_1^*(s_0, v) \leq B$. Since $C_M^\pi = \bar{C}_1^\pi(s_0, v)$ ⁵, we see that there exists a $\pi \in \Pi^D$ for which $C_M^\pi \leq B$ and so $V_M^* > -\infty$.

Since V_M^* is the value of some deterministic policy, Lemma 3 implies that $V_M^* \in \mathcal{V}$. Thus, Lemma 5 implies that $V_1^\pi(s_0, V_M^*) \geq V_M^*$ and $C_1^\pi(s_0, V_M^*) \leq C_1^*(s_0, V_M^*) \leq B$. Consequently, running π with initial state $\bar{s}_0 = (s_0, V_M^*)$ is an optimal solution to (CON). In either case, Algorithm 1 is correct. \square

D Proofs for Section 4

Definition 10. We define the exact partial sum,

$$\sigma_{h,\mathbf{v}}^{s,a}(t, u) \stackrel{\text{def}}{=} u + \sum_{s'=t}^S P_h(s' | s, a) v_{s'}. \quad (16)$$

Observation 1. We observe that both σ and $\hat{\sigma}$ can be computed recursively. Specifically, $\sigma_{h,\mathbf{v}}^{s,a}(S+1, u) = u$ and $\sigma_{h,\mathbf{v}}^{s,a}(t, u) = \sigma_{h,\mathbf{v}}^{s,a}(t, u + P_h(t | s, a)v_t)$. Similarly, $\hat{\sigma}_{h,\mathbf{v}}^{s,a}(S+1, u) = u$ and $\hat{\sigma}_{h,\mathbf{v}}^{s,a}(t, u) = \sigma_{h,\mathbf{v}}^{s,a}(t, \lfloor u + P_h(t | s, a)v_t \rfloor_{\mathcal{G}})$.

For completeness, and to assist with other arguments, we also prove the exact recursion we presented in Definition 3 is correct using Lemma 6.

Lemma 6. For any $t \in [S+1]$ and $u \in \mathbb{R}$, we have that,

$$g_{h,v}^{s,a}(t, u) = \min_{\mathbf{v} \in \mathcal{V}^{S-t+1}} g_{h,\mathbf{v}}^{s,a}(t) \quad (17)$$

$$\text{s.t. } u + \sum_{s'=t}^S P_h(s' | s, a) v_{s'} \geq v.$$

Moreover, $\bar{C}_h^*(s, v) = \min_{a \in \mathcal{A}} c_h(s, a) + g_{h,v}^{s,a}(1, r_h(s, a))$.

D.1 Proof of Lemma 6

Proof. We proceed by induction on t .

Base Case. For the base case, we consider $t = S+1$. Since $\sum_{s'=S+1}^S P_h(s' | s, a) v_{s'} = 0$ is the empty sum, the condition $u + \sum_{s'=S+1}^S P_h(s' | s, a) v_{s'} \geq v$ is true iff $u \geq v$. Also, for any \mathbf{v} , $g_{h,\mathbf{v}}^{s,a}(S+1) = 0$ by definition. Thus, the minimum defining $g_{h,\mathbf{v}}^{s,a}(S+1, u)$ is 0 when $u \geq v$ and is ∞ due to infeasibility otherwise. In symbols, $g_{h,v}^{s,a}(S+1, u) = \chi_{\{u \geq v\}}$ as was to be shown.

⁵We can view $\bar{C}(\bar{V})$ as the extension of $C(V)$ needed to formally evaluate memory-augmented policies. Since we consider deterministic policies, it is trivial to convert any memory-augmented policy into a history-dependent policy that is defined in the original environment M .

Inductive Step. For the inductive step, we consider $t \leq S$. We see that $g_{h,v}^{s,a}(t, u)$

$$\begin{aligned}
&= \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t+1} \\ u + \sum_{s'=t}^S P_h(s'|s, a) v_{s'} \geq v}} g_{h,\mathbf{v}}^{s,a}(t) \\
&= \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t+1} \\ u + \sum_{s'=t}^S P_h(s'|s, a) v_{s'} \geq v}} \alpha \left(\beta \left(P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t) \right), g_{h,\mathbf{v}}^{s,a}(t+1) \right) \\
&= \min_{v_t \in \mathcal{V}} \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t} \\ (u + P_h(t|s, a)v_t) + \sum_{s'=t+1}^S P_h(s'|s, a)v_{s'} \geq v}} \alpha \left(\beta \left(P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t) \right), g_{h,\mathbf{v}}^{s,a}(t+1) \right) \\
&= \min_{v_t \in \mathcal{V}} \alpha \left(\beta \left(P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t) \right), \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t} \\ (u + P_h(t|s, a)v_t) + \sum_{s'=t+1}^S P_h(s'|s, a)v_{s'} \geq v}} g_{h,\mathbf{v}}^{s,a}(t+1) \right) \\
&= \min_{v_t \in \mathcal{V}} \alpha \left(\beta \left(P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t) \right), g_{h,v}^{s,a}(t+1, u + P_h(t | s, a)v_t) \right)
\end{aligned}$$

The second lined used **(SR)**. The third line split the optimization into the first decision and the remaining decisions and decomposed the sum in the constraint. The fourth line used the fact that α is a non-decreasing function of both its arguments and the fact that the second optimization only concerns the second argument. The last line used the induction hypothesis.

The observation that $\min_{a \in \mathcal{A}} c_h(s, a) + g_{h,v}^{s,a}(1, r_h(s, a)) = \bar{C}_h^*(s, v)$ then follows from the definition of $\bar{\mathcal{A}}_h(s, v)$ and **(BU)**:

$$\begin{aligned}
\min_{a \in \mathcal{A}} c_h(s, a) + g_{h,v}^{s,a}(1, r_h(s, a)) &= \min_{a \in \mathcal{A}} c_h(s, a) + \min_{\substack{\mathbf{v} \in \mathcal{V}^S \\ r_h(s, a) + \sum_{s'} P_h(s'|s, a)v_s \geq v}} g_{h,\mathbf{v}}^{s,a}(1) \\
&= \min_{a \in \mathcal{A}} \min_{\substack{\mathbf{v} \in \mathcal{V}^S \\ r_h(s, a) + \sum_{s'} P_h(s'|s, a)v_s \geq v}} c_h(s, a) + g_{h,\mathbf{v}}^{s,a}(1) \\
&= \min_{(a, \mathbf{v}) \in \bar{\mathcal{A}}_h(s, v)} c_h(s, a) + g_{h,\mathbf{v}}^{s,a}(1) \\
&= \bar{C}_h^*(s, v).
\end{aligned}$$

□

D.2 Proof of Lemma 1

Proof. We proceed by induction on t .

Base Case. For the base case, we consider $t = S + 1$. By definition, $\hat{\sigma}_{h,\hat{\mathbf{v}}}^{s,a}(S + 1, u) = u$ so the constraint is satisfied iff $u \geq v$. Since for any $\hat{\mathbf{v}}$, $\hat{g}_{h,\hat{\mathbf{v}}}^{s,a}(S + 1) = 0$ by definition, the minimum defining $\hat{g}_{h,\hat{\mathbf{v}}}^{s,a}(S + 1, u)$ is 0 when $u \geq v$ and is ∞ due to infeasibility otherwise. In symbols, $\hat{g}_{h,v}^{s,a}(S + 1, u) = \chi_{\{u \geq v\}}$ as was to be shown.

Inductive Step. For the inductive step, we consider $t \leq S$. We see that,

$$\begin{aligned}
\hat{g}_{h,v}^{s,a}(t,u) &= \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t+1} \\ \hat{\sigma}_{h,\mathbf{v}}^{s,a}(t,u) \geq v}} g_{h,\mathbf{v}}^{s,a}(t) \\
&= \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t+1} \\ \hat{\sigma}_{h,\mathbf{v}}^{s,a}(t,u) \geq v}} \alpha \left(\beta \left(P_h(t \mid s, a), \bar{C}_{h+1}^*(t, v_t) \right), g_{h,\mathbf{v}}^{s,a}(t+1) \right) \\
&= \min_{v_t \in \mathcal{V}} \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t} \\ \hat{\sigma}_{h,\mathbf{v}}^{s,a}(t+1, \lfloor u + P_h(t \mid s, a) v_t \rfloor_{\mathcal{G}}) \geq v}} \alpha \left(\beta \left(P_h(t \mid s, a), \bar{C}_{h+1}^*(t, v_t) \right), g_{h,\mathbf{v}}^{s,a}(t+1) \right) \\
&= \min_{v_t \in \mathcal{V}} \alpha \left(\beta \left(P_h(t \mid s, a), \bar{C}_{h+1}^*(t, v_t) \right), \min_{\substack{\mathbf{v} \in \mathcal{V}^{S-t} \\ \hat{\sigma}_{h,\mathbf{v}}^{s,a}(t+1, \lfloor u + P_h(t \mid s, a) v_t \rfloor_{\mathcal{G}}) \geq v}} g_{h,\mathbf{v}}^{s,a}(t+1) \right) \\
&= \min_{v_t \in \mathcal{V}} \alpha \left(\beta \left(P_h(t \mid s, a), \bar{C}_{h+1}^*(t, v_t) \right), \hat{g}_{h,v}^{s,a}(t+1, \lfloor u + P_h(t \mid s, a) v_t \rfloor_{\mathcal{G}}) \right)
\end{aligned}$$

The second lined used (SR). The third line split the optimization into the first decision and the remaining decisions and used the recursive definition of $\hat{\sigma}$ in the constraint. The fourth line used the fact that α is a non-decreasing function of both its arguments and the fact that the second optimization only concerns the second argument. The last line used the induction hypothesis. \square

D.3 Proof of Proposition 4

Proof. The runtime guarantee is easily seen since Algorithm 4 consists of nested loops. The fact that it computes an optimal solution for M absent rounding or lower bounding follows immediately from Lemma 6. \square

E Proofs for Section 5

E.1 Helpful Technical Lemmas (Additive)

The following claims all assume Definition 7.

Observation 2. For any $v \in \mathbb{R}$,

$$v - \delta \leq \lfloor v \rfloor_{\mathcal{G}} \leq v. \quad (18)$$

Lemma 7. For any $h \in [H]$, $s \in \mathcal{S}$, $a \in \mathcal{A}$, $\mathbf{v} \in \mathbb{R}^S$, $u \in \mathbb{R}$, and $t \in [S+1]$, we have,

$$\sigma_{h,\mathbf{v}}^{s,a}(t,u) - (S-t+1)\delta \leq \hat{\sigma}_{h,\mathbf{v}}^{s,a}(t,u) \leq \sigma_{h,\mathbf{v}}^{s,a}(t,u). \quad (19)$$

Lemma 8 (Cost). For any $h \in [H+1]$ and $(s, v) \in \bar{\mathcal{S}}$, $\hat{C}_h^*(s, \lfloor v \rfloor_{\mathcal{G}}) \leq \bar{C}_h^*(s, v)$.

Lemma 9 (Approximation). Suppose that $\pi \in \Pi^D$. For all $h \in [H+1]$ and $(s, \hat{v}) \in \hat{\mathcal{S}}$, if $\hat{C}_h^\pi(s, \hat{v}) < \infty$, then $\hat{V}_h^\pi(s, \hat{v}) \geq \hat{v} - \delta(S+1)(H-h+1)$.

E.2 Helpful Technical Lemmas (Relative)

The following claims all assume Definition 8.

Observation 3. For any $v \in \mathbb{R}$,

$$v(1-\delta) \leq \lfloor v \rfloor_{\mathcal{G}} \leq v. \quad (20)$$

Lemma 10. For any $h \in [H]$, $s \in \mathcal{S}$, $a \in \mathcal{A}$, $\mathbf{v} \in \mathbb{R}_{\geq 0}^S$, $u \in \mathbb{R}_{\geq 0}$, and $t \in [S+1]$, we have,

$$\sigma_{h,\mathbf{v}}^{s,a}(t,u)(1-\delta)^{S-t+1} \leq \hat{\sigma}_{h,\mathbf{v}}^{s,a}(t,u) \leq \sigma_{h,\mathbf{v}}^{s,a}(t,u). \quad (21)$$

Lemma 11 (Cost). Suppose all rewards are non-negative. For any $h \in [H+1]$ and $(s, v) \in \bar{\mathcal{S}}$, $\hat{C}_h^*(s, \lfloor v \rfloor_{\mathcal{G}}) \leq \bar{C}_h^*(s, v)$.

Lemma 12 (Approximation). Suppose all rewards are non-negative and $\pi \in \Pi^D$. For all $h \in [H+1]$ and $(s, \hat{v}) \in \hat{\mathcal{S}}$, if $\hat{C}_h^\pi(s, \hat{v}) < \infty$, then $\hat{V}_h^\pi(s, \hat{v}) \geq \hat{v}(1-\delta)^{(S+1)(H-h+1)}$.

E.3 Proof of Observation 2

Proof. Using properties of the floor function, we can infer that,

$$\lfloor v \rfloor_{\mathcal{G}} = \left\lfloor \frac{v}{\delta} \right\rfloor \delta \leq \frac{v}{\delta} \delta = v,$$

and,

$$\lfloor v \rfloor_{\mathcal{G}} = \left\lfloor \frac{v}{\delta} \right\rfloor \delta \geq \left(\left\lceil \frac{v}{\delta} \right\rceil - 1 \right) \delta = \left\lceil \frac{v}{\delta} \right\rceil \delta - \delta \geq v - \delta.$$

□

E.4 Proof of Lemma 7

Proof. We proceed by induction on t .

Base Case. For the base case, we consider $t = S + 1$. By definition, we have $\hat{\sigma}_{h,\hat{v}}^{s,a}(S + 1, u) = u = \sigma_{h,\hat{v}}^{s,a}(S + 1, u)$.

Inductive Step. For the inductive step, we consider $t \leq S$. We first see that,

$$\begin{aligned} \hat{\sigma}_{h,\hat{v}}^{s,a}(t, u) &= \hat{\sigma}_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) \\ &\leq \sigma_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) \\ &= \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}} + \sum_{s'=t+1}^S P_h(s' | s, a) \hat{v}_t \\ &\leq u + \sum_{s'=t}^S P_h(s' | s, a) \hat{v}_t \\ &= \sigma_{h,\hat{v}}^{s,a}(t, u). \end{aligned}$$

The first inequality used the induction hypothesis and the second inequality used the fact that $\lfloor x \rfloor_{\mathcal{G}} \leq x$.

We also see that,

$$\begin{aligned} \hat{\sigma}_{h,\hat{v}}^{s,a}(t, u) &= \hat{\sigma}_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) \\ &\geq \sigma_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) - \delta(S - t) \\ &= \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}} + \sum_{s'=t+1}^S P_h(s' | s, a) \hat{v}_t - \delta(S - t) \\ &\geq u + \sum_{s'=t}^S P_h(s' | s, a) \hat{v}_t - \delta(S - t + 1) \\ &= \sigma_{h,\hat{v}}^{s,a}(t, u) - \delta(S - t + 1). \end{aligned}$$

The first inequality used the induction hypothesis and the second inequality used the fact that $\lfloor x \rfloor_{\mathcal{G}} \geq x - \delta$. □

E.5 Proof of Lemma 8

Proof. We proceed by induction on h . Let $(s, v) \in \bar{\mathcal{S}}$ be arbitrary.

Base Case. For the base case, we consider $h = H + 1$. Since $\lfloor v \rfloor_{\mathcal{G}} \leq v$, we immediately see,

$$\hat{C}_{H+1}^*(s, \lfloor v \rfloor_{\mathcal{G}}) = \chi_{\{\lfloor v \rfloor_{\mathcal{G}} \leq 0\}} \leq \chi_{\{v \leq 0\}} = \bar{C}_{H+1}^*(s, v).$$

Inductive Step. For the inductive step, we consider $h \leq H$. If $\bar{C}_h^*(s, v) = \infty$, then trivially $\hat{C}_h^*(s, \lfloor v \rfloor_{\mathcal{G}}) \leq \bar{C}_h^*(s, v)$. Instead, suppose that $\bar{C}_h^*(s, v) < \infty$. Let π be a solution to the optimality equations for \bar{M} so that $\bar{C}_h^\pi(s, v) = \bar{C}_h^*(s, v) < \infty$. Since $\bar{C}_h^*(s, v) < \infty$, we know that $(a^*, \mathbf{v}^*) = \pi_h(s, v) \in \hat{\mathcal{A}}_h(s, v)$. By the definition of $\hat{\mathcal{A}}_h(s, v)$, we know that,

$$\sigma_{h, \mathbf{v}^*}^{s, a^*}(1, r_h(s, a^*)) = r_h(s, a^*) + \sum_{s'} P_h(s' | s, a^*) v_{s'}^* \geq v \geq \lfloor v \rfloor_{\mathcal{G}}.$$

For each $s' \in \mathcal{S}$, define $\hat{v}_{s'}^* \stackrel{\text{def}}{=} \lfloor v_{s'}^* \rfloor_{\mathcal{G}}$ and recall that $v_{s'}^* \in \mathcal{V}$. We first observe that,

$$\begin{aligned} \sigma_{h, \hat{\mathbf{v}}^*}^{s, a^*}(1, r_h(s, a^*)) &= r_h(s, a^*) + \sum_{s'} P_h(s' | s, a) \lfloor v_{s'}^* \rfloor_{\mathcal{G}} \\ &\geq r_h(s, a^*) + \sum_{s'} P_h(s' | s, a) (v_{s'} - \delta) \\ &= r_h(s, a^*) + \sum_{s'} P_h(s' | s, a) v_{s'} - \delta \\ &= \sigma_{h, \mathbf{v}^*}^{s, a^*}(1, r_h(s, a^*)) - \delta. \end{aligned}$$

Then by [Lemma 7](#),

$$\begin{aligned} \hat{\sigma}_{h, \hat{\mathbf{v}}^*}^{s, a^*}(1, r_h(s, a^*)) &\geq \sigma_{h, \hat{\mathbf{v}}^*}^{s, a^*}(1, r_h(s, a^*)) - \delta S \\ &\geq \sigma_{h, \mathbf{v}^*}^{s, a^*}(1, r_h(s, a^*)) - \delta(S + 1) \\ &\geq \lfloor v \rfloor_{\mathcal{G}} - \delta(S + 1) \\ &= \kappa(\lfloor v \rfloor_{\mathcal{G}}). \end{aligned}$$

Thus, $(a^*, \hat{\mathbf{v}}^*) \in \hat{\mathcal{A}}_h(s, \lfloor v \rfloor_{\mathcal{G}})$.

Since $v_{s'}^* \in \mathcal{V}$, the induction hypothesis implies that $\hat{C}_{h+1}^*(s', \hat{v}_{s'}^*) \leq \bar{C}_{h+1}^*(s', v_{s'}^*) = \bar{C}_{h+1}^\pi(s', v_{s'}^*)$. The optimality equations for \hat{M} then imply that,

$$\begin{aligned} \hat{C}_h^*(s, \lfloor v \rfloor_{\mathcal{G}}) &= \min_{(a, \hat{\mathbf{v}}) \in \hat{\mathcal{A}}_h(s, v)} c_h(s, a) + f \left(\left(P_h(s' | s, a), \hat{C}_{h+1}^*(s', \hat{v}_{s'}^*) \right)_{s' \in P_h(s, a)} \right) \\ &\leq c_h(s, a^*) + f \left(\left(P_h(s' | s, a^*), \hat{C}_{h+1}^*(s', \hat{v}_{s'}^*) \right)_{s' \in P_h(s, a^*)} \right) \\ &\leq c_h(s, a^*) + f \left(\left(P_h(s' | s, a), \bar{C}_{h+1}^\pi(s', v_{s'}^*) \right)_{s' \in P_h(s, a^*)} \right) \\ &= \bar{C}_h^\pi(s, v) \\ &= \bar{C}_h^*(s, v). \end{aligned}$$

The first inequality used the fact that $(a^*, \mathbf{v}^*) \in \hat{\mathcal{A}}_h(s, v)$. The second inequality relied on f being non-decreasing and the induction hypothesis. The penultimate equality used [\(TR\)](#). This completes the proof. \square

E.6 Proof of [Lemma 9](#)

Proof. We proceed by induction on h . Let $(s, \hat{v}) \in \hat{\mathcal{S}}$ be arbitrary.

Base Case. For the base case, we consider $h = H + 1$. By definition and assumption, $\hat{C}_{H+1}^\pi(s, \hat{v}) = \chi_{\{\hat{v} \leq 0\}} < \infty$. Thus, it must be the case that $\hat{v} \leq 0$ and so by definition $\hat{V}_{H+1}^\pi(s, \hat{v}) = 0 \geq \hat{v}$.

Inductive Step. For the inductive step, we consider $h \leq H$. As in the proof of [Lemma 5](#), we know that $\pi_h(s, v) = (a, \hat{\mathbf{v}}) \in \hat{\mathcal{A}}_h(s, \hat{v})$ and for any $s' \in \mathcal{S}$ with $P_h(s' | s, a) > 0$ that $\hat{C}_{h+1}^\pi(s', v_{s'}) < \infty$. Thus, the induction hypothesis implies that $\hat{V}_{h+1}^\pi(s', \hat{v}_{s'}) \geq \hat{v}_{s'} - \delta(S + 1)(H - h)$ for any such

$s' \in \mathcal{S}$. By the policy evaluation equations, we see that,

$$\begin{aligned}
\hat{V}_h^\pi(s, \hat{v}) &= r_h(s, a) + \sum_{s'} P_h(s' | s, a) \hat{V}_{h+1}^\pi(s', \hat{v}_{s'}) \\
&\geq r_h(s, a) + \sum_{s'} P_h(s' | s, a) \hat{v}_{s'} - \delta(S+1)(H-h) \\
&= \sigma_{h, \hat{v}}^{s, a}(1, r_h(s, a)) - \delta(S+1)(H-h) \\
&\geq \hat{\sigma}_{h, \hat{v}}^{s, a}(1, r_h(s, a)) - \delta(S+1)(H-h) \\
&\geq \hat{v} - \delta(S+1) - \delta(S+1)(H-h) \\
&= \hat{v} - \delta(S+1)(H-h+1).
\end{aligned}$$

The first inequality used the induction hypothesis. The second inequality used [Lemma 7](#). The third inequality used the fact that by definition of $\hat{\mathcal{A}}_h(s, \hat{v})$ and κ , $\hat{\sigma}_{h, \hat{v}}^{s, a}(1, r_h(s, a)) \geq \kappa(\hat{v}) = \hat{v} - \delta(S+1)$. This completes the proof. \square

E.7 Proof of [Theorem 2](#)

Proof.

Correctness. If $\hat{C}_1^*(s_0, v) > B$ for all $\hat{v} \in \hat{\mathcal{V}}$, then $C_M^* > B$ since otherwise we would have $\hat{C}_1^*(s_0, \lfloor v \rfloor_{\mathcal{G}}) \leq \bar{C}_1^*(s_0, v) \leq C_M^* \leq B$ by [Lemma 8](#). Thus, if [Algorithm 5](#) outputs ‘‘infeasible’’ it is correct.

On the other hand, suppose that there exists some $\hat{v} \in \hat{\mathcal{V}}$ for which $\hat{C}_1^*(s_0, \hat{v}) \leq B$. By standard MDP theory, we know that since $\pi \in \Pi^D$ is a solution to \hat{M} , it must satisfy the optimality equations. In particular, $\hat{C}_1^\pi(s_0, \hat{v}) = \hat{C}_1^*(s_0, \hat{v}) \leq B$. As in the proof of [Theorem 1](#), since $C_M^\pi = \hat{C}_1^\pi(s_0, \hat{v})$, we see that there exists a $\pi \in \Pi^D$ for which $C_M^\pi \leq B$ and so $V_M^* > -\infty$.

Since V_M^* is the value of some deterministic policy, [Lemma 3](#) implies that $V_M^* \in \mathcal{V}$. Thus, [Lemma 9](#) implies that $\hat{V}_1^\pi(s_0, \lfloor V_M^* \rfloor_{\mathcal{G}}) \geq \lfloor V_M^* \rfloor_{\mathcal{G}} - \delta(S+1)H \geq V_M^* - \delta(1 + (S+1)H) = V_M^* - \epsilon$ and $\hat{C}_1^\pi(s_0, V_M^*) \leq C_1^*(s_0, V_M^*) \leq B$. Consequently, running π with initial state $\bar{s}_0 = (s_0, \lfloor V_M^* \rfloor_{\mathcal{G}})$ is an optimal solution to [\(CON\)](#). In either case, [Algorithm 5](#) is correct.

Complexity. For the complexity claim, we observe that the running time of [Algorithm 5](#) is $O(HS^2A|\hat{\mathcal{V}}|^2|\hat{\mathcal{U}}|)$. To bound $|\hat{\mathcal{V}}|$, we observe that the number of integer multiples of δ required to capture the range $[-Hr_{max}, Hr_{max}]$ is at most $O(\frac{Hr_{max}}{\delta}) = O(H^2Sr_{max}/\epsilon)$ by definition of δ . Moreover, $|\hat{\mathcal{U}}| = O(|\hat{\mathcal{V}}| + S) = O(|\hat{\mathcal{V}}|)$ for sufficiently large $\frac{r_{max}}{\epsilon}$.

In particular, we see that the range of the rounded sums defining $\hat{\mathcal{U}}$ is at widest $[-2Hr_{max} - \delta S, 2Hr_{max}]$ since for any $t+1$ the rounded input is,

$$\lfloor [r_h(s, a) + P_h(1 | s, a)\hat{v}_1]_{\mathcal{G}} + \dots + P_h(t | s, a)\hat{v}_t \rfloor_{\mathcal{G}} \leq r_h(s, a) + \sum_{s'=1}^t P_h(s' | s, a)\hat{v}_{s'},$$

which is at most $2Hr_{max}$, and,

$$\lfloor [r_h(s, a) + P_h(1 | s, a)\hat{v}_1]_{\mathcal{G}} + \dots + P_h(t | s, a)\hat{v}_t \rfloor_{\mathcal{G}} \geq r_h(s, a) + \sum_{s'=1}^t P_h(s' | s, a)\hat{v}_{s'} - \delta t,$$

which is at least $-2Hr_{max} - \delta S$. Overall, we see that $O(|\hat{\mathcal{V}}|^2|\hat{\mathcal{U}}|) = O(|\hat{\mathcal{V}}|^3) = O(H^6S^3r_{max}^3/\epsilon^3)$ implying that the total run time is $O(H^7S^5Ar_{max}^3/\epsilon^3)$ as claimed. \square

E.8 Proof of [Observation 3](#)

Proof. Using properties of the floor function, we can infer that,

$$\lfloor v \rfloor_{\mathcal{G}} = v^{min} \left(\frac{1}{1-\delta} \right) \left\lfloor \log_{\frac{1}{1-\delta}} \frac{v}{v^{min}} \right\rfloor \leq v^{min} \left(\frac{1}{1-\delta} \right)^{\log_{\frac{1}{1-\delta}} \frac{v}{v^{min}}} = \frac{v}{v^{min}} v^{min} = v,$$

and,

$$\lfloor v \rfloor_{\mathcal{G}} = v^{\min} \left(\frac{1}{1-\delta} \right)^{\left\lceil \log_{\frac{1}{1-\delta}} \frac{v}{v^{\min}} \right\rceil} \geq v^{\min} \left(\frac{1}{1-\delta} \right)^{\log_{\frac{1}{1-\delta}} \frac{v}{v^{\min}} - 1} = v(1-\delta).$$

□

E.9 Proof of Lemma 10

Proof. We proceed by induction on t .

Base Case. For the base case, we consider $t = S + 1$. By definition, we have $\hat{\sigma}_{h,\hat{v}}^{s,a}(S + 1, u) = u = \sigma_{h,\hat{v}}^{s,a}(S + 1, u)$.

Inductive Step. For the inductive step, we consider $t \leq S$. We first see that,

$$\begin{aligned} \hat{\sigma}_{h,\hat{v}}^{s,a}(t, u) &= \hat{\sigma}_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) \\ &\leq \sigma_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) \\ &= \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}} + \sum_{s'=t+1}^S P_h(s' | s, a) \hat{v}_t \\ &\leq u + \sum_{s'=t}^S P_h(s' | s, a) \hat{v}_t \\ &= \sigma_{h,\hat{v}}^{s,a}(t, u). \end{aligned}$$

The first inequality used the induction hypothesis and the second inequality used the fact that $\lfloor x \rfloor_{\mathcal{G}} \leq x$.

We also see that,

$$\begin{aligned} \hat{\sigma}_{h,\hat{v}}^{s,a}(t, u) &= \hat{\sigma}_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) \\ &\geq \sigma_{h,\hat{v}}^{s,a}(t + 1, \lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}}) (1 - \delta)^{S-t} \\ &= \left(\lfloor u + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}} + \sum_{s'=t+1}^S P_h(s' | s, a) \hat{v}_t \right) (1 - \delta)^{S-t} \\ &\geq \left((1 - \delta)u + (1 - \delta) \sum_{s'=t}^S P_h(s' | s, a) \hat{v}_t \right) (1 - \delta)^{S-t} \\ &= \sigma_{h,\hat{v}}^{s,a}(t, u) (1 - \delta)^{S-t+1}. \end{aligned}$$

The first inequality used the induction hypothesis and the second inequality used the fact that $\lfloor x \rfloor_{\mathcal{G}} \geq x - \delta$ and the fact that all rewards and values are non-negative allowing us to add a $(1 - \delta)$ -factor to the other value demands. □

E.10 Proof of Lemma 11

Proof. We proceed by induction on h . Let $(s, v) \in \bar{\mathcal{S}}$ be arbitrary.

Base Case. For the base case, we consider $h = H + 1$. Since $\lfloor v \rfloor_{\mathcal{G}} \leq v$, we immediately see,

$$\hat{C}_{H+1}^*(s, \lfloor v \rfloor_{\mathcal{G}}) = \chi_{\{\lfloor v \rfloor_{\mathcal{G}} \leq 0\}} \leq \chi_{\{v \leq 0\}} = \bar{C}_{H+1}^*(s, v).$$

Inductive Step. For the inductive step, we consider $h \leq H$. If $\bar{C}_h^*(s, v) = \infty$, then trivially $\hat{C}_h^*(s, \lfloor v \rfloor_{\mathcal{G}}) \leq \bar{C}_h^*(s, v)$. Instead, suppose that $\bar{C}_h^*(s, v) < \infty$. Let π be a solution to the optimality

equations for \bar{M} so that $\bar{C}_h^\pi(s, v) = \bar{C}_h^*(s, v) < \infty$. Since $\bar{C}_h^*(s, v) < \infty$, we know that $(a^*, \mathbf{v}^*) = \pi_h(s, v) \in \bar{\mathcal{A}}_h(s, v)$. By the definition of $\bar{\mathcal{A}}_h(s, v)$, we know that,

$$\sigma_{h, \mathbf{v}^*}^{s, a^*}(1, r_h(s, a^*)) = r_h(s, a^*) + \sum_{s'} P_h(s' | s, a^*) v_{s'}^* \geq v \geq \lfloor v \rfloor_{\mathcal{G}}.$$

For each $s' \in \mathcal{S}$, define $\hat{v}_{s'}^* \stackrel{\text{def}}{=} \lfloor v_{s'}^* \rfloor_{\mathcal{G}}$ and recall that $v_{s'}^* \in \mathcal{V}$. We first observe that,

$$\begin{aligned} \sigma_{h, \hat{\mathbf{v}}^*}^{s, a^*}(1, r_h(s, a^*)) &= r_h(s, a^*) + \sum_{s'} P_h(s' | s, a) \lfloor v_{s'}^* \rfloor_{\mathcal{G}} \\ &\geq r_h(s, a^*) + \sum_{s'} P_h(s' | s, a) v_{s'} (1 - \delta) \\ &\geq \left(r_h(s, a^*) + \sum_{s'} P_h(s' | s, a) v_{s'} \right) (1 - \delta) \\ &= \sigma_{h, \mathbf{v}^*}^{s, a^*}(1, r_h(s, a^*)) (1 - \delta). \end{aligned}$$

The second inequality used the fact that all rewards are non-negative. Then by [Lemma 10](#),

$$\begin{aligned} \hat{\sigma}_{h, \hat{\mathbf{v}}^*}^{s, a^*}(1, r_h(s, a^*)) &\geq \sigma_{h, \hat{\mathbf{v}}^*}^{s, a^*}(1, r_h(s, a^*)) (1 - \delta)^S \\ &\geq \sigma_{h, \mathbf{v}^*}^{s, a^*}(1, r_h(s, a^*)) (1 - \delta)^{S+1} \\ &\geq \lfloor v \rfloor_{\mathcal{G}} (1 - \delta)^{S+1} \\ &= \kappa(\lfloor v \rfloor_{\mathcal{G}}). \end{aligned}$$

Thus, $(a^*, \hat{\mathbf{v}}^*) \in \hat{\mathcal{A}}_h(s, \lfloor v \rfloor_{\mathcal{G}})$.

Since $v_{s'}^* \in \mathcal{V}$, the induction hypothesis implies that $\hat{C}_{h+1}^*(s', \hat{v}_{s'}^*) \leq \bar{C}_{h+1}^*(s', v_{s'}^*) = \bar{C}_{h+1}^\pi(s', v_{s'}^*)$. The optimality equations for \hat{M} then imply that,

$$\begin{aligned} \hat{C}_h^*(s, \lfloor v \rfloor_{\mathcal{G}}) &= \min_{(a, \hat{\mathbf{v}}) \in \hat{\mathcal{A}}_h(s, v)} c_h(s, a) + f \left(\left(P_h(s' | s, a), \hat{C}_{h+1}^*(s', \hat{v}_{s'}^*) \right)_{s' \in P_h(s, a)} \right) \\ &\leq c_h(s, a^*) + f \left(\left(P_h(s' | s, a^*), \hat{C}_{h+1}^*(s', \hat{v}_{s'}^*) \right)_{s' \in P_h(s, a^*)} \right) \\ &\leq c_h(s, a^*) + f \left(\left(P_h(s' | s, a), \bar{C}_{h+1}^\pi(s', v_{s'}^*) \right)_{s' \in P_h(s, a^*)} \right) \\ &= \bar{C}_h^\pi(s, v) \\ &= \bar{C}_h^*(s, v). \end{aligned}$$

The first inequality used the fact that $(a^*, \mathbf{v}^*) \in \hat{\mathcal{A}}_h(s, v)$. The second inequality relied on f being non-decreasing and the induction hypothesis. The penultimate equality used [\(TR\)](#).

This completes the proof. \square

E.11 Proof of [Lemma 12](#)

Proof. We proceed by induction on h . Let $(s, \hat{v}) \in \hat{\mathcal{S}}$ be arbitrary.

Base Case. For the base case, we consider $h = H+1$. By definition and assumption, $\hat{C}_{H+1}^\pi(s, \hat{v}) = \chi_{\{\hat{v} \leq 0\}} < \infty$. Thus, it must be the case that $\hat{v} \leq 0$ and so by definition $\hat{V}_{H+1}^\pi(s, \hat{v}) = 0 \geq \hat{v}$.

Inductive Step. For the inductive step, we consider $h \leq H$. As in the proof of [Lemma 5](#), we know that $\pi_h(s, v) = (a, \hat{\mathbf{v}}) \in \hat{\mathcal{A}}_h(s, \hat{v})$ and for any $s' \in \mathcal{S}$ with $P_h(s' | s, a) > 0$ that $\hat{C}_{h+1}^\pi(s', v_{s'}) < \infty$. Thus, the induction hypothesis implies that $\hat{V}_{h+1}^\pi(s', \hat{v}_{s'}) \geq \hat{v}_{s'} (1 - \delta)^{(S+1)(H-h)}$ for any such

$s' \in \mathcal{S}$. By the policy evaluation equations, we see that,

$$\begin{aligned}
\hat{V}_h^\pi(s, \hat{v}) &= r_h(s, a) + \sum_{s'} P_h(s' | s, a) \hat{V}_{h+1}^\pi(s', \hat{v}_{s'}) \\
&\geq r_h(s, a) + \sum_{s'} P_h(s' | s, a) \hat{v}_{s'} (1 - \delta)^{(S+1)(H-h)} \\
&\geq \sigma_{h, \hat{v}}^{s, a}(1, r_h(s, a)) (1 - \delta)^{(S+1)(H-h)} \\
&\geq \hat{\sigma}_{h, \hat{v}}^{s, a}(1, r_h(s, a)) (1 - \delta)^{(S+1)(H-h)} \\
&\geq \hat{v} (1 - \delta)^{S+1} (1 - \delta)^{(S+1)(H-h)} \\
&= \hat{v} (1 - \delta)^{(S+1)(H-h+1)}.
\end{aligned}$$

The first inequality used the induction hypothesis. The second inequality used the fact that the rewards are non-negative. The third inequality used [Lemma 10](#). The fourth inequality used the fact that by definition of $\hat{\mathcal{A}}_h(s, \hat{v})$ and κ , $\hat{\sigma}_{h, \hat{v}}^{s, a}(1, r_h(s, a)) \geq \kappa(\hat{v}) = \hat{v} (1 - \delta)^{S+1}$.

This completes the proof. \square

E.12 Proof of [Theorem 3](#)

Proof.

Correctness. If $\hat{C}_1^*(s_0, v) > B$ for all $\hat{v} \in \hat{\mathcal{V}}$, then $C_M^* > B$ since otherwise we would have $\hat{C}_1^*(s_0, \lfloor v \rfloor_{\mathcal{G}}) \leq \bar{C}_1^*(s_0, v) \leq C_M^* \leq B$ by [Lemma 11](#). Thus, if [Algorithm 5](#) outputs ‘‘infeasible’’ it is correct.

On the other hand, suppose that there exists some $\hat{v} \in \hat{\mathcal{V}}$ for which $\hat{C}_1^*(s_0, \hat{v}) \leq B$. By standard MDP theory, we know that since $\pi \in \Pi^D$ is a solution to \hat{M} , it must satisfy the optimality equations. In particular, $\hat{C}_1^\pi(s_0, \hat{v}) = \hat{C}_1^*(s_0, \hat{v}) \leq B$. As in the proof of [Theorem 1](#), since $C_M^\pi = \hat{C}_1^\pi(s_0, \hat{v})$, we see that there exists a $\pi \in \Pi^D$ for which $C_M^\pi \leq B$ and so $V_M^* > -\infty$.

Since V_M^* is the value of some deterministic policy, [Lemma 3](#) implies that $V_M^* \in \mathcal{V}$. Thus, [Lemma 12](#) implies that $\hat{V}_1^\pi(s_0, \lfloor V_M^* \rfloor_{\mathcal{G}}) \geq \lfloor V_M^* \rfloor_{\mathcal{G}} (1 - \delta)^{(S+1)H} \geq V_M^* (1 - \delta)^{(S+1)H+1} = V_M^* (1 - \frac{\epsilon}{(S+1)H+1})^{(S+1)H+1} \geq V_M^* (1 - \epsilon)$ and $\hat{C}_1^\pi(s_0, V_M^*) \leq C_1^*(s_0, V_M^*) \leq B$. Consequently, running π with initial state $\bar{s}_0 = (s_0, \lfloor V_M^* \rfloor_{\mathcal{G}})$ is an optimal solution to [\(CON\)](#). In either case, [Algorithm 5](#) is correct.

Complexity. For the complexity claim, we observe that the running time of [Algorithm 5](#) is $O(HS^2A|\hat{\mathcal{V}}|^2|\hat{\mathcal{U}}|)$. To bound $|\hat{\mathcal{V}}|$, we observe that the number of v_{min} -scaled powers of $1/(1 - \delta)$ required to capture the range $[0, Hr_{max}]$ is at most one plus the largest power needed, which is

$$\begin{aligned}
O(\log_{1/(1-\delta)}(\frac{Hr_{max}}{v_{min}})) &= O(\log(\frac{Hr_{max}}{v_{min}}) / \log(1/(1 - \delta))) \\
&= O(\log(\frac{Hr_{max}}{v_{min}}) / \delta) \\
&= O(\log(HS \frac{Hr_{max}}{p_{min}^H r_{min}}) / \epsilon) \\
&= O(H^2 S \log(\frac{r_{max}}{p_{min} r_{min}}) / \epsilon),
\end{aligned}$$

by definition of δ and the fact that $\log(\frac{1}{1-\delta}) = -\log(1 - \delta) \geq -\log(e^{-\delta}) = \delta$. Moreover, $|\hat{\mathcal{U}}| = O(|\hat{\mathcal{V}}|)$.

We see that the range of the rounded sums is at widest $[0, 2Hr_{max}]$ since for any $t + 1$ rounding non-negative sums is at least 0 and,

$$\lfloor [r_h(s, a) + P_h(1 | s, a) \hat{v}_1]_{\mathcal{G}} + \dots + P_h(t | s, a) \hat{v}_t \rfloor_{\mathcal{G}} \leq r_h(s, a) + \sum_{s'=1}^t P_h(s' | s, a) \hat{v}_{s'},$$

which is at most $2Hr_{max}$. Then, the same analysis from before shows that the number of scaled powers of $1/(1-\delta)$ needed to cover this interval is $O(|\hat{\mathcal{V}}|)$. Thus, we see that $O(|\hat{\mathcal{V}}|^2|\hat{\mathcal{U}}|) = O(|\hat{\mathcal{V}}|^3) = O(H^6 S^3 \log(\frac{r_{max}}{p_{min}r_{min}})^3/\epsilon^3)$ implying that the total run time is $O(H^7 S^5 A \log(\frac{r_{max}}{p_{min}r_{min}})^3/\epsilon^3)$ as claimed. \square

F Extensions

F.1 Stochastic Costs

Suppose each cost $c_h(s, a)$ is replaced with a cost distribution $C_h(s, a)$. Here, we consider finitely supported cost distributions whose supports are at most $m \in \mathbb{N}$. Then, instead of the agent occurring cost $c_h(s, a)$ upon taking action a in state s at time h , the agent occurs a random cost $c_h \sim C_h(s, a)$. Generally, this necessitates histories be cost dependent, and so the policy evaluation equations become,

$$V_h^\pi(\tau_h) = r_h(s, a) + \sum_{c', s'} C_h(c' | s, a) P_h(s' | s, a) V_{h+1}^\pi(\tau_h, a, c', s'). \quad (\text{CPE})$$

Cover MDP. This implicitly changes the definition of \mathcal{V} since the histories considered in the definition must now include cost history. Since the cost distributions are finitely supported, \mathcal{V} remains a finite set. The main difference for \bar{M} is that the future value demands must depend on both the immediate cost and the next state. This slightly changes the action space:

$$\bar{\mathcal{A}}_h(s, v) \stackrel{\text{def}}{=} \left\{ (a, \mathbf{v}) \in \mathcal{A} \times \mathcal{V}^{m \times S} \mid r_h(s, a) + \sum_{c', s'} C_h(c' | s, a) P_h(s' | s, a) v_{c', s'} \geq v \right\}.$$

Bellman Updates. In order to solve \bar{M} using [Algorithm 4](#), we must extend the definition of TSR to also be recursive in the immediate costs. The key difference of the *TSRC* condition is that g 's recursion is now two dimensional.

Definition 11 (TSRC). We call a criterion C *time-space-cost-recursive* (TSRC) if $C_M^\pi = C_1^\pi(s_0)$ where $C_{H+1}^\pi(\cdot) = \mathbf{0}$ and for any $h \in [H]$ and $\tau_h \in \mathcal{H}_h$ letting $s = s_h(\tau_h)$ and $a = \pi_h(\tau_h)$,

$$C_h^\pi(\tau_h) = c_h(s, a) + f \left((C_h(c' | s, a), P_h(s' | s, a), C_{h+1}^\pi(\tau_h, a, c', s'))_{c', s'} \right). \quad (22)$$

In the above, $c' \in C_h(s, a)$ and $s' \in P_h(s, a)$. We now require that f be computable in $O(mS)$ time. We also require that the f term above is equal to $g_h^{\tau_h, a}(1, 1)$, where, $g_h^{\tau_h, a}(m+1, 1) = 0$, $g_h^{\tau_h, a}(k, S+1) = g_h^{\tau_h, a}(k+1, 1)$, and,

$$g_h^{\tau_h, a}(k, t) = \alpha \left(\beta \left(C_h(c_k | s, a), P_h(t | s, a), C_{h+1}^\pi(\tau_h, a, t) \right), g_h^{\tau_h, a}(k, t+1) \right). \quad (23)$$

In the above, we assume c_k is the k th supported cost of $C_h(s, a)$. Again, both α, β can be computed in $O(1)$ time, but now $\alpha(\beta(y, \cdot), x) = x$ whenever $0 \in y$.

Our examples from before also carry over to the stochastic cost setting.

Proposition 5 (TSCR examples). *The following classical constraints can be modeled by a TSCR cost constraint.*

1. (Expectation Constraints) *We capture these constraints by defining $\alpha(x, y) \stackrel{\text{def}}{=} x + y$ and $\beta(x, y, z) \stackrel{\text{def}}{=} xyz$.*
2. (Almost Sure Constraints) *We capture these constraints by defining $\alpha(x, y) \stackrel{\text{def}}{=} \max(x, y)$ and $\beta(x, y, z) \stackrel{\text{def}}{=} [x > 0 \wedge y > 0]z$.*
3. (Anytime Constraints) *We capture these constraints by defining $\alpha(x, y) \stackrel{\text{def}}{=} \max(0, \max(x, y))$ and $\beta(x, y, z) \stackrel{\text{def}}{=} [x > 0 \wedge y > 0]z$.*

We can then modify our approximate recursion from before.

Definition 12. We define, $\hat{g}_{h,v}^{s,a}(m+1, 1, u) \stackrel{\text{def}}{=} \chi_{\{u \geq v\}}$, $\hat{g}_{h,v}^{s,a}(k, S+1, u) \stackrel{\text{def}}{=} \hat{g}_{h,v}^{s,a}(k+1, 1, u)$ and for $t \leq S$,

$$\hat{g}_{h,v}^{s,a}(k, t, u) \stackrel{\text{def}}{=} \min_{v_{k,t} \in \mathcal{V}} \alpha \left(\beta (C_h(c_k | s, a), P_h(t | s, a), \bar{C}_{h+1}^*(t, v_{k,t})), \hat{g}_{h,v}^{s,a}(k, t+1, \lfloor u + C_h(c_k | s, a)P_h(t | s, a)v_{k,t} \rfloor_{\mathcal{G}}) \right). \quad (24)$$

Approximation. Lastly, our rounding now occur error over time, space, and cost. Thus, we simply need to slightly modify our rounding functions. The main change is we use $\delta \stackrel{\text{def}}{=} \frac{\epsilon}{H(mS+1)+1}$. We also further relax our lower bounds to $\kappa(v) \stackrel{\text{def}}{=} v - \delta(mS+1)$ and $\kappa \stackrel{\text{def}}{=} v(1-\delta)^{mS+1}$ respectively. Our running times correspondingly will have m^3 terms now.

F.2 Infinite Discounting

Approximations. Since we focus on approximation algorithms, the infinite discounted case can be immediately handled by using the idea of effective horizon. We can treat the problem as a finite horizon problem where the finite horizon H defined so that $\sum_{h=H}^{\infty} \gamma^{h-1} r_{max} \leq \epsilon'$. By choosing ϵ' and ϵ small enough, we can get traditional value approximations. The discounting also ensures the effective horizon H is polynomially sized implying efficient computation. We just need to assume that 0-cost actions are always available so that the policy can guarantee feasibility after the effective horizon has passed.

Hardness. We also note that all of the standard hardness results concerning deterministic policy computation carries over to the infinite discounting case even if all quantities are stationary.

F.3 Faster Approximations

We can significantly improve the running time of our FPTAS. The main guarantee is given in [Corollary 1](#). They key step is to modify [Algorithm 3](#) to use the differences instead of the sums. It is easy to see that this is equivalent since,

$$r_h(s, a) + \sum_{s'} P_h(s' | s, a)v_{s'} \geq v \iff v - \sum_{s'} P_h(s' | s, a)v_{s'} \leq r_h(s, a).$$

Since rounding down the differences make them larger, it becomes harder to be below $r_h(s, a)$. Consequently, we now interpret κ as an upper bound for $r_h(s, a)$ instead of a lower bound on v . The approximate dynamic programming method based on differences can be seen in [Definition 13](#).

Definition 13. Fix a rounding down function $\lfloor \cdot \rfloor_{\mathcal{G}}$ and upper bound function κ . For any $h \in [H]$, $s \in \mathcal{S}$, $v \in \mathcal{V}$, and $u \in \mathbb{R}$, we define, $\hat{g}_{h,v}^{s,a}(S+1, u) = \chi_{\{u \leq \kappa(r_h(s, a))\}}$ and for $t \leq S$,

$$\hat{g}_h^{s,a}(t, u) \stackrel{\text{def}}{=} \min_{v_t \in \mathcal{V}} \alpha \left(\beta (P_h(t | s, a), \bar{C}_{h+1}^*(t, v_t)), \hat{g}_h^{s,a}(t+1, \lfloor u - P_h(t | s, a)v_t \rfloor_{\mathcal{G}}) \right). \quad (\text{DIF})$$

The recursion is nearly identical to the originally, and unsurprisingly, it retains the same theoretical guarantees but in the reverse order. The guarantees can be seen in [Lemma 13](#), which is straightforward to prove following the approach in the proof of [Lemma 1](#).

Lemma 13. For any $t \in [S+1]$ and $u \in \mathbb{R}$, we have that,

$$\hat{g}_h^{s,a}(t, u) = \min_{\mathbf{v} \in \mathcal{V}^{S-t+1}} \hat{g}_{h,\hat{\mathbf{v}}}^{s,a}(t) \quad \text{s.t.} \quad \tilde{\sigma}_{h,\mathbf{v}}^{s,a}(t, u) \leq \kappa(r_h(s, a)), \quad (25)$$

where $\tilde{\sigma}_{h,\mathbf{v}}^{s,a}(t, u) \stackrel{\text{def}}{=} \lfloor [u - P_h(t | s, a)v_t]_{\mathcal{G}} - \dots - P_h(S | s, a)v_S \rfloor_{\mathcal{G}}$.

The difference version so far does not help us get faster algorithms. The key is in how we use it. Since the base case of the recursion is $r_h(s, a)$ and not v , we can compute the approximate bellman update for all v 's simultaneously. This ends up saving us a factor of $|\mathcal{V}|$ that we had in the original [Algorithm 4](#). The new algorithm is defined in [Algorithm 6](#). The inputs to the recursion are define in [Definition 14](#).

Algorithm 6 Approx Solve

Input: (\bar{M}, \bar{C})

- 1: $\hat{C}_{H+1}^*(s, v) \leftarrow \chi_{\{v \leq 0\}}$ for all $(s, v) \in \bar{\mathcal{S}}$
 - 2: **for** $h \leftarrow H$ down to 1 **do**
 - 3: **for** $s \in \mathcal{S}$ **do**
 - 4: **for** $a \in \mathcal{A}$ **do**
 - 5: $\hat{g}_h^{s,a}(S+1, u) \leftarrow \chi_{\{u \leq \kappa(r_h(s,a))\}} \forall u \in \hat{\mathcal{U}}_h^{s,a}$
 - 6: **for** $t \leftarrow S$ down to 1 **do**
 - 7: **for** $u \in \hat{\mathcal{U}}_h^{s,a}$ **do**
 - 8: $\hat{v}_{t,a} \hat{g}_h^{s,a}(t, u) \leftarrow (\text{DIF})$
 - 9: **for** $v \in \mathcal{V}$ **do**
 - 10: $a^*, \hat{C}_h^*(s, v) \leftarrow \min_{a \in \mathcal{A}} c_h(s, a) + \hat{g}_h^{s,a}(1, v)$
 - 11: $\pi_h(s, v) \leftarrow a^*$
 - 12: **return** π and \hat{C}^*
-

Definition 14. For any $h \in [H]$, $s \in \mathcal{S}$, and $a \in \mathcal{A}$, we define $\hat{\mathcal{U}}_h^{s,a}(1) \stackrel{\text{def}}{=} \mathcal{V}$ and for any $t \in [S]$,

$$\hat{\mathcal{U}}_h^{s,a}(t+1) \stackrel{\text{def}}{=} \bigcup_{v_t \in \mathcal{V}} \bigcup_{\hat{\sigma} \in \hat{\mathcal{U}}_h^{s,a}(t)} \{ \lfloor \hat{\sigma} - P_h(t \mid s, a) v_t \rfloor_{\mathcal{G}} \}. \quad (26)$$

Proposition 6. The running time of **Algorithm 6** is $O(HS^2A|\mathcal{V}|\hat{\sigma})$.

Corollary 1 (Running Time Improvements). Using **Algorithm 6**, the running time of our additive-FPTAS becomes $O(H^5S^4Ar_{max}^2/\epsilon^2)$, and the running time of our relative-FPTAS becomes $O(H^5S^4A \log(\frac{r_{max}}{r_{min}p_{min}})^2/\epsilon^2)$

Approximation Details. Although the running times our clear from removing the factor of $|\hat{\mathcal{V}}|$, we need to slightly alter our approximation schemes for this to work. First, we need to use $\kappa(r_h(s, a)) \stackrel{\text{def}}{=} r_h(s, a) + \delta$ for the additive approximation. The proof from before goes through almost identically. However, for the relative approximation, no choice of upper bound can ensure enough feasibility. Thus, we simply use $\kappa(r_h(s, a)) \stackrel{\text{def}}{=} r_h(s, a)$ and apply a different analysis. We also note that technically, differences can become negative. To deal with this the relative rounding function should simply send any negative number to 0: $\lfloor -x \rfloor_{\mathcal{G}} \stackrel{\text{def}}{=} 0$. The analysis is mostly the same, but the feasibility statement must be slightly modified.

Lemma 14. Suppose all rewards are non-negative. For any $h \in [H+1]$ and $(s, v) \in \bar{\mathcal{S}}$, $\hat{C}_h^*(s, \lfloor v(1-\delta)^{H-h+1} \rfloor_{\mathcal{G}}) \leq \bar{C}_h^*(s, v)$.

The idea is that since no fixed upper bound can capture arbitrary input values, we simply input relative values. Then, the feasibility part of **Lemma 11** goes through as before. The proofs mostly remain the same, but the rounding must again change. We must now start at the smaller v_{min} that is the original v_{min} scaled by a factor of $(1-\delta)^H$ to ensure that $\lfloor V_M^*(1-\delta)^H \rfloor$ is in $\hat{\mathcal{V}}$. This makes $\hat{\mathcal{V}}$ larger, but not by too much as we argued in previous analyses.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We believe the claims made are accurate.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We point out our assumed reward bounds and note they are unavoidable under worst-case analysis.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Our paper provides all proofs in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any experiments.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We believe our paper follows the code of ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss in the introduction and conclusions potential positive impacts of our work. We do not believe there are any direct negative impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any releasable materials.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any experiments.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper is purely theoretical and does not include any experiments.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper is purely theoretical and does not use any crowdsourcing nor human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper is purely theoretical and does not use any crowdsourcing nor human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.