Burning RED: Unlocking Subtask-Driven Reinforcement Learning and Risk-Awareness in Average-Reward Markov Decision Processes

Juan Sebastian Rojas, Chi-Guhn Lee

Keywords: Average-Reward Reinforcement Learning, Risk-Sensitive Decision-Making, CVaR

Summary

Average-reward Markov decision processes (MDPs) provide a foundational framework for sequential decision-making under uncertainty. However, average-reward MDPs have remained largely unexplored in reinforcement learning (RL) settings, with the majority of RL-based efforts having been allocated to discounted MDPs. In this work, we study a unique structural property of average-reward MDPs and utilize it to introduce *Reward-Extended Differential* (or *RED*) reinforcement learning: a novel RL framework that can be used to effectively and efficiently solve various learning objectives, or *subtasks*, simultaneously in the average-reward setting. We introduce a family of RED learning algorithms for prediction and control, including proven-convergent algorithms for the tabular case. We then showcase the power of these algorithms by demonstrating how they can be used to learn a policy that optimizes, for the first time, the well-known conditional value-at-risk (CVaR) risk measure in a fully-online manner, *without* the use of an explicit bi-level optimization scheme or an augmented state-space.

Contribution(s)

- 1. We provide a general-purpose framework and a corresponding set of prediction/control algorithms for solving an arbitrary number of learning objectives, or *subtasks*, simultaneously in the average-reward setting with only a TD error-based update, including provenconvergent algorithms for the tabular case.
 - **Context:** Our work builds on (and can be viewed as a generalization of) Wan et al. (2021), which proposed proven-convergent average-reward RL algorithms that are able to learn and/or optimize the value function and average-reward simultaneously using only the TD error. In particular, the focus in Wan et al. (2021) was on proving the convergence of such algorithms, without exploring the underlying structural properties of the average-reward MDP that made such a process possible to begin with. In this work, we formalize these underlying properties, and utilize them to show that if one modifies, or *extends*, the reward from the MDP with various learning objectives, then these objectives, or *subtasks*, can be solved simultaneously using a modified, or *reward-extended*, version of the TD error.
- 2. We utilize the framework described in Contribution 1 to derive the first family of RL algorithms that can optimize the well-known conditional value-at-risk (CVaR) risk measure in a fully-online manner without the use of an explicit bi-level optimization scheme or an augmented state-space. We perform an empirical evaluation on two toy experiments, thereby illustrating the properties and effectiveness of the algorithms, while also noting that a more comprehensive empirical study is needed to fully gauge their practical implications.

Context: Several prior works have investigated CVaR optimization in the discounted setting (e.g. Bäuerle and Ott (2011) and Chow et al. (2015)). However, no prior work has developed an algorithm for CVaR optimization that does not require either an augmented state-space or an explicit bi-level optimization, which can, for example, involve solving multiple MDPs. In the average-reward setting, Xia et al. (2023) proposed a set of algorithms for optimizing the CVaR risk measure, however their methods require the use of an augmented state-space and a sensitivity-based bi-level optimization. By contrast, our work, to the best of our knowledge, is the first to optimize CVaR in an MDP-based setting without the use of an explicit bi-level optimization scheme or an augmented state-space.

Burning RED: Unlocking Subtask-Driven Reinforcement Learning and Risk-Awareness in Average-Reward Markov Decision Processes

Juan Sebastian Rojas¹, Chi-Guhn Lee¹

juan.rojas@mail.utoronto.ca, cglee@mie.utoronto.ca

¹University of Toronto, Canada

Abstract

Average-reward Markov decision processes (MDPs) provide a foundational framework for sequential decision-making under uncertainty. However, average-reward MDPs have remained largely unexplored in reinforcement learning (RL) settings, with the majority of RL-based efforts having been allocated to discounted MDPs. In this work, we study a unique structural property of average-reward MDPs and utilize it to introduce *Reward-Extended Differential* (or *RED*) reinforcement learning: a novel RL framework that can be used to effectively and efficiently solve various learning objectives, or *subtasks*, simultaneously in the average-reward setting. We introduce a family of RED learning algorithms for prediction and control, including proven-convergent algorithms for the tabular case. We then showcase the power of these algorithms by demonstrating how they can be used to learn a policy that optimizes, for the first time, the well-known conditional value-at-risk (CVaR) risk measure in a fully-online manner, *without* the use of an explicit bi-level optimization scheme or an augmented state-space.

1 Introduction

Markov decision processes (MDPs) (Puterman, 1994) are a long-established framework for sequential decision-making under uncertainty. Discounted MDPs, which aim to optimize a potentially-discounted sum of rewards over time, have enjoyed success in recent years when utilizing reinforcement learning (RL) solution methods (Sutton and Barto, 2018) to tackle certain problems of interest in various domains. Despite this success, however, these MDP-based methods have yet to be fully embraced in real-world applications due to the various intricacies and implications of real-world operation that often outweigh the capabilities of current state-of-the-art methods (Dulac-Arnold et al., 2021). We therefore turn to the less-explored average-reward MDP, which aims to optimize the reward received per time-step, to see how its unique structural properties can be leveraged to tackle challenging problems that have eluded its discounted counterpart.

In particular, we present results that show how the average-reward MDP's unique structural properties can be leveraged to enable a more *subtask-driven* approach to reinforcement learning, where various learning objectives, or *subtasks*, are solved simultaneously (and in a fully-online manner) to help solve a larger, central learning objective. Importantly, we find a compelling case-study in the realm of risk-aware decision-making that illustrates how this subtask-driven approach can alleviate some of the computational challenges and complexities that can arise in the discounted setting.

More formally, we introduce *Reward-Extended Differential* (or *RED*) reinforcement learning: a first-of-its-kind RL framework that makes it possible to solve various subtasks simultaneously in the average-reward setting. At the heart of this framework is the novel concept of the reward-extended temporal-difference (TD) error, an extension of the celebrated TD error (Sutton, 1988), which we

derive by leveraging a unique structural property of average-reward MDPs, and utilize to solve various subtasks simultaneously. We first present the RED RL framework in a generalized way, then adopt it to successfully tackle a problem that has exceeded the capabilities of current state-of-the-art methods in risk-aware decision-making: learning a policy that optimizes the well-known conditional value-at-risk (CVaR) risk measure (Rockafellar and Uryasev, 2000) in a fully-online manner without the use of an explicit bi-level optimization scheme or an augmented state-space.

Our work is organized as follows: In Section 2, we provide a brief overview of related work. In Section 3, we give an overview of the fundamental concepts related to average-reward RL and CVaR. In Section 4, we motivate the need and opportunity for a subtask-driven approach to RL through the lens of CVaR optimization. In Section 5, we introduce the RED RL framework, including the concept of the reward-extended TD error. We also introduce a family of RED RL algorithms for prediction and control, and highlight their convergence properties (with full convergence proofs in Appendix B). In Section 6, we use the RED RL framework to derive a subtask-driven approach for CVaR optimization, and provide empirical results which show that this approach can be used to successfully learn a policy that optimizes the CVaR risk measure. Finally, in Section 7, we emphasize our framework's potential usefulness towards tackling other challenging problems outside the realm of risk-awareness, highlight some of its limitations, and suggest some directions for future research.

2 Related Work

Average-Reward Reinforcement Learning: Average-reward (or average-cost) MDPs, despite being one of the most well-studied frameworks for sequential decision-making under uncertainty (Puterman, 1994), have remained relatively unexplored in reinforcement learning (RL) settings. To date, notable works on the subject (in the context of RL) include Schwartz (1993), Tsitsiklis and Van Roy (1999), Abounadi et al. (2001), Gosavi (2004), Bhatnagar et al. (2009), and Wan et al. (2021). Most relevant to our work is Wan et al. (2021), which provided a rigorous theoretical treatment of average-reward MDPs in the context of RL, and proposed the proven-convergent 'Differential TD-learning' and 'Differential Q-learning' algorithms. Our work builds on the methods from Wan et al. (2021) to develop a theoretical framework for solving various learning objectives simultaneously.

We note that these learning objectives, or *subtasks*, as explored in our work, are different from that of hierarchical RL (e.g. Sutton et al. (1999)). In particular, in hierarchical RL, the focus is on using temporally-abstracted actions, known as 'options' (or 'skills'), such that the agent learns a policy for each option, as well as an inter-option policy. By contrast, in our work, we learn a single policy, and the subtasks are not part of the action-space. Similarly, the notion of solving multiple objectives in parallel has been widely-explored in the discounted setting (e.g. McLeod et al. (2021)). However, much of this work focuses on learning multiple state representations (or 'features'), options, policies, and/or value functions. By contrast, in our work, we learn a single policy and value function, and the subtasks are not part of the state or action-spaces. To the best of our knowledge, our work is the first to explore the solving of subtasks simultaneously in the average-reward setting.

Risk-Aware Learning and Optimization in MDPs: The notion of risk-aware learning and optimization in MDP-based settings has been long-studied, from the well-established expected utility framework (Howard and Matheson, 1972), to the more contemporary framework of coherent risk measures (Artzner et al., 1999). To date, these risk-based efforts have focused almost exclusively on the discounted setting. Importantly, optimizing the CVaR risk measure in this setting typically requires augmenting the state-space and/or having to utilize an explicit bi-level optimization scheme, which can, for example, involve solving multiple MDPs. Seminal works that have investigated CVaR optimization in the standard discounted setting include Bäuerle and Ott (2011) and Chow et al. (2015); Hau et al. (2023a). In the distributional setting, works such as Dabney et al. (2018) and Keramati et al. (2020) have proposed CVaR optimization approaches that do not require an augmented state-space or an explicit bi-level optimization; however, it was later shown in Lim and Malik (2022) that such approaches converge to neither the optimal dynamic-CVaR nor the optimal static-CVaR policies (Lim and Malik (2022) then proposed a valid approach that utilizes an

augmented state-space). Some works have investigated optimizing a time-consistent (Ruszczyński, 2010) interpretation of CVaR in the discounted setting; however, such approaches only approximate CVaR, as CVaR is generally not a time-consistent risk measure in the discounted setting (Boda and Filar, 2006). Other works have investigated optimizing similar objectives to CVaR that are more computationally tractable, such as the entropic value-at-risk (e.g. Hau et al. (2023b)).

Most similar to our work (in non-average-reward settings) is Stanko and Macek (2019), which proposed a similar CVaR update scheme to the one derived in our work. However, all of the methods proposed in Stanko and Macek (2019) require either an augmented state-space or an explicit bi-level optimization. In the average-reward setting, Xia et al. (2023) proposed a set of algorithms for optimizing the CVaR risk measure; however, their methods require the use of an augmented state-space and a sensitivity-based bi-level optimization. By contrast, our work, to the best of our knowledge, is the first to optimize the CVaR risk measure in an MDP-based setting without the use of an explicit bi-level optimization scheme or an augmented state-space. We note that other works have investigated optimizing other risk measures in the average-reward setting, such as the exponential cost (e.g. Murthy et al. (2023)), and variance (e.g. Prashanth and Ghavamzadeh (2016)).

3 Preliminaries

3.1 Average-Reward Reinforcement Learning

A finite average-reward MDP is the tuple $\mathcal{M} \doteq \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, p \rangle$, where \mathcal{S} is a finite set of states, \mathcal{A} is a finite set of actions, $\mathcal{R} \subset \mathbb{R}$ is a bounded set of rewards, and $p: \mathcal{S} \times \mathcal{A} \times \mathcal{R} \times \mathcal{S} \to [0,1]$ is a probabilistic transition function that describes the dynamics of the environment. At each discrete time step, $t=0,1,2,\ldots$, an agent chooses an action, $A_t \in \mathcal{A}$, based on its current state, $S_t \in \mathcal{S}$, and receives a reward, $R_{t+1} \in \mathcal{R}$, while transitioning to a (potentially) new state, S_{t+1} , such that $p(s',r\mid s,a) = \mathbb{P}(S_{t+1}=s',R_{t+1}=r\mid S_t=s,A_t=a)$. In an average-reward MDP, an agent aims to find a policy, $\pi:\mathcal{S} \to \mathcal{A}$, that optimizes the long-run (or limiting) average-reward, \bar{r} , which is defined as follows for a given policy, π :

$$\bar{r}_{\pi}(s) \doteq \lim_{n \to \infty} \frac{1}{n} \sum_{t=1}^{n} \mathbb{E}[R_t \mid S_0 = s, A_{0:t-1} \sim \pi].$$
 (1)

In this work, we limit our discussion to *stationary Markov* policies, which are time-independent policies that satisfy the Markov property.

When working with average-reward MDPs, it is common to simplify Equation (1) into a more workable form by making certain assumptions about the Markov chain induced by following policy π . To this end, a *unichain* assumption is typically used when doing prediction (learning) because it ensures the existence of a unique limiting distribution of states, $\mu_{\pi}(s) \doteq \lim_{t \to \infty} \mathbb{P}(S_t = s \mid A_{0:t-1} \sim \pi)$, that is independent of the initial state, thereby simplifying Equation (1) to the following:

$$\bar{r}_{\pi} = \sum_{s \in \mathcal{S}} \mu_{\pi}(s) \sum_{a \in \mathcal{A}} \pi(a \mid s) \sum_{s' \in \mathcal{S}} \sum_{r \in \mathcal{R}} p(s', r \mid s, a) r. \tag{2}$$

Similarly, a *communicating* assumption is typically used for control (optimization) because it ensures the existence of a unique optimal average-reward, $\bar{r}*$, that is independent of the initial state.

To solve an average-reward MDP, solution methods such as dynamic programming or RL can be used in conjunction with the following *Bellman* (or *Poisson*) equations:

$$v_{\pi}(s) = \sum_{a} \pi(a \mid s) \sum_{s'} \sum_{r} p(s', r \mid s, a) [r - \bar{r}_{\pi} + v_{\pi}(s')], \tag{3}$$

$$q_{\pi}(s, a) = \sum_{s'} \sum_{r} p(s', r \mid s, a) [r - \bar{r}_{\pi} + \max_{a'} q_{\pi}(s', a')], \tag{4}$$

where $v_{\pi}(s)$ is the state-value function and $q_{\pi}(s,a)$ is the state-action value function for a given policy, π . Solution methods for average-reward MDPs are typically referred to as *differential* methods because of the reward difference (i.e., $r - \bar{r}_{\pi}$) operation that occurs in Equations (3) and (4). We note that solution methods typically find the solutions to Equations (3) and (4) up to a constant. This is typically not a concern, given that the relative ordering of policies is usually what is of interest.

In the context of RL, Wan et al. (2021) proposed the tabular 'Differential TD-learning' and 'Differential Q-learning' algorithms, which are able to learn and/or optimize the value function and average-reward simultaneously using only the TD error. The 'Differential TD-learning' algorithm is shown below:

$$\delta_t \doteq R_{t+1} - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t) \tag{5a}$$

$$V_{t+1}(s) \doteq V_t(s), \quad \forall s \neq S_t$$
 (5b)

$$V_{t+1}(S_t) \doteq V_t(S_t) + \alpha_t \rho_t \delta_t \tag{5c}$$

$$\bar{R}_{t+1} \doteq \bar{R}_t + \alpha_{\bar{r},t} \rho_t \delta_t \tag{5d}$$

where $V_t: \mathcal{S} \to \mathbb{R}$ is a table of state-value function estimates, \bar{R}_t is an estimate of the averagereward, \bar{r}_{π} , δ_t is the TD error, $\rho_t \doteq \pi(A_t \mid S_t) / B(A_t \mid S_t)$ is the importance sampling ratio (with behavior policy, B), α_t is the value function step size, and $\alpha_{\bar{r},t}$ is the average-reward step size.

3.2 Conditional Value-at-Risk (CVaR)

Consider a random variable X with a finite mean on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and with a cumulative distribution function $F(x) = \mathbb{P}(X \leq x)$. The (left-tail) value-at-risk (VaR) of X with parameter $\tau \in (0,1)$ represents the τ -quantile of X, such that $\mathrm{VaR}_{\tau}(X) = \sup\{x \mid F(x) \leq \tau\}$. The (left-tail) conditional value-at-risk (CVaR) of X with parameter τ is defined as follows:

$$CVaR_{\tau}(X) = \frac{1}{\tau} \int_{0}^{\tau} VaR_{u}(X) du.$$
 (6)

When F(X) is continuous at $x = \text{VaR}_{\tau}(X)$, $\text{CVaR}_{\tau}(X)$ can be interpreted as the expected value of X conditioned on X being less than or equal to $\text{VaR}_{\tau}(X)$, such that $\text{CVaR}_{\tau}(X) = \mathbb{E}[X \mid X \leq \text{VaR}_{\tau}(X)]$.

Importantly, $\text{CVaR}_{\tau}(X)$ can be formulated as the solution to the following optimization problem (Rockafellar and Uryasev, 2000):

$$\text{CVaR}_{\tau}(X) = \sup_{y \in \mathbb{R}} \mathbb{E}[y - \frac{1}{\tau}(y - X)^{+}] = \mathbb{E}[\text{VaR}_{\tau}(X) - \frac{1}{\tau}(\text{VaR}_{\tau}(X) - X)^{+}], \tag{7}$$

where $(u)^+ = \max(u, 0)$. Existing MDP-based methods typically leverage the above formulation when optimizing for CVaR, by augmenting the state-space with a state that corresponds (either directly or indirectly) to an estimate of $\operatorname{VaR}_{\tau}(X)$ (in this case, y), and solving the following bi-level optimization:

$$\sup_{\pi} \text{CVaR}_{\tau}(X) = \sup_{\pi} \sup_{y \in \mathbb{R}} \mathbb{E}[y - \frac{1}{\tau}(y - X)^{+}] = \sup_{y \in \mathbb{R}} (y - \frac{1}{\tau} \sup_{\pi} \mathbb{E}[(y - X)^{+}]), \tag{8}$$

where the 'inner' optimization problem can be solved using standard MDP solution methods.

In discounted MDPs, the random variable X corresponds to a (potentially-discounted) sum of rewards. In average-reward MDPs, X corresponds to the limiting *per-step reward*. In other words, the natural interpretation of CVaR in the average-reward setting is that of the CVaR associated with the limiting per-step reward distribution, as shown below (for a given policy, π) (Xia et al., 2023):

$$CVaR_{\tau,\pi}(s) \doteq \lim_{n \to \infty} \frac{1}{n} \sum_{t=1}^{n} CVaR_{\tau}[R_t \mid S_0 = s, A_{0:t-1} \sim \pi].$$
 (9)

As with the average-reward (i.e., Equation (1)), a unichain assumption (or similar) makes this CVaR objective independent of the initial state. In recent years, CVaR has emerged as a popular risk measure, in part because it is a 'coherent' risk measure (Artzner et al., 1999), meaning that it satisfies key mathematical properties which can be meaningful in safety-critical and risk-related applications.

Figure 1 depicts the agent-environment interaction in an average-reward MDP, where following policy π yields a limiting average-reward and reward CVaR.

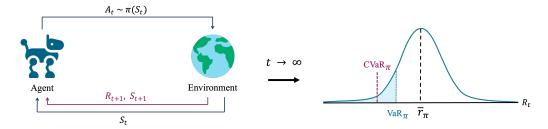


Figure 1: Illustration of the agent-environment interaction in an average-reward MDP. As $t \to \infty$, following policy π yields a limiting per-step reward distribution with an average-reward, \bar{r}_{π} , and a conditional value-at-risk, CVaR_{π} . Standard average-reward RL methods aim to optimize the average-reward, \bar{r}_{π} . By contrast, in our work, we aim to optimize CVaR_{π} .

4 A Subtask-Driven Approach

In this section, we motivate the need and opportunity for a subtask-driven approach to RL through the lens of CVaR optimization. Let us begin by considering the standard approach used by existing MDP-based methods for CVaR optimization. This approach, which is described in Equation (8), requires that we pick a wide range of guesses for the optimal value-at-risk, VaR, and that for each guess, y, we solve an MDP. Then, among all the MDP solutions, we pick the best one as our final solution (which corresponds to y = VaR). Moreover, to further compound the computational costs associated with solving several MDPs, this approach requires that the state-space be augmented with a state that corresponds (either directly or indirectly) to the VaR guess, y (e.g. see Bäuerle and Ott (2011)). Hence, this approach requires the use of both an explicit bi-level optimization scheme, and an augmented state-space. Importantly, however, this complex process would not be needed if we somehow knew what the optimal value for y (i.e., VaR) was. In fact, in the average-reward setting, if we know this optimal value, VaR, then optimizing for CVaR ultimately amounts to optimizing an average (as per Equation (7)), which can be done trivially using the standard average-reward MDP.

As such, it would appear that, to optimize CVaR, we are stuck between two extremes: a significantly computationally-expensive process if we do not know the optimal value-at-risk, VaR, and a trivial process if we do. But what if we could somehow estimate VaR along the way? That is, keep some sort of running estimate of VaR that we optimize simultaneously as we optimize CVaR. Indeed, such an approach has been proposed in the discounted setting (e.g. Stanko and Macek (2019)); however, no approach has been able to successfully remove both the augmented state-space and the explicit bi-level optimization requirements. The primary difficulty lies in *how* one updates the estimate of VaR along the way.

Critically, this is where the findings from Wan et al. (2021) become relevant. In particular, Wan et al. (2021) proposed proven-convergent algorithms for the average-reward setting that can learn and/or optimize the value function and average-reward simultaneously using only the TD error. In other words, these algorithms are able to solve two learning objectives simultaneously using only the TD error. Yet, the focus in Wan et al. (2021) was on proving the convergence of such algorithms, without exploring the underlying structural properties of the average-reward MDP that made such a process possible to begin with. In this work, we formalize these underlying properties, and utilize them to show that if one modifies, or *extends*, the reward from the MDP with various learning objectives

that satisfy certain key properties, then these objectives, or *subtasks*, can be solved simultaneously using a modified, or *reward-extended*, version of the TD error. Consequently, in terms of CVaR optimization, this allows us to develop appropriate learning updates for the VaR and CVaR estimates based solely on the TD error, such that we can optimize the CVaR risk measure in a fully-online manner without needing to augment the state-space or perform an explicit bi-level optimization.

In Section 5, we present the theoretical framework that enables the aforementioned subtask-driven approach. Then, in Section 6, we adapt this general-purpose framework for CVaR optimization.

5 Reward-Extended Differential (RED) Reinforcement Learning

In this section, we present our primary contribution: a framework for solving various learning objectives, or *subtasks*, simultaneously in the average-reward setting. We call this framework *reward-extended differential* (or *RED*) reinforcement learning. The 'differential' part of the name stems from the use of differential algorithms typically associated with average-reward MDPs. The 'reward-extended' part of the name stems from the use of the *reward-extended TD error*, a novel concept that we will introduce shortly. Through this framework, we show how the average-reward MDP's unique structural properties can be leveraged to solve (i.e., learn or optimize) any given subtask by using only a TD error-based update. We first provide a formal definition for a (generic) subtask, then proceed to derive a framework that allows us to solve any subtask that satisfies this definition. In the subsequent section, we utilize this framework to tackle the CVaR optimization problem.

Definition 5.1 (Subtask). A subtask, z_i , is any scalar prediction or control objective belonging to a corresponding bounded set $\mathcal{Z}_i \subset \mathbb{R}$, such that there exists a linear or piecewise linear subtask function, $f: \mathcal{R} \times \mathcal{Z}_1 \times \mathcal{Z}_2 \times \cdots \times \mathcal{Z}_i \times \cdots \times \mathcal{Z}_n \to \tilde{\mathcal{R}}$, where \mathcal{R} is the bounded set of observed per-step rewards from the MDP \mathcal{M} , $\tilde{\mathcal{R}} \subset \mathbb{R}$ is a bounded set of 'extended' per-step rewards whose long-run average is the primary prediction or control objective of the MDP, $\tilde{\mathcal{M}} \doteq \langle \mathcal{S}, \mathcal{A}, \tilde{\mathcal{R}}, \tilde{p} \rangle$, and $\mathcal{Z} = \{z_1 \in \mathcal{Z}_1, z_2 \in \mathcal{Z}_2, \ldots, z_n \in \mathcal{Z}_n\}$ is the set of n subtasks that we wish to solve, such that:

i) f is invertible with respect to each input given all other inputs; and

ii) each subtask $z_i \in \mathcal{Z}$ in f is independent of the states and actions, and hence independent of the observed per-step reward, $R_t \in \mathcal{R}$, such that $\mathbb{P}(S_{t+1} = s', \tilde{R}_{t+1} = f(r, z_1, \dots, z_n) \mid S_t = s, A_t = a) = \mathbb{P}(S_{t+1} = s', R_{t+1} = r \mid S_t = s, A_t = a)$, and $\mathbb{E}[f_j(R_t, z_1, z_2, \dots, z_n)] = f_j(\mathbb{E}[R_t], z_1, z_2, \dots, z_n)$, where f_j denotes the jth piecewise segment of f, and \mathbb{E} denotes any expectation taken with respect to the states and actions.

In essence, the above definition states that a subtask is some constant, z_i , that we wish to learn and/or optimize. From an algorithmic perspective, this means that we will start with some initial estimate (or guess) for the subtask, $Z_{i,t}$, then update this estimate at every time step, such that $Z_{i,t} \to z_i$ or $Z_{i,t} \to z_i^*$, depending on whether we are doing prediction or control (where z_i^* denotes the optimal subtask value). But how can we derive an appropriate update rule that accomplishes this? In the following section, we will introduce the *reward-extended TD error*, through which we can derive such an update rule for any subtask that satisfies Definition 5.1, such that $Z_{i,t} \to z_i$ when doing prediction and $Z_{i,t} \to z_i^*$ when doing control.

5.1 The Reward-Extended TD Error

In this section, we introduce and derive the *reward-extended TD error*. In particular, we derive a generic, subtask-specific, TD-like error, $\beta_{i,t}$, through which we can learn and/or optimize any subtask that satisfies Definition 5.1 via the update rule: $Z_{i,t+1} = Z_{i,t} + \alpha_{z_i,t}\beta_{i,t}$, where $Z_{i,t}$ is an estimate of subtask z_i , $\alpha_{z_i,t}$ is the step size, and $\beta_{i,t}$ is the reward-extended TD error for subtask z_i .

Importantly, we will show that the reward-extended TD error satisfies the following property: $\mathbb{E}_{\pi}[\beta_{i,t}] \to 0 \ \forall i=1,2,\ldots,n$ as $\mathbb{E}_{\pi}[\delta_t] \to 0$, where δ_t is the regular TD error, such that minimizing the regular TD error allows us to solve all subtasks simultaneously. This motivates our naming of the reward-extended TD error, given that it is intrinsically tied to the regular TD error.

Let us begin by considering the common RL update rule of the form: $NewEstimate \leftarrow OldEstimate + StepSize [Target - OldEstimate]$ (Sutton and Barto, 2018; Naik, 2024). Our aim is to find an appropriate set of subtask-specific 'targets', $\{\phi_{i,t}\}_{i=1}^n$, such that $\mathbb{E}_{\pi}[\beta_{i,t}] = \mathbb{E}_{\pi}[\phi_{i,t} - Z_{i,t}] \rightarrow 0 \ \forall i = 1, 2, \ldots, n \ \text{as} \ \mathbb{E}_{\pi}[\delta_t] \rightarrow 0$. To this end, let us consider a generic piecewise linear subtask function with m piecewise segments:

$$\tilde{R}_{t} = \begin{cases}
b_{r}^{1}R_{t} + b_{0}^{1} + b_{1}^{1}z_{1} + b_{2}^{1}z_{2} + \dots + b_{n}^{1}z_{n}, & r_{0} \leq R_{t} < r_{1} \\
b_{r}^{2}R_{t} + b_{0}^{2} + b_{1}^{2}z_{1} + b_{2}^{2}z_{2} + \dots + b_{n}^{2}z_{n}, & r_{1} \leq R_{t} < r_{2} \\
\vdots \\
b_{r}^{m}R_{t} + b_{0}^{m} + b_{1}^{m}z_{1} + b_{2}^{m}z_{2} + \dots + b_{n}^{m}z_{n}, & r_{m-1} \leq R_{t} \leq r_{m}
\end{cases} ,$$
(10)

where $r_k \in \mathcal{R} \ \forall k = 0, 1, \dots, m$, and $r_0 \leq r_1 \leq \dots \leq r_m$, such that r_0, r_m represent the lower and upper bounds of the observed per-step reward, R_t , respectively. Moreover, $b_r^j, b_0^j \in \mathbb{R}$, and $b_i^j \in \mathbb{R} \setminus \{0\}$, where b^j denotes a (predefined) constant in the jth piecewise segment of \tilde{R}_t .

Now, let us consider the TD error, δ_t , associated with (10) in the prediction setting. Let $\tilde{R}_{j,t}$ be shorthand for the jth segment of (10), such that the TD error at any time step can be expressed as:

$$\delta_{j,t} = \tilde{R}_{j,t+1} - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t)$$
(11a)

$$=b_r^j R_{t+1} + b_0^j + b_1^j Z_{1,t} + b_2^j Z_{2,t} + \dots + b_n^j Z_{n,t} - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t), \tag{11b}$$

where $V_t: \mathcal{S} \to \mathbb{R}$ denotes a table of state-value function estimates, \bar{R}_t denotes an estimate of the average-reward, \bar{r}_{π} , $Z_{i,t}$ denotes an estimate of subtask $z_i \ \forall i=1,2,\ldots,n$, and j corresponds to the piecewise condition, $r_{j-1} \leq R_{t+1} \leq r_j$, that is satisfied by the observed per-step reward, R_{t+1} .

Hence, as learning progresses, different $\tilde{R}_{j,t+1}$ values will be used to define the TD error based on which piecewise condition is satisfied at a given time step. Moreover, we know that the probability that $\delta_t = \delta_{j,t}$ is equal to the probability that $r_{j-1} \leq R_{t+1} < r_j$. This allows us to express the expected TD error associated with (10) as follows:

$$\mathbb{E}_{\pi}[\delta_{t}] = \sum_{j=1}^{m} \mathbb{P}(r_{j-1} \le R_{t+1} < r_{j}) \mathbb{E}_{\pi}[\delta_{j,t}]. \tag{12}$$

Now, let us consider the implications of $\mathbb{E}_{\pi}[\delta_t] \to 0$ as it relates to $\mathbb{E}_{\pi}[\delta_{j,t}]$. One possibility is that $\mathbb{E}_{\pi}[\delta_{j,t}] \to 0 \ \forall j=1,2,\ldots,m$. However, this may not necessarily be the case; it is possible that, for example, a pair of non-zero $\mathbb{P}(r_{j-1} \leq R_{t+1} < r_j)\mathbb{E}_{\pi}[\delta_{j,t}]$ terms cancel each other out, such that $\mathbb{E}_{\pi}[\delta_t] \to 0$ but $\mathbb{E}_{\pi}[\delta_{j,t}] \to \lambda_j \ \forall j=1,2,\ldots,m$, where $\lambda_j \in \mathbb{R}$. In such a case, what we do know is that if $\mathbb{E}_{\pi}[\delta_t] \to 0$, then the Bellman equation (3) must be satisfied, such that: $V_t(s) = \mathbb{E}_{\pi}[\tilde{R}_{t+1} - \bar{R}_t + V_t(S_{t+1}) \mid S_t = s]$. As such, we can write the following expression for λ_j , and solve for an arbitrary subtask, z_i , as follows:

$$\lambda_j = \mathbb{E}_{\pi}[\tilde{R}_{j,t+1} - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t)]$$
(13a)

$$= \mathbb{E}_{\pi} \left[\tilde{R}_{j,t+1} - \bar{R}_t + V_t(S_{t+1}) - \left(\tilde{R}_{t+1} - \bar{R}_t + V_t(S_{t+1}) \right) \right]$$
 (13b)

$$= \mathbb{E}_{\pi}[\tilde{R}_{j,t+1}] - \mathbb{E}_{\pi}[\tilde{R}_{t+1}] \tag{13c}$$

$$= \mathbb{E}_{\pi}[\tilde{R}_{j,t+1}] - \bar{r}_{\pi} \tag{13d}$$

$$= \mathbb{E}_{\pi}[b_r^j R_{t+1} + b_0^j + \ldots + b_{i-1}^j z_{i-1} + b_{i+1}^j z_{i+1} + \ldots + b_n^j z_n - \bar{r}_{\pi}] + b_i^j z_i$$
 (13e)

$$\implies z_{i} = \mathbb{E}_{\pi} \left[-\frac{1}{b_{i}^{j}} \left(b_{r}^{j} R_{t+1} + b_{0}^{j} + \ldots + b_{i-1}^{j} z_{i-1} + b_{i+1}^{j} z_{i+1} + \ldots + b_{n}^{j} z_{n} - \bar{r}_{\pi} - \lambda_{j} \right) \right]$$

$$\tag{13f}$$

$$\doteq \mathbb{E}_{\pi}[\phi_{i,j}],\tag{13g}$$

where we used the fact that z_i is independent of the states and actions to pull it out of the expectation. Here, we use $\phi_{i,j}$ to denote the expression inside the expectation in Equation (13f). Hence, to learn z_i from experience, we can utilize the common RL update rule, using the term inside the expectation in Equation (13g), $\phi_{i,j}$, as the 'target', which yields the update:

$$Z_{i,t+1} = Z_{i,t} + \alpha_{z_{i},t} \begin{cases} \phi_{i,1,t} - Z_{i,t}, & r_{0} \leq R_{t+1} < r_{1} \\ \vdots \\ \phi_{i,m,t} - Z_{i,t}, & r_{m-1} \leq R_{t+1} \leq r_{m} \end{cases}$$

$$= Z_{i,t} + \alpha_{z_{i},t} \begin{cases} (-1/b_{i}^{1}) \left(\tilde{R}_{1,t+1} - \bar{R}_{t} - \delta_{t} \right), & r_{0} \leq R_{t+1} < r_{1} \\ \vdots \\ (-1/b_{i}^{m}) \left(\tilde{R}_{m,t+1} - \bar{R}_{t} - \delta_{t} \right), & r_{m-1} \leq R_{t+1} \leq r_{m} \end{cases}$$

$$(14a)$$

$$= Z_{i,t} + \alpha_{z_{i},t} \begin{cases} (-1/b_{i}^{1}) \left(\tilde{R}_{1,t+1} - \bar{R}_{t} - \delta_{t} \right), & r_{0} \leq R_{t+1} < r_{1} \\ \vdots \\ (-1/b_{i}^{m}) \left(\tilde{R}_{m,t+1} - \bar{R}_{t} - \delta_{t} \right), & r_{m-1} \leq R_{t+1} \leq r_{m} \end{cases}$$

$$(14b)$$

$$\doteq Z_{i,t} + \alpha_{z_i,t} \beta_{i,t}, \tag{14c}$$

where $Z_{i,t}$ is the estimate of subtask z_i at time t, $\phi_{i,j,t} \doteq (-1/b_i^j)(b_r^j R_{t+1} + b_0^j + \ldots + b_{i-1}^j Z_{i-1,t} + b_0^j + \ldots + b_{i-1}^j Z_{i-1,t})$ $b_{i+1}^j Z_{i+1,t} + \ldots + b_n^j Z_{n,t} - \bar{R}_t - \delta_t$), and $\alpha_{z_i,t}$ is the step size.

As such, we now have an expression for the reward-extended TD error for subtask z_i , $\beta_{i,t}$. We will now show that this term satisfies the desired property: $\mathbb{E}_{\pi}[\beta_{i,t}] \to 0 \ \forall i = 1, 2, \dots, n \text{ as } \mathbb{E}_{\pi}[\delta_t] \to 0$, such that minimizing the regular TD error allows us to solve all the subtasks simultaneously:

Theorem 5.1. Consider an average-reward MDP with a set of reward-extended TD errors, $\{\beta_{i,t}\}_{i=1}^n$, as defined in Equation (14), corresponding to a subtask function with n subtasks that satisfy Definition 5.1. The set of reward-extended TD errors, $\{\beta_{i,t}\}_{i=1}^n$, satisfies the following property: $\mathbb{E}_{\pi}[\beta_{i,t}] \to 0 \ \forall i=1,2,\ldots,n$ as $\mathbb{E}_{\pi}[\delta_t] \to 0$, where $\beta_{i,t}$ denotes the reward-extended TD error for subtask z_i , and δ_t denotes the regular TD error.

Proof. Let us consider the reward-extended TD error associated with an arbitrary jth segment of R_t for an arbitrary ith subtask: $\beta_{i,j,t} \doteq (-1/b_i^j)(R_{j,t+1} - \bar{R}_t - \delta_t)$. As $\mathbb{E}_{\pi}[\delta_t] \to 0$, $\bar{R}_t \to \bar{r}_{\pi}$ (by Theorem 3 of Wan et al. (2021); see Remark 5.3) and $\delta_t \to \lambda_j$ for this jth segment. Hence, $\mathbb{E}_{\pi}[\beta_{i,j,t}] \to (-1/b_i^j) (\mathbb{E}_{\pi}[\tilde{R}_{j,t+1}] - \bar{r}_{\pi} - \lambda_j) = (-1/b_i^j) (\lambda_j - \lambda_j) = 0. \text{ Now, because we chose } j$ arbitrarily, we have, for all $j \in \{1, 2, \dots, m\}$, that $\mathbb{E}_{\pi}[\beta_{i,j,t}] \to 0$. As such, and because we chose i arbitrarily, we can conclude that $\mathbb{E}_{\pi}[\beta_{i,t}] = \sum_{j=1}^m \mathbb{P}(r_{j-1} \leq R_{t+1} < r_j) \mathbb{E}_{\pi}[\beta_{i,j,t}] \to 0 \ \forall i = 1, 2, \dots, n \text{ as } \mathbb{E}_{\pi}[\delta_t] \to 0. \text{ This completes the proof.}$

As such, we have derived the desired update rule that we can use to solve any given subtask in the prediction setting. The same logic can be applied in the control setting to derive equivalent updates, where we note that it directly follows from Definition 5.1 that the existence of an optimal average-reward, $\bar{r}*$, implies the existence of corresponding optimal subtask values, $z_i^* \ \forall z_i \in \mathcal{Z}$.

Remark 5.1. In the case of a (non-piecewise) linear subtask function, the expression for the reward-extended TD error can be simplified to $\beta_{i,t} = (-1/b_i)\delta_t$ by setting $\lambda = 0$ in Equation (13a), solving for the target, z_i , and applying a similar process to the one described in Equation (14).

Remark 5.2. Given Remark 5.1, it can be shown that if one treats the average-reward, \bar{r}_{π} , as a subtask, and derives the reward-extended TD error for it, the process yields the average-reward update (e.g. Equation (5d)) from the Differential algorithms proposed in Wan et al. (2021). Hence, our work can be viewed as a generalization of the work performed in Wan et al. (2021).

Remark 5.3. Strictly speaking, $\bar{R}_t \to \bar{r}_\pi + c$, $c \in \mathbb{R}$. This is because average-reward solution methods typically find the solutions to the Bellman equations (3) and (4) up to an additive constant, c. This means that, like the average-reward estimate, our subtask estimates converge to the actual subtask values, up to an additive constant. For simplicity, we omit this additive constant in our work, unless strictly necessary, given that it is commonplace to assume that solutions in the averagereward setting are correct up to an additive constant.

5.2 The RED Algorithms

In this section, we introduce the *RED RL algorithms*, which integrate the update rules derived in the previous section into the average-reward RL framework from Wan et al. (2021). The full algorithms, including algorithms that utilize function approximation, are included in Appendix A.

RED TD-learning algorithm (tabular): We update a table of estimates, $V_t : \mathcal{S} \to \mathbb{R}$ as follows:

$$\tilde{R}_{t+1} = f(R_{t+1}, Z_{1,t}, Z_{2,t}, \dots, Z_{n,t})$$
(15a)

$$\delta_t = \tilde{R}_{t+1} - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t) \tag{15b}$$

$$V_{t+1}(S_t) = V_t(S_t) + \alpha_t \rho_t \delta_t \tag{15c}$$

$$\bar{R}_{t+1} = \bar{R}_t + \alpha_{\bar{r},t} \rho_t \delta_t \tag{15d}$$

$$Z_{i,t+1} = Z_{i,t} + \alpha_{z_i,t} \rho_t \beta_{i,t}, \quad \forall z_i \in \mathcal{Z}$$
 (15e)

where R_t is the observed reward, $Z_{i,t}$ is an estimate of subtask z_i , $\beta_{i,t}$ is the reward-extended TD error for subtask z_i , \bar{R}_t is an estimate of the long-run average-reward of \tilde{R}_t , \bar{r}_{π} , δ_t is the TD error, ρ_t is the importance sampling ratio, and α_t , $\alpha_{\bar{r},t}$, and $\alpha_{z_i,t}$ are the step sizes.

Wan et al. (2021) showed for their Differential TD-learning algorithm that R_t converges to \bar{r}_{π} , and V_t converges to a solution of v in Equation (3) for a given policy, π . We now provide an equivalent theorem for our RED TD-learning algorithm, which also shows that $Z_{i,t}$ converges to $z_{i,\pi} \ \forall z_i \in \mathcal{Z}$, where $z_{i,\pi}$ denotes the subtask value induced when following policy π :

Theorem 5.2 (informal). The RED TD-learning algorithm (15) converges, almost surely, \bar{R}_t to $\bar{\tau}_\pi$, $Z_{i,t}$ to $z_{i,\pi}$ $\forall z_i \in \mathcal{Z}$, and V_t to a solution of v in the Bellman equation (3), up to an additive constant, c, if the following assumptions hold: 1) the Markov chain induced by the target policy, π , is unichain, 2) every state—action pair for which $\pi(a|s) > 0$ occurs an infinite number of times under the behavior policy, 3) the step sizes are decreased appropriately, 4) V_t is updated an infinite number of times for all states, such that the ratio of the update frequency of the most-updated state to the least-updated state is finite, and 5) the subtasks are in accordance with Definition 5.1.

RED Q-learning algorithm (tabular): We update $Q_t : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ as follows:

$$\tilde{R}_{t+1} = f(R_{t+1}, Z_{1,t}, Z_{2,t}, \dots, Z_{n,t})$$
(16a)

$$\delta_t = \tilde{R}_{t+1} - \bar{R}_t + \max_a Q_t(S_{t+1}, a) - Q_t(S_t, A_t)$$
(16b)

$$Q_{t+1}(S_t, A_t) = Q_t(S_t, A_t) + \alpha_t \delta_t \tag{16c}$$

$$\bar{R}_{t+1} = \bar{R}_t + \alpha_{\bar{r},t} \delta_t \tag{16d}$$

$$Z_{i,t+1} = Z_{i,t} + \alpha_{z_i,t}\beta_{i,t}, \quad \forall z_i \in \mathcal{Z}$$
(16e)

where R_t is the observed reward, $Z_{i,t}$ is an estimate of subtask z_i , $\beta_{i,t}$ is the reward-extended TD error for subtask z_i , \bar{R}_t is an estimate of the long-run average-reward of \tilde{R}_t , \bar{r}_{π} , δ_t is the TD error, and α_t , $\alpha_{\bar{r},t}$, and $\alpha_{z_i,t}$ are the step sizes. Wan et al. (2021) showed for their Differential Q-learning algorithm that R_t converges to $\bar{r}*$, and Q_t converges to a solution of q in Equation (4). We now provide an equivalent theorem for our RED Q-learning algorithm, which also shows that $Z_{i,t}$ converges to the corresponding optimal subtask value $z_i^* \ \forall z_i \in \mathcal{Z}$:

Theorem 5.3 (informal). The RED Q-learning algorithm (16) converges, almost surely, \bar{R}_t to $\bar{r}*$, $Z_{i,t}$ to z_i^* $\forall z_i \in \mathcal{Z}$, \bar{r}_{π_t} to $\bar{r}*$, z_{i,π_t} to z_i^* $\forall z_i \in \mathcal{Z}$, and Q_t to a solution of q in the Bellman optimality equation (4), up to an additive constant, c, where π_t is any greedy policy with respect to Q_t , if the following assumptions hold: 1) the MDP is communicating, 2) the solution of q in (4) is unique up to a constant, 3) the step sizes are decreased appropriately, 4) Q_t is updated an infinite number of times for all state-action pairs, such that the ratio of the update frequency of the most-updated state—action pair to the least-updated state—action pair is finite, and 5) the subtasks are in accordance with Definition 5.1.

See Appendix B for the formal version of these theorems, along with the full convergence proofs.

6 Case Study: RED RL for CVaR Optimization

In this section, we present a case-study which illustrates how the subtask-driven approach that was derived in Section 5 can be used to successfully optimize the CVaR risk measure, *without* the use of an explicit bi-level optimization scheme (as in Equation (8)), or an augmented state-space.

First, in order to leverage the RED RL framework for CVaR optimization, we need to derive a valid subtask function for CVaR that satisfies the requirements of Definition 5.1. It turns out that we can use Equation (7) as a basis for the subtask function. The details of the adaptation of Equation (7) into a subtask function are presented in Appendix C. Critically, as discussed in Appendix C, optimizing the long-run average of the *extended* reward (\tilde{R}_t) from this subtask function corresponds to optimizing the long-run CVaR of the *observed* reward (R_t) . Hence, we can utilize CVaR-specific versions of the RED algorithms presented in Equations (15) and (16) (or their non-tabular equivalents) to optimize VaR and CVaR, such that CVaR corresponds to the primary control objective (i.e., the \bar{r}_{π} that we want to optimize), and VaR is the (single) subtask. We call the resulting algorithms, the *RED CVaR algorithms*. These algorithms, which are shown in full in Appendix C, update CVaR in an analogous way to the average-reward (i.e., CVaR corresponds to \bar{R}_t in Equations (15) or (16)), and update VaR using a VaR-specific version of Equation (15e) or (16e) as follows:

$$VaR_{t+1} = \begin{cases} VaR_t + \alpha_{VaR,t} \left(\delta_t + CVaR_t - VaR_t \right), & R_{t+1} \ge VaR_t \\ VaR_t + \alpha_{VaR,t} \left(\left(\frac{\tau}{\tau - 1} \right) \delta_t + CVaR_t - VaR_t \right), & R_{t+1} < VaR_t \end{cases},$$
(17)

where τ is the CVaR parameter, δ_t is the TD error, R_t is the observed reward, and $\alpha_{\text{\tiny VaR}},t$ is the step size. As such, given Theorems 5.1 - 5.3, we now have a subtask-driven approach for CVaR optimization that is able to simultaneously optimize VaR and CVaR without the use of an explicit bi-level optimization scheme or an augmented state-space (see Appendix C for a more formal argument).

We now present empirical results obtained when applying the RED CVaR algorithms on two RL tasks. The first task corresponds to a two-state environment that we created to test the RED CVaR algorithms. It is called the *red-pill blue-pill* task (see Appendix E), where at every time step an agent can take either a 'red pill', which takes them to the 'red world' state, or a 'blue pill', which takes them to the 'blue world' state. Each state has its own characteristic per-step reward distribution, and in this case, for a sufficiently low CVaR parameter, τ , the red world state has a reward distribution with a lower (worse) mean but a higher (better) CVaR compared to the blue world state. As such, this task allows us to answer the following question: *can the RED CVaR algorithms successfully enable the agent to learn a policy that prioritizes optimizing the reward CVaR over the average-reward?* In particular, we would expect that the RED CVaR algorithms learn a policy that prefers to stay in the red world, and that regular, risk-neutral Differential algorithms (i.e., from Wan et al. (2021)) learn a policy that prefers to stay in the blue world. This task is illustrated in Figure 2.

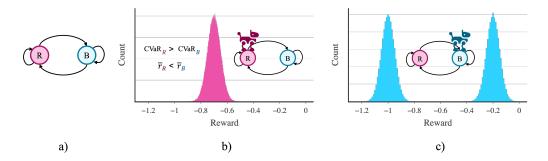


Figure 2: a) The *red-pill blue-pill* environment. b) + c) The per-step reward distributions of the b) 'red world', and c) 'blue world' states. For a sufficiently low CVaR parameter, τ , the red world state has a lower (worse) average-reward but a higher (better) reward CVaR than the blue world state.

The second task is the well-known *inverted pendulum* task, where an agent learns how to optimally balance an inverted pendulum. We chose this task because it provides us with the opportunity to test our algorithms in an environment where: 1) we must use function approximation (given the high-dimensional state-space), and 2) where the optimal CVaR policy and the optimal average-reward policy are the same policy (i.e., the policy that best balances the pendulum will yield a limiting reward distribution with both the optimal average-reward and reward CVaR). This hence allows us to directly compare the performance of our RED CVaR algorithms to that of the Differential algorithms, as well as to gauge how function approximation affects the performance of our algorithms.

In terms of empirical results, Figure 3 shows rolling averages of the average-reward and reward CVaR as learning progresses in both tasks when using a regular (risk-neutral) Differential algorithm (to optimize the average-reward) vs. a RED CVaR algorithm (to optimize the reward CVaR). As shown in the figure, in the red-pill blue-pill task, the RED CVaR algorithm successfully enables the agent to learn a policy that prioritizes maximizing the reward CVaR over the average-reward, thereby achieving a sort of *risk-awareness*. In the inverted pendulum task, both methods converge to the same policy, as expected.

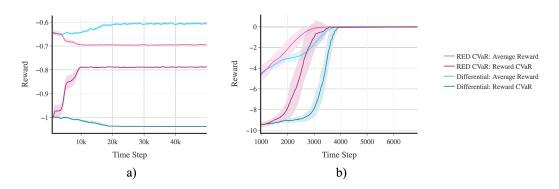


Figure 3: Rolling average-reward and reward CVaR as learning progresses when using a (risk-neutral) Differential algorithm vs. a (risk-aware) RED CVaR algorithm in the **a**) red-pill blue-pill, and **b**) inverted pendulum tasks. A solid line denotes the mean average-reward or reward CVaR, and the corresponding shaded region denotes a 95% confidence interval over **a**) 50 runs, or **b**) 10 runs. In both tasks, the RED CVaR algorithms enable the agent to learn a policy that prioritizes maximizing the reward CVaR over the average-reward, thereby achieving a sort of *risk-awareness*.

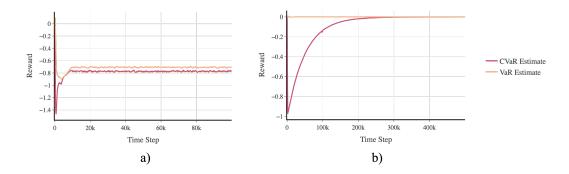


Figure 4: Typical convergence plots of the agent's VaR and CVaR estimates as learning progresses when using the RED CVaR algorithms in the **a**) red-pill blue-pill, and **b**) inverted pendulum tasks with an initial guess of 0.0 for both estimates. In both tasks, the estimates converge to the correct VaR and CVaR values, up to an additive constant, thereby yielding the optimal CVaR policy, and hence, the results shown in Figure 3.

Figure 4 shows typical convergence plots of the agent's VaR and CVaR estimates as learning progresses in both tasks when using the RED CVaR algorithms. As shown in the figure, the estimates converge in both tasks. In particular, the estimates converge to the correct VaR and CVaR values, up to an additive constant, thereby yielding the optimal CVaR policy, and hence, the results in Figure 3.

The full set of experimental details and results, including additional experiments performed, can be found in Appendix D.

7 Discussion, Limitations, and Future Work

In this work, we introduced *reward-extended differential* (or *RED*) reinforcement learning: a novel reinforcement learning framework that can be used to solve various learning objectives, or *subtasks*, simultaneously in the average-reward setting. We introduced a family of RED RL algorithms for prediction and control, and then showcased how these algorithms could be utilized to effectively and efficiently tackle the CVaR optimization problem. More specifically, we were able to use the RED RL framework to derive a set of algorithms that can optimize the CVaR risk measure without using an explicit bi-level optimization scheme or an augmented state-space, thereby alleviating some of the computational challenges and complexities that arise when performing risk-based optimization in the discounted setting. Empirically, we showed that the RED-based CVaR algorithms fared well in both tabular and linear function approximation settings.

More broadly, our work has introduced a theoretically-sound framework that allows for a subtask-driven approach to reinforcement learning, where various learning objectives (or subtasks) are solved simultaneously to help solve a larger, central learning objective. In this work, we showed (both theoretically and empirically) how this framework can be utilized to predict and/or optimize any arbitrary number of subtasks simultaneously in the average-reward setting. Central to this result is the novel concept of the reward-extended TD error, which is utilized in our framework to develop learning rules for the subtasks, and satisfies key theoretical properties that make it possible to solve any given subtask in a fully-online manner by minimizing the regular TD error. Moreover, we built upon existing results from Wan et al. (2021) to show the almost sure convergence of tabular algorithms derived from our framework. While we have only begun to grasp the implications of our framework, we have already seen some promising indications in the CVaR case study: the ability to turn explicit bi-level optimization problems into implicit bi-level optimizations that can be solved in a fully-online manner, as well as the potential to turn certain states (that meet certain conditions) into subtasks, thereby reducing the size of the state-space.

Nonetheless, while these results are encouraging, they are subject to a number of limitations. Firstly, by nature of operating in the average-reward setting, we are subject to the somewhat-strict assumptions made about the Markov chain induced by the policy (e.g. unichain or communicating). These assumptions could restrict the applicability of our framework, as they may not always hold in practice. Similarly, our definition for a subtask requires that the associated subtask function be linear or piecewise linear with respect to the subtasks, which may limit the applicability of our framework to simpler subtask functions. Finally, it remains to be seen empirically how our framework performs when dealing with multiple subtasks, when taking on more complex tasks, and/or when utilizing nonlinear function approximation. Importantly, we emphasize that the empirical evaluation performed in this work is limited, and as such, a more comprehensive empirical study is needed to fully gauge the practical implications of the proposed framework and CVaR algorithms.

Future work should look to address these limitations, as well as explore how these promising results can be extended to other domains, beyond the risk-awareness problem. In particular, we believe that the ability to optimize various subtasks simultaneously, as well as the potential to reduce the size of the state-space, by converting certain states to subtasks (where appropriate), could help alleviate significant computational challenges in other areas moving forward.

Acknowledgments

We gratefully acknowledge funding from NSERC Discovery Grant # RGPIN-2021-02760. We thank Margaret P. Chapman for insightful conversations and feedback in the early stages of this work. We thank Noah Sheridan for providing an insightful review of this work. We thank anonymous ICLR reviewer dBpb for their valuable feedback and suggestions during an earlier review process. Finally, we thank the anonymous RLC/RLJ reviewers and area chair for their useful feedback and commentary during the review process.

References

- J Abounadi, D Bertsekas, and V S Borkar. Learning algorithms for markov decision processes with average cost. *SIAM Journal on Control and Optimization*, 40(3):681–698, 2001.
- Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Math. Finance*, 9(3):203–228, July 1999.
- Andrew G Barto, Richard S Sutton, and Charles W Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Trans. Syst. Man Cybern.*, SMC-13(5):834–846, September 1983.
- Shalabh Bhatnagar, Richard S Sutton, Mohammad Ghavamzadeh, and Mark Lee. Natural actor–critic algorithms. *Automatica* (Oxf.), 45(11):2471–2482, November 2009.
- Kang Boda and Jerzy A Filar. Time consistent dynamic risk measures. *Math. Methods Oper. Res.* (*Heidelb.*), 63(1):169–186, February 2006.
- Vivek S Borkar. Stochastic Approximation: A Dynamical Systems Viewpoint. Springer, 2009.
- Nicole Bäuerle and Jonathan Ott. Markov decision processes with average-value-at-risk criteria. *Math. Methods Oper. Res.*, 74(3):361–379, December 2011.
- Yinlam Chow, Aviv Tamar, Shie Mannor, and Marco Pavone. Risk-sensitive and robust decision-making: a CVaR optimization approach. In *Advances in Neural Information Processing Systems* 28, 2015.
- Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. Implicit quantile networks for distributional reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Gabriel Dulac-Arnold, Nir Levine, Daniel J Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Mach. Learn.*, 110(9):2419–2468, September 2021.
- Abhijit Gosavi. Reinforcement learning for long-run average cost. *Eur. J. Oper. Res.*, 155(3):654–674, June 2004.
- Jia Lin Hau, Erick Delage, Mohammad Ghavamzadeh, and Marek Petrik. On dynamic programming decompositions of static risk measures in markov decision processes. In *Advances in Neural Information Processing Systems 37*, 2023a.
- Jia Lin Hau, Marek Petrik, and Mohammad Ghavamzadeh. Entropic risk optimization in discounted MDPs. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, 2023b.
- Ronald A Howard and James E Matheson. Risk-sensitive markov decision processes. *Manage. Sci.*, 18(7):356–369, March 1972.

- Ramtin Keramati, Alex Tamkin, Christoph Dann, and Emma Brunskill. Being optimistic to be conservative: Quickly learning a CVaR policy. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):4436–4443, 2020.
- Shiau Hong Lim and Ilyas Malik. Distributional reinforcement learning for risk-sensitive policies. In *Advances in Neural Information Processing Systems 35*, 2022.
- Matthew McLeod, Chunlok Lo, Matthew Schlegel, Andrew Jacobsen, Raksha Kumaraswamy, Marhta White, and Adam White. Continual auxiliary task learning. In Advances in Neural Information Processing Systems 34, 2021.
- Yashaswini Murthy, Mehrdad Moharrami, and R Srikant. Modified policy iteration for exponential cost risk sensitive MDPs. In *Proceedings of the 5th Annual Learning for Dynamics and Control Conference*, 2023.
- Abhishek Naik. Reinforcement Learning for Continuing Problems Using Average Reward. PhD thesis, University of Alberta, 2024.
- L A Prashanth and Mohammad Ghavamzadeh. Variance-constrained actor-critic algorithms for discounted and average reward MDPs. *Mach. Learn.*, 105(3):367–417, December 2016.
- Martin L Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
- R Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–41, 2000.
- Andrzej Ruszczyński. Risk-averse dynamic programming for markov decision processes. *Math. Program.*, 125(2):235–261, October 2010.
- Anton Schwartz. A reinforcement learning method for maximizing undiscounted rewards. In *International Conference on Machine Learning*, 1993.
- Silvestr Stanko and Karel Macek. Risk-averse distributional reinforcement learning: A CVaR optimization approach. In *Proceedings of the 11th International Joint Conference on Computational Intelligence*, 2019.
- Richard S Sutton. Learning to predict by the methods of temporal differences. *Mach. Learn.*, 3:944, 1988.
- Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction, 2nd edition.* MIT Press, November 2018.
- Richard S Sutton, Doina Precup, and Satinder Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artif. Intell.*, 112(1):181–211, August 1999.
- John N Tsitsiklis and Benjamin Van Roy. Average cost temporal-difference learning. *Automatica*, 35:1799, 1999.
- Yi Wan, Abhishek Naik, and Richard S Sutton. Learning and planning in average-reward markov decision processes. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.
- Li Xia, Luyao Zhang, and Peter W Glynn. Risk-sensitive markov decision processes with long-run CVaR criterion. *Prod. Oper. Manag.*, 32(12):4049–4067, December 2023.

A RED RL Algorithms

In this appendix, we provide pseudocode for our RED RL algorithms. We first present tabular algorithms, whose convergence proofs are included in Appendix B, and then provide equivalent algorithms that utilize function approximation.

Algorithm 1 RED TD-Learning (Tabular)

```
Input: the policy \pi to be evaluated, policy B to be used, piecewise linear subtask function f with
n subtasks, m piecewise segments, piecewise conditions r_{j-1} \le R < r_j such that f_j denotes the
jth segment of f that satisfies r_{j-1} \leq R < r_j, and constants b_1^j, b_2^j, \dots, b_n^j \ \forall j = 1, 2, \dots, m
Algorithm parameters: step size parameters \alpha, \alpha_{\bar{r}}, \alpha_{z_1}, \alpha_{z_2}, ..., \alpha_{z_n}
Initialize V(s) \forall s; \bar{R} arbitrarily (e.g. to zero)
Initialize subtasks Z_1, Z_2, \dots, Z_n arbitrarily (e.g. to zero)
Obtain initial S
while still time to train do
   A \leftarrow action given by B for S
  Take action A, observe R, S'
   R = f(R, Z_1, Z_2, \dots, Z_n)
   \delta = \tilde{R} - \bar{R} + V(S') - V(S)
   \rho = \pi(A \mid S) / B(A \mid S)
   V(S) = V(S) + \alpha \rho \delta
   \bar{R} = \bar{R} + \alpha_{\bar{r}} \rho \delta
   for i = 1, 2, ..., n do
      \beta_i = \sum_{j=1}^m (-1/b_i^j)(f_j - \bar{R} - \delta) \mathbb{1}\{r_{j-1} \le R < r_j\} \text{ (see Remark 5.1 for non-piecewise } f)
      Z_i = Z_i + \alpha_{z_i} \rho \beta_i
   end for
   S = S'
end while
return V
```

Algorithm 2 RED Q-Learning (Tabular)

```
Input: the policy \pi to be used (e.g., \varepsilon-greedy), piecewise linear subtask function f with n sub-
tasks, m piecewise segments, piecewise conditions r_{i-1} \le R < r_i such that f_i denotes the jth
segment of f that satisfies r_{j-1} \leq R < r_j, and constants b_1^j, b_2^j, \dots, b_n^j \ \forall j = 1, 2, \dots, m
Algorithm parameters: step size parameters \alpha, \alpha_{\bar{r}}, \alpha_{z_1}, \alpha_{z_2}, ..., \alpha_{z_n}
Initialize Q(s, a) \forall s, a; \bar{R} arbitrarily (e.g. to zero)
Initialize subtasks Z_1, Z_2, \dots, Z_n arbitrarily (e.g. to zero)
Obtain initial S
while still time to train do
   A \leftarrow action given by \pi for S
   Take action A, observe R, S'
   R = f(R, Z_1, Z_2, \dots, Z_n)
   \delta = \tilde{R} - \bar{R} + \max_a Q(S', a) - Q(S, A)
   Q(S, A) = Q(S, A) + \alpha \delta
   \bar{R} = \bar{R} + \alpha_{\bar{r}} \delta
   for i = 1, 2, ..., n do
      \beta_i = \sum_{j=1}^m (-1/b_i^j)(f_j - \bar{R} - \delta) \mathbb{1}\{r_{j-1} \le R < r_j\} (see Remark 5.1 for non-piecewise f
      Z_i = Z_i + \alpha_{z_i} \beta_i
   end for
   S = S'
end while
return Q
```

Algorithm 3 RED TD-Learning (Function Approximation)

```
Input: the policy \pi to be evaluated, policy B to be used, a differentiable state-value function
parameterization: \hat{v}(s, w), piecewise linear subtask function f with n subtasks, m piecewise
segments, piecewise conditions r_{j-1} \leq R < r_j such that f_j denotes the jth segment of f that
satisfies r_{j-1} \leq R < r_j, and constants b_1^j, b_2^j, \dots, b_n^j \ \forall j = 1, 2, \dots, m
Algorithm parameters: step size parameters \alpha, \alpha_{\bar{r}}, \alpha_{z_1}, \alpha_{z_2}, ..., \alpha_{z_n}
Initialize state-value weights \boldsymbol{w} \in \mathbb{R}^d arbitrarily (e.g. to 0)
Initialize subtasks Z_1, Z_2, \dots, Z_n arbitrarily (e.g. to zero)
Obtain initial S
while still time to train do
   A \leftarrow action given by B for S
   Take action A, observe R, S'
   R = f(R, Z_1, Z_2, \dots, Z_n)
   \delta = R - \bar{R} + \hat{v}(S', \boldsymbol{w}) - \hat{v}(S, \boldsymbol{w})
   \rho = \pi(A \mid S)/B(A \mid S)
   \boldsymbol{w} = \boldsymbol{w} + \alpha \rho \delta \nabla \hat{v}(S, \boldsymbol{w})
   \bar{R} = \bar{R} + \alpha_{\bar{r}} \rho \delta
   for i = 1, 2, ..., n do
      \beta_i = \sum_{j=1}^m (-1/b_i^j)(f_j - \bar{R} - \delta)\mathbb{1}\{r_{j-1} \le R < r_j\} (see Remark 5.1 for non-piecewise f)
      Z_i = Z_i + \alpha_{z_i} \rho \beta_i
   end for
   S = S'
end while
return w
```

Algorithm 4 RED Q-Learning (Function Approximation)

```
Input: the policy \pi to be used (e.g., \varepsilon-greedy), a differentiable state-action value function pa-
rameterization: \hat{q}(s, a, w), piecewise linear subtask function f with n subtasks, m piecewise
segments, piecewise conditions r_{i-1} \leq R < r_i such that f_i denotes the jth segment of f that
satisfies r_{j-1} \leq R < r_j, and constants b_1^j, b_2^j, \ldots, b_n^j \ \forall j = 1, 2, \ldots, m
Algorithm parameters: step size parameters \alpha, \alpha_{\bar{r}}, \alpha_{z_1}, \alpha_{z_2}, ..., \alpha_{z_n}
Initialize state-action value weights \boldsymbol{w} \in \mathbb{R}^d arbitrarily (e.g. to 0)
Initialize subtasks Z_1, Z_2, \dots, Z_n arbitrarily (e.g. to zero)
Obtain initial S
while still time to train do
   A \leftarrow action given by \pi for S
   Take action A, observe R, S'
   \tilde{R} = f(R, Z_1, Z_2, \dots, Z_n)
   \delta = R - \bar{R} + \max_{a} \hat{q}(S', a, \boldsymbol{w}) - \hat{q}(S, A, \boldsymbol{w})
   \mathbf{w} = \mathbf{w} + \alpha \delta \nabla \hat{q}(S, A, \mathbf{w})
   \bar{R} = \bar{R} + \alpha_{\bar{r}} \delta
   for i=1,2,\ldots,n do
      \beta_i = \sum_{j=1}^m (-1/b_i^j)(f_j - \bar{R} - \delta)\mathbb{1}\{r_{j-1} \le R < r_j\} (see Remark 5.1 for non-piecewise f)
      Z_i = Z_i + \alpha_{z_i} \beta_i
   end for
   S = S'
end while
return w
```

B Convergence Proofs

In this appendix, we present the full convergence proofs for the tabular RED TD-learning and tabular RED Q-learning algorithms. Our general strategy is as follows: we first show that the results from Wan et al. (2021), which show the almost sure convergence of the value function and average-reward estimates of differential algorithms, are applicable to our algorithms. We then build upon these results to show that the subtask estimates of our algorithms converge as well.

For consistency, we adopt similar notation as Wan et al. (2021) for our proofs:

- For a given vector x, let $\sum x$ denote the sum of all elements in x, such that $\sum x \doteq \sum_i x(i)$.
- Let \bar{r}_* denote the optimal average-reward.
- Let z_{i_*} denote the corresponding optimal subtask value for subtask $z_i \in \mathcal{Z}$.

B.1 Convergence Proof for the Tabular RED TD-learning Algorithm

In this section, we present the proof for the convergence of the value function, average-reward, and subtask estimates of the RED TD-learning algorithm. Similar to what was done in Wan et al. (2021), we will begin by considering a general algorithm, called *General RED TD*. We will first define General RED TD, then show how the RED TD-learning algorithm is a special case of this algorithm. We will then provide the necessary assumptions, state the convergence theorem of General RED TD, and then provide a proof for the theorem, where we show that the value function, average-reward, and subtask estimates converge, thereby showing that the RED TD-learning algorithm converges. We begin by introducing the General RED TD algorithm:

Consider an MDP $\mathcal{M} \doteq \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, p \rangle$, a behavior policy, B, and a target policy, π . Given a state $s \in \mathcal{S}$ and discrete step $n \geq 0$, let $A_n(s) \sim B(\cdot \mid s)$ denote the action selected using the behavior policy, let $R_n(s, A_n(s)) \in \mathcal{R}$ denote a sample of the resulting reward, and let $S'_n(s, A_n(s)) \sim p(\cdot, \cdot \mid s, a)$ denote a sample of the resulting state. Let $\{Y_n\}$ be a set-valued process taking values in the set of nonempty subsets of \mathcal{S} , such that: $Y_n = \{s : s \text{ component of the } |\mathcal{S}| \text{-sized table of state-value estimates, } V$, that was updated at step $n\}$. Let $\nu(n,s) \doteq \sum_{j=0}^n I\{s \in Y_j\}$, where I is the indicator function, such that $\nu(n,s)$ represents the number of times that V(s) was updated up until step n.

Now, let f be a valid subtask function (see Definition 5.1), such that $\tilde{R}_n(s,A_n(s)) \doteq f(R_n(s,A_n(s)),Z_{1,n},Z_{2,n},\ldots,Z_{k,n})$ for k subtasks $\in \mathcal{Z}$, where $\tilde{R}_n(s,A_n(s))$ is the extended reward, \mathcal{Z} is the set of subtasks, and $Z_{i,n}$ denotes the estimate of subtask $z_i \in \mathcal{Z}$ at step n. Consider an MDP with the extended reward: $\tilde{\mathcal{M}} \doteq \langle \mathcal{S}, \mathcal{A}, \tilde{\mathcal{R}}, \tilde{p} \rangle$, such that $\tilde{R}_n(s,A_n(s)) \in \tilde{\mathcal{R}}$. The update rules of General RED TD for this MDP are as follows, $\forall n \geq 0$:

$$V_{n+1}(s) \doteq V_n(s) + \alpha_{\nu(n,s)} \rho_n(s) \delta_n(s) I\{s \in Y_n\}, \quad \forall s \in \mathcal{S},$$
(B.1)

$$\bar{R}_{n+1} \doteq \bar{R}_n + \sum_{s} \alpha_{\bar{r},\nu(n,s)} \rho_n(s) \delta_n(s) I\{s \in Y_n\}, \tag{B.2}$$

$$Z_{i,n+1} \doteq Z_{i,n} + \sum_{s} \alpha_{z_i,\nu(n,s)} \rho_n(s) \beta_{i,n}(s) I\{s \in Y_n\}, \quad \forall z_i \in \mathcal{Z},$$
 (B.3)

where,

$$\delta_n(s) \doteq \tilde{R}_n(s, A_n(s)) - \bar{R}_n + V_n(S'_n(s, A_n(s))) - V_n(s)$$

$$= f(R_n(s, A_n(s)), Z_{1,n}, Z_{2,n}, \dots, Z_{k,n}) - \bar{R}_n + V_n(S'_n(s, A_n(s))) - V_n(s),$$
(B.4)

and,

$$\beta_{i,n}(s) \doteq \phi_{i,n}(s) - Z_{i,n}, \quad \forall z_i \in \mathcal{Z}.$$
 (B.5)

Here, $\rho_n(s) \doteq \pi(A_n(s) \mid s) / B(A_n(s) \mid s)$ denotes the importance sampling ratio (with behavior policy, B), \bar{R}_n denotes the estimate of the average-reward (see Equation (2)), $\delta_n(s)$ denotes the TD error, $\phi_{i,n}(s)$ denotes the (potentially-piecewise) subtask target, as defined in Section 5.1, and $\alpha_{\nu(n,s)}$, $\alpha_{\bar{r},\nu(n,s)}$, and $\alpha_{z_i,\nu(n,s)}$ denote the step sizes at time step n for state s.

We now show that the RED TD-learning algorithm is a special case of the General RED TD algorithm. Consider a sequence of experience from our MDP, $\tilde{\mathcal{M}}$: $S_t, A_t(S_t), \tilde{R}_{t+1}, S_{t+1}, \ldots$. Now recall the set-valued process $\{Y_n\}$. If we let n = time step t, we have:

$$Y_t(s) = \begin{cases} 1, s = S_t, \\ 0, \text{ otherwise,} \end{cases}$$

as well as $S'_n(S_t, A_t(S_t)) = S_{t+1}, R_n(S_t, A_t) = R_{t+1}, \text{ and } \tilde{R}_n(S_t, A_t(S_t)) = \tilde{R}_{t+1}.$

Hence, update rules (B.1), (B.2), (B.3), (B.4), and (B.5) become:

$$V_{t+1}(S_t) \doteq V_t(S_t) + \alpha_{\nu(t,S_t)} \rho_t(S_t) \delta_t \text{ and } V_{t+1}(s) \doteq V_t(s), \forall s \neq S_t,$$
 (B.6)

$$\bar{R}_{t+1} \doteq \bar{R}_t + \alpha_{\bar{r},\nu(t,S_t)} \rho_t(S_t) \delta_t, \tag{B.7}$$

$$Z_{i,t+1} \doteq Z_{i,t} + \alpha_{z_i,\nu(t,S_t)} \rho_t(S_t) \beta_{i,t}, \quad \forall z_i \in \mathcal{Z},$$
(B.8)

$$\delta_t \doteq \tilde{R}_{t+1} - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t),
= f(R_{t+1}, Z_{1,t}, Z_{2,t}, \dots, Z_{k,t}) - \bar{R}_t + V_t(S_{t+1}) - V_t(S_t),$$
(B.9)

$$\beta_{i,t} \doteq \phi_{i,t} - Z_{i,t}, \quad \forall z_i \in \mathcal{Z},$$
 (B.10)

which are RED TD-learning's update rules with $\alpha_{\nu(t,S_t)}$, $\alpha_{\bar{r},\nu(t,S_t)}$, and $\alpha_{z_i,\nu(t,S_t)}$ denoting the step sizes at time t.

We now specify the assumptions on General RED TD that are needed to ensure convergence:

Assumption B.1 (Unichain Assumption). The Markov chain induced by the policy is unichain. That is, the induced Markov chain consists of a single recurrent class and a potentially-empty set of transient states.

Assumption B.2 (Coverage Assumption). $B(a \mid s) > 0$ if $\pi(a \mid s) > 0$ for all $s \in \mathcal{S}$, $a \in \mathcal{A}$.

Assumption B.3 (Step Size Assumption). $\alpha_n > 0$, $\sum_{n=0}^{\infty} \alpha_n = \infty$, $\sum_{n=0}^{\infty} \alpha_n^2 < \infty$.

Assumption B.4 (Asynchronous Step Size Assumption 1). Let $[\cdot]$ denote the integer part of (\cdot) . For $x \in (0,1)$,

$$\sup_{i} \frac{\alpha_{[xi]}}{\alpha_i} < \infty$$

and

$$\frac{\sum_{j=0}^{[yi]} \alpha_j}{\sum_{j=0}^{i} \alpha_j} \to 1$$

uniformly in $y \in [x, 1]$.

Assumption B.5 (Asynchronous Step Size Assumption 2). There exists $\Delta > 0$ such that

$$\liminf_{n \to \infty} \frac{\nu(n,s)}{n+1} \ge \Delta,$$

a.s., for all $s \in \mathcal{S}$.

Furthermore, for all x > 0, and

$$N(n,x) = \min \left\{ m \ge n : \sum_{i=n+1}^{m} \alpha_i \ge x \right\},\,$$

the limit

$$\lim_{n \to \infty} \frac{\sum_{i=\nu(n,s)}^{\nu(N(n,x),s)} \alpha_i}{\sum_{i=\nu(n,s')}^{\nu(N(n,x),s')} \alpha_i}$$

exists a.s. for all s, s'.

Assumption B.6 (Average-Reward Step Size Assumption). The average-reward step size, $\alpha_{\bar{r},n}$, can be written as a constant fraction of the value function step size, α_n , such that $\alpha_{\bar{r},n} \doteq \eta_r \alpha_n$, where η_r is a positive scalar.

Assumption B.7 (Subtask Function Assumption). *The subtask function*, f, is 1) linear or piecewise linear, and 2) is invertible with respect to each input given all other inputs.

Assumption B.8 (Subtask Independence Assumption). Each subtask, $z_i \in \mathcal{Z}$, in f is independent of the states and actions, and hence independent of the observed reward, R_n , such that $\tilde{p}(s', f(r, z_1, \ldots, z_k)|s, a) = p(s', r|s, a)$, and $\mathbb{E}[f_j(R_n, Z_{1,n}, Z_{2,n}, \ldots, Z_{k,n})] = f_j(\mathbb{E}[R_n], Z_{1,n}, Z_{2,n}, \ldots, Z_{k,n})$, where f_j denotes the jth piecewise segment of f, and \mathbb{E} denotes any expectation taken with respect to the states and actions.

Assumption B.9 (Subtask Uniqueness Assumption). If the Bellman equation associated with $f(R_n, z_1, z_2, ..., z_k)$ admits a unique solution, then that unique solution corresponds to a unique combination of subtasks, $z_1, z_2, ..., z_k$.

Assumption B.10 (Subtask Step Size Assumptions). If the subtask function is strictly (i.e., non-piecewise) linear, the subtask step sizes, $\{\alpha_{z_i,n}\}_{i=1}^k$, can be written as constant, subtask-specific fractions of the value function step size, α_n , such that $\alpha_{z_i,n} \doteq \eta_{z_i} \alpha_n \ \forall z_i \in \mathcal{Z}$, where $\{\eta_{z_i}\}_{i=1}^k$ are positive scalars. Alternatively, if the subtask function is piecewise linear with at least two piecewise segments, the subtask step sizes satisfy the following properties: $\alpha_{z_1,n}/\alpha_n \to 0$, $\alpha_{z_1,n}/\alpha_{\bar{r},n} \to 0$, $\{\alpha_{z_i,n}/\alpha_{z_{i-1},n} \to 0\}_{i=2}^k$, and $\sum_{n=0}^{\infty} (\alpha_n^2 + \alpha_{\bar{r},n}^2 + \alpha_{z_1,n}^2 + \alpha_{z_2,n}^2 + \ldots + \alpha_{z_k,n}^2) < \infty$.

We refer the reader to Wan et al. (2021) for an in-depth discussion on Assumptions B.1-B.6. Note that Assumptions B.3-B.5 apply to the value function, average-reward, and subtask step sizes. Assumptions B.7-B.10 outline the subtask-related requirements needed to show convergence. In particular, Assumption B.7 ensures that we can explicitly write out the update (B.3), and Assumption B.8 ensures that we do not break the Markov property in the process (i.e., we preserve the Markov property by ensuring that the subtasks are independent of the states and actions, and thereby also independent of the observed reward). Assumption B.9 ensures that only a unique combination of subtasks yields the solution to a given Bellman equation. Finally, Assumption B.10 outlines additional step size requirements needed to show convergence.

Having stated the necessary assumptions, we next point out that it is easy to verify that under Assumption B.1, the following system of equations:

$$v_{\pi}(s) = \sum_{a} \pi(a \mid s) \sum_{s', \bar{r}} \tilde{p}(s', \tilde{r} \mid s, a) (\tilde{r} - \bar{r}_{\pi} + v_{\pi}(s')), \quad \forall s \in \mathcal{S}$$

$$= \sum_{a} \pi(a \mid s) \sum_{s', r} p(s', r \mid s, a) (f(r, z_{1, \pi}, z_{2, \pi}, \dots, z_{k, \pi}) - \bar{r}_{\pi} + v_{\pi}(s')), \quad \forall s \in \mathcal{S},$$
(B.11)

and,

$$\bar{r}_{\pi} - \bar{R}_0 = \eta_r \left(\sum v_{\pi} - \sum V_0 \right), \tag{B.12}$$

$$z_{i,\pi} - Z_{i,0} = \eta_i \left(\sum v_{\pi} - \sum V_0 \right), \text{ for all } z_i \in \mathcal{Z},$$
(B.13)

has a unique solution of v_{π} , where \bar{r}_{π} denotes the average-reward induced by following a given policy, π , and $z_{i,\pi}$ denotes the corresponding subtask value for subtask $z_i \in \mathcal{Z}$. Denote this unique solution of v_{π} as v_{∞} .

We are now ready to state the convergence theorem:

Theorem B.1.1 (Convergence of General RED TD). *If Assumptions B.1 – B.10 hold, then General RED TD (Equations* (B.1) – (B.5)) *converges a.s.,* \bar{R}_n *to* \bar{r}_{π} , $Z_{i,n}$ *to* $z_{i,\pi}$ $\forall z_i \in \mathcal{Z}$, and V_n to v_{∞} .

We prove this theorem in Sections B.1.1 and B.1.2. To do so, we first show that General RED TD is of the same form as *General Differential TD* from Wan et al. (2021), thereby allowing us to apply their convergence results for the value function and average-reward estimates of General Differential TD to General RED TD. We then build upon these results, using similar techniques as Wan et al. (2021), to show that the subtask estimates converge as well.

B.1.1 Proof of Theorem **B.1.1** (for *Linear* Subtask Functions)

We first provide the proof for *linear* subtask functions, where the reward-extended TD error can be expressed as a constant, subtask-specific fraction of the regular TD error, such that $\beta_{i,n}(s) = (-1/b_i)\delta_n(s)$. We consider the *piecewise linear* case in Section B.1.2.

Convergence of the value function and average-reward estimates:

Consider the increment to \bar{R}_n at each step. Given Assumption B.6, we can see from Equation (B.2) that the increment is η_r times the increment to V_n . As such, as was done in Wan et al. (2021), we can write the cumulative increment as follows:

$$\bar{R}_n - \bar{R}_0 = \eta_r \sum_{j=0}^{n-1} \sum_s \alpha_{\nu(j,s)} \rho_j(s) \delta_j(s) I\{s \in Y_j\}$$

$$= \eta_r \left(\sum_s V_n - \sum_s V_0 \right)$$

$$\Rightarrow \bar{R}_n = \eta_r \sum_s V_n - \eta_r \sum_s V_0 + \bar{R}_0 = \eta_r \sum_s V_n - c_r, \tag{B.14}$$

where
$$c_r \doteq \eta_r \sum V_0 - \bar{R}_0$$
. (B.15)

Similarly, consider the increment to $Z_{i,n}$ (for an arbitrary subtask $z_i \in \mathcal{Z}$) at each step. As per Remark 5.1, and given Assumption B.10, we can write the increment in Equation (B.3) as some constant, subtask-specific fraction of the increment to V_n . Consequently, we can write the cumulative increment as follows:

$$Z_{i,n} - Z_{i,0} = \eta_{z_i} \sum_{j=0}^{n-1} \sum_{s} \alpha_{\nu(j,s)} \rho_j(s) \beta_{i,j}(s) I\{s \in Y_j\}$$

$$= \eta_{z_i} \sum_{j=0}^{n-1} \sum_{s} \alpha_{\nu(j,s)} \rho_j(s) (-1/b_i) \delta_j(s) I\{s \in Y_j\}$$

$$= \eta_i \left(\sum_{s} V_n - \sum_{s} V_0 \right)$$

$$\implies Z_{i,n} = \eta_i \sum V_n - \eta_i \sum V_0 + Z_{i,0} = \eta_i \sum V_n - c_i, \tag{B.16}$$

where,

$$c_i \doteq \eta_i \sum V_0 - Z_{i,0}, \text{ and}$$
 (B.17)

$$\eta_i \doteq (-1/b_i)\eta_{z_i}. \tag{B.18}$$

Now consider the subtask function, f. At any given time step, the subtask function can be written as: $f_n = \tilde{R}_n(s, A_n(s)) = b_r R_n(s, A_n(s)) + b_0 + b_1 Z_{1,n} + \ldots + b_k Z_{k,n}$, where $b_r, b_0 \in \mathbb{R}$ and $b_i \in \mathbb{R} \setminus \{0\}$. Given Equation (B.16), we can write the subtask function as follows:

$$f_n = b_r R_n(s, A_n(s)) + b_0 + b_1(\eta_1 \sum V_n - c_1) + \dots + b_k(\eta_k \sum V_n - c_k)$$

$$= b_r R_n(s, A_n(s)) + \eta_f \sum V_n - c_f,$$
(B.19)

where,
$$\eta_f = \sum_{j=1}^k b_j \eta_j$$
 and $c_f = \sum_{j=1}^k b_j c_j - b_0$.

As such, we can substitute \bar{R}_n and $Z_{i,n}$ $\forall z_i \in \mathcal{Z}$ in (B.1) with (B.14) and (B.19), respectively, $\forall s \in \mathcal{S}$, which yields:

$$V_{n+1}(s) = V_n(s) + \dots$$

$$\alpha_{\nu(n,s)}\rho_n(s) \left(b_r R_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta_r \sum_{s} V_n + c_r + \eta_f \sum_{s} V_n - c_f \right) I\{s \in Y_n\}$$

$$V_{n+1}(s) = V_n(s) + \dots$$

$$\alpha_{\nu(n,s)}\rho_n(s) \left(b_r R_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta_T \sum_{s} V_n + c_T \right) I\{s \in Y_n\}$$

$$V_{n+1}(s) = V_n(s) + \dots$$

$$\alpha_{\nu(n,s)}\rho_n(s) \left(\widehat{R}_n(s, A_n(s)) + V_n(S'_n(s, A_n(s))) - V_n(s) - \eta_T \sum V_n\right) I\{s \in Y_n\},$$
(B.20)

where
$$\eta_{\scriptscriptstyle T}=\eta_{\scriptscriptstyle r}-\eta_{\scriptscriptstyle f},$$
 $c_{\scriptscriptstyle T}=c_r-c_f,$ and $\widehat{R}_n(s,A_n(s))\doteq b_rR_n(s,A_n(s))+c_{\scriptscriptstyle T}.$

Equation (B.20) is now in the same form as Equation (B.37) (i.e., General Differential TD) from Wan et al. (2021), who showed that the equation converges a.s. V_n to v_∞ as $n \to \infty$. Moreover, from this result, Wan et al. (2021) showed that \bar{R}_n converges a.s. to \bar{r}_π as $n \to \infty$. Given that General RED TD adheres to all the assumptions listed for General Differential TD in Wan et al. (2021), these convergence results apply to General RED TD.

Convergence of the subtask estimates:

Consider Equation (B.20). We can rewrite this equation, $\forall s \in \mathcal{S}$, as follows:

$$V_{n+1}(s) = V_n(s) + \dots$$

$$\alpha_{\nu(n,s)}\rho_n(s) \left(b_r R_n(s, A_n(s)) + \eta_f \sum_{s} V_n + c_r - \eta_r \sum_{s} V_n + V_n(S'_n(s, A_n(s))) - V_n(s) \right) I\{s \in Y_n\}$$

$$V_{n+1}(s) = V_n(s) + \dots$$

$$\alpha_{\nu(n,s)}\rho_n(s) \left(\hat{R}_n(s, A_n(s)) + c_T - \eta_r \sum V_n + V_n(S'_n(s, A_n(s))) - V_n(s)\right) I\{s \in Y_n\},$$
(B.21)

where,

$$\hat{R}_n(s, A_n(s)) \doteq b_r R_n(s, A_n(s)) + \eta_f \sum V_n$$
(B.22)

$$= b_r R_n(s, A_n(s)) + b_1(\eta_1 \sum V_n) + \ldots + b_k(\eta_k \sum V_n)$$
 (B.23)

$$= b_r R_n(s, A_n(s)) + b_1 \hat{Z}_{1,n} + \ldots + b_k \hat{Z}_{k,n}.$$
 (B.24)

Now consider an MDP, $\hat{\mathcal{M}}$, which has rewards, $\hat{\mathcal{R}}$, as defined in Equation (B.22), has the same state and action spaces as the MDP $\hat{\mathcal{M}}$, and has the transition probabilities defined as:

$$\hat{p}(s', \hat{r} \mid s, a) \doteq p(s', r \mid s, a) \tag{B.25}$$

$$= \tilde{p}(s', \tilde{r} \mid s, a)$$
 (by Definition 5.1), (B.26)

such that $\hat{\mathcal{M}} \doteq \langle \mathcal{S}, \mathcal{A}, \hat{\mathcal{R}}, \hat{p} \rangle$. It is easy to check that the unichain assumption holds for the MDP, $\hat{\mathcal{M}}$. Moreover, given Equation (B.21) and Assumptions B.7 and B.8, the average-reward induced by following policy π for the MDP, $\hat{\mathcal{M}}$, \hat{r}_{π} , can be written as follows:

$$\hat{\bar{r}}_{\pi} = \bar{r}_{\pi} - c_{\pi}. \tag{B.27}$$

Now, because

$$\begin{split} v_{\infty}(s) &= \sum_{a} \pi(a \mid s) \sum_{s', \tilde{r}} \tilde{p}(s', \tilde{r} \mid s, a) (\tilde{r} - \bar{r}_{\pi} + v_{\infty}(s')) \quad \text{(from (B.11))} \\ &= \sum_{a} \pi(a \mid s) \sum_{s', \tilde{r}} \tilde{p}(s', \tilde{r} \mid s, a) (\tilde{r} - (\hat{\bar{r}}_{\pi} + c_{_{T}}) + v_{\infty}(s')) \quad \text{(from (B.27))} \\ &= \sum_{a} \pi(a \mid s) \sum_{s', \tilde{r}} \tilde{p}(s', \tilde{r} \mid s, a) (\tilde{r} - c_{_{T}} - \hat{\bar{r}}_{\pi} + v_{\infty}(s')) \\ &= \sum_{a} \pi(a \mid s) \sum_{s', \tilde{r}} \tilde{p}(s', \tilde{r} \mid s, a) (\hat{r} - \hat{\bar{r}}_{\pi} + v_{\infty}(s')) \quad \text{(from (B.21))} \\ &= \sum_{a} \pi(a \mid s) \sum_{s', \tilde{r}} \hat{p}(s', \hat{r} \mid s, a) (\hat{r} - \hat{\bar{r}}_{\pi} + v_{\infty}(s')) \quad \text{(from (B.26))}, \end{split}$$

we can see that v_{∞} is a solution of not just the state-value Bellman equation for the MDP, $\tilde{\mathcal{M}}$, but also the state-value Bellman equation for the MDP, $\hat{\mathcal{M}}$.

Next, consider an arbitrary *i*th subtask. As per Equations (B.23) and (B.24), we can write the subtask value induced by following policy π for the MDP, $\hat{\mathcal{M}}$, $\hat{z}_{i,\pi}$, as follows:

$$\hat{z}_{i,\pi} = z_{i,\pi} + c_i. \tag{B.28}$$

We can then combine Equations (B.13), (B.16), and (B.28), which yields:

$$\hat{z}_{i,\pi} = \eta_i \sum v_{\infty}. \tag{B.29}$$

Next, we can combine Equation (B.16) with the result from Wan et al. (2021) which shows that $V_n \to v_\infty$, which yields:

$$Z_{i,n} \to \eta_i \sum v_{\infty} - c_i.$$
 (B.30)

Moreover, because $\eta_i \sum v_{\infty} = \hat{z}_{i,\pi}$ (Equation (B.29)), we have:

$$Z_{i,n} \to \hat{z}_{i,\pi} - c_i. \tag{B.31}$$

Finally, because $\hat{z}_{i,\pi} = z_{i,\pi} + c_i$ (Equation (B.28)), we have:

$$Z_{i,n} \to z_{i,\pi}$$
 a.s. as $n \to \infty$. (B.32)

B.1.2 Proof of Theorem **B.1.1** (for *Piecewise Linear* Subtask Functions)

We now provide the proof for *piecewise linear* subtask functions, where the reward-extended TD error can be expressed as follows:

$$\beta_{i,n}(s) = \begin{cases} (-1/b_i^1) \left(\tilde{R}_{1,n}(s, A_n(s)) - \bar{R}_n - \delta_n(s) \right), & r_0 \leq R_n(s, A_n(s)) < r_1 \\ \vdots \\ (-1/b_i^m) \left(\tilde{R}_{m,n}(s, A_n(s)) - \bar{R}_n - \delta_n(s) \right), & r_{m-1} \leq R_n(s, A_n(s)) \leq r_m \end{cases}$$

where $r_u \in \mathcal{R} \ \forall u = 0, 1, \dots, m$, and $r_0 \leq r_1 \leq \dots \leq r_m$, such that r_0, r_m represent the lower and upper bounds of the observed per-step reward, $R_n(s, A_n(s))$, respectively. Our general strategy in this case is to use a *multiple-timescales* argument, such that we leverage Theorem 2 in Section 6 of Borkar (2009), along with the results from Theorem B.3 of Wan et al. (2021).

To begin, let us consider Assumption B.10, which enables the formulation of a multiple-timescales argument. In particular, the $\alpha_{z_1,n}/\alpha_n \to 0$, $\alpha_{z_1,n}/\alpha_{\bar{r},n} \to 0$, and $\{\alpha_{z_i,n}/\alpha_{z_{i-1},n} \to 0\}_{i=2}^k$ conditions imply that the subtask step sizes, $\{\alpha_{z_i,n}\}_{n=0}^{\infty} \forall z_i \in \mathcal{Z}$, decrease to 0 at faster rates than the value function and average-reward step sizes, $\{\alpha_n\}_{n=0}^{\infty}$ and $\{\alpha_{\bar{r},n}\}_{n=0}^{\infty}$, respectively. This implies that the subtask updates move on slower timescales compared to the value function and average-reward updates. Hence, as argued in Section 6 of Borkar (2009), the (faster) value function and average-reward updates, (B.1) and (B.2), view the (slower) subtask updates, (B.3), as quasi-static, while the (slower) subtask updates view the (faster) value function and average-reward updates as nearly equilibrated (as we will show below, the results from Wan et al. (2021) imply the existence of such an equilibrium point). Similarly, the $\{\alpha_{z_i,n}/\alpha_{z_{i-1},n} \to 0\}_{i=2}^k$ condition implies that each subtask update views the other subtask updates as either quasi-static or nearly-equilibrated.

As such, having established the multiple-timescales argument, we now proceed to show the convergence of the value function, average-reward, and subtask estimates:

Convergence of the value function and average-reward estimates:

Given the multiple-timescales argument, such that the subtask estimates are viewed as quasi-static (i.e., constant), Equation (B.1) can be viewed as being of the same form as Equation (B.30) (i.e., General Differential TD) from Wan et al. (2021), who showed (via Theorem B.3) that the equation converges, almost surely, V_n to v_∞ as $n \to \infty$. Moreover, from this result, Wan et al. (2021) showed that \bar{R}_n converges, almost surely, to \bar{r}_π as $n \to \infty$. Given that General RED TD adheres to all the assumptions listed for General Differential TD in Wan et al. (2021), these convergence results apply to General RED TD.

Convergence of the subtask estimates:

Let us consider the asynchronous subtask updates (B.3). Each update in (B.3) is of the same form as Equation 7.1.2 of Borkar (2009). Accordingly, to show the convergence of the subtask estimates, we can apply the result in Section 7.4 of Borkar (2009), which shows the convergence of asynchronous updates that are of the same form as Equation 7.1.2. To apply this result, given Assumptions B.4 and B.5, we only need to show the convergence of the *synchronous* version of the subtask updates:

$$Z_{i,n+1} = Z_{i,n} + \alpha_{z_i,n} \left(g_i(Z_{i,n}) + M_{n+1}^{z_i} \right) \ \forall z_i \in \mathcal{Z}, \tag{B.33}$$

where,

$$g_{i}(Z_{i,n})(s) \doteq \sum_{a} \pi(a \mid s) \sum_{s',r} p(s',r \mid s,a) \phi_{i,n}(s) - Z_{i,n},$$

$$\begin{cases}
-\frac{1}{b_{i}^{1}} \left(b_{r}^{1} R_{n}(s,A_{n}(s)) + \dots + b_{i-1}^{1} Z_{i-1,n} + b_{i+1}^{1} Z_{i+1,n} + \dots + b_{k}^{1} Z_{k,n} - \bar{R}_{n} - \delta_{n}(s) \right), \dots \\
\dots, r_{0} \leq R_{n}(s,A_{n}(s)) < r_{1},
\end{cases}$$

$$\vdots$$

$$\begin{cases}
-\frac{1}{b_{i}^{m}} \left(b_{r}^{m} R_{n}(s,A_{n}(s)) + \dots + b_{i-1}^{m} Z_{i-1,n} + b_{i+1}^{m} Z_{i+1,n} + \dots + b_{k}^{m} Z_{k,n} - \bar{R}_{n} - \delta_{n}(s) \right), \dots \\
\dots, r_{m-1} \leq R_{n}(s,A_{n}(s)) \leq r_{m},
\end{cases}$$

$$M_{n+1}^{z_{i}}(s) \doteq \rho_{n}(s) \left(\phi_{i,n}(s) - Z_{i,n} \right) - q_{i}(Z_{i,n})(s).$$

To show the convergence of the synchronous update (B.33) under the multiple-timescales argument, we can apply the result of Theorem 2 in Section 6 of Borkar (2009) to show that $Z_{i,n} \to z_{i,\pi} \forall z_i \in \mathcal{Z}$ a.s. as $n \to \infty$. This theorem requires that 3 assumptions be satisfied. As such, we will now show, via Lemmas B.1 - B.3, that these 3 assumptions are indeed satisfied:

Lemma B.1. The value function update, $V_{n+1} = V_n + \alpha_n(h(V_n) + M_{n+1})$, where

$$h(V_n)(s) \doteq \sum_{a} \pi(a \mid s) \sum_{s', \tilde{r}} \tilde{p}(s', \tilde{r} \mid s, a)(\tilde{r} - \bar{R}_n + V_n(s') - V_n(s)),$$

$$= \sum_{a} \pi(a \mid s) \sum_{s', r} p(s', r \mid s, a)(f(r, Z_{1,n}, Z_{2,n}, \dots, Z_{k,n}) - \psi(V_n) + V_n(s') - V_n(s)),$$

$$M_{n+1}(s) \doteq \rho_n(s) \left(\tilde{R}_n(s, A_n(s)) - \psi(V_n) + V_n(S'_n(s, A_n(s))) - V_n(s) \right) - h(V_n)(s),$$

 $\psi(V_n) = \bar{R}_n$ is a 'reference function' as defined in Wan et al. (2021), and $Z_{1,n}, Z_{2,n}, ..., Z_{k,n}$ are quasi-static under the multiple-timescales argument,

has a globally asymptotically stable equilibrium, $v_{\infty}(Z_{1,n}, Z_{2,n}, \dots, Z_{k,n})$, where v_{∞} is a Lipschitz map.

Proof. This was shown in Theorem B.3 of Wan et al. (2021).

Lemma B.2. The subtask update rules (B.33) each have a globally asymptotically stable equilibrium, $z_{i,\pi}$.

Proof. Let us begin by considering the 'fastest' subtask, $Z_{1,n}$, under the multiple-timescales argument, such that all other subtasks, $Z_{2,n},\ldots,Z_{k,n}$, operate on slower timescales and can be considered quasi-static. Applying the results of Theorem B.3 of Wan et al. (2021) under the multiple-timescales argument, we have that: $h(V_n) \to 0$ and $\bar{R}_n \to \bar{r}_\pi$ as $V_n \to v_\infty$. Importantly, as argued in Section 5, $h(V_n) \to 0$ implies that $\delta_n(s) \to \lambda_j \in \mathbb{R}$ for each jth piecewise segment in $\phi_{i,n}(s)$. As such, we can interpret $Z_{1,n}$ as being the only 'moving' variable in $g_1(Z_{1,n})$, such that all other parameters in the update are static, quasi-static, or nearly-equilibrated.

Let us now consider the ODE associated with $g_1(Z_{1,n})$,

$$\dot{x}_t = g_1(x_t). \tag{B.34}$$

To show that the update rule for $Z_{1,n}$ has a globally asymptotically stable equilibrium, $z_{1,\pi}$, it suffices to show that there exists a *Lyapunov function* for the associated ODE (B.34).

To this end, we first note, given the discussion in Section 5, that $z_{1,\pi}$ is an equilibrium point for the update rule associated with $Z_{1,n}$. Moreover, given Assumptions B.1 and B.9, we know that $z_{1,\pi}$ is the unique equilibrium point.

We now show the existence of a Lyapunov function with respect to the aforementioned equilibrium point, $z_{1,\pi}$. In particular, we consider the function, L, defined by:

$$L(Z_1) = \frac{1}{2}(Z_1 - z_{1,\pi})^2.$$

To establish that L is a Lyapunov function, we must show that:

- 1. L is continuous,
- 2. $L(Z_1) = 0$ if $Z_1 = z_{1,\pi}$,
- 3. $L(Z_1) > 0$ if $Z_1 \neq z_{1,\pi}$, and
- 4. For any solution $\{x_t\}_{t\geq 0}$ of the associated ODE (B.34) and $0\leq s < t$, we have $L(x_t) < L(x_s)$ for all $x_s \neq z_{1,\pi}$.

It directly follows from the definition of L that the first three conditions are satisfied.

We now show that fourth condition is also satisfied:

Let $\{x_t\}_{t\geq 0}$ be a solution to the associated ODE (B.34), and let $0\leq s < t$. By the chain rule, we have that:

$$\frac{d}{dt}L(x_t) = \frac{\partial L}{\partial x_t} \cdot \frac{dx_t}{dt} = (x_t - z_{1,\pi}) \cdot g_1(x_t) = (x_t - z_{1,\pi}) \cdot (\mathbb{E}_{\pi}[\phi_1] - x_t) = (x_t - z_{1,\pi}) \cdot (z_{1,\pi} - x_t).$$
(B.35)

We will now analyze the sign of the term $(x_t - z_{1,\pi}) \cdot (z_{1,\pi} - x_t)$ when $x_t \neq z_{1,\pi}$. There are two cases to consider:

Case 1: $x_t > z_{1,\pi}$: We have that $(x_t - z_{1,\pi}) > 0$ and $(z_{1,\pi} - x_t) < 0$. Hence, we can conclude that $(x_t - z_{1,\pi}) \cdot (z_{1,\pi} - x_t) < 0$.

Case 2: $x_t < z_{1,\pi}$: We have that $(x_t - z_{1,\pi}) < 0$ and $(z_{1,\pi} - x_t) > 0$. Hence, we can conclude that $(x_t - z_{1,\pi}) \cdot (z_{1,\pi} - x_t) < 0$.

As such, in both cases we can conclude that $(x_t - z_{1,\pi}) \cdot (z_{1,\pi} - x_t) < 0$, which implies that $L(x_t) < L(x_s)$ for all $x_s \neq z_{1,\pi}$ and $0 \leq s < t$.

As such, we have now verified the four conditions, and can therefore conclude that L is a valid Lyapunov function. Consequently, we can conclude, under the multiple-timescales argument, that $Z_{1,n}$ has a globally asymptotically stable equilibrium, $z_{1,\pi}$.

Critically, we can leverage the above result to show, using the same techniques as above, that $Z_{2,n}$ has a globally asymptotically stable equilibrium, $z_{2,\pi}$, where V_n , \bar{R}_n , and $Z_{1,n}$ are considered to be nearly equilibrated, and all remaining subtasks are considered quasi-static. The process can then be repeated for $Z_{3,n}$ and so forth, thereby showing that the subtask update rules (B.33) each have a globally asymptotically stable equilibrium, $z_{i,\pi}$. This completes the proof.

Lemma B.3.
$$\sup_n(||V_n|| + ||Z_{1,n}||) + ||Z_{2,n}|| + \ldots + ||Z_{k,n}||) < \infty \ a.s.$$

Proof. It was shown in Theorem B.3 of Wan et al. (2021) that $\sup_n(||V_n||) < \infty$ a.s. Hence, we only need to show that $\sup_n(||Z_{i,n}||) < \infty \ \forall z_i \in \mathcal{Z}$ a.s. To this end, we can apply Theorem 7 in Section 3 of Borkar (2009). This theorem requires 4 assumptions for each $z_i \in \mathcal{Z}$:

- (A1) The function g_i is Lipschitz. That is, $||g_i(x) g_i(y)|| \le U_i ||x y||$ for some $0 < U_i < \infty$.
- (A2) The sequence $\{\alpha_{z_i,n}\}_{n=0}^{\infty}$ satisfies $\alpha_{z_i,n} > 0 \ \forall n \ge 0, \ \sum \alpha_{z_i,n} = \infty$, and $\sum \alpha_{z_i,n}^2 < \infty$.
- (A3) $\{M_n^{z_i}\}_{n=0}^{\infty}$ is a martingale difference sequence that is square-integrable.
- (A4) The functions $g_i(x)_d \doteq g_i(dx)/d$, $d \geq 1, x \in \mathbb{R}$, satisfy $g_i(x)_d \to g_i(x)_\infty$ as $d \to \infty$, uniformly on compacts for some $g_{i_\infty} \in C(\mathbb{R})$. Furthermore, the ODE $\dot{x}_t = g_i(x_t)_\infty$ has the origin as its unique globally asymptotically stable equilibrium.

Consider an arbitrary ith subtask. We note that Assumption (A1) is satisfied given that all operators in g_i are Lipschitz. Moreover, we note that Assumption B.3 satisfies Assumption (A2).

We now show that $\{M_n^{z_i}\}_{n=0}^\infty$ is indeed a martingale difference sequence that is square-integrable, such that $\mathbb{E}_B[M_{n+1}^{z_i}\mid\mathcal{F}_n]=0$ a.s., $n\geq 0$, and $\mathbb{E}_B[(M_{n+1}^{z_i})^2\mid\mathcal{F}_n]<\infty$ a.s., $n\geq 0$, where B is the behaviour policy (and π is the target policy). To this end, let $\mathcal{F}_n\doteq\sigma(Z_{i,u},M_u^{z_i},u\leq n), n\geq 0$ denote an increasing family of σ -fields. We have that, for any $s\in\mathcal{S}$:

$$\mathbb{E}_{B}[M_{n+1}^{z_{i}}(s) \mid \mathcal{F}_{n}] = \mathbb{E}_{B} \left[\rho_{n}(s) \left(\phi_{i,n}(s) - Z_{i,n} \right) - g_{i}(Z_{i,n})(s) \mid \mathcal{F}_{n} \right]$$

$$= \mathbb{E}_{\pi} \left[\phi_{i,n}(s) - Z_{i,n} \right] - g_{i}(Z_{i,n})(s)$$

$$= 0$$

Moreover, since all components of $\{M_n^{z_i}\}_{n=0}^{\infty}$ involve bounded quantities, it directly follows that $\mathbb{E}_B[||M_{n+1}^{z_i}||^2 \mid \mathcal{F}_n] \leq K$ for some finite constant K>0. Hence, Assumption (A3) is verified.

We now verify Assumption (A4). Under the multiple-timescales argument, we have that:

$$g_i(x)_{\infty} = \lim_{d \to \infty} g_i(x)_d = \lim_{d \to \infty} \frac{\mathbb{E}_{\pi} \left[\phi_i\right] - dx}{d} = 0 - x = -x.$$
 (B.36)

Clearly, $g_i(x)_{\infty}$ is continuous in every $x \in \mathbb{R}$. As such, we have that $g_i(x)_{\infty} \in C(\mathbb{R})$. Now consider the ODE $\dot{x} = g_i(x)_{\infty}$. This ODE has the origin as an equilibrium since $g_i(0)_{\infty} = 0$. Furthermore, given Lemma B.2, we can conclude that this equilibrium must be the unique globally asymptotically stable equilibrium, thereby satisfying Assumption (A4).

Assumptions (A1) - (A4) are hence verified, meaning that we can apply the results of Theorem 7 in Section 3 of Borkar (2009) to conclude that $\sup_n(||Z_{i,n}||) < \infty \ \forall z_i \in \mathcal{Z}$, almost surely, and hence, that $\sup_n(||V_n|| + ||Z_{1,n}||) + ||Z_{2,n}|| + \ldots + ||Z_{k,n}||) < \infty$, almost surely.

As such, we have now verified the 3 assumptions required by Theorem 2 in Section 6 of Borkar (2009), which means that we can apply the result of the theorem to conclude that $Z_{i,n} \to z_{i,\pi} \, \forall z_i \in \mathcal{Z}$, almost surely, as $n \to \infty$.

This completes the proof of Theorem B.1.1.

B.2 Convergence Proof for the Tabular RED Q-learning Algorithm

In this section, we present the proof for the convergence of the value function, average-reward, and subtask estimates of the RED Q-learning algorithm. Similar to what was done in Wan et al. (2021), we will begin by considering a general algorithm, called *General RED Q*. We will first define General RED Q, then show how the RED Q-learning algorithm is a special case of this algorithm. We will then provide the necessary assumptions, state the convergence theorem of General RED Q, and then provide a proof for the theorem, where we show that the value function, average-reward, and subtask estimates converge, thereby showing that the RED Q-learning algorithm converges. We begin by introducing the General RED Q algorithm:

Consider an MDP $\mathcal{M} \doteq \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, p \rangle$. Given a state $s \in \mathcal{S}$, action $a \in \mathcal{A}$, and discrete step $n \geq 0$, let $R_n(s,a) \in \mathcal{R}$ denote a sample of the resulting reward, and let $S'_n(s,a) \sim p(\cdot, \cdot \mid s,a)$ denote a sample of the resulting state. Let $\{Y_n\}$ be a set-valued process taking values in the set of nonempty subsets of $\mathcal{S} \times \mathcal{A}$, such that: $Y_n = \{(s,a) : (s,a) \text{ component of the } |\mathcal{S} \times \mathcal{A}| \text{-sized table of state-action value estimates, } Q$, that was updated at step n. Let $\nu(n,s,a) \doteq \sum_{j=0}^n I\{(s,a) \in Y_j\}$, where I is the indicator function, such that $\nu(n,s,a)$ represents the number of times that the (s,a) component of Q was updated up until step n.

Now, let f be a valid subtask function (see Definition 5.1), such that $\tilde{R}_n(s,a) \doteq f(R_n(s,a),Z_{1,n},Z_{2,n},\ldots,Z_{k,n})$ for k subtasks $\in \mathcal{Z}$, where $\tilde{R}_n(s,a)$ is the extended reward, \mathcal{Z} is the set of subtasks, and $Z_{i,n}$ denotes the estimate of subtask $z_i \in \mathcal{Z}$ at step n. Consider an MDP with the extended reward: $\tilde{\mathcal{M}} \doteq \langle \mathcal{S}, \mathcal{A}, \tilde{\mathcal{R}}, \tilde{p} \rangle$, such that $\tilde{R}_n(s,a) \in \tilde{\mathcal{R}}$. The update rules of General RED Q for this MDP are as follows, $\forall n \geq 0$:

$$Q_{n+1}(s,a) \doteq Q_n(s,a) + \alpha_{\nu(n,s,a)} \delta_n(s,a) I\{(s,a) \in Y_n\}, \quad \forall s \in \mathcal{S}, a \in \mathcal{A},$$
 (B.37)

$$\bar{R}_{n+1} \doteq \bar{R}_n + \sum_{s,a} \alpha_{\bar{r},\nu(n,s,a)} \delta_n(s,a) I\{(s,a) \in Y_n\},$$
(B.38)

$$Z_{i,n+1} \doteq Z_{i,n} + \sum_{s,a} \alpha_{z_i,\nu(n,s,a)} \beta_{i,n}(s,a) I\{(s,a) \in Y_n\}, \quad \forall z_i \in \mathcal{Z}$$
(B.39)

where,

$$\delta_n(s,a) \doteq \tilde{R}_n(s,a) - \bar{R}_n + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a)$$

$$= f(R_n(s,a), Z_{1,n}, Z_{2,n}, \dots, Z_{k,n}) - \bar{R}_n + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a),$$
(B.40)

and,

$$\beta_{i,n}(s,a) \doteq \phi_{i,n}(s,a) - Z_{i,n}, \quad \forall z_i \in \mathcal{Z}. \tag{B.41}$$

Here, \bar{R}_n denotes the estimate of the average-reward (see Equation (2)), $\delta_n(s,a)$ denotes the TD error, $\phi_{i,n}(s,a)$ denotes the (potentially-piecewise) subtask target, as defined in Section 5.1, and $\alpha_{\nu(n,s,a)}$, $\alpha_{\bar{r},\nu(n,s,a)}$, and $\alpha_{z_i,\nu(n,s,a)}$ denote the step sizes at time step n for state-action pair (s,a).

We now show that the RED Q-learning algorithm is a special case of the General RED Q algorithm. Consider a sequence of experience from our MDP, $\tilde{\mathcal{M}}$: $S_t, A_t, \tilde{R}_{t+1}, S_{t+1}, \ldots$. Now recall the set-valued process $\{Y_n\}$. If we let n = time step t, we have:

$$Y_t(s, a) = \begin{cases} 1, s = S_t \text{ and } a = A_t, \\ 0, \text{ otherwise,} \end{cases}$$

as well as $S_n'(S_t,A_t)=S_{t+1},$ $R_n(S_t,A_t)=R_{t+1},$ and $\tilde{R}_n(S_t,A_t)=\tilde{R}_{t+1}.$

Hence, update rules (B.37), (B.38), (B.39), (B.40), and (B.41) become:

$$Q_{t+1}(S_t, A_t) \doteq Q_t(S_t, A_t) + \alpha_{\nu(t, S_t, A_t)} \delta_t; \quad Q_{t+1}(s, a) \doteq Q_t(s, a), \forall s \neq S_t, a \neq A_t, \quad (B.42)$$

$$\bar{R}_{t+1} \doteq \bar{R}_t + \alpha_{\bar{r},\nu(t,S_t,A_t)} \delta_t, \tag{B.43}$$

$$Z_{i,t+1} \doteq Z_{i,t} + \alpha_{z_i,\nu(t,S_t,A_t)}\beta_{i,t}, \quad \forall z_i \in \mathcal{Z},$$
(B.44)

$$\delta_{t} \doteq \tilde{R}_{t+1} - \bar{R}_{t} + \max_{a'} Q_{t}(S_{t+1}, a') - Q_{t}(S_{t}, A_{t}),$$

$$= f(R_{t+1}, Z_{1,t}, Z_{2,t}, \dots, Z_{k,t}) - \bar{R}_{t} + \max_{a'} Q_{t}(S_{t+1}, a') - Q_{t}(S_{t}, A_{t}),$$
(B.45)

$$\beta_{i,t} \doteq \phi_{i,t} - Z_{i,t}, \quad \forall z_i \in \mathcal{Z},$$
 (B.46)

which are RED Q-learning's update rules with $\alpha_{\nu(t,S_t,A_t)}$, $\alpha_{\bar{r},\nu(t,S_t,A_t)}$, and $\alpha_{z_i,\nu(t,S_t,A_t)}$ denoting the step sizes at time t.

We now specify the assumptions specific to General RED Q that are needed to ensure convergence. We refer the reader to Wan et al. (2021) for an in-depth discussion on these assumptions:

Assumption B.11 (Communicating Assumption). The MDP has a single communicating class. That is, each state in the MDP is accessible from every other state under some deterministic stationary policy.

Assumption B.12 (State-Action Value Function Uniqueness). *There exists a unique solution of q only up to a constant in the Bellman equation* (4).

Assumption B.13 (Asynchronous Step Size Assumption 3). There exists $\Delta > 0$ such that

$$\liminf_{n \to \infty} \frac{\nu(n, s, a)}{n+1} \ge \Delta,$$

a.s., for all $s \in \mathcal{S}$, $a \in \mathcal{A}$.

Furthermore, for all x > 0, and

$$N(n,x) = \min \left\{ m > n : \sum_{i=n+1}^{m} \alpha_i \ge x \right\},\,$$

the limit

$$\lim_{n \rightarrow \infty} \frac{\sum_{i=\nu(n,s,a)}^{\nu(N(n,x),s,a)} \alpha_i}{\sum_{i=\nu(n,s',a')}^{\nu(N(n,x),s',a')} \alpha_i}$$

exists a.s. for all s, s', a, a'.

Having stated the necessary assumptions, we next point out that it is easy to verify that under Assumption B.11, the following system of equations:

$$q_{*}(s, a) = \sum_{s', \tilde{r}} \tilde{p}(s', \tilde{r} \mid s, a)(\tilde{r} - \bar{r}_{*} + \max_{a'} q_{*}(s, a)), \quad \forall s \in \mathcal{S}, a \in \mathcal{A}$$

$$= \sum_{s', r} p(s', r \mid s, a)(f(r, z_{1_{*}}, z_{2_{*}}, \dots, z_{k_{*}}) - \bar{r}_{*} + \max_{a'} q_{*}(s, a)), \quad \forall s \in \mathcal{S}, a \in \mathcal{A},$$
(B.47)

and,

$$\bar{r}_* - \bar{R}_0 = \eta_r \left(\sum q_* - \sum Q_0 \right),$$
 (B.48)

$$z_{i_*} - Z_{i,0} = \eta_i \left(\sum q_* - \sum Q_0 \right), \quad \forall z_i \in \mathcal{Z}, \tag{B.49}$$

has a unique solution for q_* , where \bar{r}_* denotes the optimal average-reward, and z_{i_*} denotes the corresponding optimal subtask value for subtask $z_i \in \mathcal{Z}$. Denote this unique solution solution of q_* as q_{∞} .

We are now ready to state the convergence theorem:

Theorem B.2.1 (Convergence of General RED Q). If Assumptions B.3, B.4, B.6, B.7, B.8, B.9, B.10, B.11, B.12, and B.13 hold, then the General RED Q algorithm (Equations B.37–B.41) converges a.s. \bar{R}_n to \bar{r}_* , $Z_{i,n}$ to z_{i_*} $\forall z_i \in \mathcal{Z}$, \bar{r}_{π_t} to \bar{r}_* , z_{i,π_t} to z_{i_*} $\forall z_i \in \mathcal{Z}$, and Q_n to q_∞ , where π_t is any greedy policy with respect to Q_t , and z_{i,π_t} denotes the subtask value induced by following policy π_t .

We prove this theorem in Sections B.2.1 and B.2.2. To do so, we first show that General RED Q is of the same form as *General Differential Q* from Wan et al. (2021), thereby allowing us to apply their convergence results for the value function and average-reward estimates of General Differential Q to General RED Q. We then build upon these results, using similar techniques as Wan et al. (2021), to show that the subtask estimates converge as well.

B.2.1 Proof of Theorem **B.2.1** (for *Linear* Subtask Functions)

We first provide the proof for *linear* subtask functions, where the reward-extended TD error can be expressed as a constant, subtask-specific fraction of the regular TD error, such that $\beta_{i,n}(s,a) = (-1/b_i)\delta_n(s,a)$. We consider the *piecewise linear* case in Section B.2.2.

Convergence of the value function and average-reward estimates:

Consider the increment to \bar{R}_n at each step. Given Assumption B.6, we can see from Equation (B.38) that the increment is η_r times the increment to Q_n . As such, as was done in Wan et al. (2021), we can write the cumulative increment as follows:

$$\bar{R}_n - \bar{R}_0 = \eta_r \sum_{j=0}^{n-1} \sum_{s,a} \alpha_{\nu(j,s,a)} \delta_j(s,a) I\{(s,a) \in Y_j\}$$

$$=\eta_r\left(\sum Q_n-\sum Q_0\right)$$

$$\implies \bar{R}_n = \eta_r \sum Q_n - \eta_r \sum Q_0 + \bar{R}_0 = \eta_r \sum Q_n - c_r, \tag{B.50}$$

where
$$c_r \doteq \eta_r \sum Q_0 - \bar{R}_0$$
. (B.51)

Similarly, consider the increment to $Z_{i,n}$ (for an arbitrary subtask $z_i \in \mathcal{Z}$) at each step. As per Remark 5.1, and given Assumption B.6, we can write the increment in Equation (B.39) as some constant, subtask-specific fraction of the increment to Q_n . Consequently, we can write the cumulative increment as follows:

$$Z_{i,n} - Z_{i,0} = \eta_{z_i} \sum_{j=0}^{n-1} \sum_{s,a} \alpha_{\nu(j,s,a)} \beta_{i,j}(s,a) I\{(s,a) \in Y_j\}$$

$$= \eta_{z_i} \sum_{j=0}^{n-1} \sum_{s,a} \alpha_{\nu(j,s,a)} (-1/b_i) \delta_j(s,a) I\{(s,a) \in Y_j\}$$

$$= \eta_i \left(\sum_{j=0}^{n} Q_n - \sum_{j=0}^{n} Q_j \right)$$

$$\implies Z_{i,n} = \eta_i \sum_{j=0}^{n} Q_n - \eta_j \sum_{j=0}^{n} Q_j + Z_{i,0} = \eta_j \sum_{j=0}^{n} Q_j - c_i,$$
(B.52)

where,

$$c_i \doteq \eta_i \sum Q_0 - Z_{i,0}, \text{ and}$$
 (B.53)

$$\eta_i \doteq (-1/b_i)\eta_{z_i}.\tag{B.54}$$

Now consider the subtask function, f. At any given time step, the subtask function can be written as: $f_n = \tilde{R}_n(s,a) = b_r R_n(s,a) + b_0 + b_1 Z_{1,n} + \ldots + b_k Z_{k,n}$, where $b_r, b_0 \in \mathbb{R}$ and $b_i \in \mathbb{R} \setminus \{0\}$. Given Equation (B.52), we can write the subtask function as follows:

$$f_n = b_r R_n(s, a) + b_0 + b_1 (\eta_1 \sum_{i=1}^n Q_n - c_1) + \dots + b_k (\eta_k \sum_{i=1}^n Q_n - c_k)$$

$$= b_r R_n(s, a) + \eta_f \sum_{i=1}^n Q_n - c_f,$$
(B.55)

where, $\eta_f = \sum_{j=1}^k b_j \eta_j$ and $c_f = \sum_{j=1}^k b_j c_j - b_0$.

As such, we can substitute \bar{R}_n and $Z_{i,n}$ $\forall z_i \in \mathcal{Z}$ in (B.37) with (B.50) and (B.55), respectively, $\forall s \in \mathcal{S}, a \in \mathcal{A}$, which yields:

$$Q_{n+1}(s,a) = Q_n(s,a) + \dots$$

$$\alpha_{\nu(n,s,a)} \left(b_r R_n(s,a) + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a) - \eta_r \sum_{a'} Q_n + c_r + \eta_f \sum_{a'} Q_n - c_f \right) I\{(s,a) \in Y_n\}$$

$$Q_{n+1}(s,a) = Q_n(s,a) + \dots$$

$$\alpha_{\nu(n,s,a)} \left(b_r R_n(s,a) + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a) - \eta_T \sum_{a'} Q_n + c_T \right) I\{(s,a) \in Y_n\}$$

$$Q_{n+1}(s,a) = Q_n(s,a) + \dots$$

$$\alpha_{\nu(n,s,a)} \left(\widehat{R}_n(s,a) + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a) - \eta_T \sum_{a'} Q_n \right) I\{(s,a) \in Y_n\},$$
(B.56)

where $\eta_{\scriptscriptstyle T}=\eta_{\scriptscriptstyle r}-\eta_{\scriptscriptstyle f},$ $c_{\scriptscriptstyle T}=c_r-c_f,$ and $\widehat{R}_n(s,a)\doteq b_rR_n(s,a)+c_{\scriptscriptstyle T}.$

Equation (B.56) is now in the same form as Equation (B.14) (i.e., General Differential Q) from Wan et al. (2021), who showed that the equation converges a.s. Q_n to q_∞ as $n \to \infty$. Moreover, from this result, Wan et al. (2021) showed that \bar{R}_n converges a.s. to \bar{r}_* as $n \to \infty$, and that \bar{r}_{π_t} converges a.s. to \bar{r}_* , where π_t is a greedy policy with respect to Q_t . Given that General RED Q adheres to all the assumptions listed for General Differential Q in Wan et al. (2021), these convergence results apply to General RED Q.

Convergence of the subtask estimates:

Consider Equation (B.56). We can rewrite this equation, $\forall s \in \mathcal{S}, a \in \mathcal{A}$, as follows:

$$Q_{n+1}(s,a) = Q_n(s,a) + \dots$$

$$\alpha_{\nu(n,s,a)} \left(b_r R_n(s,a) + \eta_f \sum_{n} Q_n + c_{r} - \eta_r \sum_{n} Q_n + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a) \right) I\{(s,a) \in Y_n\}$$

$$Q_{n+1}(s,a) = Q_n(s,a) + \dots$$

$$\alpha_{\nu(n,s,a)} \left(\hat{R}_n(s,a) + c_T - \eta_r \sum_{a'} Q_n + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a) \right) I\{(s,a) \in Y_n\},$$
(B.57)

where,

$$\hat{R}_n(s,a) \doteq b_r R_n(s,a) + \eta_f \sum Q_n \tag{B.58}$$

$$= b_r R_n(s, a) + b_1(\eta_1 \sum Q_n) + \dots + b_k(\eta_k \sum Q_n)$$
 (B.59)

$$= b_r R_n(s, a) + b_1 \hat{Z}_{1,n} + \ldots + b_k \hat{Z}_{k,n}.$$
 (B.60)

Now consider an MDP, $\hat{\mathcal{M}}$, which has rewards, $\hat{\mathcal{R}}$, as defined in Equation (B.58), has the same state and action spaces as $\hat{\mathcal{M}}$, and has the transition probabilities defined as:

$$\hat{p}(s', \hat{r} \mid s, a) \doteq p(s', r \mid s, a) \tag{B.61}$$

$$= \tilde{p}(s', \tilde{r} \mid s, a) \quad \text{(by Definition 5.1)}, \tag{B.62}$$

such that $\hat{\mathcal{M}} \doteq \langle \mathcal{S}, \mathcal{A}, \hat{\mathcal{R}}, \hat{p} \rangle$. It is easy to check that the communicating assumption holds for the MDP, $\hat{\mathcal{M}}$. Moreover, given Equation (B.57) and Assumptions B.7 and B.8, the optimal average-reward for the MDP, $\hat{\mathcal{M}}$, \hat{r}_* , can be written as follows:

$$\hat{\bar{r}}_* = \bar{r}_* - c_{\tau}. \tag{B.63}$$

Now, because

$$\begin{split} q_{\infty}(s,a) &= \sum_{s',\tilde{r}} \tilde{p}(s',\tilde{r}\mid s,a) (\tilde{r} - \bar{r}_* + \max_{a'} q_{\infty}(s',a')) \quad \text{(from (B.47))} \\ &= \sum_{s',\tilde{r}} \tilde{p}(s',\tilde{r}\mid s,a) (\tilde{r} - (\hat{\bar{r}}_* + c_T) + \max_{a'} q_{\infty}(s',a')) \quad \text{(from (B.63))} \\ &= \sum_{s',\tilde{r}} \tilde{p}(s',\tilde{r}\mid s,a) (\tilde{r} - c_T - \hat{\bar{r}}_* + \max_{a'} q_{\infty}(s',a')) \\ &= \sum_{s',\tilde{r}} \tilde{p}(s',\tilde{r}\mid s,a) (\hat{r} - \hat{\bar{r}}_* + \max_{a'} q_{\infty}(s',a')) \quad \text{(from (B.57))} \\ &= \sum_{s',\tilde{r}} \hat{p}(s',\tilde{r}\mid s,a) (\hat{r} - \hat{\bar{r}}_* + \max_{a'} q_{\infty}(s',a')) \quad \text{(from (B.62))}, \end{split}$$

we can see that q_{∞} is a solution of not just the state-action value Bellman optimality equation for the MDP, $\tilde{\mathcal{M}}$, but also the state-action value Bellman optimality equation for the MDP, $\hat{\mathcal{M}}$.

Next, consider an arbitrary *i*th subtask. As per Equations (B.59) and (B.60), we can write the optimal subtask value for the MDP, $\hat{\mathcal{M}}$, \hat{z}_{i_*} , as follows:

$$\hat{z}_{i_*} = z_{i_*} + c_i. \tag{B.64}$$

We can then combine Equations (B.49), (B.52), and (B.64), which yields:

$$\hat{z}_{i_*} = \eta_i \sum q_{\infty}. \tag{B.65}$$

Next, we can combine Equation (B.52) with the result from Wan et al. (2021) which shows that $Q_n \to q_{\infty}$, which yields:

$$Z_{i,n} \to \eta_i \sum q_\infty - c_i.$$
 (B.66)

Moreover, because $\eta_i \sum q_\infty = \hat{z}_{i_*}$ (Equation (B.65)), we have:

$$Z_{i,n} \to \hat{z}_{i_*} - c_i. \tag{B.67}$$

Finally, because $\hat{z}_{i_*} = z_{i_*} + c_i$ (Equation (B.64)), we have:

$$Z_{i,n} \to z_{i_*}$$
 a.s. as $n \to \infty$. (B.68)

We conclude by considering $z_{i,\pi_t} \ \forall z_i \in \mathcal{Z}$, where π_t is a greedy policy with respect to Q_t . Given that $Q_t \to q_\infty$ and $\bar{r}_{\pi_t} \to \bar{r}_*$, almost surely, it directly follows from Definition 5.1 that $z_{i,\pi_t} \to z_{i_*} \ \forall z_i \in \mathcal{Z}$, almost surely.

B.2.2 Proof of Theorem **B.2.1** (for *Piecewise Linear* Subtask Functions)

We now provide the proof for *piecewise linear* subtask functions, where the reward-extended TD error can be expressed as follows:

$$\beta_{i,n}(s,a) = \begin{cases} (-1/b_i^1) \left(\tilde{R}_{1,n}(s,a) - \bar{R}_n - \delta_n(s,a) \right), & r_0 \le R_n(s,a) < r_1 \\ \vdots & \vdots \\ (-1/b_i^m) \left(\tilde{R}_{m,n}(s,a) - \bar{R}_n - \delta_n(s,a) \right), & r_{m-1} \le R_n(s,a) \le r_m \end{cases},$$

where $r_u \in \mathcal{R} \ \forall u = 0, 1, \dots, m$, and $r_0 \leq r_1 \leq \dots \leq r_m$, such that r_0, r_m represent the lower and upper bounds of the observed per-step reward, $R_n(s, a)$, respectively. Our general strategy in this case is to use a *multiple-timescales* argument, such that we leverage Theorem 2 in Section 6 of Borkar (2009), along with the results from Theorems B.1 and B.2 of Wan et al. (2021).

To begin, let us consider Assumption B.10, which enables the formulation of a multiple-timescales argument. In particular, the $\alpha_{z_1,n}/\alpha_n \to 0$, $\alpha_{z_1,n}/\alpha_{\bar{r},n} \to 0$, and $\{\alpha_{z_i,n}/\alpha_{z_{i-1},n} \to 0\}_{i=2}^k$ conditions imply that the subtask step sizes, $\{\alpha_{z_i,n}\}_{n=0}^{\infty} \forall z_i \in \mathcal{Z}$, decrease to 0 at faster rates than the value function and average-reward step sizes, $\{\alpha_n\}_{n=0}^{\infty}$ and $\{\alpha_{\bar{r},n}\}_{n=0}^{\infty}$, respectively. This implies that the subtask updates move on slower timescales compared to the value function and average-reward updates. Hence, as argued in Section 6 of Borkar (2009), the (faster) value function and average-reward updates, (B.37) and (B.38), view the (slower) subtask updates, (B.39), as quasistatic, while the (slower) subtask updates view the (faster) value function and average-reward updates as nearly equilibrated (as we will show below, the results from Wan et al. (2021) imply the existence of such an equilibrium point). Similarly, the $\{\alpha_{z_i,n}/\alpha_{z_{i-1},n} \to 0\}_{i=2}^k$ condition implies that each subtask update views the other subtask updates as either quasi-static or nearly-equilibrated.

As such, having established the multiple-timescales argument, we now proceed to show the convergence of the value function, average-reward, and subtask estimates:

Convergence of the value function and average-reward estimates:

Given the multiple-timescales argument, such that the subtask estimates are viewed as quasi-static (i.e., constant), Equation (B.37) can be viewed as being of the same form as Equation (B.4) (i.e., General Differential Q) from Wan et al. (2021), who showed (via Theorem B.2) that the equation converges a.s. Q_n to q_∞ as $n\to\infty$. Moreover, from this result, Wan et al. (2021) showed that \bar{R}_n converges a.s. to \bar{r}_* as $n\to\infty$, and that \bar{r}_{π_t} converges a.s. to \bar{r}_* , where π_t is a greedy policy with respect to Q_t . Given that General RED Q adheres to all the assumptions listed for General Differential Q in Wan et al. (2021), these convergence results apply to General RED Q.

Convergence of the subtask estimates:

Let us consider the asynchronous subtask updates (B.39). Each update in (B.39) is of the same form as Equation 7.1.2 of Borkar (2009). Accordingly, to show the convergence of the subtask estimates, we can apply the result in Section 7.4 of Borkar (2009), which shows the convergence of asynchronous updates that are of the same form as Equation 7.1.2. To apply this result, given Assumptions B.4 and B.13, we only need to show the convergence of the *synchronous* version of the subtask updates:

$$Z_{i,n+1} = Z_{i,n} + \alpha_{z_i,n} \left(g_i(Z_{i,n}) + M_{n+1}^{z_i} \right) \ \forall z_i \in \mathcal{Z},$$
(B.69)

where,

$$g_{i}(Z_{i,n})(s,a) \doteq \sum_{s',r} p(s',r \mid s,a)\phi_{i,n}(s,a) - Z_{i,n},$$

$$\begin{pmatrix} -\frac{1}{b_{i}^{1}} \left(b_{r}^{1}R_{n}(s,a) + \dots + b_{i-1}^{1}Z_{i-1,n} + b_{i+1}^{1}Z_{i+1,n} + \dots + b_{k}^{1}Z_{k,n} - \bar{R}_{n} - \delta_{n}(s,a) \right), \dots \\ \dots, r_{0} \leq R_{n}(s,a) < r_{1}, \\ \vdots \\ -\frac{1}{b_{i}^{m}} \left(b_{r}^{m}R_{n}(s,a) + \dots + b_{i-1}^{m}Z_{i-1,n} + b_{i+1}^{m}Z_{i+1,n} + \dots + b_{k}^{m}Z_{k,n} - \bar{R}_{n} - \delta_{n}(s,a) \right), \dots \\ \dots, r_{m-1} \leq R_{n}(s,a) \leq r_{m}, \\ M_{n+1}^{z_{i}}(s,a) \doteq (\phi_{i,n}(s,a) - Z_{i,n}) - g_{i}(Z_{i,n})(s,a). \end{pmatrix}$$

To show the convergence of the synchronous update (B.69) under the multiple-timescales argument, we can apply the result of Theorem 2 in Section 6 of Borkar (2009) to show that $Z_{i,n} \to z_{i_*} \forall z_i \in \mathcal{Z}$ a.s. as $n \to \infty$. This theorem requires that 3 assumptions be satisfied. As such, we will now show, via Lemmas B.4 - B.6, that these 3 assumptions are indeed satisfied:

Lemma B.4. The value function update, $Q_{n+1} = Q_n + \alpha_n(h(Q_n) + M_{n+1})$, where

$$\begin{split} h(Q_n)(s,a) &\doteq \sum_{s',\tilde{r}} \tilde{p}(s',\tilde{r}\mid s,a)(\tilde{r}-\bar{R}_n + \max_{a'} Q_n(s',a') - Q_n(s,a)), \\ &= \sum_{s',r} p(s',r\mid s,a)(f(r,Z_{1,n},Z_{2,n},\ldots,Z_{k,n}) - \psi(Q_n) + \max_{a'} Q_n(s',a') - Q_n(s,a)), \\ M_{n+1}(s,a) &\doteq \left(\tilde{R}_n(s,a) - \psi(Q_n) + \max_{a'} Q_n(S'_n(s,a),a') - Q_n(s,a)\right) - h(Q_n)(s,a), \\ \psi(Q_n) &= \bar{R}_n \text{ is a 'reference function' as defined in Wan et al. (2021), and} \\ Z_{1,n}, Z_{2,n}, \ldots, Z_{k,n} \text{ are quasi-static under the multiple-timescales argument,} \end{split}$$

has a globally asymptotically stable equilibrium, $q_{\infty}(Z_{1,n}, Z_{2,n}, \dots, Z_{k,n})$, where q_{∞} is a Lipschitz map.

Proof. This was shown in Theorem B.2 of Wan et al. (2021).

Lemma B.5. The subtask update rules (B.69) each have a globally asymptotically stable equilibrium, z_{i_*} .

Proof. Let us begin by considering the 'fastest' subtask, $Z_{1,n}$, under the multiple-timescales argument, such that all other subtasks, $Z_{2,n},\ldots,Z_{k,n}$, operate on slower timescales and can be considered quasi-static. Applying the results of Theorems B.1 and B.2 of Wan et al. (2021) under the multiple-timescales argument, we have that: $h(Q_n) \to 0$ and $\bar{R}_n \to \bar{r}_*$ as $Q_n \to q_\infty$. Importantly, as argued in Section 5, $h(Q_n) \to 0$ implies that $\delta_n(s,a) \to \lambda_j \in \mathbb{R}$ for each jth piecewise segment in $\phi_{i,n}(s,a)$. As such, we can interpret $Z_{1,n}$ as being the only 'moving' variable in $g_1(Z_{1,n})$, such that all other parameters in the update are static, quasi-static, or nearly-equilibrated.

Let us now consider the ODE associated with $g_1(Z_{1,n})$,

$$\dot{x}_t = g_1(x_t). \tag{B.70}$$

To show that the update rule for $Z_{1,n}$ has a globally asymptotically stable equilibrium, z_{1_*} , it suffices to show that there exists a *Lyapunov function* for the associated ODE (B.70).

To this end, we first note, given the discussion in Section 5, that z_{1_*} is an equilibrium point for the update rule associated with $Z_{1,n}$. Moreover, given Assumptions B.9, B.11, and B.12, we know that z_{1_*} is the unique equilibrium point.

We now show the existence of a Lyapunov function with respect to the aforementioned equilibrium point, z_{1_*} . In particular, we consider the function, L, defined by:

$$L(Z_1) = \frac{1}{2}(Z_1 - z_{1_*})^2.$$

To establish that L is a Lyapunov function, we must show that:

- 1. L is continuous,
- 2. $L(Z_1) = 0$ if $Z_1 = z_1$,
- 3. $L(Z_1) > 0$ if $Z_1 \neq z_{1}$, and
- 4. For any solution $\{x_t\}_{t\geq 0}$ of the associated ODE (B.70) and $0\leq s< t$, we have $L(x_t)< L(x_s)$ for all $x_s\neq z_{1_*}$.

It directly follows from the definition of L that the first three conditions are satisfied.

We now show that fourth condition is also satisfied:

Let $\{x_t\}_{t\geq 0}$ be a solution to the associated ODE (B.70), and let $0\leq s < t$. By the chain rule, we have that:

$$\frac{d}{dt}L(x_t) = \frac{\partial L}{\partial x_t} \cdot \frac{dx_t}{dt} = (x_t - z_{1_*}) \cdot g_1(x_t) = (x_t - z_{1_*}) \cdot (\mathbb{E}_{\pi}[\phi_1] - x_t) = (x_t - z_{1_*}) \cdot (z_{1_*} - x_t). \tag{B.71}$$

We will now analyze the sign of the term $(x_t - z_{1_*}) \cdot (z_{1_*} - x_t)$ when $x_t \neq z_{1_*}$. There are two cases to consider:

Case 1: $x_t > z_{1_*}$: We have that $(x_t - z_{1_*}) > 0$ and $(z_{1_*} - x_t) < 0$. Hence, we can conclude that $(x_t - z_{1_*}) \cdot (z_{1_*} - x_t) < 0$.

Case 2: $x_t < z_{1_*}$: We have that $(x_t - z_{1_*}) < 0$ and $(z_{1_*} - x_t) > 0$. Hence, we can conclude that $(x_t - z_{1_*}) \cdot (z_{1_*} - x_t) < 0$.

As such, in both cases we can conclude that $(x_t - z_{1_*}) \cdot (z_{1_*} - x_t) < 0$, which implies that $L(x_t) < L(x_s)$ for all $x_s \neq z_{1_*}$ and $0 \leq s < t$.

As such, we have now verified the four conditions, and can therefore conclude that L is a valid Lyapunov function. Consequently, we can conclude, under the multiple-timescales argument, that $Z_{1,n}$ has a globally asymptotically stable equilibrium, $z_{1,n}$.

Critically, we can leverage the above result to show, using the same techniques as above, that $Z_{2,n}$ has a globally asymptotically stable equilibrium, z_{2*} , where Q_n , \bar{R}_n , and $Z_{1,n}$ are considered to be nearly equilibrated, and all remaining subtasks are considered quasi-static. The process can then be repeated for $Z_{3,n}$ and so forth, thereby showing that the subtask update rules (B.69) each have a globally asymptotically stable equilibrium, z_{i*} . This completes the proof.

Lemma B.6.
$$\sup_n(||Q_n|| + ||Z_{1,n}||) + ||Z_{2,n}|| + \ldots + ||Z_{k,n}||) < \infty$$
 a.s.

Proof. It was shown in Theorem B.2 of Wan et al. (2021) that $\sup_n(||Q_n||) < \infty$ a.s. Hence, we only need to show that $\sup_n(||Z_{i,n}||) < \infty \ \forall z_i \in \mathcal{Z}$ a.s. To this end, we can apply Theorem 7 in Section 3 of Borkar (2009). This theorem requires 4 assumptions for each $z_i \in \mathcal{Z}$:

- (A1) The function g_i is Lipschitz. That is, $||g_i(x) g_i(y)|| \le U_i ||x y||$ for some $0 < U_i < \infty$.
- (A2) The sequence $\{\alpha_{z_i,n}\}_{n=0}^{\infty}$ satisfies $\alpha_{z_i,n} > 0 \ \forall n \ge 0, \ \sum \alpha_{z_i,n} = \infty, \ \text{and} \ \sum \alpha_{z_i,n}^2 < \infty.$
- (A3) $\{M_n^{z_i}\}_{n=0}^{\infty}$ is a martingale difference sequence that is square-integrable.
- (A4) The functions $g_i(x)_d \doteq g_i(dx)/d$, $d \geq 1, x \in \mathbb{R}$, satisfy $g_i(x)_d \to g_i(x)_\infty$ as $d \to \infty$, uniformly on compacts for some $g_{i_\infty} \in C(\mathbb{R})$. Furthermore, the ODE $\dot{x}_t = g_i(x_t)_\infty$ has the origin as its unique globally asymptotically stable equilibrium.

Consider an arbitrary ith subtask. We note that Assumption (A1) is satisfied given that all operators in g_i are Lipschitz. Moreover, we note that Assumption B.3 satisfies Assumption (A2).

We now show that $\{M_n^{z_i}\}_{n=0}^\infty$ is indeed a martingale difference sequence that is square-integrable, such that $\mathbb{E}_{\pi}[M_{n+1}^{z_i} \mid \mathcal{F}_n] = 0$ a.s., $n \geq 0$, and $\mathbb{E}_{\pi}[(M_{n+1}^{z_i})^2 \mid \mathcal{F}_n] < \infty$ a.s., $n \geq 0$. To this end, let $\mathcal{F}_n \doteq \sigma(Z_{i,u}, M_u^{z_i}, u \leq n), n \geq 0$ denote an increasing family of σ -fields. We have that, for any $s \in \mathcal{S}$ and $a \in \mathcal{A}$: $\mathbb{E}_{\pi}[M_{n+1}^{z_i}(s,a) \mid \mathcal{F}_n] = \mathbb{E}_{\pi}\left[\phi_{i,n}(s,a) - Z_{i,n}\right] - g_i(Z_{i,n})(s,a) = 0$. Moreover, since all components of $\{M_n^{z_i}\}_{n=0}^\infty$ involve bounded quantities, it directly follows that $\mathbb{E}[||M_{n+1}^{z_i}||^2 \mid \mathcal{F}_n] \leq K$ for some finite constant K > 0. Hence, Assumption (A3) is verified.

We now verify Assumption (A4). Under the multiple-timescales argument, we have that:

$$g_i(x)_{\infty} = \lim_{d \to \infty} g_i(x)_d = \lim_{d \to \infty} \frac{\mathbb{E}_{\pi} \left[\phi_i\right] - dx}{d} = 0 - x = -x.$$
 (B.72)

Clearly, $g_i(x)_{\infty}$ is continuous in every $x \in \mathbb{R}$. As such, we have that $g_i(x)_{\infty} \in C(\mathbb{R})$. Now consider the ODE $\dot{x} = g_i(x)_{\infty}$. This ODE has the origin as an equilibrium since $g_i(0)_{\infty} = 0$. Furthermore, given Lemma B.5, we can conclude that this equilibrium must be the unique globally asymptotically stable equilibrium, thereby satisfying Assumption (A4).

Assumptions (A1) - (A4) are hence verified, meaning that we can apply the results of Theorem 7 in Section 3 of Borkar (2009) to conclude that $\sup_n(||Z_{i,n}||) < \infty \ \forall z_i \in \mathcal{Z}$, almost surely, and hence, that $\sup_n(||Q_n|| + ||Z_{1,n}||) + ||Z_{2,n}|| + \ldots + ||Z_{k,n}||) < \infty$, almost surely. \square

As such, we have now verified the 3 assumptions required by Theorem 2 in Section 6 of Borkar (2009), which means that we can apply the result of the theorem to conclude that $Z_{i,n} \to z_{i_*} \, \forall z_i \in \mathcal{Z}$ a.s. as $n \to \infty$.

Finally, as was done in the proof for linear subtask functions, we conclude the proof by considering $z_{i,\pi_t} \ \forall z_i \in \mathcal{Z}$, where π_t is a greedy policy with respect to Q_t . Given that $Q_t \to q_\infty$ and $\bar{r}_{\pi_t} \to \bar{r}_*$, almost surely, it directly follows from Definition 5.1 that $z_{i,\pi_t} \to z_{i_*} \ \forall z_i \in \mathcal{Z}$, almost surely.

This completes the proof of Theorem B.2.1.

C Leveraging the RED RL Framework for CVaR Optimization

This appendix contains details regarding the adaptation of the RED RL framework for CVaR optimization. We first derive an appropriate subtask function, then use it to adapt the RED RL algorithms (see Appendix A) for CVaR optimization. In doing so, we arrive at the *RED CVaR algorithms*, which are presented in full at the end of this appendix. These RED CVaR algorithms allow us to optimize CVaR (and VaR) without the use of an augmented state-space or an explicit bi-level optimization. We also provide a convergence proof for the tabular RED CVaR Q-learning algorithm, which shows that the VaR and CVaR estimates converge to the optimal long-run VaR and CVaR, respectively.

C.1 A Subtask-Driven Approach for CVaR Optimization

In this section, we use the RED RL framework to derive a subtask-driven approach for CVaR optimization that does not require an augmented state-space or an explicit bi-level optimization. To begin, let us consider Equation (7), which is displayed below as Equation (C.1) for convenience:

$$CVaR_{\tau}(R_t) = \sup_{y \in \mathcal{R}} \mathbb{E}[y - \frac{1}{\tau}(y - R_t)^+]$$
 (C.1a)

$$= \mathbb{E}[\operatorname{VaR}_{\tau}(R_t) - \frac{1}{\tau}(\operatorname{VaR}_{\tau}(R_t) - R_t)^+], \tag{C.1b}$$

where $\tau \in (0,1)$ denotes the CVaR parameter, and R_t denotes the observed per-step reward.

We can see from Equation (C.1) that CVaR can be interpreted as an expectation (or average) of sorts, which suggests that it may be possible to leverage the average-reward MDP to optimize this expectation, by treating the reward CVaR as the average-reward, \bar{r}_{π} , that we want to optimize. However, this requires that we know the optimal value of the scalar, y, because the expectation in Equation (C.1b) only holds for this optimal value (which corresponds to the per-step reward VaR). Unfortunately, this optimal value is typically not known beforehand, so in order to optimize CVaR, we also need to optimize y.

Importantly, we can utilize RED RL framework to turn the optimization of y into a subtask, such that CVaR is the primary control objective (i.e., the \bar{r}_{π} that we want to optimize), and VaR (y in Equation (C.1)), is the (single) subtask. This is in contrast to existing MDP-based methods, which typically leverage Equation (C.1) when optimizing for CVaR by augmenting the state-space with a state that corresponds (either directly or indirectly) to an estimate of VaR $_{\tau}(R_t)$ (in this case, y), and solving the bi-level optimization shown in Equation (8), thereby increasing computational costs.

To utilize the RED RL framework, we first need to derive a valid subtask function for CVaR that satisfies the requirements of Definition 5.1. To this end, let us consider Equation (C.1). In particular, we can see that if we treat the expression inside the expectation in Equation (C.1) as our subtask function, f (see Definition 5.1), then we have a piecewise linear subtask function that is invertible with respect to each input given all other inputs, where the subtask, VaR, is independent of the observed per-step reward. Hence, we can adapt Equation (C.1) as our subtask function (given that is satisfies Definition 5.1), as follows:

$$\tilde{R}_t = \text{VaR} - \frac{1}{\tau} (\text{VaR} - R_t)^+, \tag{C.2}$$

where R_t is the observed per-step reward, \tilde{R}_t is the extended per-step reward, VaR is the valueat-risk of the observed per-step reward, and τ is the CVaR parameter. Importantly, this is a valid subtask function with the following properties: the average (or expected value) of the *extended* reward corresponds to the CVaR of the *observed* reward, and the optimal average of the *extended* reward corresponds to the optimal CVaR of the *observed* reward. This is formalized as Corollaries C.1 - C.4 below: **Corollary C.1.** *The function presented in Equation* (C.2) *is a valid subtask function.*

Proof. The function presented in Equation (C.2) is clearly a piecewise linear function that is invertible with respect to each input given all other inputs. Moreover, the subtask, VaR, is independent of the observed per-step reward. Hence, this function satisfies Definition 5.1 for the subtask, VaR. \Box

Corollary C.2. If the subtask, VaR (from Equation (C.2)) is estimated, and such an estimate is equal to the long-run VaR of the observed per-step reward, then the average (or expected value) of the extended reward, \tilde{R}_t , from Equation (C.2) is equal to the long-run CVaR of the observed per-step reward.

Proof. This follows directly from Equation (C.1b).

Corollary C.3. If the subtask, VaR (from Equation (C.2)) is estimated, and the resulting average of the extended reward from Equation (C.2) is equal to the long-run CVaR of the observed per-step reward, then the VaR estimate is equal to the long-run VaR of the observed per-step reward.

Proof. This follows directly from Equation (C.1b).

Corollary C.4. A policy that yields an optimal long-run average of the extended reward, \tilde{R}_t , from Equation (C.2) is a CVaR-optimal policy. In other words, the optimal long-run average of the extended reward corresponds to the optimal long-run CVaR of the observed reward.

Proof. For a given policy, we know from Equation (C.1a) that, across a range of VaR estimates, the best possible long-run average of the extended reward for that policy corresponds to the long-run CVaR of the observed reward for that same policy. Hence, the best possible long-run average of the extended reward that can be achieved across various policies and VaR estimates, corresponds to the optimal long-run CVaR of the observed reward.

As such, we now have a subtask function with a single subtask, VaR, and an extended reward whose average, when optimized, corresponds to the optimal CVaR of the observed reward. We are now ready to apply the RED RL framework. First, we can derive the reward-extended TD error update for our subtask, VaR, using the methodology outlined in Section 5.1, where, in this case, we have a piecewise linear subtask function with two segments. To this end, if we assume that the CDF of the per-step reward distribution is continuous at $VaR_{\tau}(R_{t+1})$, the resulting subtask update is as follows:

$$\operatorname{VaR}_{t+1} = \begin{cases} \operatorname{VaR}_{t} + \alpha_{\operatorname{VaR},t} \left(\delta_{t} + \operatorname{CVaR}_{t} - \operatorname{VaR}_{t} \right), & R_{t+1} \ge \operatorname{VaR}_{t} \\ \operatorname{VaR}_{t} + \alpha_{\operatorname{VaR},t} \left(\left(\frac{\tau}{\tau - 1} \right) \delta_{t} + \operatorname{CVaR}_{t} - \operatorname{VaR}_{t} \right), & R_{t+1} < \operatorname{VaR}_{t} \end{cases}, \tag{C.3}$$

where δ_t is the regular TD error, and $\alpha_{\text{VaR},t}$ is the step size.

Note that to derive the update (C.3), we used the fact that when $R_{t+1} < \text{VaR}_{\tau}(R_{t+1})$ and the CDF of the per-step reward distribution is continuous at $\text{VaR}_{\tau}(R_{t+1})$, we have that

$$\mathbb{E}[\operatorname{VaR}_{\tau}(R_{t+1}) - \frac{1}{\tau}(\operatorname{VaR}_{\tau}(R_{t+1}) - R_{t+1}) \mid R_{t+1} < \operatorname{VaR}_{\tau}(R_{t+1})]$$
 (C.4a)

$$= \text{VaR}_{\tau}(R_{t+1}) - \frac{1}{\tau} \left(\text{VaR}_{\tau}(R_{t+1}) - \mathbb{E}[R_{t+1} \mid R_{t+1} < \text{VaR}_{\tau}(R_{t+1})] \right) \tag{C.4b}$$

$$= \operatorname{VaR}_{\tau}(R_{t+1}) - \frac{1}{\tau} (\operatorname{VaR}_{\tau}(R_{t+1}) - \operatorname{CVaR}_{\tau}(R_{t+1})). \tag{C.4c}$$

With this update, we now have all the components needed to utilize the RED algorithms in Appendix A to optimize CVaR (where CVaR corresponds to the \bar{r}_{π} that we want to optimize). We call these CVaR-specific algorithms, the *RED CVaR algorithms*. The full algorithms are included at the end of this appendix.

We now present the tabular *RED CVaR Q-learning* algorithm, along with a corresponding theorem and convergence proof which shows that the VaR and CVaR estimates converge to the optimal long-run VaR and CVaR of the observed reward, respectively:

RED CVaR Q-learning algorithm (tabular): We update a table of estimates, $Q_t : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ as follows:

$$\tilde{R}_{t+1} = \text{VaR}_t - \frac{1}{\tau} (\text{VaR}_t - R_{t+1})^+$$
 (C.5a)

$$\delta_t = \tilde{R}_{t+1} - \text{CVaR}_t + \max_{a} Q_t(S_{t+1}, a) - Q_t(S_t, A_t)$$
 (C.5b)

$$Q_{t+1}(S_t, A_t) = Q_t(S_t, A_t) + \alpha_t \delta_t \tag{C.5c}$$

$$Q_{t+1}(s,a) = Q_t(s,a), \quad \forall s, a \neq S_t, A_t \tag{C.5d}$$

$$CVaR_{t+1} = CVaR_t + \alpha_{CVaR,t}\delta_t$$
 (C.5e)

$$VaR_{t+1} = \begin{cases} VaR_t + \alpha_{VaR,t} \left(\delta_t + CVaR_t - VaR_t \right), & R_{t+1} \ge VaR_t \\ VaR_t + \alpha_{VaR,t} \left(\left(\frac{\tau}{\tau - 1} \right) \delta_t + CVaR_t - VaR_t \right), & R_{t+1} < VaR_t \end{cases},$$
 (C.5f)

where R_t denotes the observed reward, VaR_t denotes the VaR estimate, $CVaR_t$ denotes the CVaR estimate, δ_t denotes the TD error, and α_t , α_{CVaR} , t, and α_{VaR} , t denote the step sizes.

Theorem C.1.1. The RED CVaR Q-learning algorithm (C.5) converges, almost surely, CVaR_t to CVaR*, VaR_t to VaR*, CVaR_{π_t} to CVaR*, VaR_{π_t} to VaR*, and Q_t to a solution of q in the Bellman Equation (4), up to an additive constant, c, where π_t is any greedy policy with respect to Q_t, if the following assumptions hold: 1) the MDP is communicating, 2) the solution of q in (4) is unique up to a constant, 3) the step sizes are in accordance with Assumptions B.3, B.4, B.6, B.10, and B.13, 4) Q_t is updated an infinite number of times for all state-action pairs, such that the ratio of the update frequency of the most-updated state—action pair to the least-updated state—action pair is finite, 5) the subtask function outlined in Equation (C.2) is in accordance with Definition 5.1, 6) there is a unique solution to the CVaR optimization (C.1), and 7) the CDF of the limiting per-step reward distribution is continuous at VaR_T(R_{t+1}).

Proof. By definition, the RED CVaR Q-learning algorithm (C.5) is of the form of the generic RED Q-learning algorithm (16), where CVaR_t corresponds to \bar{R}_t and VaR_t corresponds to $Z_{i,t}$ for a single subtask. We also know from Corollary C.1 that the subtask function used is valid. Hence, Theorem 5.3 applies, such that:

- i) CVaR_t and CVaR_{π_t} converge a.s. to the optimal long-run average, $\bar{r}*$, of the extended reward from the subtask function (i.e., the optimal long-run average of \tilde{R}_t),
- ii) VaR_t and VaR_{π_t} converge a.s. to the corresponding optimal subtask value, z*, and
- iii) Q_t converges to a solution of q in the Bellman Equation (4),

all up to an additive constant, c.

Hence, to complete the proof, we need to show that $\bar{r}* = \text{CVaR}^*$ and $z* = \text{VaR}^*$:

From Corollary C.4 we know that the optimal long-run average of the extended reward corresponds to the optimal long-run CVaR of the observed reward, hence we can conclude that $\bar{r}*=\text{CVaR}^*$. Finally, from Corollary C.3 we can deduce that since CVaR_t converges a.s. to CVaR^* , then z* must correspond to VaR^* . This completes the proof.

As such, with the RED CVaR Q-learning algorithm, we now have a way to optimize the long-run CVaR (and VaR) of the observed reward without the use of an augmented state-space, or an explicit bi-level optimization. See Section 6 and Appendix D for empirical results obtained when using the RED CVaR algorithms.

C.2 Additional Commentary

We now provide additional commentary on the subtask-driven approach for CVaR optimization:

Remark C.1. In the above analysis, we made the assumption that the CDF of the limiting per-step reward distribution is continuous at $VaR_{\tau}(R_{t+1})$, thereby yielding a simpler update (i.e., Equation (C.3)) for the subtask, VaR. We note that this assumption was made for simplicity and convenience, and that an equivalent analysis could be conducted for a more generic VaR update (that does not make the aforementioned assumption) to yield the same theoretical conclusion as Theorem C.1.1. Moreover, as shown empirically in Section 6 and Appendix D, we found that using the simplified version of the VaR update (C.3) yielded strong performance.

Remark C.2. A natural question to ask could be whether we can extend the above convergence results to the prediction case. In other words, can we show that a tabular RED CVaR TD-learning algorithm will converge to the long-run VaR and CVaR of the observed reward induced by following a given policy? It turns out that, because we are not optimizing the expectation in Equation (C.1a) when doing prediction (we are only learning it), we cannot guarantee that we will eventually find the optimal VaR estimate, which implies that we may not recover the CVaR value (since Equation (C.1b) only holds to the optimal VaR value). However, this is not to say that a RED CVaR TD-learning algorithm has no use; in fact, we do use such an algorithm as part of an actor-critic architecture for optimizing CVaR in the inverted pendulum experiment (see Appendix D). Empirically, as discussed in Section 6, we find that this actor-critic approach is able to find the optimal CVaR policy.

Remark C.3. It should be noted that in the risk measure literature, risk measures are typically classified into two categories: static or dynamic. Static risk measures are most commonly characterized as being interpretable, but lacking in time consistency (Boda and Filar, 2006). Conversely, dynamic risk measures are typically characterized as being difficult to interpret, but capable of time consistency. Curiously, in our case, the CVaR risk measure that we aim to optimize can be characterized as being both interpretable and time-consistent as $t \to \infty$ (while noting that there is some time inconsistency before $t \to \infty$). We note that this does not affect the significance of our results, but rather suggests that the current categorization of risk measures may be insufficient to capture such nuances that occur in the average-reward setting.

C.3 RED CVaR Algorithms

Below is the pseudocode for the RED CVaR algorithms.

Algorithm 5 RED CVaR Q-Learning (Tabular)

```
Input: the policy \pi to be used (e.g., \varepsilon-greedy)
Algorithm parameters: step size parameters \alpha, \alpha_{\text{\tiny CVaR}}, \alpha_{\text{\tiny VaR}}; CVaR parameter \tau
Initialize Q(s, a) \forall s, a (e.g. to zero)
Initialize CVaR arbitrarily (e.g. to zero)
Initialize VaR arbitrarily (e.g. to zero)
Obtain initial S
while still time to train do
    A \leftarrow action given by \pi for S
   Take action A, observe R, S'
   \tilde{R} = \text{VaR} - \frac{1}{\tau} \max{\{\text{VaR} - R, 0\}}
   \delta = R - \text{CVaR} + \max_a Q(S', a) - Q(S, A)
   Q(S, A) = Q(S, A) + \alpha \delta
   \text{CVaR} = \text{CVaR} + \alpha_{\text{\tiny CVaR}} \delta
   if R \geq \text{VaR} then
       VaR = VaR + \alpha_{VaR}(\delta + CVaR - VaR)
       	ext{VaR} = 	ext{VaR} + lpha_{	ext{VaR}} \left( \left( rac{	au}{	au - 1} 
ight) \delta + 	ext{CVaR} - 	ext{VaR} 
ight)
   end if
   S = S'
end while
return Q
```

Algorithm 6 RED CVaR Actor-Critic

```
Input: a differentiable state-value function parameterization \hat{v}(s, w); a differentiable policy pa-
rameterization \pi(a \mid s, \boldsymbol{\theta})
Algorithm parameters: step size parameters \alpha, \alpha_{\pi}, \alpha_{\text{CVaR}}, \alpha_{\text{VaR}}; CVaR parameter \tau
Initialize state-value weights w \in \mathbb{R}^d and policy weights \theta \in \mathbb{R}^{d'} (e.g. to 0)
Initialize CVaR arbitrarily (e.g. to zero)
Initialize VaR arbitrarily (e.g. to zero)
Obtain initial S
while still time to train do
    A \sim \pi(\cdot \mid S, \boldsymbol{\theta})
   Take action A, observe R, S'
    \tilde{R} = \text{VaR} - \frac{1}{\tau} \max{\{\text{VaR} - R, 0\}}
    \delta = \tilde{R} - \text{CVaR} + \hat{v}(S', \boldsymbol{w}) - \hat{v}(S, \boldsymbol{w})
    \boldsymbol{w} = \boldsymbol{w} + \alpha \delta \nabla \hat{v}(S, \boldsymbol{w})
    \theta = \theta + \alpha_{\pi} \delta \nabla \ln \pi (A \mid S, \theta)
    \text{CVaR} = \text{CVaR} + \alpha_{\text{CVaR}} \delta
   if R \geq \text{VaR} then
        \mathrm{VaR} = \mathrm{VaR} + \alpha_{\mathrm{vaR}} (\delta + \mathrm{CVaR} - \mathrm{VaR})
    else
        	extsf{VaR} = 	extsf{VaR} + lpha_{	extsf{VaR}} \left( \left( rac{	au}{	au - 1} 
ight) \delta + 	extsf{CVaR} - 	extsf{VaR} 
ight)
    end if
    S = S'
end while
return w, \theta
```

D Numerical Experiments

This appendix contains details regarding the numerical experiments performed as part of this work. We discuss the experiments performed in the *red-pill blue-pill* environment (see Appendix E for more details on the red-pill blue-pill environment), as well as the experiments performed in the well-known *inverted pendulum* environment.

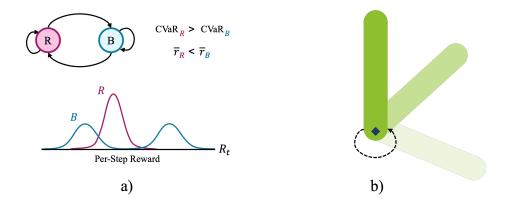


Figure D.1: An illustration of the a) red-pill blue-pill, and b) inverted pendulum environments.

The aim of the experiments was to contrast and compare the RED CVaR algorithms (see Appendix C) with the Differential algorithms from Wan et al. (2021) in the context of risk-aware decision-making. In particular, we aimed to show how the RED CVaR algorithms could be utilized to optimize for CVaR (without the use of an augmented state-space or an explicit bi-level optimization scheme), and contrast the results to those of the Differential algorithms, which served as a sort of 'baseline' to illustrate how our *risk-aware* approach contrasts a *risk-neutral* approach. In other words, we aimed to show whether our algorithms could successfully enable a learning agent to act in a risk-aware manner instead of the usual risk-neutral manner.

In terms of the algorithms used, we used the RED CVaR Q-learning algorithm (Algorithm 5) in the red-pill blue-pill experiment, and the RED CVaR Actor-Critic algorithm (Algorithm 6) in the inverted pendulum experiment. In terms of the (risk-neutral) Differential algorithms used for comparison (see Appendix D.3 for the full algorithms), we used the Differential Q-learning algorithm (Algorithm 7) in the red-pill blue-pill experiment, and the Differential Actor-Critic algorithm (Algorithm 8) in the inverted pendulum experiment.

D.1 Red-Pill Blue-Pill Experiment

In the first experiment, we consider a two-state environment that we created for the purposes of testing our algorithms. It is called the *red-pill blue-pill* environment (see Appendix D), where at every time step an agent can take either a 'red pill', which takes them to the 'red world' state, or a 'blue pill', which takes them to the 'blue world' state. Each state has its own characteristic per-step reward distribution, and in this case, for a sufficiently low CVaR parameter, τ , the red world state has a per-step reward distribution with a lower (worse) mean but a higher (better) CVaR compared to the blue world state. As such, this task allows us to answer the following question: *can a RED CVaR algorithm successfully enable the agent to learn a policy that prioritizes optimizing the reward CVaR over the average-reward?* In particular, we would expect that the RED CVaR Q-learning algorithm enables the agent to learn a policy that prefers to stay in the red world, and that the (risk-neutral) Differential Q-learning algorithm (from Wan et al. (2021)) enables the agent to learn a policy that prefers to stay in the blue world. This task is illustrated in Figure D.1a).

For this experiment, we ran both algorithms using various combinations of step sizes for each algorithm. We used an ε -greedy policy with a fixed epsilon of 0.1, and a CVaR parameter, τ , of 0.25. We set all initial guesses to zero. We ran the algorithms for 100k time steps.

For the Differential Q-learning algorithm, we tested every combination of the value function step size, $\alpha \in \{1/n, 2\text{e-4}, 2\text{e-3}, 2\text{e-2}, 2\text{e-1}\}$ (where 1/n refers to a step size sequence that decreases the step size according to the time step, n), with the average-reward step size, $\eta\alpha$, where $\eta \in \{1\text{e-4}, 1\text{e-3}, 1\text{e-2}, 1\text{e-1}, 1.0, 2.0\}$, for a total of 30 unique combinations. Each combination was run 50 times using different random seeds, and the results were averaged across the runs. The resulting (averaged) average-reward over the last 1,000 time steps is displayed in Figure D.2. As shown in the figure, a value function step size of 2e-4 and an average-reward η of 1.0 resulted in the highest average-reward in the final 1,000 time steps in the red-pill blue-pill task. These are the parameters used to generate the results displayed in Figure 3a).

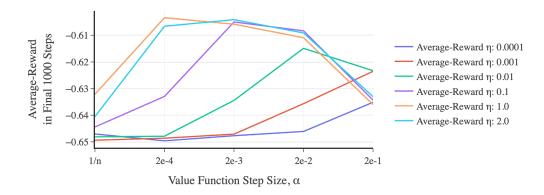


Figure D.2: Step size tuning results for the red-pill blue-pill task when using the Differential Q-learning algorithm. The average-reward in the final 1,000 steps is displayed for various combinations of value function and average-reward step sizes.

For the RED CVaR Q-learning algorithm, we tested every combination of the value function step size, $\alpha \in \{1/n, 2\text{e-4}, 2\text{e-3}, 2\text{e-2}, 2\text{e-1}\}$, with the average-reward (in this case CVaR) step size, $\alpha_{\text{CVaR}} \doteq \eta_{\text{CVaR}} \alpha$, where $\eta_{\text{CVaR}} \in \{1\text{e-4}, 1\text{e-3}, 1\text{e-2}, 1\text{e-1}, 1.0, 2.0\}$, and the VaR step size, $\alpha_{\text{VaR}} \doteq \eta_{\text{VaR}} \alpha$, where $\eta_{\text{VaR}} \in \{1\text{e-4}, 1\text{e-3}, 1\text{e-2}, 1\text{e-1}, 1.0, 2.0\}$, for a total of 180 unique combinations. Each combination was run 50 times using different random seeds, and the results were averaged across the runs. A value function step size of 2e-2, an average-reward (CVaR) η of 1e-1, and a VaR η of 1e-1 yielded the best results and were used to generate the results displayed in Figures 3a) and 4a).

Follow-up Experiment: Varying the CVaR Parameter

Given the results shown in Figure 3a), we can see that, with proper hyperparameter tuning, the tabular RED CVaR Q-learning algorithm is able to reliably find the optimal CVaR policy for a CVaR parameter, τ , of 0.25. In the context of the red-pill blue-pill environment, this means that the agent learns to stay in the red world state because the state has a characteristic reward distribution with a better (higher) CVaR compared to the blue world state. By contrast, the risk-neutral Differential Q-learning algorithm yields an average-reward optimal policy that keeps the agent in the blue world state because the state has a better (higher) average-reward compared to the red world state.

Now consider what would happen if we used the RED CVaR Q-learning algorithm with a τ of 0.99. By definition, a CVaR corresponding to a $\tau \approx 1.0$ is equivalent to the average-reward. Hence, with a τ of 0.99, we would expect that the optimal CVaR policy corresponds to staying in the blue world state (since it has the better average-reward). This means that for some τ between 0.25 and 0.99, there is a critical point where the CVaR-optimal policy changes from staying in the red world (let us call this the *red policy*) to staying in the blue world state (let us call this the *blue policy*).

Importantly, we can estimate this critical point using simple Monte Carlo (MC). We are able to use MC in this case because both policies effectively stay in a single state (the red or blue world state), such that the CVaR of the policies can be estimated by sampling the characteristic reward distribution of each state, while accounting for the exploration ε . Figure D.3 shows the MC estimate of the CVaR of the red and blue policies for a range of CVaR parameters, assuming an exploration ε of 0.1. Note that we used the same distribution parameters listed in Appendix E for the red-pill blue-pill environment. As shown in Figure D.3, this critical point occurs somewhere around $\tau \approx 0.8$.

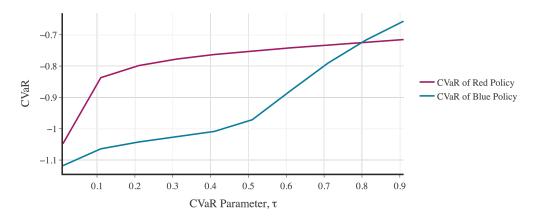


Figure D.3: Monte Carlo estimates of the CVaR of the red and blue policies for a range of CVaR parameters in the red-pill blue-pill environment.

As such, one way that we can further validate the tabular RED CVaR Q-learning algorithm is by re-running the red-pill blue-pill experiment for different CVaR parameters, and seeing if the optimal CVaR policy indeed changes at a $\tau \approx 0.8$. Importantly, this allows us to empirically validate whether the algorithm actually optimizes at the desired risk level. When running this experiment, we used the same hyperparameters used to generate the results in Figure 3a). We ran the experiment for $\tau \in \{0.1, 0.25, 0.5, 0.75, 0.85, 0.9\}$. For each τ , we performed 50 runs using different random seeds, and the results were averaged across the runs.

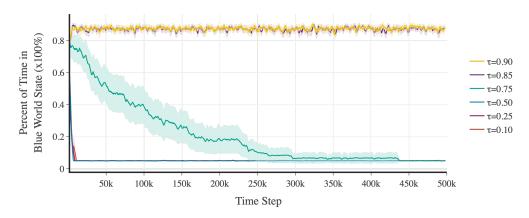


Figure D.4: Rolling percent of time that the agent stays in the blue world state as learning progresses when using the RED CVaR Q-learning algorithm in the red-pill blue-pill environment for a range of CVaR parameters. A solid line denotes the mean percent of time spent in the blue world state, and the corresponding shaded region denotes a 95% confidence interval over 50 runs.

Figure D.4 shows the results of this experiment. In particular, the figure shows a rolling percent of time that the agent stays in the blue world state as learning progresses (note that we used an exploration ε of 0.1). From the figure, we can see that for $\tau \in \{0.1, 0.25, 0.5, 0.75\}$, the agent learns to stay in the red world state, and for $\tau \in \{0.85, 0.9\}$, the agent learns to stay in the blue world state. This is consistent with what we would expect, given that the critical point is $\tau \approx 0.8$. Hence, these results further validate that our algorithm is able to optimize at the desired risk level.

D.2 Inverted Pendulum Experiment

In the second experiment, we consider the well-known *inverted pendulum* environment, where an agent is tasked with learning how to optimally balance an inverted pendulum. We chose this task because it provides us with the opportunity to test our algorithms in an environment where: 1) we must use function approximation (given the high-dimensional state-space), and 2) where the optimal CVaR policy and the optimal average-reward policy are the same policy (i.e., the policy that best balances the pendulum will yield a limiting reward distribution with both the optimal average-reward and reward CVaR). This hence allows us to directly compare the performance of our RED CVaR algorithms to that of the regular Differential algorithms, as well as to gauge how function approximation affects the performance of our algorithms. For this task, we utilized a simple actor-critic architecture (Barto et al., 1983; Sutton and Barto, 2018) as this allowed us to compare the performance of a (non-tabular) RED CVaR TD-learning algorithm with a (non-tabular) Differential TD-learning algorithm. This task is illustrated in Figure D.1b).

For this experiment, we ran both algorithms using various combinations of step sizes for each algorithm. We used a fixed CVaR parameter, τ , of 0.1. We set all initial guesses to zero. We ran the algorithms for 100k time steps. For simplicity, we used tile coding (Sutton and Barto, 2018) for both the value function and policy parameterizations, where we parameterized a softmax policy. For each parameterization, we used 32 tilings, each with 8 X 8 tiles.

By using a linear function approximator (i.e., tile coding), the gradients for the value function and policy parameterizations can be simplified as follows:

$$\nabla \hat{v}(s, \boldsymbol{w}) = \boldsymbol{x}(s),\tag{D.1}$$

$$\nabla \ln \pi(a \mid s, \boldsymbol{\theta}) = \boldsymbol{x}_h(s, a) - \sum_{u \in \mathcal{A}} \pi(u \mid s, \boldsymbol{\theta}) \boldsymbol{x}_h(s, u), \tag{D.2}$$

where $s \in \mathcal{S}$, $a \in \mathcal{A}$, x(s) is the state feature vector, and $x_h(s, a)$ is the softmax preference vector.

For the Differential Actor-Critic algorithm, we tested every combination of the value function step size, $\alpha \in \{1/n, 2\text{e-4}, 2\text{e-3}, 2\text{e-2}\}$ (where 1/n refers to a step size sequence that decreases the step size according to the time step, n), with η 's for the average-reward and policy step sizes, $\eta \alpha$, where $\eta \in \{1\text{e-3}, 1\text{e-2}, 1\text{e-1}, 1.0, 2.0\}$, for a total of 100 unique combinations. Each combination was run 10 times using different random seeds, and the results were averaged across the runs. A value function step size of 2e-3, a policy η of 2.0, and an average-reward η of 1e-2 yielded the best results and were used to generate the results displayed in Figure 3b).

For the RED CVaR Actor-Critic algorithm, we tested every combination of the value function step size, $\alpha \in \{1/n, 2\text{e-4}, 2\text{e-3}, 2\text{e-2}\}$, with η 's for the average-reward (CVaR), VaR, and policy step sizes, $\eta \alpha$, where $\eta \in \{1\text{e-3}, 1\text{e-2}, 1\text{e-1}, 1.0, 2.0\}$, for a total of 500 unique combinations. Each combination was run 10 times using different random seeds, and the results were averaged across the runs. A value function step size of 2e-3, a policy η of 1.0, an average-reward (CVaR) η of 1e-2, and a VaR η of 1e-3 were used to generate the results displayed in Figures 3b) and 4b).

D.3 Risk-Neutral Differential Algorithms

This section contains the pseudocode for the risk-neutral Differential algorithms used for comparison in our experiments.

Algorithm 7 Differential Q-Learning (Tabular)

```
Input: the policy \pi to be used (e.g., \varepsilon-greedy)

Algorithm parameters: step size parameters \alpha, \eta
Initialize Q(s,a) \forall s,a (e.g. to zero)

Initialize \bar{R} arbitrarily (e.g. to zero)

Obtain initial S

while still time to train do

A \leftarrow action given by \pi for S

Take action A, observe R,S'

\delta = R - \bar{R} + \max_a Q(S',a) - Q(S,A)

Q(S,A) = Q(S,A) + \alpha\delta

\bar{R} = \bar{R} + \eta\alpha\delta

S = S'

end while

return Q
```

Algorithm 8 Differential Actor-Critic

return w, θ

```
Input: a differentiable state-value function parameterization \hat{v}(s, \boldsymbol{w}); a differentiable policy parameterization \pi(a \mid s, \boldsymbol{\theta})

Algorithm parameters: step size parameters \alpha, \eta_{\pi}, \eta_{\bar{R}}

Initialize state-value weights \boldsymbol{w} \in \mathbb{R}^d and policy weights \boldsymbol{\theta} \in \mathbb{R}^{d'} (e.g. to 0)

Initialize \bar{R} arbitrarily (e.g. to zero)

Obtain initial S

while still time to train do

A \sim \pi(\cdot \mid S, \boldsymbol{\theta})

Take action A, observe R, S'

\delta = R - \bar{R} + \hat{v}(S', \boldsymbol{w}) - \hat{v}(S, \boldsymbol{w})

\boldsymbol{w} = \boldsymbol{w} + \alpha \delta \nabla \hat{v}(S, \boldsymbol{w})

\boldsymbol{\theta} = \boldsymbol{\theta} + \eta_{\pi} \alpha \delta \nabla \ln \pi(A \mid S, \boldsymbol{\theta})

\bar{R} = \bar{R} + \eta_{\bar{R}} \alpha \delta

S = S'

end while
```

E Red-Pill Blue-Pill Environment

This appendix contains a Python implementation of the red-pill blue-pill environment introduced in this work. The environment consists of a two-state MDP, where at every time step an agent can take either a 'red pill', which takes them to the 'red world' state, or a 'blue pill', which takes them to the 'blue world' state. Each state has its own characteristic per-step reward distribution, and in this case, for a sufficiently low CVaR parameter, τ , the red world state has a per-step reward distribution with a lower (worse) mean but higher (better) CVaR compared to the blue world state. More specifically, the red world state reward distribution is characterized as a gaussian distribution with a mean of -0.7 and a standard deviation of 0.05. Conversely, the blue world state is characterized by a mixture of two gaussian distributions with means of -1.0 and -0.2, and standard deviations of 0.05. We assume all rewards are non-positive. The Python implementation of the environment is provided below:

```
import pandas as pd
import numpy as np
class EnvironmentRedPillBluePill:
  def __init__(self, dist_2_mix_coefficient = 0.5):
    # set distribution parameters
    self.dist_1 = {mean': -0.7, 'stdev': 0.05}
    self.dist_2a = {'mean': -1.0, 'stdev': 0.05}
    self.dist_2b = {'mean': -0.2, 'stdev': 0.05}
    self.dist_2_mix_coefficient = dist_2_mix_coefficient
    # start state
    self.start_state = np.random.choice(['redworld', 'blueworld'])
  def env_start(self, start_state=None):
    # return initial state
    if pd. isnull(start state):
      return self.start_state
    else:
      return start_state
  def env_step(self, state, action, terminal=False):
    if action == 'red_pill':
      next_state = 'redworld'
    elif action == 'blue_pill':
      next_state = 'blueworld'
    if state == 'redworld':
      reward = np.random.normal(loc=self.dist 1['mean'],
                                 scale=self.dist_1['stdev'])
    elif state == 'blueworld':
      dist = np.random.choice(['dist2a', 'dist2b'],
                              p=[self.dist_2_mix_coefficient,
                                 1 - self.dist_2_mix_coefficient])
      if dist == 'dist2a':
        reward = np.random.normal(loc=self.dist_2a['mean'],
                                   scale=self.dist_2a['stdev'])
      elif dist == 'dist2b':
        reward = np.random.normal(loc=self.dist_2b['mean'],
                                   scale=self.dist_2b['stdev'])
    return min(0, reward), next_state, terminal
```