

Defense Against Reward Poisoning Attacks in Reinforcement Learning

Kiarash Banihashem*
University of Maryland

kiarash@umd.edu

Adish Singla
Max Planck Institute for Software Systems

adishs@mpi-sws.org

Goran Radanovic
Max Planck Institute for Software Systems

gradanovic@mpi-sws.org

Reviewed on OpenReview: <https://openreview.net/forum?id=goPsLn3RVo>

Abstract

We study defense strategies against reward poisoning attacks in reinforcement learning. As a threat model, we consider cost-effective targeted attacks—these attacks minimally alter rewards to make the attacker’s target policy uniquely optimal under the poisoned rewards, with the optimality gap specified by an attack parameter. Our goal is to design agents that are robust against such attacks in terms of the worst-case utility w.r.t. the true, unpoisoned, rewards while computing their policies under the poisoned rewards. We propose an optimization framework for deriving optimal defense policies, both when the attack parameter is known and unknown. For this optimization framework, we first provide characterization results for generic attack cost functions. These results show that the functional form of the attack cost function and the agent’s knowledge about it are critical for establishing lower bounds on the agent’s performance, as well as for the computational tractability of the defense problem. We then focus on a cost function based on ℓ_2 norm, for which we show that the defense problem can be efficiently solved and yields defense policies whose expected returns under the true rewards are lower bounded by their expected returns under the poison rewards. Using simulation-based experiments, we demonstrate the effectiveness and robustness of our defense approach.

1 Introduction

One of the key challenges in designing trustworthy AI systems is ensuring that they are technically robust and resilient to security threats European Commission (2019). Amongst many requirements that an AI system ought to satisfy in order to be deemed trustworthy is robustness to adversarial attacks Hamon et al. (2020). Standard approaches to reinforcement learning (RL) Sutton & Barto (2018) have shown to be susceptible to adversarial attacks which manipulate the feedback that an agent receives from its environment. These attacks broadly fall under two categories: a) *test-time* attacks, which manipulate an agent’s input data at test-time without changing its policy Huang et al. (2017); Lin et al. (2017); Tretschk et al. (2018); Kumar et al. (2021), and b) *training-time* attacks that manipulate an agent’s input data at training-time, influencing the agent’s learned policy Zhang & Parkes (2008); Ma et al. (2019); Huang & Zhu (2019); Rakhsha et al. (2020; 2021); Zhang et al. (2020b); Sun et al. (2020); Liu & Lai (2021); Rangi et al. (2022). In this paper, we focus on training-time attacks, and more specifically, on *targeted reward poisoning* attacks that modify (i.e., *poison*) rewards to force an agent into adopting a *target* policy Ma et al. (2019); Rakhsha et al. (2021); Rangi et al. (2022). Prior work on reward poisoning attacks on RL primarily focuses on designing optimal attacks. In this paper,

*This work was done as a part of an internship project at Max Planck Institute for Software Systems.

we take a different perspective on *targeted* reward poisoning attacks, and focus on designing *defense strategies* against such attacks. This is challenging, given that the attacker is typically unconstrained in poisoning the rewards to force the target policy, while the agent’s performance is measured under the true reward function, which is unknown. The key idea that we exploit in our work is that the poisoning attacks have an underlying *structure* arising from the attacker’s objective to minimize the cost of the attack needed to force the target policy. We therefore ask the following question:

Can we design an effective defense strategy against targeted reward poisoning attacks by exploiting the cost-effective nature of these attacks?

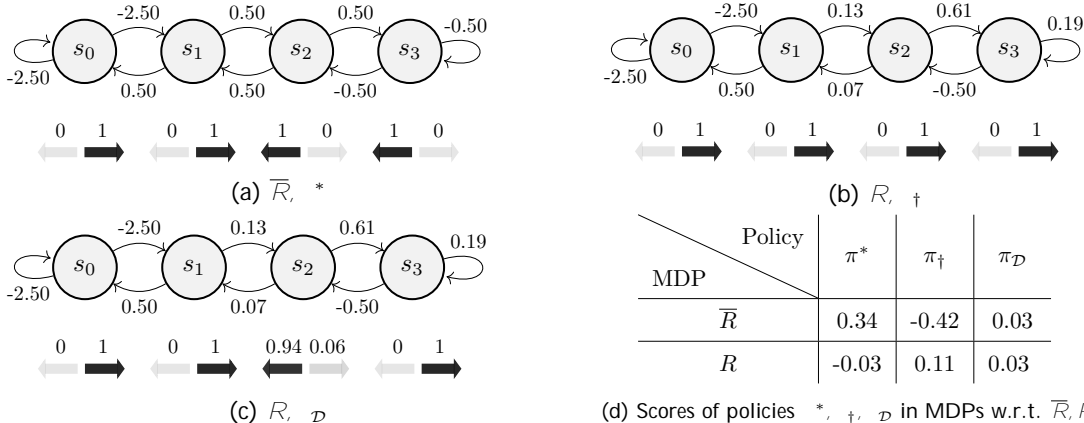


Figure 1: A simple chain environment with 4 states and two possible actions: *left* and *right*. s_0 is the initial state. The agent goes in the direction of its action with probability 90%, and otherwise the next state is selected uniformly at random from the other 3 states. Weights on edges indicate rewards for the action taken. For example in Fig. 1a, if the agent takes *left* in state s_1 , it receives 0.5. We denote the true rewards by \bar{R} , the poisoned rewards by R , the optimal policy under \bar{R} by π^* , the target policy (which is uniquely optimal under R) by π_{\dagger} , and the defense policy (which is derived from our framework) by $\pi_{\mathcal{D}}$. (a) shows \bar{R} and π^* . In particular, the numbers above the arrows and the different shades of gray show the probabilities of taking actions *left* and *right* under π^* . (b) shows R and π_{\dagger} . (c) shows $\pi_{\mathcal{D}}$ that our optimization framework derived from R , and by reasoning about the goal of the attack (π_{\dagger}). In particular, our optimization framework maximizes the worst-case performance under \bar{R} : while the optimization procedure does not know \bar{R} , it can constrain the set of plausible candidates for \bar{R} using R . (d) Table. 1d): Each entry in the table indicates the score of a (policy, reward function) pair, where the score is a scaled version of the total discounted return (see Section 3). For example, the score of policy π_{\dagger} equals -0.42 and 0.11 under R and \bar{R} respectively. Our defense policy significantly improves upon this and achieves a score of 0.03 . For comparison, the score of π^* equals 0.34 . Moreover, unlike for the target policy π_{\dagger} , the score of our defense policy $\pi_{\mathcal{D}}$ under \bar{R} is always at least as high as its score under R , as predicted by our results (see Theorem 5.1). The results are obtained with parameters $\epsilon_{\dagger} = 0.1$, $\epsilon_{\mathcal{D}} = 0.2$ and $\gamma = 0.99$ (see Section 3).

In this paper, we study this question in depth and under different assumption on the agent’s knowledge about the attack cost function. Perhaps surprisingly, the answer to this question is sometimes affirmative. While an agent only has access to the poisoned rewards, it may still be able infer some information about the true reward function, using the fact that the attack is cost-effective. By maximizing the worst-case utility over the set of plausible candidates for the true reward function, the agent can substantially limit the influence of the attack. The approach we study can be understood from Figure 1 which uses the chain environment from Rakhsha et al. (2021) to demonstrate the main ideas.

Contributions. We formalize this reasoning, and characterize the utility of our novel framework for designing defense policies. In summary, the key contributions include:

- We formalize the problem of designing defense policies against targeted and cost-effective reward poisoning attacks, which minimally modify the original reward function to achieve their goal (force a target policy).

- ^ We introduce a novel optimization framework for finding optimally robust defense policies this framework focuses on optimizing the agent's worst-case utility among the set of reward functions that are plausible candidates of the true reward function.
- ^ We provide characterization results that establish feasibility and computational complexity of finding optimally robust defense policies for different classes of attack cost functions. These results show that the functional form of the attack cost function and the agent's knowledge about it play a critical role in deriving optimally robust defense policies with provable performance guarantees.
- ^ Focusing on a cost function based on ℓ_2 norm, we show that optimally robust defense policies can be efficiently computed. We further establish lower bounds on the true performance of defense policies derived from our framework and computable from the poisoned rewards.
- ^ We empirically demonstrate the effectiveness and robustness of our approach using numerical simulations.

To our knowledge, this is the first framework for studying this type of defenses against reward poisoning attacks that try to force a target policy at a minimal cost.

2 Related Work

While this paper is broadly related to the literature on adversarial machine learning (e.g., Huang et al. (2011)), we recognize four themes in supervised learning (SL) and reinforcement learning (RL) that closely connect to our work.

Poisoning attacks in SL and RL. This paper is closely related to data poisoning attacks, first introduced and extensively studied supervised learning Biggio et al. (2012); Xiao et al. (2012); Mei & Zhu (2015); Xiao et al. (2015); Li et al. (2016); Koh & Liang (2017); Biggio & Roli (2018). These attacks are also called training-time attacks, and unlike test time attacks Szegedy et al. (2014); Pinto et al. (2017); Behzadan & Munir (2017); Zhang et al. (2020a); Moosavi-Dezfooli et al. (2016); Nguyen et al. (2015); Madry et al. (2018), which attack an already trained agent, they change data points during the training phase, which in turn affects the parameters of the learned model. Data poisoning attacks have also been studied in the bandits literature Jun et al. (2018); Ma et al. (2018); Liu & Shro (2019) and in RL (see Section 1).

Defenses against poisoning attacks in SL. In supervised learning, defenses against data poisoning attacks are often based on data sanitization that removes outliers from the training set Cretu et al. (2008); Paudice et al. (2018), trusted data points that support robust learning Nelson et al. (2008); Zhang et al. (2018), or robust estimation Charikar et al. (2017); Diakonikolas et al. (2019). Recently, Wu et al. (2022) have considered aggregation based defenses that can certify an RL agent's policy against a limited number of changes in the training dataset. While such defenses can mitigate some attack strategies, they are in general susceptible to data poisoning attacks Steinhardt et al. (2017); Koh et al. (2018).

Robustness to model uncertainty. There is a rich literature that studies robustness to uncertainties in reward functions McMahan et al. (2003); Regan & Boutilier (2010), and transition models Nilim & El Ghaoui (2005); Iyengar (2005); Bagnell et al. (2001) for MDP models. Typically, these works consider settings in which instead of knowing the exact parameters of the MDP, the agent has access to a set of possible parameters (uncertainty set). These works design policies that perform well in the worst case. More recent works have proposed ways to scale up these approaches via function approximation Tamar et al. (2014), as well as utilize them in online settings Lim et al. (2013). While our work uses the same principles of robust optimization, we do not assume that the uncertainty set, i.e., the set of all possible rewards, is directly given. Instead, we show how to derive it from the poisoned reward function.

Robustness to corrupted episodes. Another important line of work is the literature on robust learners that receive corrupted input during their training phase. Such learners have recently been designed for bandits and experts settings Lykouris et al. (2018); Gupta et al. (2019); Bogunovic et al. (2020); Amir et al. (2020), and episodic reinforcement learning Lykouris et al. (2019); Zhang et al. (2021). Typically, these works consider an attack model in which the adversary can arbitrarily corrupt a limited number of episodes. As we operate in the non-episodic setting and do not assume a limit in the attacker's poisoning budget, these works are orthogonal to the aspects we study in this paper. Instead, we utilize the structure of the attack in order to design a defense algorithm.

3 Formal Setting

In this section, we describe our formal setting, and identify relevant background details on reward poisoning attacks, as well as our problem statement. The problem formulation specifies our objectives that we establish and formally analyze in the next sections.

3.1 Preliminaries

We consider a standard reinforcement learning setting in which the environment is described by a discrete-time discounted Markov Decision Processes (MDP) Puterman (1994), denoted as $M = (S; A; R; P; \gamma; s^0)$, where: S is the state space, A is the action space, $R : S \times A \rightarrow \mathbb{R}$ is the reward function, $P : S \times A \times S \rightarrow [0; 1]$ is the transition model with $P(s; a; s^0)$ denoting the probability of transitioning to state s^0 by taking action a in state s , $\gamma \in [0; 1)$ is the discount factor, and s^0 is the initial state distribution. We consider state and action spaces, i.e., S and A , that are finite and discrete, and due to this we can adopt a vector notation for quantities dependent on states or state-action pairs. W.l.o.g., we assume that $|A| \geq 2$.

A generic (stochastic) policy is denoted by π , and it is a mapping $\pi : S \rightarrow \mathcal{P}(A)$, where $\mathcal{P}(A)$ is the probability simplex over action space A . We use $\pi(s; a)$ to denote the probability of taking action a in state s . While deterministic policies are a special case of stochastic policies, when explicitly stating that a policy is deterministic, we assume that it is a mapping from states to actions, i.e., $\pi : S \rightarrow A$. We denote the set of all policies by Π and the set of all deterministic policies by Π^{det} . For policy π , we define its score, as $V(\pi) = \mathbb{E} \left(\sum_{t=1}^{\infty} \gamma^{t-1} R(s_t; a_t) \right)$, where states s_t is sampled from the initial state distribution s^0 , and then subsequent states s_t are obtained by executing policy π in the MDP. The score of a policy is therefore its total expected return scaled by a factor of $(1 - \gamma)$.

Finally, we consider occupancy measures. We denote the state-action occupancy measure in the Markov chain induced by policy π by $\mu(s; a) = \mathbb{E} \left(\sum_{t=1}^{\infty} \gamma^{t-1} \mathbb{1}[s_t = s; a_t = a] \right)$. Given the MDP M , the set of realizable state-action occupancy measures under any (stochastic) policy $\pi \in \Pi$ is denoted by \mathcal{M} . Score $V(\pi)$ and μ satisfy $V(\pi) = \sum_{s, a} \mu(s; a) R(s; a)$, where $h \cdot i$ computes the dot product between two vectors of size $|S| \cdot |A|$. We denote by $\mu(s) = \mathbb{E} \left(\sum_{t=1}^{\infty} \gamma^{t-1} \mathbb{1}[s_t = s] \right)$ the state occupancy measure in the Markov chain induced by policy $\pi \in \Pi$. State-action occupancy measure $\mu(s; a)$ and state occupancy measure $\mu(s)$ satisfy $\mu(s; a) = \mu(s) \pi(s; a)$. We focus on ergodic MDPs, which in turn implies that $\mu(s) > 0$ for all $s \in S$ and $\pi \in \Pi$ (Puterman (1994)). This is a standard assumption in this line of work (e.g, see Rakhsha et al. (2021)) and is used to ensure the feasibility of the attacker's optimization problem.

3.2 Reward Poisoning Attacks

We consider reward poisoning attacks on an offline learning agent that optimally change the original reward function with the goal of deceiving the agent to adopt a deterministic policy $\pi^* \in \Pi^{\text{det}}$, called target policy. This type of attack has been extensively studied in the literature, and here we utilize the attack formulation based on the works of Ma et al. (2019); Rakhsha et al. (2020; 2021); Zhang et al. (2020b). In the following, we introduce the necessary notation, the attacker's model, and the agent's model (without defense).

Notation. We use \bar{M} to denote the true or original MDP with true, unpoisoned, reward function \bar{R} , i.e., $\bar{M} = (S; A; \bar{R}; P; \gamma; s^0)$. We use M to denote the modified or poisoned MDP with poisoned reward function R , i.e., $M = (S; A; R; P; \gamma; s^0)$. Note that only the reward function R changes across these MDPs. Quantities that depend on reward functions have analogous notation. For example, the score of policy π under \bar{R} is denoted by $V(\pi; \bar{R})$, whereas its score under R is denoted by $V(\pi; R)$. We denote an optimal policy under \bar{R} by π^* , i.e., $\pi^* = \arg \max_{\pi \in \Pi} V(\pi; \bar{R})$.

Attack model. The attacker we consider in this paper has full knowledge of \bar{M} . It can be modeled by a function $A(c; R^0; \gamma; \pi^*)$ that returns a set of poisoned rewards functions for a given attack cost function c , reward function R^0 , target policy π^* , and a desired attack parameter γ .¹ In particular, the attack problem is defined by the following optimization problem:

$$\min_R c(R; R^0) \quad \text{s.t.} \quad V(\pi^*; R) + \gamma \geq 2 \gamma V(\pi^*; \bar{R}) \quad (P1)$$

¹For cost functions c_p defined by (1) with finite $p > 1$, this set has a single element.

A common class of cost functions are ℓ_p -norms of manipulations Ma et al. (2019); Rakhsha et al. (2020; 2021), i.e.,

$$c(R; R^0) = c_p(R; R^0) = \|R - R^0\|_p; \quad (1)$$

with $p \geq 1$. As shown by Rakhsha et al. (2021), the attack problem (P1) is feasible for this class of cost functions and ergodic MDPs. Furthermore, instead of considering all deterministic policies, it is sufficient to consider policies that differ from π_y in a single action. Using $\pi_y^f; s; a$ to denote a policy that follows a $\pi_y(s)$ in state s and $\pi_y^f(s)$ in states $s \in \mathcal{S}$, (P1) can be rewritten as follows:

$$\min_R c(R; R^0) \quad \text{s.t.} \quad \pi_y^f; s; a + \pi_y^f; s; a \geq \pi_y(s); \quad (P1')$$

The equivalence between (P1) and (P1') is shown by Rakhsha et al. (2021), but we also provide further details in Appendix. Intuitively, (P1) and (P1') are equivalent because in order for a policy to be optimal it is sufficient (and necessary) that the policy is better than any of its neighbor policies. This fact allows us to reduce the number of constraints. Whereas the number of constraints in (P1) is exponential in $|\mathcal{S}|$ and $|\mathcal{A}|$, the number of constraint in (P1') is polynomial in $|\mathcal{S}|$ and $|\mathcal{A}|$, which in turn implies that (P1') is tractable (for a fixed p).

To better understand the optimization problem (P1'), we can consider ℓ_2 attack cost (i.e., c_2) defined as the Euclidean distance between R^0 and R . By solving this problem, i.e., setting $R \in \mathcal{A}(c_2; \bar{R}; \pi_y; \pi_y)$, the attacker finds the closest reward function to \bar{R} for which π_y is a uniquely optimal policy (with attack parameter π_y).

Remark 3.1. Note that the optimization problem (P1) may not be feasible if we lift the assumption that underlying MDP is ergodic. This can be seen from the constraints of the optimization problem (P1'): if target policy π_y does not visit a certain state s , then it has the same score as its neighbor policies $\pi_y^f; s; a$. Hence, the primary reason for assuming ergodicity is to make the attack problem feasible.

Agent without defense. The agent receives the poisoned MDP $\mathcal{M} := (\mathcal{S}; A; \bar{R}; P; \gamma)$ where the underlying true reward function \bar{R} (unknown to the agent) has been poisoned to \bar{R} . In the existing works on reward poisoning attacks, an agent naively optimizes score v (score w.r.t. \bar{R}). Because of this, the agent ends up adopting policy π_y .

3.3 Problem Statement

Perhaps unsurprisingly, the agent without defense, could perform arbitrarily badly under the true reward function \bar{R} . Our goal is to design a robust agent that derives its policy using the poisoned MDP $\mathcal{M} := (\mathcal{S}; A; \bar{R}; P; \gamma)$, but has provable worst-case guarantees w.r.t. \bar{R} . This agent has access to the poisoned reward vector $R \in \mathcal{A}(c; \bar{R}; \pi_y; \pi_y)$, but \bar{R} , π_y , and π_y are not given to the agent. Figure 2 illustrates a generic problem setting studied in this paper. In general, the agent does not know \bar{R} , but is given a class of cost functions \mathcal{C} that contains c , i.e., $c \in \mathcal{C}$. \mathcal{C} represents the agent's knowledge about the attack cost function; in a special case where \mathcal{C} contains only one element, the agent knows the cost function. Notice that π_y is obtainable by solving the optimization problem $\arg \max_b v$ as π_y is uniquely optimal in \mathcal{M} . On the other hand, \bar{R} is unknown to the agent. In terms of π_y , we will focus on two cases, the case when π_y is known to the agent, and the case when it is not.

In the first case, we can formulate the following optimization problem of maximizing the worst case performance of the agent, given that \bar{R} is unknown:

$$\max_{R; c \in \mathcal{C}} \min_{R} v \quad \text{s.t.} \quad R \in \mathcal{A}(c; R; \pi_y; \pi_y); \quad (P2a)$$

In other words, we calculate the set of all possible reward functions \mathcal{R} such that an attack on R could lead to the solution \bar{R} , i.e., $\bar{R} \in \mathcal{A}(c; R; \pi_y; \pi_y)$. We then find a policy π_y for which the worst-case score, i.e. $\min_{R; c} v$ is maximized, where the minimum is over all reward functions R calculated previously and all cost functions $c \in \mathcal{C}$.

Figure 2: The problem setting studied in this paper. The attack A modifies the original (true) reward function \bar{R} to force y while minimizing its cost. The defense D aims to optimize the agent's score under \bar{R} , but it only sees the poisoned reward R . Nevertheless, it can (in principle) find the set of all reward functions consistent with R and A , and search for a policy that maximizes the worst-case score of the agent over this consistent set.

Knowledge about attack cost	Guarantees on the value	Complexity
General C (e.g., s.t. $c_{\text{const}} \geq 2C$) $C = \text{fk } R \quad R^{\text{fk}_p} \text{ s.t. } p \in [1; 1)$	No for any R , Proposition 4.2 Yes, Theorem 4.4	NP-hard, Theorem 4.5
$C = \text{fk } R \quad R^{\text{fk}_1} g$	No for some R , Theorem 4.3 Yes, Section 5	NP-hard, Appendix Convex, Section 5
$C = \text{fk } R \quad R^{\text{fk}_2} g$	Yes, Appendix	Convex, Appendix

Table 1: Characterization results for different levels of the agent's knowledge about the attack cost function, expressed through C . In general, if C can be arbitrary, the optimization problem (P2a) may be unbounded from below regardless of R . For some classes C , e.g., that contain p -norm attack costs ($p \in [1; 1)$), (P2a) has the optimal solution, but this solution may be computationally hard. When the attack cost function is known, properties of problem (P2a) depend on the functional form of the attack cost function, as indicated by the 1 -norm, 2 -norm, and 1 -norm attack costs.

For the case when the agent does not know y , we use the following optimization problem:

$$\max_{R; c \geq 2C} \min_{D} \quad \text{s.t. } R \in A(c; R; y;) \text{ and } 0 < D \quad (\text{P2b})$$

where the agent uses D as an upper bound on y . We denote solutions to the optimization problems (P2a) and (P2b) by D , and it will be clear from the context which optimization problem we are referring to with D .

Remark 3.2. We note that some structural assumptions on the attack model are needed to guarantee robustness. For example, prior work typically considers untargeted attacks with budget constraints, e.g., that put an upper limit on the cost of the attack or the number of episodes in which the attacker can attack Lykouris et al. (2019); Zhang et al. (2021). This paper focuses on targeted attacks that minimize the cost of the attack. Given the strategic nature of the attacker, the structural assumptions on the attack model that (P2a) and (P2b) rely on are fairly natural.

4 Characterization Results for Generic Attack Cost

In this section, we provide characterization results showing the importance of the agent's knowledge about the attack cost function. The overview of the characterization results is shown in Table 1. To corresponding proofs can be found in Appendix.

Remark 4.1. The results presented in this section are stated for the optimization problem (P2a). However, the same results also hold for the optimization problem (P2b).

4.1 General Attack Cost

We start by stating what is perhaps an expected result: if the attack cost function can be arbitrary, then no defense can achieve any provable guarantee. It is relatively easy to see why this claim should hold. If the

agent believes that the attack cost function can be constant, then from the agent's perspective \bar{R} can be any reward function. Since rewards are not bounded, this in turn implies that no matter which policy the agent selects, no worst-case guarantees are possible. More formally, we obtain the following claim.

Proposition 4.2. Let $c_{\text{const}}(R; R^0)$ be a constant cost function, and assume that $c_{\text{const}} \geq 2C$. Then the optimization problem (P2a) is unbounded from below.

Given this result, it is clear that for provable defenses: a) the attack cost function cannot not be arbitrary in that R has to be informative about \bar{R} ; b) the agent should have some knowledge about the attack cost function. In the next subsection, we consider cost functions based on ℓ_p norms, i.e., c_p , commonly adopted by prior work on reward poisoning attacks Ma et al. (2019); Rakhsha et al. (2020; 2021).

4.2 ℓ_p Attack Cost

In contrast to constant cost functions, the strategic nature of the attacker is more apparent when it optimizes c_p , so the agent may infer some information about the true reward function \bar{R} from the poisoned rewards R which it can access. Our first result shows that for $p = 1$ the success of such an inference procedure depends on R . This is formally captured by the following theorem.

Theorem 4.3. There exists an instance of the problem setting, i.e., $\text{MDPM} = (S; A; R; P; \gamma)$ for which the optimization problem (P2a) is unbounded from below when $\alpha_1 \geq 2C$.

This theorem paints a relatively bleak picture for the possibility of achieving provable guarantees. Note two important observations. First, the impossibility result is a weaker variant of the result stated in Proposition 4.2 as it holds only for some MDPs. Second, α_1 is measuring the maximum modification of reward function \bar{R} , so critical information about \bar{R} may be lost this also provides intuition behind the impossibility results. In contrast, when $p \neq 1$, c_p is affected by all the modifications of \bar{R} . In fact, when p is restricted to take values in $[1; 1)$, a lower bound on the optimal value of (P2a) can always be derived from R .

Theorem 4.4. Consider any policy π_{γ} s.t. $\pi_{\gamma}(\gamma(s)|s) = 0$. For $C = f c_p$ s.t. $p \in [1; 1)$, the optimal value of the optimization problem (P2a) is bounded from below by $b^{\pi_{\gamma}}$.

A direct consequence of Theorem 4.4 is that the optimal solution to (P2a) always exists, and its worst case performance under \bar{R} is at least $b^{\pi_{\gamma}}$. Note that we can easily find a (deterministic) policy π_{γ} that maximizes the lower bound by solving $\max_{\pi} \sum_{s \in \mathcal{S}} \pi(s) b^{\pi(s)}$.² Furthermore, for any policy π_{γ} s.t. $\pi_{\gamma}(\gamma(s)|s) = 0$ we have that $b^{\pi_{\gamma}} \leq b^{\pi_{\gamma}}$. This means that we can efficiently find a defense policy with provable performance guarantees. However, such a defense policy may not be optimally robust in that its performance lower bound would not match the optimal one. We now turn to computational complexity challenges in deriving optimally robust defense policies: the next theorem provides a hardness result for $C = f c_p$ s.t. $p \in [1; 1)$, which admits guarantees on the optimal solution to (P2a).

Theorem 4.5. For $C = f c_p$ s.t. $p \in [1; 1)$, it is NP-hard to determine whether the optimal value of the optimization problem (P2a) is greater than or equal to b^{γ} .

Given that even for this natural choice of cost functions, $C = f c_p$ s.t. $p \in [1; 1)$, the problem of finding optimally robust policies is computationally hard, in the next section we focus on ℓ_2 attack cost that is known to the agent, i.e., $C = f c_2$. In Appendix, we provide similar analysis for ℓ_1 attack cost, i.e., $C = f c_1$.

5 Characterization Results for ℓ_2 Attack Cost

In this section, we consider the attack cost function c_2 , and assume that it is known to the agent ($C = f c_2$). In the first part, we focus on characterization results for the case when the attack parameter γ is known to the agent, i.e., the optimization problem (P2a). In the second part, we focus on the optimization problem (P2b), and generalize the results from the first part to the unknown attack parameter setting. The proofs of our formal results can be found in Appendix. Figure 3 provides intuition behind the results of this section.

²Stochastic policies are not necessary in this case since we can think of this problem as searching for an optimal policy over a truncated actions space (because actions $\pi_{\gamma}(s)$ are not admissible), so an optimal deterministic policy always exists.

In particular, it illustrates attack and defense strategies for a single-state MDP with action set $\{a_1, \dots, a_7\}$ and the ℓ_2 cost function. We refer the reader to Appendix for more details.

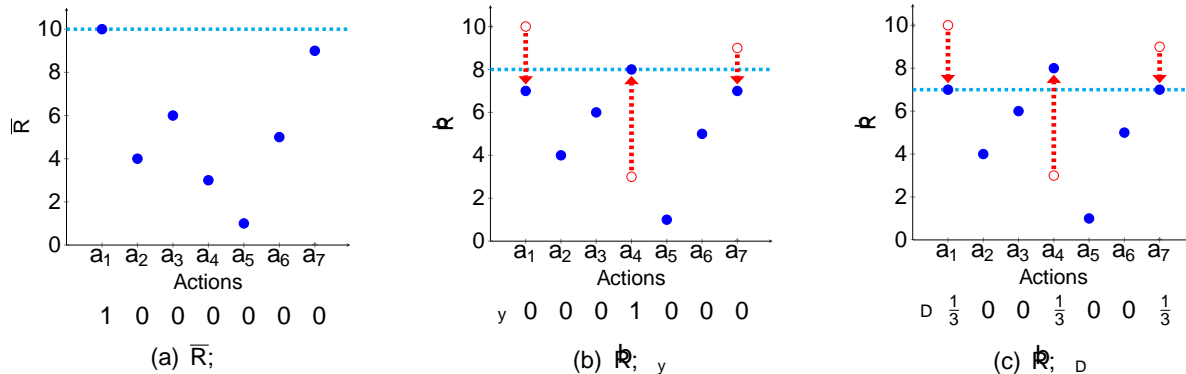


Figure 3: A single-state environment environment with 7 actions and states. In each figure, the denoted policy is uniform over actions on or above the dashed line (a) shows \bar{R} and π^* . Here, the optimal policy selects action a_1 . (b) shows \hat{R} and target policy \hat{y} with $\hat{y}_4 = 1$. Here, the target policy selects action a_4 . \hat{R} is obtained by solving (P1) (or (P1')) with the ℓ_2 cost function and $R^0 = \bar{R}$. In this case, the attack only modifies the rewards of three actions, a_1 , a_4 , and a_7 . (c) shows \hat{R} and \hat{D} with $\hat{D}_4 = 2$. The defense strategy only sees poisoned rewards \hat{R} . It first calculates the optimal action and the set of all second-best actions under \hat{R} , in this case $\{a_1, a_7\}$. If the reward of the second-best actions are no worse than \hat{D} , they form the set $\mathcal{S} = \{s; a_1\}; \{s; a_7\}$ or simply $\mathcal{S} = \{a_1, a_7\}$. We show in Appendix that the defense strategy, i.e., the solution to (P2b), selects an action uniformly at random from the set \mathcal{S} ($\hat{y}(s) = \frac{1}{2}$ for $s = a_1, a_7$). As we show in this section, the expected reward of the defense policy under \hat{R} is at least as much as its expected reward under \bar{R} , which can be easily verified from the figures: in this case, both are equal to $2\frac{2}{3}$.

5.1 Known Parameter Setting

We begin by analyzing the optimization problem (P2a). Denote by \mathcal{S} state-action pairs $(s; a)$ for which the difference between b^y and $b^{y \setminus \{s; a\}}$ is equal to ϵ , i.e., $\mathcal{S} = \{(s; a) : b^{y \setminus \{s; a\}} - b^y = \epsilon\}$.³ For the results of this section, with $\epsilon = \frac{1}{2}$ plays a critical role as we show in our formal analysis, it characterizes the feasible set of the optimization problem (P2a). In particular, in our analysis we show that \hat{R} is the solution to the attack problem for an underlying reward function R , i.e., $\hat{R} = A(R; y; \frac{1}{2})$, if and only if R can be expressed as

$$R = \hat{R} + \sum_{(s; a) \in \mathcal{S}} \alpha_{s; a} (b^{y \setminus \{s; a\}} - b^y);$$

with $\alpha_{s; a} \geq 0$. To see the importance of this result, let us instantiate $R = \bar{R}$ and calculate \hat{R} :

$$\hat{R} = \bar{R} + \sum_{(s; a) \in \mathcal{S}} \alpha_{s; a} (b^{y \setminus \{s; a\}} - b^y);$$

When the occupancy measure of \hat{D} is positively aligned with vectors $(b^{y \setminus \{s; a\}} - b^y)$, the performance of \hat{D} under the original reward function \bar{R} is at least $b^{\hat{D}} = \sum_{(s; a) \in \mathcal{S}} \alpha_{s; a} \hat{D}_{s; a}$. Hence, constraining \hat{D} to satisfy $\hat{D}_{s; a} \geq 0$ for all $(s; a) \in \mathcal{S}$ yields a guarantee on the score $\hat{V}^{\hat{D}}$, i.e., $\hat{V}^{\hat{D}} \geq b^{\hat{D}}$. These insights are formalized by Theorem 5.1, which also describes a procedure for solving (P2a). In Appendix, we provide intuition behind our analysis using a special case of our setting.

³In practice, \mathcal{S} should be calculated with some tolerance due to numerical imprecision (See Section 6).

Theorem 5.1. Consider the following optimization problem parameterized by:

$$\max_{\gamma} \sum_{s \in \mathcal{S}} \gamma(s) \sum_{a \in \mathcal{A}(s)} \left(R(s, a) - \sum_{s' \in \mathcal{S}} P(s'|s, a) \gamma(s') \right) \quad \text{s.t.} \quad \gamma(s) \geq 0 \quad \forall s \in \mathcal{S} \quad (P3)$$

For $\gamma = \gamma^*$, this optimization problem is always feasible, and its optimal solution γ^* specifies an optimal solution to the optimization problem (P2a) for $C = f_{C2g}$ with

$$D(\gamma) = \frac{\sum_{s \in \mathcal{S}} \gamma(s) \sum_{a \in \mathcal{A}(s)} \left(R(s, a) - \sum_{s' \in \mathcal{S}} P(s'|s, a) \gamma(s') \right)}{\sum_{s \in \mathcal{S}} \gamma(s)} \quad (2)$$

The score of $D(\gamma)$ is lower bounded by $\bar{D} - b^D$.

In addition to providing a characterization of the solution to (P2a), the above theorem provides an efficient algorithm for finding this solution using linear programming. As we discuss in Appendix D.3, the set of vectors γ and $\gamma^{f_{S;ag}}$ can be precomputed since the agent is given a poisoned model (and hence, knows the transition probabilities); for any policy π , the state-action occupancy measure satisfies $\gamma(s; a) = \pi(s) \sum_{a'} \pi(a'|s) \gamma(s; a')$ where the state occupancy measure π is the unique solution to the following Bellman identity

$$\pi(s) = (1 - \sum_{a \in \mathcal{A}(s)} \pi(s; a)) \pi(s) + \sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}(s')} P(s|s', a) \pi(s'; a)$$

In addition, the set of valid occupancy measures γ is the set of all vectors $\gamma \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$ satisfying the following Bellman flow constraints

$$\sum_{a \in \mathcal{A}(s)} \gamma(s; a) = (1 - \sum_{a \in \mathcal{A}(s)} \gamma(s; a)) \gamma(s) + \sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}(s')} P(s|s', a) \gamma(s'; a)$$

$$\gamma(s; a) \geq 0 \quad \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$$

Therefore, since occupancy measures $\gamma^{f_{S;ag}}$ and γ can be precomputed, the optimization problem (P3) can be efficiently solved using linear programming. In other words, computing optimally robust defense policy is computationally tractable. In order to calculate each γ^* , we need to solve $\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}(s)} \gamma(s; a)$ system of linear equations, which requires $|\mathcal{S}|^3$ operations in the worst case. Therefore, finding each γ^* for $\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}(s)} \gamma(s; a) \in [0, 1]$ requires at most $|\mathcal{S}|^4 \times |\mathcal{A}|$ operations. The linear program (P3) takes $O(|\mathcal{S}|^4 \times |\mathcal{A}|)$ time in the worst case as it has $|\mathcal{S}| \times |\mathcal{A}|$ constraints and variables. Therefore, the overall complexity of the approach is $O(|\mathcal{S}|^4 \times |\mathcal{A}|)$ in the worst case.

Theorem 5.1 also provides a performance guarantee of the defense policy w.r.t. the true reward function, i.e., $\bar{D} - b^D$. Such a bound is important in practice since it provides a certificate of the worst-case performance under the true reward function \bar{R} , even though the agent can only optimize over \mathcal{R} . In contrast to the lower bound in Theorem 4.4, the lower bound in Theorem 5.1 is optimal.

5.2 Unknown Parameter Setting

In this subsection, we focus on the optimization problem (P2b). First, note the structural difference between (P2a) and (P2b). In the former case, γ is given, and hence, the defense can infer possible values of \bar{R} by solving an inverse problem to the attack problem (P1). In particular, we know that the original reward function \bar{R} has to be in the set $\mathcal{R} = \{R : \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}(s)} \gamma(s; a) (R(s, a) - \sum_{s' \in \mathcal{S}} P(s'|s, a) \gamma(s')) \geq 0\}$. In the latter case, γ is not known, and instead we use parameter \bar{D} as an upper bound on γ . We distinguish two cases:

- **Overestimating Attack Parameter.** If $\bar{D} > \bar{D}$, then we know that \bar{R} is in the set $\mathcal{R} = \{R : \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}(s)} \gamma(s; a) (R(s, a) - \sum_{s' \in \mathcal{S}} P(s'|s, a) \gamma(s')) \geq 0\}$. Note that this set is a super-set of $\mathcal{R} = \{R : \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}(s)} \gamma(s; a) (R(s, a) - \sum_{s' \in \mathcal{S}} P(s'|s, a) \gamma(s')) \geq 0\}$, which means that it is less informative about \bar{R} .
- **Underestimating Attack Parameter.** If $\bar{D} < \bar{D}$, then the set $\mathcal{R} = \{R : \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}(s)} \gamma(s; a) (R(s, a) - \sum_{s' \in \mathcal{S}} P(s'|s, a) \gamma(s')) \geq 0\}$ will have only a single element, i.e., \bar{R} . In other words, this set typically contains no information about \bar{R} .

We analyze these two cases separately, first focusing on the former one.

5.2.1 Overestimating Attack Parameter

When $D \geq \gamma$, our formal analysis builds on the one presented in Section 5.1, and we highlight the main differences. Given that γ is not exactly known, we cannot directly operate on the set \mathcal{Y} . However, since D upper bounds γ , the defense can utilize the procedure from the previous section (Theorem 5.1) with appropriately chosen \bar{C} to solve (P2b) as we show in the following theorem.

Theorem 5.2. Assume that $D \geq \gamma$, and define $b = \min_{s; a \in \mathcal{Y}(s)} b^y - b^{y; s; a; g}$. Then, the optimization problem (P3) with $\bar{C} = \min_{D; b; g}$ is feasible and its optimal solution π_{\max} identifies an optimal policy π_D for the optimization problem (P2b) with $C = f_{C_2; g}$ via Equation (2). This policy π_D satisfies $V^{\pi_D} \leq b^D$.

To interpret the bounds, let us consider three cases:

- $\bar{R} \in \mathcal{R}$: If the attack indeed poisoned \bar{R} , then the smallest $\gamma \in (0; D]$ such that $\gamma \in \mathcal{R}$; corresponds to γ . In this case, it turns out that $\gamma = b$, and somewhat surprisingly, the defense policies of (P2a) and (P2b) coincide. (Note that this analysis assumes that $D \geq \gamma$.)
- $\bar{R} = \mathcal{R}$ and $D < b$: This corresponds to the case when the attack did not poison \bar{R} and there is no $\gamma \in (0; D]$ such that $\gamma \in \mathcal{R}$. In this case, it turns out that the optimal solution to the optimization problem (P2b) is $\pi_D = \pi_\gamma$ (indeed π_γ is uniquely optimal under \bar{R}).
- $\bar{R} = \mathcal{R}$ and $D \geq b$: This corresponds to the case when the attack did not poison \bar{R} and there is $\gamma \in (0; D]$ such that $\gamma \in \mathcal{R}$. In fact, b is the smallest such γ . In this case, it turns out that, in general, the optimal solution to the optimization problem (P2b) is $\pi_D \in \mathcal{Y}$, even though π_γ is uniquely optimal under \bar{R} .

These three cases also showcase the importance of choosing \bar{C} that is a good upper bound on γ . When $\bar{R} = \mathcal{R}$, the agent should select D that is strictly smaller than b . On the other hand, when $\bar{R} \in \mathcal{R}$, the agent should select $D \geq b$, as it will be apparent from the result of the next subsection, in particular Theorem 5.3. While the agent knows b , it does not know if $\bar{R} = \mathcal{R}$ or $\bar{R} \in \mathcal{R}$.

5.2.2 Underestimating Attack Parameter

In this subsection, we analyze the case when $D < \gamma$. We first state our result, and then discuss its implications.

Theorem 5.3. If $\gamma > D$, then $\pi_D = \pi_\gamma$ is the unique solution of the optimization problem (P2b) with $C = f_{C_2; g}$.

Therefore, together with Theorem 5.2, Theorem 5.3 is showing the importance of having a good prior knowledge about the attack parameter γ . In particular:

- When the attack did not poison the reward function (i.e., $\mathcal{R} = \bar{R}$), overestimating γ implies that π_D might not be equal to π_γ for larger values of D , even though π_γ is uniquely optimal under \bar{R} . This can have a detrimental effect since in this case $V^{\pi_D} < V^{\pi_\gamma} = V^{\pi^*}$.
- When the attack did poison the reward function \bar{R} (i.e., $\mathcal{R} \in \bar{R}$), underestimating γ implies $\pi_D = \pi_\gamma$, but π_γ might be suboptimal. In this case, the defense policy does not limit the negative influence of the attack at all, i.e., $V^{\pi_D} = V^{\pi_\gamma} = V^{\pi^*}$.

We further discuss nuances to selecting D in Section 7.

6 Experimental Evaluation

In this section we evaluate our defense strategy in an experimental setting in order to better understand its efficacy and robustness. We focus on the setting from Section 5.2: C_2 attack cost functions, which is known to the agent, i.e., $C = f_{C_2; g}$, with an unknown attack parameter γ .

Given the results in Section 5 (Theorems 5.2 and 5.3), we use the linear programming formulation (P3) together with the CVXPY solver Diamond & Boyd (2016); Agrawal et al. (2018) for calculating the solution to the defense optimization problem (P2b). In the experiments, due to limited numerical precision, V^{π_D} is calculated

with a tolerance parameter, set to 10^{-4} by default.⁴. In other words, $\| \mathbf{v} - \mathbf{f}(\mathbf{s}; \mathbf{a}) \|_1 \leq \epsilon$.




Figure 4: Experimental environments: Figures(a) and (b) show the Navigation and Grid world environment respectively while (c) and (d) show the target policy in these environments. For comparison, in the navigation environment, $\epsilon = 0.26$ and $\delta = 0.45$ while in the grid world environment, $\epsilon = 1.75$ and $\delta = 0.70$.

(a) Navigation, PreAttack (b) Navigation, PostAttack (c) Grid world, PreAttack (d) Grid world, PostAttack

Figure 5: Robustness of the defense policy against random perturbation. Results are based on average of 100 runs for each data point. Error bars around the data points indicate standard error.

Navigation environment. Our first environment, shown in Figure 4a is the Navigation environment taken from Rakhsha et al. (2021). The environment has 9 states and 2 possible actions. The reward function is action independent and has the following values: $\bar{R}(s_0; \cdot) = \bar{R}(s_1; \cdot) = \bar{R}(s_2; \cdot) = \bar{R}(s_3; \cdot) = -2.5$, $\bar{R}(s_4; \cdot) = \bar{R}(s_5; \cdot) = 1$ and $\bar{R}(s_6; \cdot) = \bar{R}(s_7; \cdot) = \bar{R}(s_8; \cdot) = 0$. When the agent takes an action, it will successfully navigate in the direction shown by the arrows with probability 0.9; otherwise, the next state will be sampled uniformly at random. The bold arrows in the figure indicate the attacker's target policy. The initial state is s_0 and the discounting factor equals 0.99.

Grid world environment. For our second environment, shown in Figure 4b, we use the grid world environment from Ma et al. (2019) with slight modifications in order to ensure ergodicity we add a 10% failure probability to each action, sampling the next state randomly in case of failure. The environment has 18 states and 4 actions: up, down, right and left. The white, gray and blue cells in the figure represent the states and the black cells represent walls. In the white and gray states, the agent will attempt to go in the direction specified by its action if there is a neighboring state in that direction. If there is no such state, the agent will attempt to stay in its own place. In the blue state G, the agent will attempt to stay in its own place regardless of the action taken. In all states, each attempt will succeed with probability 0.9; with probability 0.1, the next state will be sampled uniformly at random. In the gray and white states, the agent's reward is a function of the state it is attempting to visit. Attempting to visit a gray, white and blue state will yield a reward of -1, 1 and 2 respectively. If the agent is in a blue state, it will always receive a reward of 0. The bold arrows in the figure specify the attacker's target policy. The initial state is S and the discounting factor equals 0.9.

Policy score for different values of parameters. We first analyze the score of our defense policy in

⁴The value was chosen because the CVXPY solver uses a precision of 10^{-5} .

both environments with different values of γ and β . For comparison, we also report the scores of the target policy (π^*) and the optimal policy (π^*). The results are shown in Figures 4c and 4d. As seen in the figures, as long as $\beta > \gamma$, our defense policy significantly improves the agent's score compared to γ .

Robustness to perturbations. We now analyze our algorithm's robustness towards uncertainties in the reward functions used by the attacker and the defender. For our first experiment, which we call PreAttack, we randomly perturb the attacker's input. In particular, the input to the defender's optimization problem \mathcal{R} is sampled from $A(c_2; \bar{R} + N(0; \sigma^2 I); \gamma; \gamma)$ where I is the identity matrix, N denotes the multivariate normal distribution and σ is the perturbation parameter varied in the experiment. For our second experiment, called PostAttack, we randomly perturb the reward vector after the attack, sampling the defender's input from $A(c_2; \bar{R}; \gamma; \gamma) + N(0; \sigma^2 I)$. In both experiments we use $\gamma = 0.1$ and $\beta = 1$. As explained below, when calculating π^* , we also experiment with a larger tolerance parameter of 0.1 , denoting the defense policy in this case with π_{D+} .

The results can be seen in Figure 5. As seen in the figures, our defense policy consistently improves on the baseline obtained with no defense (i.e., γ). It is also clear that the PostAttack perturbations have a greater negative impact on our defense strategy's score. Results for π_{D+} indicate that this is due to random perturbations prohibiting our algorithm from identifying all of the elements in \mathcal{R} . While having a higher tolerance parameter helps with robustness, it can also lead to a lower performance when there is no noise because π_{D+} would falsely include additional elements. We leave choosing the tolerance parameter in a more systematic way for future work.

7 Concluding Discussions

In this paper, we introduced an optimization framework for designing defense strategies against reward poisoning attacks that change an agent's reward structure in order to steer the agent to adopt a target policy. We analyzed the utility of using such defense strategies, providing characterization results and provable guarantees on their performance. Moving forward we see several interesting future research directions.

Beyond the worst-case utility. In this paper, we defined the defense objective as the maximization of the agent's worst-case utility. While this is a sensible objective, there are other objectives that one could analyze. For example, instead of focusing on the absolute performance, one can try to optimize performance relative to the target policy. Notice that this is a somewhat different, and possibly a weaker goal, given that the target policy can have arbitrarily bad utility under \bar{R} . Additionally, one could study the agent's sub-optimality gap, i.e., the difference between its score and the score of the optimal policy π^* , and compare it to the sub-optimality gap of the target policy. Going back to the example in Figure 3, we can see that the defense strategy can significantly reduce the suboptimality gap relative to not having a defense (in this case, from $10 - 3 = 7$ to $10 - 22 = -3$ (2:7)). However, when the original reward function already forces the target policy, i.e., when $\bar{R} = \mathcal{R}$, the suboptimality gap of the target policy is equal to 0. In Figure 3, this happens if the target policy is action a_1 . Nevertheless, in this case, the suboptimality gap of the defense strategy would also be low, equal to $1 - 2$ (since the defense would randomly select a_1 or a_7), and generally upper bounded by $\beta - \gamma = 2$ if the rewards of sub-optimal actions $\{a_2, \dots, a_7\}$ were different. Providing a full theoretical treatment for the general setting is an interesting research direction.

Informed prior. We did not model prior knowledge that an agent might have about the attacker or the underlying reward function. In practice, the agent may have some information about the underlying true reward function. Incorporating such considerations calls for a Bayesian approach that could improve the agent's defense by, e.g., ruling out implausible candidates for \bar{R} in the agent's inference of \bar{R} given \mathcal{R} .

Selecting β and non-oblivious attacks. The results in Section 5.2 indicate that choosing good β is important for having a functional defense. In practice, a selection procedure for β should take into account the cost that the attacker has for different choices of γ , as well as game-theoretic considerations: attacks might not be oblivious in that the strategy for selecting γ might depend on the strategy for selecting β . Namely, a direct consequence of Theorem 5.2 is that the attack optimization problem (P1) can successfully achieve its goal if it chooses a large value of γ to large enough values. However, the cost of the attack also grows with γ , so the attack (if strategic) also needs to reason about β when selecting γ . We leave the full game-theoretic characterization of the parameter selection problem for the future work.

Unknown-model and scalability. Following prior work, we focused on attacks and defenses that have

access to an accurate transition model and operate in tabular settings. Given Theorems 5.1, 5.2 and 5.3, the defense problem can be solved using the linear program (P3), which is similar in size to the attack optimization problem (P1'). As such, we expect our method to be computationally scalable as long as the attack optimization problem can be solved.

For many RL problems however, the transition and reward functions are not known and the MDP is too large to be modeled in tabular settings. RL solutions for such problems typically rely on function approximation. For future work, it would be interesting to study the problem, for both attack and defense, in this more realistic scenario. An immediate question is how the attack problem would need to change in order to generalise to this setting. In terms of defense, while the new setting would likely pose new challenges, we expect the general min-max formulation in Section 3, as well as the techniques used for solving the optimization problems (P2a) and (P2b), to remain useful.

Acknowledgments

The authors would like to thank the anonymous reviewers for their valuable comments and suggestions. This research was, in part, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) project number 467367360.

References

- Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
- Idan Amir, Idan Attias, Tomer Koren, Roi Livni, and Yishay Mansour. Prediction with corrupted expert advice. *CoRR*, abs/2002.10286, 2020.
- J Andrew Bagnell, Andrew Y Ng, and Je G Schneider. Solving uncertain markov decision processes. Technical report, Carnegie Mellon University, 2001.
- Kiarash Banihashem, Adish Singla, Jiarui Gan, and Goran Radanovic. Admissible policy teaching through reward design. *CoRR*, abs/2201.02185, 2022.
- Vahid Behzadan and Arslan Munir. Whatever does not kill deep reinforcement learning, makes it stronger. *CoRR*, abs/1712.09344, 2017.
- Dimitri P Bertsekas. Control of uncertain systems with a set-membership description of the uncertainty PhD thesis, Massachusetts Institute of Technology, 1971.
- Battista Biggio and Fabio Roli. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84:317–331, 2018.
- Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In *ICML*, 2012.
- Ilija Bogunovic, Arpan Losalka, Andreas Krause, and Jonathan Scarlett. Stochastic linear bandits robust to adversarial attacks. *CoRR*, abs/2007.03285, 2020.
- Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data. In *STOC*, pp. 47–60, 2017.
- Gabriela F Cretu, Angelos Stavrou, Michael E Locasto, Salvatore J Stolfo, and Angelos D Keromytis. Casting out demons: Sanitizing training data for anomaly sensors. *IEEE Symposium on Security and Privacy*, pp. 81–95. IEEE, 2008.
- Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Jacob Steinhardt, and Alistair Stewart. Sever: A robust meta-algorithm for stochastic optimization. In *ICML*, pp. 1596–1606, 2019.

- Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research* 17(1):2909–2913, 2016.
- Christos Dimitrakakis, David C Parkes, Goran Radanovic, and Paul Tylkin. Multi-view decision processes: The helper-ai problem. In *NeurIPS*, pp. 5443–5452, 2017.
- European Commission. Ethics Guidelines for Trustworthy Artificial Intelligence. URL: <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>, 2019. [Online; accessed 15-January-2021].
- Ahana Ghosh, Sebastian Tschiatschek, Hamed Mahdavi, and Adish Singla. Towards deployment of robust cooperative ai agents: An algorithmic framework for learning adaptive policies. In *AAMAS*, pp. 447–455, 2020.
- Anupam Gupta, Tomer Koren, and Kunal Talwar. Better algorithms for stochastic bandits with adversarial corruptions. In *COLT*, pp. 1562–1578, 2019.
- Ronan Hamon, Henrik Junklewitz, and Ignacio Sanchez. Robustness and explainability of artificial intelligence. *Publications Office of the European Union*, 2020.
- Ling Huang, Anthony D Joseph, Blaine Nelson, Benjamin IP Rubinstein, and J Doug Tygar. Adversarial machine learning. In *ACM workshop on Security and artificial intelligence*, pp. 43–58, 2011.
- Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel. Adversarial attacks on neural network policies. *CoRR*, abs/1702.02284, 2017.
- Yunhan Huang and Quanyan Zhu. Deceptive reinforcement learning under adversarial manipulations on cost signals. In *GameSec*, pp. 217–237, 2019.
- Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research* 30(2):257–280, 2005.
- Kwang-Sung Jun, Lihong Li, Yuzhe Ma, and Xiaojin Zhu. Adversarial attacks on stochastic bandits. In *NeurIPS*, pp. 3644–3653, 2018.
- Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *ICML*, pp. 1885–1894. PMLR, 2017.
- Pang Wei Koh, Jacob Steinhardt, and Percy Liang. Stronger data poisoning attacks break data sanitization defenses. *CoRR*, abs/1811.00741, 2018.
- Aounon Kumar, Alexander Levine, and Soheil Feizi. Policy smoothing for provably robust reinforcement learning. *arXiv preprint arXiv:2106.11420*, 2021.
- Bo Li, Yining Wang, Aarti Singh, and Yevgeniy Vorobeychik. Data poisoning attacks on factorization-based collaborative filtering. In *NeurIPS*, pp. 1885–1893, 2016.
- Shiau Hong Lim, Huan Xu, and Shie Mannor. Reinforcement learning in robust markov decision processes. In *NeurIPS*, pp. 701–709, 2013.
- Yen-Chen Lin, Zhang-Wei Hong, Yuan-Hong Liao, Meng-Li Shih, Ming-Yu Liu, and Min Sun. Tactics of adversarial attack on deep reinforcement learning agents. In *IJCAI*, pp. 3756–3762, 2017.
- Fang Liu and Ness B. Shroff. Data poisoning attacks on stochastic bandits. In *ICML*, pp. 4042–4050, 2019.
- Guanlin Liu and Lifeng Lai. Provably efficient black-box action poisoning attacks against reinforcement learning. *Advances in Neural Information Processing Systems* 34, 2021.
- Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *STOC*, pp. 114–122, 2018.
- Thodoris Lykouris, Max Simchowitz, Aleksandrs Slivkins, and Wen Sun. Corruption robust exploration in episodic reinforcement learning. *CoRR*, abs/1911.08689, 2019.

- Yuzhe Ma, Kwang-Sung Jun, Lihong Li, and Xiaojin Zhu. Data poisoning attacks in contextual bandits. In *GameSec* pp. 186 204, 2018.
- Yuzhe Ma, Xuezhou Zhang, Wen Sun, and Jerry Zhu. Policy poisoning in batch reinforcement learning and control. In *NeurIPS*, pp. 14543 14553, 2019.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.
- H Brendan McMahan, George J Gordon, and Avrim Blum. Planning in the presence of cost functions controlled by an adversary. In *ICML*, pp. 536 543, 2003.
- Shike Mei and Xiaojin Zhu. Using machine teaching to identify optimal training-set attacks on machine learners. In *AAAI*, pp. 2871 2877, 2015.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *CVPR*, pp. 2574 2582, 2016.
- Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D Joseph, Benjamin IP Rubinstein, Udam Saini, Charles A Sutton, J Doug Tygar, and Kai Xia. Exploiting machine learning to subvert your spam filter. *LEET*, 8:1 9, 2008.
- Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *CVPR*, pp. 427 436, 2015.
- Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780 798, 2005.
- Andrea Paudice, Luis Muñoz-González, Andras Gyorgy, and Emil C Lupu. Detection of adversarial training examples in poisoning attacks through anomaly detection. *CoRR*, abs/1802.03041, 2018.
- Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *ICML*, pp. 2817 2826, 2017.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming* John Wiley & Sons, Inc., 1994.
- Goran Radanovic, Rati Devidze, David Parkes, and Adish Singla. Learning to collaborate in markov decision processes. In *ICML*, pp. 5261 5270, 2019.
- Amin Rakhsha, Goran Radanovic, Rati Devidze, Xiaojin Zhu, and Adish Singla. Policy teaching via environment poisoning: Training-time adversarial attacks against reinforcement learning. In *ICML*, 2020.
- Amin Rakhsha, Goran Radanovic, Rati Devidze, Xiaojin Zhu, and Adish Singla. Policy teaching in reinforcement learning via environment poisoning attacks. *Journal of Machine Learning Research*, 22(210):1 45, 2021.
- Anshuka Rangi, Haifeng Xu, Long Tran-Thanh, and Massimo Franceschetti. Understanding the limits of poisoning attacks in episodic reinforcement learning. *arXiv preprint arXiv:2208.13663*, 2022.
- Kevin Regan and Craig Boutilier. Robust policy computation in reward-uncertain mdps using nondominated policies. In *AAAI*, volume 24, 2010.
- Jacob Steinhardt, Pang Wei Koh, and Percy Liang. Certified defenses for data poisoning attacks. In *NeurIPS*, pp. 3520 3532, 2017.
- Yanchao Sun, Da Huo, and Furong Huang. Vulnerability-aware poisoning mechanism for online rl with unknown dynamics. In *International Conference on Learning Representations*, 2020.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

- Umar Syed, Michael Bowling, and Robert E Schapire. Apprenticeship learning using linear programming. In ICML , pp. 1032 1039, 2008.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In ICLR , 2014.
- Csaba Szepesvári. The asymptotic convergence-rate of q-learning. *NeurIPS*, volume 10, pp. 1064 1070, 1997.
- Aviv Tamar, Shie Mannor, and Huan Xu. Scaling up robust mdps using function approximation. In ICML , pp. 181 189, 2014.
- Edgar Tretschk, Seong Joon Oh, and Mario Fritz. Sequential attacks on agents for long-term adversarial goals. CoRR, abs/1805.12487, 2018.
- Fan Wu, Linyi Li, Chejian Xu, Huan Zhang, Bhavya Kailkhura, Krishnaram Kenthapadi, Ding Zhao, and Bo Li. Copa: Certifying robust policies for online reinforcement learning against poisoning attacks. arXiv preprint arXiv:2203.08398, 2022.
- Han Xiao, Huang Xiao, and Claudia Eckert. Adversarial label flips attack on support vector machines. In ECAI , pp. 870 875, 2012.
- Huang Xiao, Battista Biggio, Gavin Brown, Giorgio Fumera, Claudia Eckert, and Fabio Roli. Is feature selection secure against training data poisoning? In ICML , pp. 1689 1698, 2015.
- Haoqi Zhang and David C. Parkes. Value-based policy teaching with active indirect elicitation. In AAAI , 2008.
- Huan Zhang, Hongge Chen, Chaowei Xiao, Bo Li, Duane Boning, and Cho-Jui Hsieh. Robust deep reinforcement learning against adversarial perturbations on observations. CoRR, abs/2003.08938, 2020a.
- Xuezhou Zhang, Xiaojin Zhu, and Stephen Wright. Training set debugging using trusted items. In AAAI , volume 32, 2018.
- Xuezhou Zhang, Yuzhe Ma, Adish Singla, and Xiaojin Zhu. Adaptive reward-poisoning attacks against reinforcement learning. In ICML , 2020b.
- Xuezhou Zhang, Yiding Chen, Xiaojin Zhu, and Wen Sun. Robust policy gradient against strong data corruption. CoRR, abs/2102.05800, 2021.

A List of Appendices

In this section we provide a brief description of the content provided in the appendices of the paper.

- ^ Appendix B provides an intuition of our results for the ℓ_2 attack cost using special MDPs in which the agent's actions do not affect the transition dynamics. The proofs of the results presented in this Appendix can be found in Appendix I.
- ^ Appendix C provides additional details regarding the experiments.
- ^ Appendix D contains some background on reward poisoning attacks, and a brief overview of the MDP properties that are important for proving our formal results.
- ^ Appendix E contains characterization results for the attack optimization problem (P1).
- ^ Appendix F contains proofs of the formal results in Section 5.
 - The proof of Theorem 5.1 is in Section F.1.
 - The proof of Theorem 5.2 is in Section F.2.
 - The proof of Theorem 5.3 is in Section F.3.

Appendix G contains the proofs for the results in Section 4 relating to Guarantees on values. The appendix also includes an optimization framework for solving the defense optimization problem (P2a) for the ℓ_1 norm as well as additional characterization results for the defense optimization problem for more general cost functions which are used for proving the complexity results in Section 4.

The characterization result for the defense optimization problem with for the ℓ_1 norm is provided in Section G.1.

Proof of Theorem 4.4 is provided in Section G.2.

Additional characterization results for the defense optimization problem are provided in Section G.3.

Proof of Proposition 4.2 is provided in Section G.4

Proof of Theorem 4.3 is provided in Section G.5.

Appendix H contains the proofs of the formal results in Section 4 relating to the computational complexity of the defense optimization problem as well as the computational complexity result for the ℓ_1 attack cost.

The hardness result for the ℓ_1 norm is provided in Section H.1.

Proof of Theorem 4.5 is provided in Section H.2

Appendix I contains a formal treatment of the results presented in Appendix B

B Intuition of Results using Special MDPs

In this Appendix, we describe characterization results on the ℓ_2 attack cost for special MDPs, in which the agent's actions do not affect the transitions, that is, we assume that

$$P(s; a; s^0) = P(s; a^0; s^0) \quad \forall s; a; a^0; s^0. \quad (3)$$

Variants of the above condition have been studied in the literature (e.g., Szepesvári (1997); Dimitrakakis et al. (2017); Sutton & Barto (2018); Radanovic et al. (2019); Ghosh et al. (2020)). Note that this assumption implies that any two policies π^0 and π^1 have equal state occupancy measures, so we simplify the notation by denoting $\mu = \mu^0 = \mu^1$.

While the results from the previous sections incorporate this special case, we study this setting because: i) the optimal solutions to the defense problem have a simple form, enabling us to provide intuitive explanations of our main results from the previous sections, ii) using this setting, we show a tightness result for Theorem 5.2.

A more formal exposition of our results for this setting including the proofs can be found in Appendix I.

B.1 Optimal Defense Policy

In this subsection, we provide the intuition behind defense policies for the unknown parameter setting with $\mathcal{D} = \mathcal{Y}$ (Section 5.2.1). The key point about the assumption in Equation (3) is that it allows us to consider each state separately in the defense optimization problems. In particular, it can be shown that the optimization problem (P3) is equivalent to solving jSj optimization problems of the form

$$\max_{(jS) \subseteq \mathcal{P}(A)} \min_{(a) \in \mathcal{A}(jS)} \left(\sum_{s \in jS} \mu(s) \left(R(s; a) - \max_{y \in \mathcal{Y}(s)} R(s; y) \right) \right) \quad (P3b)$$

where $\mu_s = \sum_{a \in \mathcal{A}(jS)} \mu(s; a)$.⁵ If we instantiate Theorem 5.2 for special MDPs by putting $\mu_s = \min_{D; \mathcal{B}} \mu_s$, the set \mathcal{S}_s has an intuitive description: it is the set of all second-best actions (w.r.t R) in state s such that their poisoned reward is greater than or equal to $R(s; y(s)) - \frac{\mu_s}{\mu(s)}$. It turns out that the defense policy for states selects an action uniformly at random from the set $\mathcal{S}_s = \{a \in \mathcal{A}(jS) \mid R(s; a) \geq R(s; y(s)) - \frac{\mu_s}{\mu(s)}\}$. In other words, the defense policy D is given by:

$$D(a|jS) = \begin{cases} \frac{1}{|\mathcal{S}_s|} & \text{if } a \in \mathcal{S}_s \\ 0 & \text{otherwise} \end{cases}$$

⁵See Lemma 8 in Banihashem et al. (2022).

To see why, note that the objective in (P3b) only improves as we put more probability on selecting $y(s)$ (since $y(s)$ is optimal under \mathcal{R}). However, the constraints in (P3b) require that the selection probability of any action in \mathcal{S} has to be at least as high as the selection probability of $y(s)$, which in turn give us the uniform at random selection rule. Figure 3 illustrates attack and defense policies for special MDPs using a single-state MDP with action set $\{a_1, \dots, a_7\}$. To obtain defense policy π_D , we can solve the optimization problem (P3b), which implies that π_D should select an action uniformly at random from the set $\{y(s) \mid \mathcal{S} = \{a_1, a_4, a_7\}\}$.

C Additional Details Regarding Experiments

In this section we provide additional details regarding the experiments. The source code for our experiments, as well as instructions for replicating our results can be found in the Supplementary Material.

C.1 Implementation details

Both the attacker optimization problem (P1), and the defense optimization problems (P2a) and (P2b) are convex since, by Theorems 5.1 and 5.2, the defense optimization problem reduces to the linear program (P3). We use CVXPY to calculate their solutions. The code for solving these optimization problems can be found in the file `MDP.py`. The specific functions used for solving these problems are as follows:

- ^ The function `attack` implements the attacker's optimization problem (P1). As explained in the main text, this is done by solving (P1') since (P1') is equivalent to (P1).
- ^ The function `defend_known` implements the optimization problem (P2a). As explained in Section 6, the tolerance parameter is set to 10^{-4} (default value).
- ^ The function `defend_unknown` implements the optimization problem (P2b).

C.2 Running time

Following prior work Rakhsha et al. (2021), to test the running times, we use the chain environment from Rakhsha et al. (2021), but with different number of states (additional states are added between s_2 and s_3 , and the corresponding transitions and rewards are defined analogously to those for \mathcal{S}_2). The attack and defense parameters are set to $\gamma = 0.1$ and $\beta = 0.2$. Table 2 shows the average running times (across 10 runs) of the attack optimization problem (P1') and the defense optimization problem (P2b) for different sizes of the chain environment.

It should be noted that the attack and defense optimization problems are similar in size, both solve a problem with at most $|S|(|A| - 1)$ constraints on $R^{j|S_i|A_j}$. However, solving the defense problem takes more time, partly because γ , β and ϵ need to be identified before (P3) can be solved.

The machine used for obtaining these results is a Macbook Pro personal computer with 4 Gigabytes of memory and a 2.4 GHz Intel Core i5 processor.

D Background and Additional MDP Properties

In this section we briefly outline the background and MDP properties that we utilize in our proofs.

D.1 Reward Poisoning Attacks

In this section, we provide some background on the cost-efficient reward poisoning attacks, focusing on the results from Rakhsha et al. (2021).

The setting studied in Rakhsha et al. (2021) incorporates both the average and the discounted reward optimality criteria in a discrete-time Markov Decision Process (MDP), with finite state and action spaces. Our MDP setting is equivalent to their MDP setting under the discounted reward optimality criteria. This criteria can be specified by score v . As defined in the main text, score v of policy π is the total expected

S	Problem		Attack		Defense	
4			0:01s	0:5ms	0:05s	1:6ms
10			0:01s	0:2ms	0:09s	1:5ms
20			0:01s	0:1ms	0:17s	4:8ms
30			0:02s	2:0ms	0:27s	9:7ms
50			0:04s	6:8ms	0:56s	34:6ms
70			0:07s	3:0ms	1:02s	69:7ms
100			0:13s	5:4ms	1:83s	91:2ms

Table 2: Run time of the attack and defense optimization problems for the chain environment with varied number of states $|S|$. Reported numbers are average of 10 runs; standard error is shown with .

return scaled by factor 1 :

$$= E \left(\sum_{t=1}^{\infty} \gamma^t R(s_t; a_t) \right);$$

where the state s_1 is sampled from the initial state distribution , and subsequent states s_t are obtained by executing policy in the MDP. Actions a_t are sampled from policy .

As explained in the main text, the following result is important for our analysis, since it allows us to simplify the optimization problem (P1) into the optimization problem (P1').

Lemma D.1. (Lemma 1 in Rakhsha et al. (2021)) The score of a policy y is at least γ greater than all other deterministic policies if and only if its score is at least γ greater than the score of any policy y^f ; ag. In other words,

$$8 \ 2 \ \det \ n^f \ y^g : \ y \ + \ y \ () \ 8s; a \in \ y(s) : \ y \ y^f s; ag + \ y :$$

Remark D.2. As explained in Rakhsha et al. (2021), this lemma implies that the optimization problem (P1) is equivalent to (P1'). Furthermore, the optimization problem is always feasible since any policy can be made optimal with sufficient perturbation of the reward function as formally shown by Rakhsha et al. (2021) and Ma et al. (2019).

D.2 Overview of Important Quantities

Next, we provide an overview of standard MDP quantities and the quantities introduced in the main text that are important for our analysis.

In addition to score , we consider state-action value function, or Q-value function, defined as

$$Q(s; a) = E \left(\sum_{t=1}^{\infty} \gamma^t R(s_t; a_t) \mid s_1 = s; a_1 = a \right);$$

In other words, $Q(s; a)$ is the total expected return when the first state is s , the first action is a , while subsequent states s_t and actions a_t are obtained by executing policy in the MDP.

We consider two occupancy measures. By we denote the state-action occupancy measure in the Markov chain induced by policy :

$$(s; a) = E \left(\sum_{t=1}^{\infty} \gamma^t \mathbb{1}[s_t = s; a_t = a] \right);$$

Given MDP M , the set of realizable occupancy measures under any (stochastic) policy π is denoted by \mathcal{O} . Note that the following holds:

$$h_{i,j} = \sum_{s \in \mathcal{S}} \mu(s) R_i(s) \quad (4)$$

where $h_{i,j}$ in the above equation computes a dot product between two vectors of size $|\mathcal{S}|$ (i.e., two vectors in $\mathbb{R}^{|\mathcal{S}|}$). We also denote by μ the state occupancy measure in the Markov chain induced policy π , i.e.:

$$\mu(s) = \mathbb{E} \left[\sum_{t=1}^{\infty} \gamma^{t-1} \mathbb{1}[s_t = s] \right]; \quad \mu \in \mathcal{O}$$

Note that

$$\sum_{s,a} \mu(s; a) = \sum_s \mu(s) = 1$$

State-action occupancy measure and state occupancy measure satisfy

$$\mu(s; a) = \mu(s) \pi(a|s); \quad (5)$$

which for deterministic π is equivalent to

$$\mu(s; a) = \mathbb{1}[\pi(a|s) = 1] \mu(s); \quad (6)$$

Apart from the standard MDP quantities mentioned above, we also mention quantities introduced in the main text. We denote by \mathcal{A} state-action pairs $(s; a)$ for which the margin between b^y and $b^{y^f; s; ag}$ is equal to γ , i.e.:

$$\mathcal{A} = \{(s; a) : b^{y^f; s; ag} - b^y = \gamma\}; \quad (7)$$

which can be expressed through reward function R using state-action occupancy measures:

$$\mathcal{A} = \{(s; a) : \sum_{y \in \mathcal{Y}} \gamma y R(s; a) = \gamma\};$$

Finally, quantity $f^{s; ag}(\mu)$ measures how well the occupancy measure of μ is aligned with $\gamma^{f; s; ag}$ relative to γ :

$$f^{s; ag}(\mu) = \sum_{y \in \mathcal{Y}} \gamma y R(s; a); \quad (8)$$

D.3 Occupancy Measures as Linear Constraints

In this subsection, we introduce the Bellman flow linear constraints that characterize \mathcal{O} and \mathcal{A} . In order to characterize \mathcal{O} , we require the following constraints:

$$\sum_a \mu(s; a) = \sum_{s'} \mu(s') P(s; a; s); \quad (9)$$

$$\mu(s; a) \geq 0; \quad (10)$$

The importance of these constraints is reflected in the following lemma.

Lemma D.3. (Theorem 2 in Syed et al. (2008)) Let μ be a vector that satisfies the Bellman flow constraints (9) and (10). Define policy π as

$$\pi(a|s) = \frac{\mu(s; a)}{\sum_a \mu(s; a)}; \quad (11)$$

Then μ is the state-action occupancy measure of π , in other words $\mu = \mu^\pi$. Conversely, if π is a policy with state-action occupancy measure (i.e., $\mu = \mu^\pi$) then μ satisfies the Bellman flow constraints (9) and (10), as well as Equation(11).

As for π , it is well-known (e.g., see Rakhsha et al. (2021)) that a vector π is the state occupancy measure for policy π (i.e., $\pi = \pi$), if and only if

$$\pi(s) = (1 - \gamma) \pi(s) + \sum_{s', a} \gamma P(s, a | s', a) \pi(s') \quad (12)$$

E Attack Characterization Results

E.1 Characterization results for the ℓ_2 attack cost

In this section we provide characterization results for the attack optimization problem (P1) for the ℓ_2 norm, i.e., $C = \ell_2$. We will later use these results for proving the formal results presented in Section 5.1 and Section 5.2. In addition, these results provide intuition for our results about the more general ℓ_p norms, which we will discuss in the next sections. In particular, the main result of this appendix is a set of Karush Kuhn Tucker (KKT) conditions that characterize the solution to the optimization problem (P1). As we focus on the cost function $c = \ell_2$ in this section, we will drop the dependence on c in $A(c; R; \gamma; \gamma)$.

To compactly express the KKT characterization results, let us introduce state occupancy difference matrix $\mathbf{R}^{\text{diff}} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ as a matrix with rows consisting of the vectors $\pi(s') - \pi(s)$ for all neighboring policies $s', s \in \mathcal{S}$. Additionally, for all $s, a \in \mathcal{S} \times \mathcal{A}$, we use $\pi(s, a)$ to denote the transpose of the row of \mathbf{R}^{diff} corresponding to (s, a) . Note that $\pi(s, a)$ is a column vector. In this notation, given Remark D.2 and Equation (4), the optimization problem (P1) is equivalent to

$$\begin{aligned} \min_{\mathbf{R}} \quad & \frac{1}{2} \|\mathbf{R}\|_F^2 \\ \text{s.t.} \quad & \mathbf{R} \mathbf{1} \leq \mathbf{0} \end{aligned} \quad (\text{P1}')$$

where $\mathbf{1}$ is a $|\mathcal{S}| \times 1$ vector whose each element equal to 1, and \leq specifies that the left hand side is element-wise less than or equal to the right hand side. Given this notation, the following lemma states the KKT conditions for a reward function \mathbf{R} (i.e., an $|\mathcal{S}| \times |\mathcal{S}|$ vector) to be an optimal solution to the optimization problem (P1).

Lemma E.1. (KKT characterization) \mathbf{R} is a solution to the optimization problem (P1) if and only if there exists an $|\mathcal{S}| \times 1$ vector λ such that

$$\begin{aligned} (\mathbf{R} - \mathbf{R}^{\text{diff}}) + \lambda^T &= \mathbf{0} && \text{stationarity;} \\ \mathbf{R} + \lambda \mathbf{1} &\leq \mathbf{0} && \text{primal feasibility;} \\ \lambda &\geq \mathbf{0} && \text{dual feasibility;} \\ \sum_{s, a \in \mathcal{S} \times \mathcal{A}} \pi(s, a) (\pi(s, a)^T \mathbf{R} + \lambda) &= \mathbf{0} && \text{complementary slackness} \end{aligned}$$

where $\mathbf{0}$ denotes an $|\mathcal{S}| \times |\mathcal{S}|$ vector whose each element equal to 0, and likewise, $\mathbf{1}$ denotes an $|\mathcal{S}| \times 1$ vector whose each element equal to 1.

Proof. Since (P1) is always feasible (Remark D.2) and all of the constraints are linear, strong duality holds. Now, the Lagrangian of the optimization problem is equal to

$$L = \frac{1}{2} \|\mathbf{R}\|_F^2 + \lambda^T (\mathbf{R} + \lambda \mathbf{1});$$

and taking the gradient with respect to \mathbf{R} gives us

$$\nabla_{\mathbf{R}} L = (\mathbf{R} - \mathbf{R}^{\text{diff}}) + \lambda^T = \mathbf{0};$$

The statement then follows by applying the standard KKT conditions. \square

Remark E.2. (Uniqueness) The solution to the optimization problem (P1) is unique since the objective $\frac{1}{2} \|\mathbf{R}\|_F^2$ is strongly convex.

The above lemma, has the following important consequence.

Lemma E.3. Reward function R satisfies $\bar{R} = A(R; \gamma; \gamma)$ if and only if there exists some $s; a \geq 0$ such that

$$R = \bar{R} + \sum_{(s;a) \in \mathcal{Y}} X_{(s;a)} \gamma^{f(s;a)} \gamma^i$$

Proof. To prove the statement, we use Lemma E.1. The primal feasibility condition in the lemma always holds as $\bar{R} \in A(\bar{R}; \gamma; \gamma)$. Therefore $\bar{R} \in A(R; \gamma; \gamma)$ if and only if there exists $s; a \geq 0$ such that the other three conditions hold. Note that the complementary slackness condition is equivalent to

$$\delta(s; a \in \mathcal{Y}(s)) : (s; a) = 0 \implies (s; a)^T (R + \gamma - \bar{R}) = 0 \iff \delta(s; a) \geq \gamma : (s; a) = 0$$

Therefore from dual feasibility, stationarity and complementary slackness it follows that $\bar{R} \in A(R; \gamma; \gamma)$ if and only if there exists $s; a \geq 0$ such that

$$R = \bar{R} + \sum_{(s;a)} X_{(s;a)} \gamma^{f(s;a)} \gamma^i < 0; \\ \delta(s; a) \geq \gamma : (s; a) = 0$$

The Lemma therefore follows by setting $s; a = \delta(s; a)$ since setting $(s; a) = 0$ for all $(s; a) \geq \gamma$ is equivalent to not summing over the terms corresponding to $(s; a) \geq \gamma$ in the stationarity condition. \square

A direct consequence of this lemma is the following result.

Corollary E.4. Assume that $\bar{R} = A(R; \gamma; \gamma)$ and $\bar{R} \in \mathcal{R}$. It follows that

$$b = \gamma;$$

where

$$b = \min_{s; a \in \mathcal{Y}(s)} \max_{i \in \mathcal{H}} b_{\gamma^{f(s;a)}}^i$$

Proof. Assume to the contrary that $b \in \mathcal{Y}$. Given the primal feasibility condition in Lemma E.1, $b \in \mathcal{Y}$. Therefore $b > \gamma$. It follows that

$$\delta(s; a \in \mathcal{Y}(s)) : b_{\gamma^{f(s;a)}} > \gamma \implies \delta(s; a) = 0$$

Given Lemma E.3, this implies that $R = \bar{R}$, which contradicts the initial assumption $R \in \mathcal{R}$. \square

F Proofs of Section 5.1

This section of the appendix contains the proofs of the formal results presented in Section 5.

F.1 Proof of Theorem 5.1

Before proving the theorem we prove some results that we need for the proof of this theorem, as well as for the results in later sections.

Lemma F.1. Consider policy π with state-action occupancy measure μ . Solution π^* to the following optimization problem:

$$\min_{\pi} \quad \text{s.t.} \quad \bar{R} = A(R; \gamma; \gamma); \tag{P4}$$

satisfies:

$$\min = \begin{cases} b & \text{if } \delta_{s;a} \geq \gamma : f_{s;ag}(\cdot) \geq 0 \\ 1 & \text{otherwise} \end{cases}$$

Proof. We separately analyze the two cases: the case when $f_{s;ag}(\cdot) \geq 0$ for all $(s;a) \in \mathcal{S} \times \mathcal{A}$ holds, and the case when it does not.

Case 1: If $f_{s;ag}(\cdot) \geq 0$ for all $(s;a) \in \mathcal{S} \times \mathcal{A}$, then by using Equation (4) and Lemma E.3 we obtain that

$$b = \min_{\mathbf{R}} \sum_{(s;a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E}_{\mathbf{y}} [R_{s;a} D_{s;a} + \mathbb{E}_{\mathbf{y}} [y f_{s;ag}(\cdot) R_{s;a}]] \geq 0$$

Therefore, $\min = b$. Furthermore, from Lemma E.3, we know that $\mathbf{R} = \mathbf{R}^*$ satisfies the constraint in the optimization problem (P4), so the score of the optimal solution to (P4) is $\min = b$.

Case 2: Now, consider the case when $f_{s;ag}(\cdot) < 0$ for a certain state-action pair $(s;a) \in \mathcal{S} \times \mathcal{A}$. Let $\epsilon_{s;a}$ be an arbitrary positive number. From Lemma E.3, we know that

$$\mathbf{R} = \mathbf{R}^* + \epsilon_{s;a} \mathbf{D}_{s;a} + \mathbb{E}_{\mathbf{y}} [y f_{s;ag}(\cdot) \epsilon_{s;a}]$$

satisfies the constraint in the optimization problem (P4), and hence is a solution to (P4). Moreover, by using this solution together with Equation (4), we obtain

$$b = \min_{\mathbf{R}} \sum_{(s;a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E}_{\mathbf{y}} [R_{s;a} D_{s;a} + \mathbb{E}_{\mathbf{y}} [y f_{s;ag}(\cdot) R_{s;a}]] = \sum_{(s;a) \in \mathcal{S} \times \mathcal{A}} \epsilon_{s;a} f_{s;ag}(\cdot) \quad (13)$$

Since $\epsilon_{s;a}$ can be arbitrarily large and $f_{s;ag}(\cdot) < 0$, while b is fixed, \min can be arbitrarily small. Hence, the score of the optimal solution to (P4) is unbounded from below, i.e., $\min = -\infty$. \square

We can now prove Theorem 5.1, that is the following statement.

Statement: Consider the following optimization problem parameterized by:

$$\begin{aligned} \max_{\mathbf{R}} \quad & \sum_{(s;a) \in \mathcal{S} \times \mathcal{A}} \mathbb{E}_{\mathbf{y}} [R_{s;a} D_{s;a} + \mathbb{E}_{\mathbf{y}} [y f_{s;ag}(\cdot) R_{s;a}]] \\ \text{s.t.} \quad & \mathbb{E}_{\mathbf{y}} [y f_{s;ag}(\cdot) R_{s;a}] \geq 0 \quad \forall (s;a) \in \mathcal{S} \times \mathcal{A} \end{aligned} \quad (P3)$$

For $\mathbf{R} = \mathbf{R}^*$, this optimization problem is always feasible, and its optimal solution \mathbf{R}_{\max} specifies an optimal solution to the optimization problem (P2a) with

$$D_{\max}(a|s) = \frac{\max_{a^0} D_{\max}(s; a)}{\max_{a^0} D_{\max}(s; a^0)} \quad (14)$$

The score of $D_{\max}(a|s)$ is lower bounded by $\frac{b}{D_{\max}}$.

Proof. The feasibility of the problem follows from Theorem 4.4⁶. Note that \mathbf{R}_{\max} always exists since (P3) is maximizing a continuous function over a closed and bounded set. Concretely, the constraints $\mathbb{E}_{\mathbf{y}} [y f_{s;ag}(\cdot) R_{s;a}] \geq 0$ and Equations (9) and (10) each define closed sets, and since $\sum_j \pi_j = 1$, the set is bounded.

In order to see why \mathbf{R}_{\max} specifies an optimal solution to (P2a), note that we can rewrite (P2a) as

$$\max_{\mathbf{R}} \min_{\mathbf{D}}$$

⁶The proof of Theorem 4.4 does not rely on this result.

where μ_{\min} is the solution to the optimization problem (P4). Due to Lemma F.1, this could be rewritten as

$$\begin{aligned} \max & b \\ \text{s.t.} & f_{s;a}^{\mu} \geq 0 \quad \forall (s;a) \in \mathcal{S} \end{aligned}$$

Namely, maximizing a function $f(x)$ subject to constraint $x \in X$ (where $X \subseteq \mathbb{R}$) is equivalent to maximizing $\tilde{f}(x)$, where

$$\tilde{f}(x) = \begin{cases} f(x) & \text{if } x \in X \\ 1 & \text{o.w.} \end{cases}$$

Due to (4) and (8), the constrained optimization problem above can be rewritten as

$$\begin{aligned} \max & \mu \\ \text{s.t.} & \mu \leq f_{s;a}^{\mu} \quad \forall (s;a) \in \mathcal{S} \end{aligned}$$

Therefore, given Lemma D.3, μ_{\max} specifies a solution to (P2a) via (2).

Finally, note that the constraints of the optimization problem (P3) ensure that a policy whose occupancy measure is equal to μ_{\max} will have $f_{s;a}^{\mu} \geq 0$ in other words, $f_{s;a}^{\mu_{\max}}$ is non-negative for all $(s;a) \in \mathcal{S}$. Due to Lemma F.1, we know that such policy will have the worst case utility equal to b . Therefore, $\mu_{\max} = b$. \square

Remark F.2. Given Lemma D.3, the constraint (2) can equivalently be replaced with constraints (9) and (10), making the optimization problem (P3) a linear program.

F.2 Proof of Theorem 5.2

The proof of the theorem is similar to the proof of Theorem 5.1 and builds on two lemmas which we introduce in this section.

Lemma F.3. Set $b = \min_{(s;a) \in \mathcal{S}} f_{s;a}^{\mu}$, where

$$b = \min_{(s;a) \in \mathcal{S}} \left(\sum_{y \in \mathcal{Y}} b_{y|s;a}^{\mu} \right)$$

Reward function R satisfies $R = A(R; \mu; \cdot)$ for some $\mu \in (0; b]$ if and only if

$$R = \mu + \sum_{(s;a) \in \mathcal{S}} \mu_{s;a} f_{s;a}^{\mu}$$

for some $\mu_{s;a} \geq 0$.

Proof. We divide the proof into two parts, respectively proving the sufficiency and the necessity of the condition.

Part 1 (Necessity): Assume that $R = A(R; \mu; \cdot)$ for some $\mu \in (0; b]$. From the stationarity and dual feasibility conditions in Lemma E.1, we deduce

$$\forall (s;a) \in \mathcal{S} : R = \mu + \sum_{y \in \mathcal{Y}} \mu_{y|s;a} (f_{s;a}^{\mu}(y) - \mu) \tag{15}$$

We claim that $\mu_{y|s;a} = 0$ for all $(s;a) \in \mathcal{S}$. Note that this would imply the lemma's statement by setting $\mu_{s;a} = \sum_{y \in \mathcal{Y}} \mu_{y|s;a}$, since the terms corresponding to $(s;a) \in \mathcal{S}$ could be skipped in the summation of (15).

To see why the claim holds, assume that $(s; a) \notin \mathcal{D}$ for some $(s; a)$ where $a \in \mathcal{Y}(s)$. From complementary slackness, we know that $(s; a)^T R + \tilde{\nu} = 0$, which implies that

$$b = \min_{(s; a) \in \mathcal{Y}(s)} \left((s; a)^T R \right) \quad (s; a)^T R = \tilde{\nu} \tag{16}$$

However, $\tilde{\nu} = b$ holds by primal feasibility. Therefore, all the inequalities are equalities, which implies $\tilde{\nu} = b$. Since $\tilde{\nu} \in \mathcal{D}$, we conclude that $\tilde{\nu} = \min_{\mathcal{D}} b; \mathbf{g} = \dots$. Since all of the inequalities in (16) are indeed equalities, we conclude

$$(s; a)^T R = \dots = (s; a)^2 \dots;$$

which proves the claim.

Part 2 (Sufficiency): Assume that

$$R = \mathbb{R}^+ \sum_{(s; a) \in \mathcal{Y}(s)} X_{s; a} \dots \mathcal{Y}(s; a) \dots \mathcal{Y} \dots;$$

for some $\dots \geq 0$. Set $\tilde{\nu} = \dots$ and note that $\tilde{\nu} \in \mathcal{D}$ by definition. Set

$$(s; a) = \begin{cases} \dots & \text{if } (s; a) \in \mathcal{D} \\ 0 & \text{o.w.} \end{cases} \dots;$$

We now verify all the conditions of Lemma E.1 hold. Stationarity and dual feasibility hold because $R = \mathbb{R}^+ \sum_{(s; a) \in \mathcal{Y}(s)} (s; a) \dots (s; a)$ and $\dots < 0$. Primal feasibility holds because $\tilde{\nu} = \min_{\mathcal{D}} b; \mathbf{g} = b$. Finally, complementary slackness holds because

$$(s; a) \notin \mathcal{D} \Rightarrow (s; a)^2 = \dots \Rightarrow (s; a)^T R + \tilde{\nu} = 0:$$

□

Lemma F.4. Let \dots_{\min} be the solution to the following optimization problem

$$\min_{\mathcal{R}} \dots \text{ s.t. } \mathcal{R} = A(\mathcal{R}; \mathcal{Y}; \dots) \wedge 0 < \dots \in \mathcal{D} \tag{P5}$$

Then

$$\dots_{\min} = \begin{cases} b & \text{if } \dots \mathcal{D}(s; a)^2 \dots : \dots \mathcal{Y}(s; a) \dots \mathcal{Y} \dots 0 \dots; \\ 1 & \text{o.w.} \end{cases}$$

where $\dots = \min_{\mathcal{D}} b; \mathbf{g}$, and

$$b = \min_{(s; a) \in \mathcal{Y}(s)} \dots \mathcal{Y} \dots \mathcal{Y}(s; a) \dots \mathcal{Y} \dots;$$

Proof. The proof is similar to the proof of Lemma F.1. We separately analyze the two cases: the case when $\dots \mathcal{Y}(s; a) \dots \mathcal{Y} \dots 0$ for all $(s; a) \in \mathcal{D}$ holds, and the case when it does not.

Case 1: If $\dots \mathcal{Y}(s; a) \dots \mathcal{Y} \dots 0$ for all $(s; a) \in \mathcal{D}$, then by using Equation (4) and Lemma F.3 we obtain that

$$b = \dots \mathcal{D}; \mathcal{R} \dots \mathcal{E} \dots \mathcal{X} \dots \mathcal{D}; \mathcal{Y} \dots \mathcal{Y}(s; a) \dots \mathcal{Y} \dots \mathcal{E} \dots 0;$$

Therefore, $\dots = b$. Furthermore, from Lemma F.3, we know that $R = \mathcal{R}$ satisfies the constraint in the optimization problem (P5), so the score of the optimal solution to (P5) is $\dots_{\min} = b$.

Case 2: Now, consider the case when $f^{s;a}(y) < 0$ for a certain state-action pair $(s; a) \in \mathcal{S} \times \mathcal{A}$. Let $\epsilon_{s;a}$ be an arbitrary positive number. From Lemma F.3, we know that

$$R = \mathbb{R} + \epsilon_{s;a} \cdot \mathbb{1}_{(s;a)}; \quad y \in \mathcal{Y}^{f^{s;a}}$$

satisfies the constraint in the optimization problem (P5), and hence is a solution to (P5). Moreover, by using this solution together with Equation (4), we obtain

$$b = \mathbb{R} + \epsilon_{s;a} \cdot \mathbb{1}_{(s;a)}; \quad y \in \mathcal{Y}^{f^{s;a}}; \quad b = \mathbb{R} + \epsilon_{s;a} \cdot \mathbb{1}_{(s;a)} \cdot f^{s;a}(y)$$

Since $\epsilon_{s;a}$ can be arbitrarily large and $f^{s;a}(y) < 0$, while b is fixed, $\mathbb{1}_{(s;a)} \cdot f^{s;a}(y)$ can be arbitrarily small. Hence, the score of the optimal solution to (P5) is unbounded from below, i.e., $b_{\min} = -\infty$. \square

We are now ready to prove Theorem 5.2.

Statement: Assume that $\mathcal{D} \subseteq \mathcal{Y}$, and define $b = \min_{(s;a) \in \mathcal{S} \times \mathcal{A}} b^{s;a}$. Then, the optimization problem (P3) with $\mathcal{D} = \mathcal{D}$ is feasible and its optimal solution π_{\max} identifies an optimal policy $\pi_{\mathcal{D}}$ for the optimization problem (P2b) via Equation (2). This policy $\pi_{\mathcal{D}}$ satisfies $V^{\pi_{\mathcal{D}}} \geq b$.

Proof. The proof is divided into two parts, respectively proving the first and the second claim in the theorem statement.

Part 1 (Solution to (P2b)): We prove that the optimization problem (P3) is feasible, its optimal solution π_{\max} identifies an optimal solution to (P2b) via Equation (2), and satisfies $V^{\pi_{\max}} \geq b$.

The feasibility of the problem follows from Theorem 4.4. Note that π_{\max} always exists since (P3) is maximizing a continuous function over a closed and bounded set. Concretely, the constraints $y \in \mathcal{Y}^{f^{s;a}}; \quad 0 \leq \pi_j \leq 1$ and Equations (9) and (10) each define closed sets and since $\sum_j \pi_j = 1$, the set \mathcal{D} is bounded.

In order to see why π_{\max} specifies an optimal solution to (P2b), note that we can rewrite (P2b) as

$$\max_{\pi} \min_{(s;a)} b^{s;a};$$

where $\min_{(s;a)}$ is the solution to the optimization problem (P5). Due to Lemma F.4, this could be rewritten as

$$\begin{aligned} & \max b \\ & \text{s.t. } f^{s;a}(y) \geq 0 \quad \forall (s;a) \in \mathcal{S} \times \mathcal{A}; \end{aligned}$$

where $b = \min_{(s;a)} b^{s;a}$. Namely, maximizing a function $f(x)$ subject to constraint $x \in X$ (where $X \subseteq \mathbb{R}^n$) is equivalent to maximizing $\tilde{f}(x)$, where

$$\tilde{f}(x) = \begin{cases} f(x) & \text{if } x \in X \\ 1 & \text{o.w.} \end{cases};$$

Due to (4) and (8), the constrained optimization problem above can be rewritten as

$$\begin{aligned} & \max_{\pi} \mathbb{R} + \epsilon_{s;a} \cdot \mathbb{1}_{(s;a)} \\ & \text{s.t. } y \in \mathcal{Y}^{f^{s;a}}; \quad \mathbb{R} + \epsilon_{s;a} \cdot \mathbb{1}_{(s;a)} \cdot f^{s;a}(y) \geq b; \end{aligned}$$

Therefore, given Lemma D.3, π_{\max} specifies a solution to (P2b) via (2). Finally, given Lemma F.4, b satisfies the constraints of (P5) and therefore b is a lower bound on $V^{\pi_{\max}}$. \square

F.3 Proof of Theorem 5.3

Statement: If $\gamma > D$, then γ is the unique solution of the optimization problem (P2b), hence $D = \gamma$.

Proof. As in Theorem 5.2, set $\gamma = \min_{\mathbf{s}; \mathbf{a} \in \mathcal{Y}(\mathbf{s})} f(\mathbf{s}; \mathbf{a})$ where

$$f(\mathbf{s}; \mathbf{a}) = \min_{\mathbf{b} \in \mathcal{B}(\mathbf{s})} \mathbf{b}^T \mathbf{y} - \mathbf{b}^T \mathbf{y}^f(\mathbf{s}; \mathbf{a});$$

From the feasibility of the attack, we have that

$$\exists \mathbf{s}; \mathbf{a} \in \mathcal{Y}(\mathbf{s}) : \mathbf{b}^T \mathbf{y} - \mathbf{b}^T \mathbf{y}^f(\mathbf{s}; \mathbf{a}) > D \quad \Rightarrow \quad \gamma = D;$$

Therefore, given Lemma F.3, the constraint in the optimization problem (P2b) is satisfied only for $R = \mathbb{R}$. This reduces the optimization problem (P2b) to $\max_{\mathbf{b}} \mathbf{b}^T \mathbf{y}$, which has a unique optimal solution: γ . \square

G Proofs of Section 4

In this section, we provide the proofs of the results of Section 4 as well as additional results that characterize the defense optimization problem. While our results are stated for the defense optimization problem (P2a), all results hold for (P2b) as well with $\gamma = \min_{\mathbf{s}; \mathbf{a} \in \mathcal{Y}(\mathbf{s})} f(\mathbf{s}; \mathbf{a})$.

G.1 Solution to the defense optimization problem for the ℓ_1 attack cost

In this section, we present a convex optimization framework for solving the defense optimization problem (P2a) for $C = f_{C_1}$. Our main result is the following theorem, the proof of which is presented in Section G.3.

Theorem G.1. Let $\mathbf{R} \in \mathbb{R}^{|\mathcal{S}| \times (|\mathcal{A}| + 1)}$ be a matrix with rows consisting of the vectors $\mathbf{y}^f(\mathbf{s}; \mathbf{a})$ as in Section E and let $\mathbf{R}_{\mathcal{Y}}$ be the sub-matrix of \mathbf{R} consisting of the rows corresponding to \mathcal{Y} . Define the set $\mathcal{U} \subseteq \mathcal{S} \times \mathcal{A}$ as the set of state action pairs $(\mathbf{s}; \mathbf{a})$ for which the following optimization problem is feasible.

$$\exists \mathbf{s}; \mathbf{a} : \mathbf{R}^T(\mathbf{s}; \mathbf{a}) = 1 \text{ and } \mathbf{R}^T(\mathbf{s}; \mathbf{a}) = 1 \text{ and } < 0;$$

where $<$ denotes coordinate-wise inequality. Consider the following optimization problem:

$$\max_{\mathbf{b} \in \mathbb{R}^n} \mathbf{b}^T \mathbf{y} \quad \text{s.t.} \quad (\mathbf{a}; \mathbf{s}) \in \mathcal{U} \quad \forall \mathbf{s}; \mathbf{a} \in \mathcal{Y}(\mathbf{s}) \quad (17)$$

The optimization problem (17) is always feasible and its solution is a solution to the defense optimization problem (P2a) with $C = f_{C_1}$.

G.2 Proof of Theorem 4.4

Statement: Consider any policy $\pi_{\mathcal{Y}}$ s.t. $\pi_{\mathcal{Y}}(\mathbf{y}(\mathbf{s})|\mathbf{s}) = 0$. For $C = f_{C_p}$ s.t. $p \in [1; 1)$, the optimal value of problem (P2a) is bounded below by $\gamma_{\mathcal{Y}}$.

Proof. We claim that $\mathbb{R}(\mathbf{s}; \mathbf{a}) \geq R(\mathbf{s}; \mathbf{a})$ for all R satisfying $\mathbb{R} \in \mathcal{A}(R)$ and all $\mathbf{s}; \mathbf{a} \in \mathcal{Y}(\mathbf{s})$. To see why, assume that if this not the case for some $\mathbf{s}; \mathbf{a}$ and define \mathbb{R} as

$$\mathbb{R}(\mathbf{s}; \mathbf{a}) = \begin{cases} R(\mathbf{s}; \mathbf{a}) & \text{if } \mathbf{s}; \mathbf{a} = \mathbf{s}; \mathbf{a} \\ \mathbb{R}(\mathbf{s}; \mathbf{a}) & \text{o.w.} \end{cases}$$

It is clear that $\mathbb{R} \in \mathcal{A}(R)$ for all $p \in [1; 1)$. We claim that \mathbb{R} is feasible for the attack optimization problem (P1) with parameters $C_p; R; \gamma_{\mathcal{Y}}$. This would contradict the assumption, $\mathbb{R} \in \mathcal{A}(R)$ as it would mean \mathbb{R} is not an optimal solution to (P1), finishing the proof.

⁷The results related to computational hardness are discussed separately in Appendix H.

To prove the claim, since (P1) is equivalent to (P1'), we need to show that

$$R_{y^f s; ag} \leq R_y \quad y;$$

for all $s; a \in \mathcal{Y}(s)$. By definition of R however, $R_{y^f s; ag} = D_{y^f s; ag}$ as

$$R_{y^f s; ag} - D_{y^f s; ag} = R - R; \quad y^f s; ag = 0; \tag{18}$$

where the first equality follows from Equation (4), and the second equality follows from the fact that R only differs from D in $(s; a)$, $R_{y^f s; ag}(s; a) = 0$. Similarly $R_y \leq D_y$ as

$$R_y - D_y = R - R; \quad y = 0;$$

where the second inequality follows from the fact that R is never less than D . Therefore,

$$R_{y^f s; ag} \leq R_y \leq D_{y^f s; ag} \leq D_y \leq y;$$

where the second inequality follows from the assumption that $R \geq A(R)$. We have therefore shown (18), finishing the proof. \square

G.3 Characterization of the inner minimization problem in (P2a)

We begin by providing characterisation results for the inner minimization problem in (P2a) for different known cost functions. These results can be seen as extensions of Lemma F.1. Formally, for $x \in \mathcal{X}$ consider the following optimization problem

$$\min_R \quad \text{s.t.} \quad R = A(c_p; R; y; y); \tag{P4}$$

The following lemmas characterize the value of the above optimization problem for different values of p . In stating these lemmas, we use $\mathcal{R}^{j^S(j^A-1)j^Sj^A}$ be a matrix with rows consisting of the vectors $D_{y^f s; ag} \quad y$ as in Section E and let \mathcal{R} be the sub-matrix of $\mathcal{R}^{j^S(j^A-1)j^Sj^A}$ consisting of the rows corresponding to y .

Lemma G.2. Let π be a fixed policy and assume that $1 < p < \infty$ is a fixed number. Define the function $u_p : \mathcal{R} \rightarrow \mathcal{R}$ as $u_p(x) = \text{sgn}(x) |x|^{p-1}$ where $\text{sgn}(x) = 1 [x > 0] - 1 [x < 0]$ and let $u_p : \mathcal{R}^n \rightarrow \mathcal{R}^n$ be its coordinate-wise extension to \mathcal{R}^n , i.e., $u_p(x)_i = u_p(x_i)$.

The solution to (P4) equals b if

$$; u_p(\mathcal{R}^T) \geq 0 \quad \mathcal{R} < 0;$$

and equals 1 otherwise.

Lemma G.3. Let π be a fixed policy. Define the function $u_1 : \mathcal{R} \rightarrow \mathcal{R}$ as

$$u_1(x) = \begin{cases} 1 & \text{if } x \geq 0; \\ -1 & \text{o.w.} \end{cases};$$

and let $u_1 : \mathcal{R}^n \rightarrow \mathcal{R}^n$ be its coordinate-wise extension to \mathcal{R}^n , i.e., $u_1(x)_i = u_1(x)_i$. The solution to (P4) equals b if

$$; u_1(\mathcal{R}^T) \geq 0 \quad \mathcal{R} < 0 \quad \text{s.t.} \quad \mathcal{R} \geq 0;$$

and equals 1 otherwise.

Lemma G.4. Let π be a fixed policy and let $U \subseteq \mathcal{S} \times \mathcal{A}$ be the set of state action pairs $(s; a)$ for which the following optimization problem is feasible.

$$\mathcal{R}(s; a) \geq \mathcal{R}^T(s; a) - 1 \quad \text{and} \quad \mathcal{R}^T(s; a) = 1 \quad \text{and} \quad \mathcal{R} < 0; \tag{19}$$

Then the solution to (P4) is 1 if $(\mathcal{R}^T(s; a) > 0$ for some $s; a \in U$ and is b otherwise.

When it is clear from context, we will drop the dependence on u_p . The proof of Lemmas G.2, G.3 and G.4 are provided below. Note that Lemma G.4 immediately implies Theorem G.1 since the feasibility of (17) already follows from Theorem 4.4.

Proof of Lemma G.2. Throughout the proof, we will drop the dependence on c_p , y and y in A . The proof follows a similar structure as the results for the ℓ_2 norm; namely, Lemmas E.1, E.3 and F.1. We begin by analyzing the constraint $A(R) = \bar{R}$ using the KKT conditions. Since $1 < p < \infty$, we can change the objective to $\frac{1}{p} \|R - \bar{R}\|_p^p$ for convenience. Of the four KKT conditions, primal feasibility holds if and only if $\|R - \bar{R}\|_p \leq \delta$. For the stationarity condition, forming the Lagrangian of (P1), we obtain

$$L = \frac{1}{p} \|R - \bar{R}\|_p^p + \lambda^T (R - y) = \int \frac{1}{p} (R(s; a) - \bar{R}(s; a))^p + \lambda^T R - y^T \lambda$$

Taking the gradient,

$$\begin{aligned} \nabla_R L = 0 \iff & \text{sgn}(R(s; a) - \bar{R}(s; a)) |R(s; a) - \bar{R}(s; a)|^{p-1} + \lambda^T(s; a) = 0 \quad \forall s; a \\ \iff & \lambda^T(s; a) = \text{sgn}(\bar{R}(s; a) - R(s; a)) |\bar{R}(s; a) - R(s; a)|^{p-1} \quad \forall s; a \\ \iff & \lambda^T(s; a)^{\frac{1}{p-1}} \text{sgn}(\lambda^T(s; a)) = \bar{R}(s; a) - R(s; a) \quad \forall s; a \\ \iff^{(i)} & u(\lambda^T) = \bar{R} - R; \end{aligned}$$

where for (i) we have used the definition of u . The complementary slackness condition states that

$$(s; a) = 0 \iff \delta(s; a) \geq \epsilon;$$

Finally, dual feasible states that $\lambda < 0$.

Returning to the condition $A(R) = \bar{R}$, primal feasibility always holds as $\bar{R} = A(\bar{R})$. We therefore conclude that a vector R satisfies $A(R) = \bar{R}$ if and only if the following problem is feasible

$$\begin{aligned} R &= \bar{R} + u(\lambda^T) \\ \lambda &< 0 \\ (s; a) &= 0 \iff \delta(s; a) \geq \epsilon \end{aligned}$$

By definition of u , this means that

$$A(R) = \bar{R} \iff R = \bar{R} + u(\lambda^T) \text{ for some } \lambda < 0;$$

Returning back to (P4), the problem can be rewritten as

$$\begin{aligned} \min_{R; \lambda} & \quad h; R; \lambda \\ \text{s.t.} & \quad \lambda < 0 \\ & \quad R = u(\lambda^T) + \bar{R}; \end{aligned}$$

or equivalently,

$$\min_{\lambda < 0} \quad h; u(\lambda^T) + \bar{R}; \tag{20}$$

Now, assume that $u(\lambda^T) < 0$ for some $\lambda < 0$. Observe that $u(c^{-1/p} x) = c^{1/p} x$ for positive constants $c > 0$. Therefore,

$$u(c^{-1/p} \lambda^T) = c^{1/p} u(\lambda^T); \quad u(c^{-1/p} \lambda^T) < 0$$

Since $1 < p < \infty$, letting $c \rightarrow \infty$, the value of $u(c^{-1/p} \lambda^T)$ can be made arbitrarily low. Therefore, the value of (20) equals h . Conversely, assume that $u(\lambda^T) \geq 0$ for all $\lambda < 0$. It follows that the value of (20) is bigger than equal to b . Since $\lambda = 0$ is feasible for (20), the value is exactly b . \square

Before we prove Lemma G.3, we will state and prove the following Lemma which characterizes the subgradient of the ℓ_1 norm.

Lemma G.5. The vectors $w; z \in \mathbb{R}^n$ satisfy $w \in \partial \|z\|_1$ if and only if (a) $z = 0$ and $\|w\|_1 = 1$ or (b) $\|w\|_1 = 1$ and $z \in \Theta(w)$ where

$$\Theta(w) := \left\{ x \in \mathbb{R}^n : x_i \begin{cases} \geq c & \text{if } w_i > 0 \\ \leq -c & \text{if } w_i < 0 \\ \in [-c, c] & \text{if } w_i = 0 \end{cases} \text{ for some } c \in [0, 1] \right\}$$

Proof. By Proposition A.22 in Bertsekas (1971),

$$\partial \|z\|_1 = \text{conv}\{w \text{ s.t. } \|w\|_1 = 1; z^T w = \|z\|_1\} = \text{conv}\{w \text{ s.t. } \|w\|_1 = 1; z^T w = \|z\|_1\} \quad (21)$$

analyzing the above result, note that

$$\begin{aligned} z^T w &= \sum_i z_i w_i \\ &\stackrel{(a)}{=} \sum_i |z_i| |w_i| \\ &= \|z\|_1 \|w\|_1 \\ &\stackrel{(b)}{=} \|z\|_1 \end{aligned}$$

Since $z^T w = \|z\|_1$, equality holds in (b), implying $\|w\|_1 = 1$ or $z = 0$ and in (a), implying $z_i w_i = |z_i| |w_i|$ for all i . This means that if $|z_i| < \|z\|_1$, then $w_i = 0$, and if $|z_i| = \|z\|_1$, then $w_i = 0$ if $z_i = 0$ and $w_i = 0$ if $z_i = 0$. In other words (this time conditioning on w), if $\|w\|_1 < 1$, then $z = 0$. Otherwise, if $w_i > 0$ then $z_i = \|z\|_1$, if $w_i < 0$, then $z_i = -\|z\|_1$ and if $w_i = 0$ then z_i can be any value in $[-\|z\|_1; \|z\|_1]$. Therefore, the proof follows by setting $c = \|z\|_1$. (the last condition is obviously true but it is important to make explicit as will become clear later). This brings us to the following result. \square

We can now prove Lemma G.3.

Proof of Lemma G.3. The main ideas of the KKT analysis in the proof of Theorem G.2 still hold and the only condition that changes is the stationarity slackness condition. Formally, the lagrangian equals

$$L(\bar{R}, \lambda) = \bar{R}^T y + \lambda (\bar{R} - R)$$

Therefore, the stationarity condition can be written as

$$0 \leq \lambda + \bar{R} - R \leq 0 \quad \lambda \geq 0 \quad \bar{R} - R = 0 \quad (22)$$

By Lemma G.5, and given the fact that $\partial \|x\|_1 = \Theta(x)$, the above condition holds if and only if

$$\bar{R} \in \Theta(R) \wedge \lambda + \bar{R} - R = 0 \quad \lambda \geq 0 \quad \bar{R} - R = 0$$

Combining with the complementary slackness and dual feasibility, (P4) can be rewritten as

$$\begin{aligned} \min_{\lambda, \bar{R}} \quad & \lambda + \bar{R}^T y \\ \text{s.t.} \quad & \lambda \geq 0 \\ & \bar{R} \in \Theta(R) \wedge \lambda + \bar{R} - R = 0 \end{aligned}$$

Now, observe that $\partial \|c x\|_1 = \Theta(x)$ for all $c > 0$. Therefore, the $\bar{R} \in \Theta(R)$ inside the first clause of the last constraint is equivalent to $\bar{R} \in \Theta(0)$ since if $\bar{R} \in \Theta(0)$ and $\bar{R} \in \Theta(R)$, then $\bar{R} = 0$ can be satisfied by

replacing with $\frac{1}{k^T k_1}$. Similarly, the τ_1 inside the second clause is redundant. Therefore, the last constraint can be rewritten as

$$v^2 \theta(\tau)^T \epsilon_0 - v = 0$$

Now, observe that $\tau^T \epsilon_0$ is equivalent to ϵ_0 as the rows of τ are independent; this is because the only row with nonzero $(s; a \epsilon_y(s))$ is the one corresponding to $(s; a \epsilon_y(s))$. Therefore, the optimization problem can be rewritten as

$$\begin{aligned} \min_{v_i} & b + h; v_i \\ \text{s.t.} & < 0 \\ & v^2 \theta(\tau)^T \epsilon_0 - v = 0 \end{aligned}$$

Assume that there exists ϵ_0 satisfying < 0 such that $u(\tau)^T; < 0$. Then for any $c > 0$, $c u(\tau)^T \geq \theta(\tau)$ and therefore, the value of the optimization problem is 1. Otherwise, the value is lower bounded by b and the bound is attained with $v = 0$, concluding the proof. \square

Proof of Lemma G.4. We begin by analyzing (P1) as before. Forming the Lagrangian,

$$L = R \bar{R}_1 + \tau^T (R y)$$

Therefore, the stationarity condition states that

$$\tau^T \geq \theta R \bar{R}_1$$

As before, the primal feasibility condition holds, dual feasibility states that < 0 and complementary slackness states that $(s; a) = 0$ for all $(s; a) \neq y$. Therefore, $A(R) = \mathbb{R}$ if and only if

$$\tau^T \geq \theta \mathbb{R} R_1 \tag{23}$$

Note however that vectors $z; w$ satisfy $w \geq \theta z k_1$ if and only if

$$\begin{aligned} z_i & \geq 1 \text{ if } z_i > 0; & z_i & \geq 0 \text{ if } w_i = 1 \\ w_i & = 1 \text{ if } z_i < 0; & z_i & > 0 \text{ if } w_i = 1 \\ z_i & \in [1; 1] \text{ if } z_i = 0 & z_i & = 0 \text{ if } w_i \in (-1; 1) \end{aligned}$$

Now, assume that $(ajs) > 0$ for some $(s; a)$ for which (19) is feasible and θ be the vector satisfying (19). By (23), we conclude that the vector $R(s; a) = \mathbb{R}(s; a) - t^{-1} [(s; a) = (-s; a)]$ satisfies $A(R) = \mathbb{R}$ for any $t > 0$. Therefore, if $(ajs) > 0$, the solution to (P4) is 1. Conversely, if (19) is not feasible for any such $s; a$, then it follows that $R(s; a) = \mathbb{R}(s; a)$ for all $R; s; a$ satisfying $A(R) = \mathbb{R}$ and $(ajs) > 0$. Therefore, the value of the optimization problem is lower bounded by b . Since $A(\mathbb{R}) = \mathbb{R}$, the lower bound is attainable which proves the lemma. \square

G.4 Proof of Proposition 4.2

Statement: Let $c_{\text{const}}(R; R^0)$ be a constant cost function, and assume that $c_{\text{const}} \geq C$. Then problem (P2a) is unbounded from below.

Proof. We will show that for any ϵ , the value of the inner optimization problem of (P2a) equals 1, thereby proving the result.

Let π be an arbitrary policy and consider an arbitrary vector R . By definition of c_{const} and the fact that \mathbb{R} is feasible for (P1), $\mathbb{R} \geq A(c_{\text{const}}; R; y; y)$. Therefore, $R; c_{\text{const}}$ are feasible for (P2a). Since R was arbitrary, it can be chosen such that the ϵ is arbitrarily low, proving the claim. \square

G.5 Proof of Theorem 4.3

Statement: There exists MDP $M = (S; A; R; P; \gamma)$ for which problem (P2a) is unbounded from below when $c_1 \geq 2C$.

Proof. Let π be an arbitrary policy. We will build an example where the value of the inner minimization problem of (P2a) is $1 - \epsilon$.

Consider a single-state MDP with the reward function $R = (1; 0; 0; 0)$ where a_1 is the first action and $\gamma = 1$. Since R is symmetric with respect to $a_2; a_3; a_4$, we assume without loss of generality that $a_2 < a_3 < a_4$. For any x , consider the reward function $R^0 = (1 - x; x; x; x)$ with actions $a_1; a_2; a_3; a_4$. We claim that $R^0 \notin A(c_1; R^0; \gamma; \gamma)$. To see why, consider the attack optimization problem (P1). The cost of the attack optimization problem is at least x because $R^0(a_1) - R^0(a_2) = 1 - 2x$ and any feasible R needs to satisfy $R(a_1) - R(a_2) \geq 1$, which implies

$$\begin{aligned} R^0 \in A(c_1; R^0; \gamma; \gamma) & \iff \max_{\pi} \sum_{a \in A} \pi(a) R^0(a) - R^0(a_2) \geq c_1 \\ & \iff \frac{1}{2} (R^0(a_1) - R^0(a_2) + R^0(a_2) - R^0(a_2)) \geq c_1 \\ & \iff \frac{1}{2} (2x - 1 + 1) \geq c_1 \end{aligned}$$

Since $R^0 \notin A(c_1; R^0; \gamma; \gamma) = x$, we conclude $R^0 \notin A(c_1; R^0; \gamma; \gamma)$ as claimed.

We now consider the score of the policy π under the reward function R^0 and show that, with the proper choice of x , it is unbounded from below. Since $R^0 \notin A(c_1; R^0; \gamma; \gamma)$, this would show that the value of the inner minimization in (P2a) equals $1 - \epsilon$, proving the theorem.

By definition, R^0 can be written as

$$R^0 = (a_1) + x((a_2) - (a_1) - (a_3) - (a_4)) = (a_1) + x((a_2) - (1 - (a_2))) = (a_1) + x(2(a_2) - 1):$$

Since $(a_2) < (a_3) < (a_4)$ however, we conclude that $(a_2) < 0.5$. Therefore, tending x to infinity finishes the proof. \square

H Computational hardness results

In this section, we provide computational hardness results for different choices of C , showing that the defense optimization problem (P2a) is NP-hard in different cases. While our results are stated for the defense optimization problem (P2a), all results hold for (P2b) as well with $\epsilon = \min_{D; b; g}$. Our proofs rely on the results of Appendix G and we therefore refer to this and rely on notation introduced in this Appendix; namely, the notation π and u_p as introduced in Lemma G.2.

H.1 Hardness result for $p = 1$

4.5 We begin by considering the case of c_1 . By reducing the 3SAT problem to the defense optimization problem (P2a), we will show that (P2a) is NP-hard. More formally, we prove the following theorem.

Theorem H.1. For $C = f(c_1, g)$, it is NP-hard to determine whether the optimal value of problem (P2a) is greater than or equal to $b - \epsilon$.

We assume without loss of generality that the clauses do not contain duplicate variables as the 3SAT instance remains NP-hard in this case.

Before proving the theorem, we prove a weaker result which essentially proves hardness assuming we can set the matrix A and vector y can be set arbitrarily, without the restriction that they correspond to an MDP.

Proposition H.2. Assume we are given an instance of 3SAT with clauses c_1, \dots, c_m and variables (x_1, \dots, x_n) . Define $k = 2n + 1$ and $\ell = 4n + m + 1$. It is possible to build, in polynomial time, a matrix $M \in \mathbb{R}^{k \times \ell}$, a vector $e \in \mathbb{R}^k$ such that e is strictly positive and

^ If the 3SAT instance is satisfiable, then there exists a $u \in \mathbb{R}^k$ satisfying $\langle u, e \rangle < 0$ such that

$$\langle u, M^T \rangle; e < \frac{1}{4}.$$

In addition, M^T does not contain any zero entries.

^ If the 3SAT instance is not satisfiable, then for any $u \in \mathbb{R}^k$ satisfying $\langle u, e \rangle < 0$,

$$\langle u, M^T \rangle; e > \frac{1}{4}.$$

Proof. We begin by an empty matrix and vector and in each step, add a value to e and add a column of size k to M .

We will let $(a_0; a_1; \dots; a_n; b_1; \dots; b_n)$ be placeholder names for the coordinates of M^T . Since the optimization problem we are considering involves the matrix product $M^T u$, for ease of notation, we will specify each column of M by the linear map forming the corresponding coordinate of $M^T u$. As an example, $a_0 + 4a_3 + 5b_3$ is a column that has value 1 in a_0 , 4 in a_3 , 5 in b_3 and 0 everywhere else.

Defining $c := 500k$, we first add the columns below.

$$c \begin{pmatrix} a_0 + \sum_{i=1}^n x_i^2 \\ a_i + \sum_{i=1}^n x_i^2 \\ b_i \end{pmatrix} \quad (24)$$

$$0.9a_0 + (a_i + b_i) \quad \text{for } i \in [n] \quad (25)$$

$$1.1a_0 + (a_i + b_i) \quad \text{for } i \in [n] \quad (26)$$

$$4(0.9a_0 + (a_i - b_i)) \quad \text{for } i \in [n] \quad (27)$$

$$4(0.9a_0 + (b_i - a_i)) \quad \text{for } i \in [n] \quad (28)$$

The value of e for (24), (25), (26), (27) and (28) is $(6n^2 + m \frac{1}{2})$, $3n$, $3n$, 1 and 1 respectively. Note that (24) is the all-ones vector. The constant c in (24) and the constant 4 in (27) and (28) will not matter for the proof, but will be useful for later.

Next for each clause with variables $(x_i; x_j; x_k)$, we will add the following column where d_i denotes a_i if the variable x_i appears as positive in the clause and $-a_i$ if the variable appears as negative.

$$0.95 a_0 + (d_i + d_j + d_k) \quad (29)$$

The value of e in the above is 1.

We now prove that $e \perp M^T u$ have the mentioned properties.

If the 3SAT instance is satisfiable, this is easy to show. Set $a_0 = 1$ and set $(a_i; b_i)$ to $(1; 0)$ if x_i is set to true in the satisfiable arrangement and to $(0; 1)$ if x_i is set to 0. It is clear that the coordinates of $u(M^T)$ corresponding to (24), (25) and (26) equal 1, -1 and -1 respectively. Furthermore, since $f(a_i - b_i; b_i - a_i) = f(1; -1)$, exactly half of the coordinates of $u(M^T)$ corresponding to (27) and (28) equal 1, while the other half equal -1 . Furthermore, all of the coordinates corresponding to (29) equal -1 . Therefore, the inner product $\langle u, M^T \rangle; e$ equals

$$(6n^2 + m \frac{1}{2}) - n - 3n - n - 3n + 0 - m = \frac{1}{2}$$

which proves the claim.

Conversely, assume that the 3SAT instance is not satisfiable. Let s_1, \dots, s_n denote the coordinates of e with s_1 corresponding to (24) and let $u[s]$ denote $u(\mathbb{M}^T)(s)$. Let $\epsilon > 0$ satisfying $\epsilon < 0$ be an arbitrary vector. First note that since $\epsilon > 0$, $\mathbb{M}^T(s_1)$ is strictly positive and therefore $u[s_1] = 1$. Since $e(s_1) = 6n^2 + m - \frac{1}{2}$, we need to show that

$$\sum_{s \in S_1} e(s) u[s] > 6n^2 + m + \frac{1}{2} + \frac{1}{4} \quad (30)$$

$$\sum_{s \in S_1} \frac{e(s) (u[s] + 1)}{2} > \frac{1}{2} (6n^2 + m + \frac{3}{4} + 2n - 3n + 2n - 1 + m) \quad (31)$$

$$\sum_{s \in S_1} e(s) \mathbb{1}[u[s] = 1] > n + \frac{3}{8} \quad (32)$$

Now note that since the value of e in coordinates corresponding to (25) and (26) is $> n + \frac{3}{8}$, if $u[s]$ equals one in any of these coordinates, then the claim is proved.

Therefore, we assume w.l.o.g that for any i ,

$$0.9a_0 \leq a_i + b_i \leq 1.1a_0$$

This implies that if any of a_i, b_i is strictly positive, so is a_0 . Since at least one of a_0, a_i, b_i needs to be strictly positive by the assumption $\epsilon > 0$, we conclude that a_0 is strictly positive.

Now note that $u[s]$ is 1 in more than half of the coordinates corresponding to (27) and (28), then (32) will hold. We can therefore assume w.l.o.g that at least half of these coordinates equal 1. We note however that if $u[s]$ equals 1 for both (27) and (28) for some $1 \leq i \leq n$, then

$$0.9a_0 \leq a_i + b_i \leq 0.9a_0 \Rightarrow a_0 = 0$$

which is not possible. We can therefore conclude that for any i , $u(\mathbb{M}^T)$ equals 1 in exactly one of the two coordinates corresponding to (27) and (28). Therefore, either $a_i + b_i \leq 0.9a_0$ or $b_i - a_i \leq 0.9a_0$. Since $a_i + b_i \geq 0.9a_0$, we get that either $a_i \leq 0.9a_0$ or $b_i \leq 0.9a_0$. Since $a_i + b_i \leq 1.1a_0$, the bigger one is $\geq 0.9a_0$ and the smaller one is $\leq 0.2a_0$.

Finally, note that if $u[s]$ equals 1 for any of the coordinates corresponding to (29), then (32) would hold since $u[s]$ was already 1 for half of the coordinates corresponding to (27) and (28). We can therefore assume that $u[s]$ is 0 for all the coordinates corresponding to (29). Note however that if $d_i + d_j + d_k \leq 0.95a_0$, then at least one of the d_i must have been > 0.2 . This means that all of the clauses must hold true (in the 3SAT sense) with x_i set to

$$x_i = \begin{cases} 1; & \text{if } a_i \leq 0.5; \\ 0; & \text{otherwise} \end{cases}$$

This is not possible however as we assumed that the 3SAT instance was not satisfiable. □

Claim H.3. Define \mathbb{M} as above and let s_1 denote the column of the matrix corresponding to (24). For all rows v in \mathbb{M} , we have

$$v(s_1) \leq \sum_{i=1}^5 |v(s_i)|$$

which further implies $\sum_{i=1}^5 |v(s_i)| \leq (1 + 1 + 1 + 1 + 1)v(s_1)$.

Proof. The claim follows trivially from the fact that as $v(s_1) = c$ and all other entries in v are less than 4. □

Next, we prove the following lemma which essentially states that for any desired value of γ , we can find a plausible reward function R such that γ equals this value.

Lemma H.4. Let M be an ergodic MDP with unspecified reward function and let π be a policy in this MDP. For any value of $\epsilon > 0$ and any set of state-action pairs $S \times A$ satisfying $\forall (s; a) : s \in S, a \in A$, there is a reward function R such that

1. R is feasible for the attack problem (P1), i.e., $\exists R \in \mathcal{R}^k$ such that $\|R\|_1 \leq \epsilon$.
2. $\|v^* - v^\pi\|_1 \leq \epsilon$.

Proof. Consider the following reward function,

$$R(s; a) = \begin{cases} 0 & \text{if } a = \pi(s) \\ \frac{\epsilon}{2} & \text{if } a \neq \pi(s) \\ 0 & \text{o.w.} \end{cases}$$

It is clear that $\|R\|_1 = \epsilon$ and

$$v^* - v^\pi = \begin{cases} \frac{\epsilon}{2} & \text{if } (s; a) \in S \times A \\ 0 & \text{o.w.} \end{cases}; \text{ for all } (s; a) \in S \times A$$

which proves the claim. \square

We now use the above results to construct an MDP, formally proving Theorem H.1. Given an instance of 3SAT, let M ; e denote the values specified in Proposition H.2. Recall that $M \in \mathbb{R}^{k \times k}$ and $e \in \mathbb{R}^k$ where $k = 2n + 1$ and $n = 4n + m + 1$. Let $\epsilon > 0$; $\delta \in [0, 1]$ be parameters to be specified later. (see Claims H.11, H.7, and H.5.) Intuitively, we need ϵ to be close to zero and δ to be close to one.

Given these values, we will build an MDP with reward vector R with $k + 3$ states and $k + 1$ actions for which $\|R\|_1 = \epsilon$ and the 3SAT instance is satisfiable if and only if

$$\exists R \in \mathbb{R}^k : \|R\|_1 \leq \epsilon; \exists v^* < 0; \quad (33)$$

In our construction, the states s_1, \dots, s_k will each correspond to the columns of M while the states $s_{k+2}, \dots, s_{k+3}; s_0$ will be new. Furthermore, in state s_0 , each of the actions $a \in A(s_0)$ will correspond to a row of M .

Let M denote the submatrix of M with only columns corresponding to $(s; \pi(s))$. We first observe that (33) can be simplified because $v^*(s; a)$ is 0 for all $(s; a) \in S \times A$. Therefore, (33) is equivalent to

$$\exists R \in \mathbb{R}^k : \|R\|_1 \leq \epsilon; \exists v^* < 0; \quad (34)$$

In order to specify this MDP, we first specify the transition probabilities of π . In state s_0 , following π leads to s_i for $i \in \{1, \dots, k\}$ with probability δ and leads to s_0 with the probability $(1 - \delta)b(s_i)$ where

$$b = \frac{e}{\|e\|_1}; \quad (35)$$

In states $s \in S$, following π will lead back to s with probability 1. The initial distribution of the MDP is chosen as $(s_0) = 1$ and $(s_i) = 0$ for $i \in \{1, \dots, k\}$.

It is straightforward to see that

$$v^*(s_0) = (1 - \delta); \quad \forall i \in \{1, \dots, k\} : v^*(s_i) = \frac{\delta}{2}; \quad \forall i \in \{1, \dots, k\} : v^*(s_i) = (1 - \delta) \frac{e(s_i)}{\|e\|_1};$$

Now, for each of the rows in M like e , we add an action a to state s_0 with the following transition probabilities.

$$P(s_0; a; s_j) = \begin{cases} \delta & \text{if } j = 0 \\ \frac{\delta}{4} \left(\sum_i e(s_i) \right) & \text{if } j \in \{1, \dots, k\} \\ \frac{\delta}{2} e(s_j) & \text{if } j = 1 \end{cases} \quad (36)$$

where ϵ is taken to be large enough such that for all i ,

$$\min_i b(s_i) \geq \frac{1}{k} \epsilon k_1$$

Note that the value of ϵ is the same for all the rows of matrix; it is set to the maximum of $\frac{1}{\min_i b_i} \frac{1}{k} \epsilon k_1$ across all rows of M .

Claim H.5. If $\epsilon \geq \frac{1}{10}$, the transition probabilities specified in (35) are valid, i.e., they are in the range $[0, 1]$ and $\sum_{s_0} P(s_0; a; s_0) = 1$:

Proof. In order to make sure these transition probabilities are valid, they need to sum to one and they all need to be non-negative. They sum to one by definition. As for being non-negative, it holds trivially for $j = 0$ and holds for $j \neq 0$ by definition of ϵ . Formally, for $j = 1$,

$$\begin{aligned} P(s_0; y(s_0); s_0) &= \frac{1}{4} \sum_i v(s_i) = \frac{1}{2} \frac{1}{4} \sum_i \frac{P_i v(s_i)}{(\min_i b_i) \frac{1}{k} \epsilon k_1} \\ &= \left(\frac{1}{2} \frac{1}{4} \right) > 0 \end{aligned}$$

and for $j = 1$,

$$\begin{aligned} P(s_0; y(s_0); s_j) + \frac{1}{2} \frac{v(s_j)}{k} &= P(s_0; y(s_0); s_j) + \frac{1}{2} \frac{j v(s_j)}{k} \\ &= (1 - \frac{1}{k}) \min_i b_i + \frac{1}{2} \frac{v(s_j)}{k} \geq \min_i b_i \\ &= (\min_i b_i) \left(1 - \frac{3}{2} \right) > 0 \end{aligned}$$

□

Now, using Lemma H.4, set the reward for the MDP such that r consists of the added actions in s_0 . Note that there are multiple actions in the states $s_i \neq 0$ as well; however, given Lemma H.4, their transition probabilities are not important and can be set arbitrarily as long as the MDP remains ergodic.

Claim H.6. For all $a \in \mathcal{A}(s_0)$,

$$(M^T)^{-1}(a) = \frac{1}{4} \left(\sum_j v(s_j) \right); \quad \frac{1}{4} \left(\sum_j v(s_j) \right); \quad (1 - \frac{1}{k}); \quad \frac{1}{2} v$$

where v is the row corresponding to a and we have abused notation by using $\frac{1}{2} v$ to denote the last k entries of the vector.

Proof. As before, it is straightforward to see that $\sum_{y \in \mathcal{A}(s_0)} P(s_0; y; s_0) = (1 - \frac{1}{k})$ and

$$\sum_{y \in \mathcal{A}(s_i)} P(s_0; y; s_i) = \frac{1}{2} \frac{1}{4} \left(\sum_j v(s_j) \right);$$

and

$$\sum_{y \in \mathcal{A}(s_i)} P(s_0; y; s_i) = (1 - \frac{1}{k}) b(s_i) + \frac{1}{2} v(s_i);$$

□

Claim H.7. If $\max_j \sum_{i=1}^n g_i \leq \frac{1}{4kek_1}$; then

$$\frac{kek_1}{(1-\epsilon)} \sum_{j=2}^n \sum_{i=1}^n x_i^0 y(s_j) A_{ij} \leq \frac{1}{8}$$

Proof.

$$\frac{kek_1}{(1-\epsilon)} \sum_{j=2}^n \sum_{i=1}^n x_i^0 y(s_j) A_{ij} = \frac{1 + (1-\epsilon)}{(1-\epsilon)} kek_1 \tag{37}$$

$$\frac{1}{4} kek_1 (1 + (1-\epsilon)) \leq \frac{1}{8} \tag{38}$$

□

Claim H.8. The 3SAT instance is satisfiable if and only if (34) holds.

Proof. Given Claim H.6, the sub-matrix of M corresponding to columns s_{j-1} equals

$$\frac{1}{2} \hat{M}$$

Let $u_1[s]$ denote $u_1^T (M^T)(s)$. For any s ,

$$\begin{aligned} \sum_{j=2}^n y(s_j) u_1[s_j] &= \sum_{j=2}^n y(s_j) u_1[s_j] + \frac{(1-\epsilon)}{kek_1} \epsilon u_1 \left(\frac{1}{2} \hat{M}^T \right) \\ &= \sum_{j=2}^n y(s_j) u_1[s_j] + \frac{(1-\epsilon)}{kek_1} \epsilon u_1 (\hat{M}^T) \end{aligned}$$

Multiplying both sides by $\frac{kek_1}{(1-\epsilon)}$, it follows that

$$\frac{kek_1}{(1-\epsilon)} \sum_{j=2}^n y(s_j) u_1[s_j] \leq \epsilon u_1 (\hat{M}^T) \leq \frac{kek_1}{(1-\epsilon)} \sum_{j=2}^n \sum_{i=1}^n x_i^0 y(s) u_1[s_j] A_{ij} \tag{39}$$

$$\frac{kek_1}{(1-\epsilon)} \sum_{j=2}^n \sum_{i=1}^n x_i^0 y(s) A_{ij} \tag{40}$$

$$\frac{1}{8} \tag{41}$$

Therefore, by Lemma H.4, if the 3SAT is satisfiable, then there exists a u_1 such that

$$\frac{kek_1}{(1-\epsilon)} \sum_{j=2}^n y(s_j) u_1[s_j] \leq \frac{1}{4} + \frac{1}{8} < 0$$

If the 3SAT is not satisfiable, then by Lemma H.4, for all feasible u_1 ,

$$\frac{kek_1}{(1-\epsilon)} \sum_{j=2}^n y(s_j) u_1[s_j] \geq \frac{1}{4} - \frac{1}{8} > 0$$

which proves the claim. □

H.2 Proof of Theorem 4.5

Statement: For $C = f_{C_p}$ s.t. $p \in [1; 1]^g$, it is NP-hard to determine whether the optimal value of problem (P2b) is greater than or equal to γ . We will prove the result using the same MDP as the previous Section.

We need to show the 3SAT instance is satisfiable if and only if

$$p \in [1; 1]^g; \exists R^k: \exists \alpha \wedge u_p(\alpha^T); \gamma < 0:$$

As before, let M denote the submatrix of α with only columns corresponding to $(s; \gamma(s))$. As in the previous section we note that because $\gamma(s; a)$ is 0 for all $s; a \in \gamma(s)$, the above condition is equivalent to

$$p \in [1; 1]^g; \exists R^k: \exists \alpha \wedge u_p(M^T); \gamma < 0: \quad (42)$$

We split the proof into two lemmas.

Lemma H.9. If the 3SAT instance is satisfiable, then (42) holds.

Proof. In this case, then the same proof as before basically holds. Formally, consider the used before in the proof and consider the vector M^T . We need to show that there exist p for which $\gamma; u_p(M^T)$ is negative. Note however that

$$\gamma; u_p(M^T) = \sum_{i=1}^k \gamma(s_i) u_p(M^T)(s_i)$$

For fixed $x \in \mathbb{R}$, $u_p(x)$ is a continuous function of p and $\lim_{p \rightarrow 1} u_p(x) = u_1(x)$. If we show that all of the coordinates in M^T are non-zero, this would imply that $\gamma; u_p(M^T)$ converges to $\gamma; u_1(M^T) < 0$ for large enough p which proves the claim. Therefore, $\gamma; u_p(M^T)$ is continuous in p and letting p be large enough proves the claim.

It remains to verify that all of the coordinates in M^T used in the above proof were non-zero. Since all of the coordinates of α were non-zero by our construction of α , we conclude that M^T is non-zero on s_1 . For $i \in \{2; \dots; g\}$, given claim H.3, we have $\alpha(s_i) > 0$, for all rows α of M , which means M^T is strictly negative on these states. Finally, the entry corresponding to s_0 equals $\gamma(s_0)$ in all rows of M , which implies $M^T(s_0)$ is strictly negative, proving the claim. \square

Lemma H.10. If (42) holds, then the 3SAT instance is satisfiable.

Proof. We assume that $\epsilon > 0$ without loss of generality; if $\epsilon = 0$ then $\gamma; u_p(M^T) = 0$ which is not negative.

Letting $u_p[s]$ denote $u_p(M^T)(s_i)$ and $d_p[s]$ denote $\gamma(s) u_p[s]$, then $\gamma; u_p(M^T)$ can be rewritten as

$$\begin{aligned} \gamma; u_p(M^T) &= \sum_{i=1}^k \gamma(s_i) u_p[s_i] & (43) \\ &= \sum_{i=1}^k d_p[s_i] + \sum_{(24)} d_p[s] + \sum_{(25)} d_p[s] + \sum_{(26)} d_p[s] + \sum_{(27)} d_p[s] + \sum_{(28)} d_p[s] + \sum_{(29)} d_p[s] & (44) \end{aligned}$$

where in the above, we have broken the sum in different parts, depending on what s_i corresponds to and we have abused the notation by using the set (24) denote all states that correspond to Equation (24). Note that the sum corresponding to (24) consists of a single state. We will also use s_{24} to denote this state.

We begin by stating some properties of u_p and d_p .

Claim H.11. Let ϵ be close enough to 1 such that

$$\frac{\epsilon}{4} (1 - \epsilon):$$

Then for all rows v from M ,

$$8s \in s_1 : |u_p[s]| \leq u_p[s_1] \quad (45)$$

Proof. We need to show that for all rows v in M , $|v(s_1) - v(s)| \leq |v(s) - v(s_0)|$ for all $s \in s_1$. For $i = 1$, this follows immediately from Claim H.3 and Claim H.6.

As for s_0 , we need to show that

$$v(s_1) - v(s_0) \leq \frac{1}{2} |v(s_1) - v(s_0)| \quad (1)$$

Note however that

$$\frac{1}{2} |v(s_1) - v(s_0)| \leq \frac{1}{4} |v(s_1) - v(s_0)| \quad (1)$$

by the assumption on ϵ . □

Claim H.12. For any $i \in [n]$, letting s_i and s_i^0 denote the states corresponding to (27) and (28) respectively.

$$d_p[s_i] + d_p[s_i^0] \leq 0:$$

Proof. To prove this, observe that

$$0:9a_0 (a_i - b_i) + 0:9a_0 (b_i - a_i) = 1:8a_0 \leq 0$$

There are therefore two possibilities: (a) Both $0:9a_0 (a_i - b_i)$ and $0:9a_0 (b_i - a_i)$ are non-negative. In this case, both $d_p[s_i]$ and $d_p[s_i^0]$ are non-negative and the claim follows. (b) $0:9a_0 (a_i - b_i) \leq 0$ and $0:9a_0 (b_i - a_i) \geq 0$ or vice versa. Assume w.l.o.g that $0:9a_0 (a_i - b_i) \leq 0$, i.e, $d_p[s_i] \leq 0$. In this case, $|0:9a_0 (a_i - b_i)| \leq |0:9a_0 (b_i - a_i)|$ and therefore, since $v(s_i) = v(s_i^0)$, we conclude that $d_p[s_i] \leq |d_p[s_i^0]|$ and therefore the claim follows. □

We split the proof into several cases.

Case 1: There exists $s \in (25) \cup (26)$ such that $d_p[s] \leq 0$.

Let \bar{s} be such a state. In this case, define $\mathcal{C}_p[s]$ as follows

$$\mathcal{C}_p[s] = \begin{cases} 0 & \text{if } s = \bar{s} \\ 0 & \text{if } s \in (27) \cup (28) \\ d_p[s] & \text{if } s = s_1 \\ |d_p[s]| & \text{o.w.} \end{cases}$$

We first claim that $\sum_{s \in (25) \cup (26)} d_p[s] \leq \sum_{s \in (27) \cup (28)} \mathcal{C}_p[s]$. To prove this, note that $d_p[s] \leq \mathcal{C}_p[s]$ for all $s \in (27) \cup (28)$. We therefore need to prove that

$$\sum_{s \in (27) \cup (28)} d_p[s] \leq 0;$$

which follows from Claim H.12. by summing over all $i \in [n]$.

Defining $\mathcal{S} := \mathcal{S}_n(f, s_1; \text{sg} [(27) [(28)])$, we note that

$$\begin{aligned} \sum_s \mathbb{E}_p[s] &= \sum_{s \in \mathcal{S}} d_p[s_1] \sum_{s \in \mathcal{S}} |d_p[s_1]| \\ &= \sum_{s \in \mathcal{S}} y(s_1) u_p[s_1] \sum_{s \in \mathcal{S}} y(s) |u_p[s_1]| \\ &\stackrel{(45)}{=} \sum_{s \in \mathcal{S}} y(s_1) u_p[s_1] \sum_{s \in \mathcal{S}} y(s) |u_p[s_1]| \\ &= |u_p[s_1]| \sum_{s \in \mathcal{S}} y(s_1) \sum_{s \in \mathcal{S}} y(s) A \end{aligned}$$

Defining $\mathcal{S}_+ := \mathcal{S}_n(f, s_2; s_1; s_0, g)$,

$$\begin{aligned} e(s_1) \sum_{s \in \mathcal{S}_+} e(s) &= 6n^2 + m \frac{1}{2} (2n - 1) 3n - m \\ &= 6n^2 + m \frac{1}{2} (6n^2 + 3n - m) \\ &= 3n \frac{1}{2} > \frac{1}{4}; \end{aligned}$$

Claim H.7 now gives a contradiction as the inner product in (42) becomes positive.

Case 2: $d_p[s] < 0$ for all $s \in \mathcal{S}$ [(26)] and $|a_j - b_j| \geq 0.9a_0$ for some $j \in [n]$.

In this case, we conclude that

$$a_i + b_i \in [0.9; 1.1]a_0 \quad \forall i \in [1; n]$$

Since either a_i or b_i must be strictly positive for some i (because $a_0 > 0$), we conclude that $a_0 > 0$.

Similar to case 1, we introduce a new vector \mathbb{E}_p such that $\sum_s d_p[s] = \sum_s \mathbb{E}_p[s]$.

Let \mathfrak{s}_j and \mathfrak{s}_j^0 denote the states corresponding to (27) and (28) respectively for $i = j$. Note that by assumption, $u_p[\mathfrak{s}_j]; u_p[\mathfrak{s}_j^0] \geq 0$. Assume without loss of generality that $u_p[\mathfrak{s}_j] \geq u_p[\mathfrak{s}_j^0]$. Define the vector \mathbb{E}_p as

$$\mathbb{E}_p[s] = \begin{cases} \geq 0 & \text{if } s \in \mathcal{S}_n(f, \mathfrak{s}_j; g) \\ \geq d_p[s] & \text{if } s \in \mathcal{S}_n(f, s_1; \mathfrak{s}_j; g) \\ \geq \sum_j |d_p[s_j]| & \text{o.w.} \end{cases}$$

As before, $\sum_s d_p[s] = \sum_s \mathbb{E}_p[s]$. More formally, for (27) and (28), we use Claim H.12 and for all the other states, we have $d_p[s] \leq \mathbb{E}_p[s]$.

We now claim that

$$\sum_{s \in \mathcal{S}_n(f, \mathfrak{s}_j; g)} \mathbb{E}_p[s] \geq u_p[\mathfrak{s}_j] \quad (46)$$

This is because

$$(M^T)(\mathfrak{s}_j) = \frac{(M^T)(\mathfrak{s}_j) + (M^T)(\mathfrak{s}_j^0)}{2} = 3.6a_0;$$

while $(M^T)(s) \geq 3.5a_0$ for all $s \in \mathcal{S}_n(f, \mathfrak{s}_j; g)$ [(26)] [(29)]. This is because $a_i + b_i \in [0.9; 1.1]a_0$.

Defining $\mathcal{E} := (25) \cup (26) \cup (29)$, we conclude that

$$\begin{aligned} & \sum_s d_p[s] \sum_s \mathcal{E}_p[s] \\ &= \sum_{s_1} y(s_1) u_p[s_1] + \sum_{\mathcal{E}_1} y(\mathcal{E}_1) u_p[\mathcal{E}_1] + \sum_{s_2 \in \mathcal{E}} y(s_2) u_p[s_2] + \sum_{i=2}^X y(s_i) u_p[s_i] \\ & \quad \sum_{s_1} y(s_1) u_p[s_1] + \sum_{\mathcal{E}_1} y(\mathcal{E}_1) u_p[\mathcal{E}_1] + \sum_{s_2 \in \mathcal{E}} y(s_2) u_p[s_2] + \sum_{i=2}^X y(s_i) u_p[s_i] \\ &= \sum_{i=2}^X y(s_i) u_p[s_i] + \sum_{\mathcal{E}_1} y(\mathcal{E}_1) u_p[\mathcal{E}_1] + \sum_{s_2 \in \mathcal{E}} y(s_2) u_p[s_2] \\ & \quad \sum_{i=2}^X y(s_i) u_p[s_i] + \sum_{\mathcal{E}_1} y(\mathcal{E}_1) u_p[\mathcal{E}_1] + \sum_{s_2 \in \mathcal{E}} y(s_2) u_p[s_2] \\ &= u_p[\mathcal{E}_1] \sum_{i=2}^X y(s_i) + \sum_{s_2 \in \mathcal{E}} y(s_2) u_p[s_2] \end{aligned}$$

Here, for the final inequality we have used the fact that $\sum_{i=2}^X y(s_i) \geq 0$; which follows from Claim H.7.

As before, note that

$$e(s_1) + e(\mathcal{E}_1) \sum_{s_2 \in \mathcal{E}} e(s) = 6n^2 + m \left(\frac{1}{2} + 1 \right) - 2n - 3n - m \left(\frac{1}{2} \right)$$

Claim H.7 now gives a contradiction as the inner product in (42) becomes positive.

Case 3: $d_p[s] < 0$ for all $s \in (25) \cup (26)$ and $|a_i - b_j| \geq 9a_0$ for all $1 \leq i \leq n$.

Similar as before, define \mathcal{E}_p as

$$\mathcal{E}_p[s] = \begin{cases} 0 & \text{if } s \in (27) \cup (28) \\ d_p[s] & \text{if } s = s_1 \\ \min_j d_p[s_j] & \text{if } s \in (29) \\ j d_p[s_j] & \text{o.w.} \end{cases}$$

As before, $\sum_s d_p[s] \sum_s \mathcal{E}_p[s]$. Therefore,

$$\begin{aligned} & \sum_s d_p[s] \sum_s \mathcal{E}_p[s] \\ &= \sum_{s_1} y(s_1) u_p[s_1] \sum_{s_2 \in (25) \cup (26)} y(s_2) u_p[s_2] + \sum_{s_2 \in (29)} y(s_2) \min_j u_p[s_j] + \sum_{i=2}^X y(s_i) u_p[s_i] \\ & \quad \sum_{s_1} y(s_1) u_p[s_1] \sum_{s_2 \in (25) \cup (26)} y(s_2) u_p[s_2] + \sum_{s_2 \in (29)} y(s_2) u_p[s_2] \mathbb{1}[u_p[s] < 0] + \sum_{i=2}^X y(s_i) u_p[s_i] \\ &= u_p[s_1] \sum_{s_2 \in (25) \cup (26)} y(s_2) + \sum_{s_2 \in (29)} y(s_2) \mathbb{1}[u_p[s] < 0] + \sum_{i=2}^X y(s_i) u_p[s_i] \end{aligned}$$

Assume for the sake of contradiction that the 3SAT instance is not satisfiable. There are two possibilities. The first is that for at least one $s \in (29)$, we have $u_p[s] = 0$. In that case,

$$e(s_1) \sum_{s_2 \in (25) \cup (26)} e(s) + \sum_{s_2 \in (29)} e(s) \mathbb{1}[u_p[s] < 0] \geq \frac{1}{2}$$

As before, Claim H.7 gives a contradiction. Otherwise, $u_p[s] < 0$ for all s (29). In that case, we have

$$u_\infty(M^T \lambda), \mu = 6n^2 + m - 1/2 - 2n \cdot 3n - m < -1/2,$$

which contradicts Proposition H.2 as $\lambda = 0$. \square

I Proofs of Appendix B

In this section, we provide a more formal treatment of the results in Appendix B, formally stating and proving these results.

Proposition I.1. *Assume that condition (3) holds. Set ϵ as*

$$\epsilon = \min_{s, a \neq \pi^\dagger(s)} \rho^\dagger - \rho^{\dagger\{s; a\}}.$$

Consider the following policy

$$\pi_{\mathcal{D}}(a/s) = \frac{1[a \in \Theta_s \setminus \{\pi^\dagger(s)\}]}{|\Theta_s| + 1}. \quad (47)$$

Equation (47) characterizes the solution to the optimization problems (P2a) and (P2b) with parameters $\epsilon = \epsilon_\dagger$ and $\epsilon = \min\{\epsilon, \epsilon_{\mathcal{D}}\}$ respectively. Furthermore, in both cases $\bar{\rho}^D = \rho^D$

Proof. Given Theorem 5.1, Theorem 5.2, and Lemma D.3, it suffices to show that if $\epsilon \leq \epsilon$, the solution to the optimization problem

$$\begin{aligned} \max_{\psi \in \Psi} \quad & \psi, R \\ \text{s.t.} \quad & \psi^{\dagger\{s; a\}} - \psi^\dagger, \psi \geq 0 \quad s, a \in \Theta, \end{aligned} \quad (P4)$$

corresponds to the occupancy measure of policy $\pi_{\mathcal{D}}$ defined by Equation (47). Namely, the optimization problems (P2a) and (P2b) correspond to the optimization problem (P3) with parameters $\epsilon = \epsilon_\dagger$ and $\epsilon = \min\{\epsilon, \epsilon_{\mathcal{D}}\}$ respectively. Since $\epsilon \leq \epsilon_\dagger$, the primal feasibility condition in Lemma E.1 implies that the solution to the above optimization problem characterizes both cases ((P2a) and (P2b)).

Now, due to Lemma D.3, we have

$$\psi \in \Psi \quad \psi < 0 \quad s : \quad \psi(s, a) = (1 - \gamma)\sigma(s) + \gamma \sum_{\tilde{s}, \tilde{a}} P(\tilde{s}, \tilde{a}, s) \psi(\tilde{s}, \tilde{a}).$$

Since $P(\tilde{s}, \tilde{a}, s)$ is independent of \tilde{a} , the second condition is equivalent to

$$s : \quad \psi(s, a) = (1 - \gamma)\sigma(s) + \gamma \sum_{\tilde{s}} P(\tilde{s}, \pi^\dagger(\tilde{s}), s) (\sum_{\tilde{a}} \psi(\tilde{s}, \tilde{a})),$$

which, due to (12), is equivalent to

$$\psi(s, a) = \mu(s).$$

Furthermore, given the independence of the transition distributions from policies, we have the following

$$\psi^{\dagger\{s; a\}} - \psi^\dagger(\tilde{s}, \tilde{a}) = \begin{cases} \mu(s) & \text{if } (\tilde{s}, \tilde{a}) = (s, a) \\ -\mu(s) & \text{if } (\tilde{s}, \tilde{a}) = (s, \pi^\dagger(s)) \\ 0 & \text{o.w} \end{cases}. \quad (48)$$

Therefore, the constraint $\rho^{\tau\{s;a\}} - \rho^{\tau}, \psi \geq 0$ is equivalent to $\psi(s, a) \geq \psi(s, \pi_{\dagger}(s))$. Furthermore, note that

$$\begin{aligned} (s, a) \in \Theta \quad & \rho^{\tau\{s;a\}} - \rho^{\tau} = -\epsilon_{\dagger} \\ & \psi^{\tau\{s;a\}} - \psi^{\tau}, R \geq -\epsilon_{\dagger} \\ & R(s, a) - R(s, \pi_{\dagger}(s)) = -\frac{\epsilon_{\dagger}}{\mu(s)} \\ & a \in \Theta_s. \end{aligned}$$

Putting it all together, the optimization problem (P3) is equivalent to

$$\begin{aligned} \max \quad & R, \psi \\ \text{s.t.} \quad & \psi(s, \pi_{\dagger}(s)) \geq \psi(s, a) \quad s, a \in \Theta_s \\ & \psi(s, a) = \mu(s) \quad s \in S \\ & \psi(s, a) \geq 0 \quad (s, a). \end{aligned}$$

Note that the maximization is now over all vectors $\psi \in \mathbb{R}^{|\mathcal{S}| \cdot |\mathcal{A}|}$ as the constraint $\psi \in \Psi$ has been made explicit. Furthermore, given Lemma D.3 and Equation (5), any vector ψ satisfying the last two constraints (the bellman constraints) corresponds to a policy π through

$$\pi(a/s) = \frac{\psi(s, a)}{\mu(s)}.$$

In other words, probability of choosing a in state s is proportional to $\psi(s, a)$.

Now, let us analyze the solution to this optimization problem which we will denote by ψ_{\max} . This solution ψ_{\max} exists, since the optimization problem is maximizing a continuous function on a closed and bounded set.

We first claim that if $a \notin \Theta_s \setminus \{\pi_{\dagger}(s)\}$, then $\psi_{\max}(s, a) = 0$. If this is not the case, then ψ_{\max} is not optimal. Concretely, consider the following vector ψ

$$\psi(\tilde{s}, \tilde{a}) = \begin{cases} \psi_{\max}(\tilde{s}, \tilde{a}) + \frac{1}{|\Theta_s|+1} \psi_{\max}(s, a) & \text{if } \tilde{s} = s \quad \tilde{a} \in \Theta_s \setminus \{\pi_{\dagger}(s)\} \\ 0 & \text{if } \tilde{s} = s \quad \tilde{a} = a \\ \psi_{\max}(\tilde{s}, \tilde{a}) & \text{o.w.} \end{cases}.$$

In other words, we uniformly spread the probability of choosing action a in state s over the set $\Theta_s \setminus \{\pi_{\dagger}(s)\}$. The vector ψ still satisfies the constraints: if $\tilde{a} \in \Theta_s$, $\psi(s, \pi_{\dagger}(s)) - \psi(s, \tilde{a}) = \psi_{\max}(s, \pi_{\dagger}(s)) - \psi_{\max}(s, \tilde{a})$ and the objective has strictly improved because

$$\rho^{\tau\{s;a\}} - \rho^{\tau} = -\epsilon \quad -\epsilon = R(s, a) - R(s, \pi_{\dagger}(s)) - \frac{\epsilon}{\mu(s)}.$$

Since $a \notin \Theta_s$, the inequality is strict and therefore

$$\tilde{a} \in \Theta_s \setminus \{\pi_{\dagger}(s)\} : R(s, \tilde{a}) > R(s, a).$$

This means that ψ was not optimal, contradicting the initial assumption.

Now note that if $\psi_{\max}(s, a) > \psi_{\max}(s, \pi_{\dagger}(s))$ for some $a \in \Theta_s$, then again ψ_{\max} isn't optimal as we could replace it with

$$\psi(\tilde{s}, \tilde{a}) = \begin{cases} \psi_{\max}(\tilde{s}, \tilde{a}) + \frac{\psi_{\max}(s, a) - \psi_{\max}(s, \pi_{\dagger}(s))}{|\Theta_s|+1} & \text{if } \tilde{s} = s \quad \tilde{a} \in \Theta_s \setminus \{\pi_{\dagger}(s)\} \setminus \{a\} \\ \psi_{\max}(\tilde{s}, \tilde{a}) - \frac{|\Theta_s|(\psi_{\max}(s, a) - \psi_{\max}(s, \pi_{\dagger}(s)))}{|\Theta_s|+1} & \text{if } \tilde{s} = s \quad \tilde{a} = a \\ \psi_{\max}(\tilde{s}, \tilde{a}) & \text{o.w.} \end{cases}.$$

Intuitively, since the action a was being chosen with strictly higher probability than action $\pi_{\dagger}(s)$, we have uniformly spread this excess probability among the set $\Theta_s \setminus \{\pi_{\dagger}(s)\}$. This vector would still be feasible as $\psi(s, a) = \psi(s, \pi_{\dagger}(s))$ and would be strictly better in terms of utility as $R(s, \pi_{\dagger}(s)) > R(s, a)$. This contradicts our initial assumption and therefore $\psi_{\max}(s, a) = \psi_{\max}(s, \pi_{\dagger}(s))$ for all $a \in \Theta_s$.

Since the occupancy measure ψ_{\max} satisfies $\psi_{\max}(s, a) = 0$ for all $a \notin \Theta_s \setminus \{\pi_{\dagger}(s)\}$ and $\psi_{\max}(s, a) = \psi(s, \pi_{\dagger}(s))$ for all $a \in \Theta_s$, we conclude that it is the occupancy measure for the policy $\pi_{\mathcal{D}}$ as defined in Equation (47).

In order to prove $\bar{\rho}^{\mathcal{D}} = \rho^{\mathcal{D}}$, first note that for $(s, a) \in \Theta$

$$\begin{aligned} \psi^{\dagger\{s;a\}} - \psi^{\dagger, \psi^{\mathcal{D}}} &= \mu(s)(\psi^{\mathcal{D}}(s, a) - \psi^{\mathcal{D}}(s, \pi_{\dagger}(s))) \\ &= \mu(s)^2(\pi_{\mathcal{D}}(a/s) - \pi_{\mathcal{D}}(\pi_{\dagger}(s)/s)) = 0, \end{aligned}$$

where we used Equation (48) and Equation (47). Therefore

$$\begin{aligned} \bar{\rho}^{\mathcal{D}} - \rho^{\mathcal{D}} &= \bar{R} - R, \psi^{\mathcal{D}} \\ &\stackrel{(i)}{=} \sum_{(s,a) \in \Theta} \alpha_{s,a} \psi^{\dagger\{s;a\}} - \psi^{\dagger, \psi^{\mathcal{D}}} \\ &= \sum_{(s,a) \in \Theta} \alpha_{s,a} \cdot 0 \\ &= 0, \end{aligned}$$

where (i) follows from Lemma E.3 in the known parameter case and Lemma F.3 in the unknown parameter case. \square

J Additional experiments

In this section, we provide the worst-case score of the policy from the defender's perspective, i.e., the value of the inner minimization in (P2b), for our policy. Given the results in Section 5.2 (Theorems 5.3 and 5.2), by construction of $\pi_{\mathcal{D}}$, this minimum is achieved for $R = \bar{R}$ and therefore the worst-case score equals $\rho^{\mathcal{D}}$. As mentioned in Section 5.2, $\rho^{\mathcal{D}}$ is a *certificate* in that we are guaranteed $\bar{\rho}^{\mathcal{D}} \geq \rho^{\mathcal{D}}$ as long as $\epsilon_{\mathcal{D}} \leq \epsilon_{\dagger}$.

The results are shown in Figure 6. As shown in the Figure, the value of $\rho^{\mathcal{D}}$ is larger when $\epsilon_{\mathcal{D}} < \epsilon_{\dagger}$. This is in line with results of Section 5.2 as in this case, $\pi_{\mathcal{D}} = \pi_{\dagger}$ and π_{\dagger} is the optimal policy in R . In addition, when $\epsilon_{\mathcal{D}} < \epsilon_{\dagger}$, the value of $\rho^{\mathcal{D}}$ increases with ϵ_{\dagger} . This is because when we increase the attack parameter ϵ_{\dagger} , the policy π_{\dagger} needs to become optimal in R with a larger margin, which causes its score to increase.

We note that for both environments, when no defense is employed, i.e., when using the policy π_{\dagger} , we will obtain the same worst-case value when $\epsilon_{\mathcal{D}} < \epsilon_{\dagger}$ as $\pi_{\mathcal{D}} = \pi_{\dagger}$, and we will obtain the worst-case score ρ^{\dagger} when $\epsilon_{\mathcal{D}} \geq \epsilon_{\dagger}$.




Figure 6: The value of the worst-case score from the defense perspective, i.e., $\rho^{\mathcal{D}}$ under different values of $\epsilon_{\dagger}, \epsilon_{\mathcal{D}}$ for the Navigation (left) and Grid world (right) environments.