# Unpacking Reward Shaping: Understanding the Benefits of Reward Engineering on Sample Complexity 

Abhishek Gupta*<br>University of Washington<br>abhgupta@cs.washington.edu

Aldo Pacchiano*<br>Microsoft Research, NYC<br>apacchiano@microsoft.com

Yuexiang Zhai<br>UC Berkeley, EECS<br>simonzhai@berkeley.edu

Sham M. Kakade<br>Harvard University<br>sham@seas.harvard.edu

Sergey Levine<br>UC Berkeley, EECS<br>svlevine@eecs.berkeley.edu


#### Abstract

Reinforcement learning provides an automated framework for learning behaviors from high-level reward specifications, but in practice the choice of reward function can be crucial for good results - while in principle the reward only needs to specify what the task is, in reality practitioners often need to design more detailed rewards that provide the agent with some hints about how the task should be completed. The idea of this type of "reward-shaping" has been often discussed in the literature, and is often a critical part of practical applications, but there is relatively little formal characterization of how the choice of reward shaping can yield benefits in sample complexity. In this work, we build on the framework of novelty-based exploration to provide a simple scheme for incorporating shaped rewards into RL along with an analysis tool to show that particular choices of reward shaping provably improve sample efficiency. We characterize the class of problems where these gains are expected to be significant and show how this can be connected to practical algorithms in the literature. We confirm that these results hold in practice in an experimental evaluation, providing an insight into the mechanisms through which reward shaping can significantly improve the complexity of reinforcement learning while retaining asymptotic performance.


## 1 Introduction

Reinforcement learning (RL) in its most general form presents a very difficult optimization problem: when there are no constraints on the reward function or dynamics, a learning algorithm may need to exhaustively explore the entire state space to discover high-reward regions. Naïve algorithms that rely entirely on random exploration are known to be exponentially expensive [26,3], and much of the theoretical work on efficient RL algorithms has focused on smarter exploration strategies that aim to more efficiently cover state space, typically in time polynomial in the state space cardinality [6, $8,24,55,27,32,52]$. Much of this work is based on upper confidence bound (UCB) principles and prescribes some kind of exploration bonus to prioritize exploration of rarely visited regions. Analogous strategies have also been employed in a number of practical RL algorithms [41, 40, 10,

[^0]$12,22,36,43,42,37]$. However, perhaps surprisingly, much of the empirical work on reinforcement learning does not make use of explicit exploration bonuses or other dedicated exploration strategies, despite numerous theoretical results showing them to be essential to attain tractable sample complexity. Instead, practitioners often incorporate prior knowledge of each task into designing or shaping the reward function $[31,38,39,44,5,11,50,54]$, preferring this heuristic approach over the more principled exploration strategies. At this point, one may wonder why is reward shaping often practically preferable to dedicated exploration?

A likely answer to this question lies in the fact that even the best general-purpose exploration strategies still require visiting every state in the MDP at least once in the worst case. This is of course unavoidable without further assumptions. However, in practice, sample complexity that is polynomial in the size of the entire state space might still not be practical, and hence prior knowledge in the form of reward shaping is required to render such tasks tractable. Surprisingly, despite the widespread popularity of reward shaping in RL applications, the analysis of reward shaping has remained limited to proving policy invariance [31] or largely empirical observations, often relegating reward shaping to folk knowledge. In this work, we take a step towards studying the effect of reward shaping on the efficiency of RL algorithms, by asking the following question:

## Can we theoretically justify the sample complexity benefits that reward shaping from prior domain knowledge can provide for reinforcement learning?

We aim to provide a set of tools that formally analyze how reward shaping can improve the complexity of tabula-rasa RL and better direct exploration. To perform this analysis, we first propose a simple modification to standard RL algorithms -"UCBVI-Shaped" that incorporates shaped rewards into optimism based exploration. We use this algorithm instantiation to then provide a regret analysis framework that studies how this introduction of shaped rewards can (in certain cases) provide much more directed optimism than uninformed exploration, while maintaining asymptotic performance.
To approach our analysis, we specifically consider problems where the reward shaping is provided through a term $\widetilde{V}$, a (potentially suboptimal) approximation of the optimal value function $V^{\star}$. In particular, we assume that the shaping $\widetilde{V}$ is a multiplicatively bounded approximation of the optimal value function $V^{\star}$, i.e., $\widetilde{V}(s) \leq V^{\star}(s) \leq \beta \widetilde{V}(s), \forall s$, for a finite multiplicative factor $\beta$. This type of shaped reward function $\widetilde{V}$ can be incorporated into a standard RL algorithm like UCBVI [8] through two channels: (1) bonus scaling - simply reweighting a standard, decaying count-based bonus $\frac{1}{\sqrt{N_{h}(s, a)}}$ by the per-state reward shaping and (2) value projection - adaptively projecting learned value functions into ranges of value functions derived from the reward shaping.
We show that this relatively simple algorithmic instantiation lends itself to an analysis that shows significant sample complexity benefits with shaping. Intuitively, the key pieces in our complexity analysis of UCBVI-Shaped are: (1) A multiplicative sandwich condition (via $\beta$ ) between $\widetilde{V}$ and $V^{\star}$ allows for the bonus scaling to depend on $\beta \widetilde{V}$ instead of a coarse approximation of $V^{\star}$ such as $H$. This allows for a reduction of complexity from a horizon $H$ dependent term to one scaling with $\beta \widetilde{V} \leq \beta V^{\star}$ by allowing for a faster decay of the exploration bonuses while still providing enough optimism. (2) A projection of the value function prevents "over optimism" by hastening the convergence of the empirical $\widehat{Q}$ functions during value iteration, thus allowing for faster detection of sub-optimal actions. This results in the ability to prune out large parts of the state space, as we also validate empirically.
To summarize, the key contribution of this work is to characterize how reward shaping can provably improve sample efficiency by providing gains in both $|\mathcal{S}|$ and $H$ dependent terms. We do so by analyzing the gains from reward shaping in two different ways: bonus scaling and value projection. We show that the "quality" (determined by $\beta$ ) of the provided shaping can significantly improve the sample efficiency of the resulting reinforcement learning algorithm, provide a set of analysis techniques to understand improvements in sample complexity from shaping, and confirm our findings with numerical experiments.

## 2 Related Work

Regret Analysis in Finite Horizon Episodic Tabular MDPs. Recent research on regret analysis has studied both model-based and model-free RL methods. Model-based methods [23, 4, 8, 20, 17, 35] first learn a model from previous experience and use the learned model for future decision making. In contrast, model-free methods [24, 9, 55, 46, 30, 27] aim to learn the value function without the
model estimation stage and use the learned value function for decision making. Our analysis lies in the model-based framework and is similar to the setting of [8] but with additional assumptions on knowing $\widetilde{V}$ as a multiplicatively bounded approximation of $V^{\star}$. The main difference between our results and the aforementioned works is that we consider the reward shaping setting with a truncated interval assumption. As a result, we reduce the state dependency $|\mathcal{S}|$ to the some smaller "effective state space". Our work is also closely related to [21]. We discuss this in Section 3.

Regret Analysis with Linear Function Approximation. A recent line of literature has investigated the linear function approximation setting by assuming the transition kernels or the value functions can be represented by $d$-dimensional linear features [25,53, 7, 56, 13, 48, 47, 45] or even general function classes [49]. With the aforementioned assumptions, the regret analysis will only depend on the ambient dimension $d$ (or other intrinsic complexity measure), instead of $S, A$ in the tabular setting [45, 48, 13, 56, 7, 52, 25], which could greatly decrease the complexity of learning. Instead of posing structural assumptions on the function class representing the MDP's values or dynamics, we ask the question of whether having approximate knowledge of $V^{\star}$ can improve the speed of learning.

Practical Reward Design and Reward Shaping. Ng et al. [31] proposed a potential-based shaping function $F$ that ensures policy invariance under transformation. However, unbiased potential based reward shaping is rarely used in practice. In most applications, heuristic reward design is carefully performed with potentially biased reward functions. In some large-scale practical RL tasks [11, $50,51]$, the reward functions are heavily handcrafted based on prior knowledge. Besides reward engineering, a distinct line of work applies uninformed exploration algorithms like count-based reward shaping or intrinsic rewards to encourage greater state visitation [41, 40, 10, 12, 22, 36, 43, 42, 37, 28]. Importantly, these methods optimize for the worst case, as they try to cover all states since the exact location of the reward is unknown. In contrast, we look at the problem of incorporating domain knowledge via reward shaping into the exploration process. Closely related to our work is that of Cheng et al. [14], which studies how to incorporate shaping (heuristics) into the RL process via reducing the effective horizon. This work provides gains by reducing the horizon factor while our work provides gains by reducing the size of the effective state space that needs to be searched through.

## 3 Overview

We consider an episodic Markov Decision process $\mathcal{M}=\left(\mathcal{S}, \mathcal{A}, \mathbb{P}^{\star}, r, H\right)$ where $\mathcal{S}$ corresponds to the state space, $\mathcal{A}$ is the action space, $\mathbb{P}^{\star}$ is the transition operator, $r: \mathcal{S} \times \mathcal{A} \rightarrow[0,1]$ is the reward function and $H$ is the problem horizon. We use $|\mathcal{S}|,|\mathcal{A}|$ to denote the number of state and actions, and we use $\mathbb{P}^{\star}(\cdot \mid s, a)$ to denote the transition probability of state action pair $(s, a)$. The value function $V^{\pi}\left(s_{0}\right)$ of a policy $\pi$ starting at an initial state $s_{0}$ is defined as $V^{\pi}\left(s_{0}\right):=\mathbb{E}_{\pi}\left[\sum_{h=0}^{H} r\left(s_{h}, a_{h}\right)\right]$, where $\mathbb{E}_{\pi}$ denote the transition dynamic of $\mathcal{M}$ under $\pi$. Similarly, $V^{\star}\left(s_{0}\right)$ is the value function of the optimal policy $\pi^{\star}$. We consider a sequential interaction between a learner and the MDP $\mathcal{M}$ occurring in rounds indexed by $t \in \mathbb{N}$. At the start of round $t$ the learner selects a policy $\pi_{t}$ that is used to gather a sample trajectory from $\mathcal{M}$. As is standard in the literature, we measure the learner's performance up to round $T$ by $\operatorname{Regret}(T):=\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)$.

Our goal is to show that knowing a reward shaping term $\widetilde{V}$ allows for significantly more sample efficient learning, which with high probability has a sublinear regret upper bound. This bound has a leading term that depends on an "effective state space" of $\mathcal{M}$, determined by the quality of a reward shaping term $\widetilde{V}$ and the nature of the particular MDP being solved. This pruned effective state space can be much smaller than $|\mathcal{S}|$. Additionally, we will show an improved horizon dependence as the shaping term $\widetilde{V}$ allows for the bonus terms to be smaller and therefore decay faster thus replacing horizon factors of $H$ with $\beta \widetilde{V}$ ones.

Intuitively, reward shaping allows for the consideration of a reduced effective state space. Consider the maze environment in Fig. 1 with agent starting in the middle. Without knowing where the goal is, an uninformed exploration algorithm needs to explore nearly the entire maze. However, an effectively incorporated reward shaping term $\widetilde{V}$ would allow the algorithm to prune out the entire left half of the maze, effectively halving the effective state space size. This can happen in two ways: firstly, if the shaping can directly indicate that certain actions are sub-optimal, then entire regions of the space can be eliminated from search. Secondly, even if states are not eliminated, if their corresponding bonuses are scaled according to their shaped value $\widetilde{V}$ from the shaping rather than uniformly with $H$, it limits unnecessary exploration of suboptimal states.

Based on this intuition, we propose modifications to a standard reinforcement algorithm that allows us to perform this analysis effectively. To aid our analysis, we introduce UCBVI-Shaped, a modification of the UCBVI algorithm [8] that uses an additional reward shaping term $\widetilde{V}$ in two ways: (1) Bonus scaling: $\widetilde{V}$ is used to provide an exploration bonus that combines inverse state visitation counts with reward shaping. This allows for the reduction of overoptimism and the dropping of the horizon $H$ dependence (2) Value projection: the shaping term is used to clip the empirical value function $\widehat{V}_{h}(s)$, which enables pruning unnecessary elements of the state space and allows for a bound that depends only on the effective state space size. Informally, our main result for UCBVI-Shaped can be summarized as follows:
Theorem 3.1 (Main Informal). With probability at least $1-\delta$, the regret of UCBVI - Shaped satisfies ${ }^{2}$

$$
\operatorname{Regret}(T)=\mathcal{O}\left(B(\widetilde{V}, \mathcal{M}) \log (T / \delta)+\operatorname{Regret}-\operatorname{UCBVI}\left(\mathcal{S}_{\text {remain }}, \mathcal{A}, T\right)\right)
$$

Here, $\mathcal{S}_{\text {remain }}$ corresponds to the set of states in $\mathcal{M}$ where the information contained in $\widetilde{V}$ was not enough to rule them out, and Regret $-\operatorname{UCBVI}\left(\mathcal{S}_{\text {remain }} \mathcal{A}, T\right)$ is a UCBVI regret upper bound function evaluated on $\mathcal{S}_{\text {remain }}$ and $\mathcal{A} . B(\widetilde{V}, \mathcal{M})$ is a time-agnostic complexity measure that depends solely on the $\mathcal{M}$ and the quality of the shaping $\widetilde{V}$.

We will define $\mathcal{S}_{\text {remain }}$ and $B(\widetilde{V}, \mathcal{M})$ and describe their relationship to $\widetilde{V}$ precisely in Section 5. If $\mathcal{S}_{\text {remain }} \ll \mathcal{S}$, the regret of UCBVI - Shaped can be substantially smaller than Regret $-\operatorname{UCBVI}(\mathcal{S}, \mathcal{A})$, particularly since the first term of the regret grows logarithmically while the second scales with $\sqrt{T}$. We state our assumptions next:
Assumption 1. The quality of the shaping term $\widetilde{V}$ is described by a parameter $\beta$. We assume access to a shaping "value function" estimators $\widetilde{V}_{h}: \mathcal{S} \rightarrow \mathbb{R}$ such that $V_{h}^{\star}(s) \leq \beta \widetilde{V}_{h}(s)$, for all $s \in \mathcal{S}, h \in[H]$ and for some $\beta \geq 1$.

Instead of absorbing $\beta$ into the definition of $\widetilde{V}_{h}$ we allow $\widetilde{V}_{h}(s)<V_{h}^{\star}(s)$ for some states $s \in \mathcal{S}$. We'll show that learning a value of $\beta$ that turns this assumption true can be performed online. This assumption is intimately related to the optimistic $\widetilde{Q}$ assumptions of [21]. A thorough comparison with this work can be found in Appendix B.2.
Assumption 2. We assume the reward functions satisfy $r(s, a) \in[0,1], \forall(s, a) \in \mathcal{S} \times \mathcal{A}$.
Assumption 3. The states $\mathcal{S}$ are $h$-indexed, i.e., the states reachable at time $h \in[H]$ are disjoint from the states reachable at time $h^{\prime} \in[H]$ when $h \neq h^{\prime}$.

Contributions. Our main conceptual innovation is to introduce the notions of pseudosuboptimal and path-pseudosuboptimal states to quantify the "effective" size of the state space as a function of the quality of the shaping term $\widetilde{V}$ and use these notions to show how UCBVI-Shaped can attain significantly improved regret rates. In contrast with for example [21], where the regret rates will always scale at least with the number of states, our regret guarantees depend on an effective state size that may be orders of magnitude smaller. We believe this captures the real complexity improvement that reward shaping may yield, namely, avoid exploration of unnecessary areas of the state space. Other approaches such as [21] may (in general) at most yield an improvement in the dependence on the effective size of the action space. We further show that incorporating shaping into the exploration bonus term improves the horizon-dependence in the bound when the shaping is good enough, allowing us to replace a leading $H$ term with $\max _{s} \beta \widetilde{V}(s)$.

[^1]
## 4 The UCBVI-Shaped Algorithm

To perform analysis of reward shaping, we build on the framework of the UCBVI algorithm [8]. UCBVI is an exploration algorithm based on the upper confidence bound, described in detail in Appendix D. This forms a convenient base algorithm for incorporating shaped rewards in a way that admits faster learning while maintaining asymptotic performance (as we will discuss in Section 5). Our algorithm, UCBVI-Shaped, is a combination of two changes to the upper confidence bound algorithm. First, we modify the bonus scaling to depend on $\widetilde{V}$. Second, we introduce a projection subroutine into value iteration, implemented as a standard clipping operation.

```
Algorithm 1 UCBVI - Shaped
    Input reward function \(r\) (assumed to be known), confidence parameters
    for \(t=1, \ldots, T\)
        Compute \(\widehat{P}_{t}\) using all previous empirical transition data as \(\widehat{\mathbb{P}}_{t}\left(s^{\prime} \mid s, a\right):=\frac{N_{h}^{t}\left(s, a, s^{\prime}\right)}{N_{h}^{t}(s, a)}, \forall h, s, a, s^{\prime}\).
        Compute reward bonus \(b_{h}^{t}(s, a)\) from Eqn. 1 (roughly of order \(\frac{\widetilde{V}}{\sqrt{N(s, a)}}\) ) \(\triangleright\) Bonus scaling
        Run Value-Iteration with Projection (Algorithm 2).
        Set \(\pi_{t}\) as the returned policy of VI.
    End for
```

As in standard UCBVI, we define $N_{h}^{t}(s, a)$ to be the visitation count for the state-action pair $(s, a)$ at iteration $t-1$ for horizon $h: N_{h}^{t}(s, a):=\sum_{i=0}^{t-1} \mathbf{1}\left\{\left(s_{h}^{i}, a_{h}^{i}\right)=(s, a)\right\}$. Similarly to $N_{h}^{t}(s, a)$, we use $N_{h}^{t}\left(s, a, s^{\prime}\right):=\sum_{i=0}^{t-1} \mathbf{1}\left\{\left(s_{h}^{i}, a_{h}^{i}, s_{h+1}^{i}\right)=\left(s, a, s^{\prime}\right)\right\}$ as the visitation count of state-action pair $(s, a)$ and the subsequent state $s^{\prime}$. We then use $\widehat{\mathbb{P}}_{t}\left(s^{\prime} \mid s, a\right):=\frac{N_{h}^{t}\left(s, a, s^{\prime}\right)}{N_{h}^{t}(s, a)},{ }^{3} \quad \forall h, s, a, s^{\prime}$ to denote the empirical transition kernels at iteration $t$. UCBVI uses value iteration with the empirical transition function $\widehat{\mathbb{P}}_{t}$ and a reward function augmented with an exploration bonus, given by $r_{h}+b_{h}^{t}$. This is defined as a dynamic programming procedure that starts at $H$ and then proceeds backward in time to $h=0$, updating according to Algorithm 1 and 2. The key algorithmic changes between UCBVI and UCBVIShaped are highlighted in blue: (1) scaling bonus $b_{h}^{t}$ by the shaping term $\widetilde{V}$ and (2) projecting the empirical value function $\widehat{V}_{h}^{t}(s)$ by an upper bound based on the sandwiched shaping, $\beta \widetilde{V}(s, a)$. As we discuss in Section 6, we also show how the sandwich factor $\beta$ does need to be provided beforehand but instead can be inferred through a straightforward technique for online model selection. While the approximate order of the bouns term is $\frac{\widetilde{V}}{\sqrt{N(s, a)}}$, a more detailed description can be found in Section 5.

```
Algorithm 2 Value Iteration with Projection
    Input \(\left\{\widehat{P}_{t}, r+b_{h}^{t}\right\}_{h=0}^{H-1}\).
    \(\widehat{V}_{H}^{t}(s)=0, \forall s, \widehat{Q}_{h}^{t}(s, a)\)
    While not converged
        \(\widehat{Q}_{h}^{t}(s, a)=\min \left\{r_{h}(s, a)+b_{h}^{t}(s, a)+\widehat{\mathbb{P}}_{t}(\cdot \mid s, a) \cdot \widehat{V}_{h+1}^{t}, H\right\}\)
        \(\widehat{V}_{h}^{t}(s)=\min \left(\max _{a} \widehat{Q}_{h}^{t}(s, a), \beta \widetilde{V}(s)\right) \quad \triangleright\) Value projection
    6: \(\pi_{h}^{t}(s)=\arg \max _{a} \widehat{Q}_{h}^{t}(s, a), \forall h, s, a\).
```


## 5 Analyzing UCBVI-Shaped

In this section, we will derive our main result on the sample complexity of the UCBVI-Shaped algorithm, and along the way introduce pseudosuboptimal and path-pseudosuboptimal states as a tool for deriving bounds that depend only on the "effective" state space as determined by the provided

[^2]reward shaping. We will first introduce some notation to use in our analysis. As described in the previous section, UCBVI-Shaped proceeds in rounds. At the beginning of round $t$, the learner has access to an empirical model $\widehat{\mathbb{P}}_{t}$ built from the data collected up to iteration $t-1$. The bonus terms we consider are built by taking the empirical second moment of $\widetilde{V}$. This is related to the definition of bonus $_{2}$ in Azar et al. [8]. Importantly, the empirical value functions $\widehat{V}_{h}^{t}$ are clipped above by $\beta \widetilde{V}_{h}$ :
\[

$$
\begin{equation*}
b_{h}^{t}(s, a)=\min \left(16 \beta \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{2|\mathcal{S}||\mathcal{A}|}{\delta}}{N_{h}^{t}(s, a)}}+\frac{12 \beta \widetilde{V}^{\max }}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}, 2 \beta \widetilde{V}^{\max }\right), \tag{1}
\end{equation*}
$$

\]

where $\widetilde{V}_{h}^{\max }=\max _{s^{\prime}} \widetilde{V}_{h}\left(s^{\prime}\right)$ and $\widetilde{V}^{\max }=\max _{s^{\prime}, h^{\prime}} \widetilde{V}_{h^{\prime}}\left(s^{\prime}\right)$.
Despite clipping and a modified bonus term, the $\widehat{Q}$ and $\widehat{V}$ values of UCBVI-Shaped are optimistic:
Lemma 5.1. With probability at least $1-\delta$ we have

$$
\begin{equation*}
\widehat{V}_{0}^{t}\left(s_{0}\right) \geq V_{0}^{\star}\left(s_{0}\right), \forall s_{0} \in \mathcal{S} ; \quad \text { and } \quad \widehat{Q}_{h}^{t}(s, a) \geq Q_{h}^{\star}(s, a), \quad \forall(s, a) \in \mathcal{S} \times \mathcal{A} \tag{2}
\end{equation*}
$$

for all $t, h \in \mathbb{N} \times[H]$, where $\widehat{V}_{h}^{t}$ is computed via Algorithm 2.
The proof of Lemma 5.1 can be found in Appendix A.1. Optimism (Lemma 5.1) and the simulation Lemma [2] imply that:

$$
\begin{equation*}
V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq \widehat{V}_{1}^{t}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)=\mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H} b_{h}^{t}\left(s_{h}, a_{h}\right)+\left(\widehat{\mathbb{P}}_{h}^{t}\left(\cdot \mid s_{h}, a_{h}\right)-\mathbb{P}^{\star}\left(\cdot \mid s_{h}, a_{h}\right)\right) \cdot \widehat{V}_{h+1}^{\pi_{t}}\right] . \tag{3}
\end{equation*}
$$

We now turn our attention to characterize the information contained in $\widetilde{V}$.

### 5.1 Pruning of the State Space

Consider the following "surrogate" Q functions, $\widetilde{Q}_{h}^{u}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ induced by $\widetilde{V}$ via:

$$
\widetilde{Q}_{h}^{u}(s, a)=\mathbb{E}_{s^{\prime} \sim \mathbb{P}(\cdot \mid s, a)}\left[r(s, a)+\beta \tilde{V}_{h+1}\left(s^{\prime}\right)\right]
$$

By Assumption 1, we can bound $Q^{\star}$ via $Q_{h}^{\star}(s, a) \leq \widetilde{Q}_{h}^{u}(s, a)$. The basis of our main results is the following observation. If an action $a$ satisfies $\widetilde{Q}_{h}^{u}(s, a)<V_{h}^{\star}(s)$ for state $s$, then $a$ is a suboptimal action for state $s$. The projection (Step 5 of Algorithm 2) ensures the 'empirical' $Q$-functions $\widehat{Q}_{h}^{t}(s, a)$ of Algorithm 1 will quickly converge to values upper bounded by $\widetilde{Q}_{h}^{u}(s, a)$. Since optimism guarantees that $\widehat{Q}_{h}^{t}\left(s, \pi_{*}(s)\right) \geq V^{\star}(s)$, and the policy executed by UCBVI-Shaped is greedy w.r.t $\widehat{Q}_{h}^{t}$, actions belonging to state-action pairs such that $\widetilde{Q}_{h}^{u}(s, a)<V_{h}^{\star}(s)$ will quickly stopped being played by the algorithm. Moreover, all states that are only accessible through state-action


Figure 2: PathPseudoSub ${ }_{\Delta} \times \mathcal{A}$ is split from $\operatorname{Support}\left(d^{\pi^{*}}(s, a)\right)$ by BoundaryPseudoSub $_{\Delta}$. UCBVIShaped can avoid exploring over PathPseudoSub ${ }_{\Delta} \times \mathcal{A}$. pairs of this kind will also stop being visited by the algorithm after only a few iterations. In the subsequent discussion we will call PseudoSub to the set of state action pairs that are quickly 'prunned out' by UCBVI-Shaped, the set of states only accessible through these state action pairs as PathPseudoSub and the set of state-action pairs in PseudoSub that are not in PathPseudoSub $\times \mathcal{A}$ as BoundaryPseudoSub. Once all states in BoundaryPseudoSub have been visited enough, no states in PathPseudoSub will be visited again. Provided this happens sufficiently fast, we can show a regret bound that is independent on the size of PathPseudoSub. The subsequent discussion is aimed at formalizing this and culminates with our main result (Theorem 5.2).
Given a parameter $\Delta>0$, we say that an action $a$ is $\Delta$-pseudosuboptimal ${ }^{4}$ for state $s$ if $V_{h}^{\star}(s) \geq$ $\Delta+\widetilde{Q}_{h}^{u}(s, a)$. We denote the set of $\Delta-$ pseudosuboptimal state action pairs as:

$$
\operatorname{PseudoSub}_{\Delta}=\left\{(s, a) \in \mathcal{S} \times \mathcal{A} \text { s.t. } V_{h}^{\star}(s) \geq \Delta+\widetilde{Q}_{h}^{u}(s, a)\right\} .
$$

[^3]The intuition we want to capture is that all states $\tilde{s}$ that can be accessed only through traversing a state action pair in $\mathrm{PseudoSub}_{\Delta}$, can be safely ignored because we can determine their suboptimality as soon as the state action pairs in PseudoSub $\Delta$ that lead to $\tilde{s}$ are visited enough.

Now we define the set of $\Delta$-path-pseudosuboptimal states as the states that can be reached only by traversing a $\Delta-$ pseudosuboptimal state action pair:

PathPseudoSub $_{\Delta}=\left\{s \in \mathcal{S}\right.$ s.t. all feasible paths from initial states to $s$ intersect PseudoSub $\left.{ }_{\Delta}\right\}$.
Notice that there may be state action pairs in PathPseudoSub ${ }_{\Delta} \times \mathcal{A}$ that may not be in PseudoSub ${ }_{\Delta}$. In fact, there may exist states $s$ in $\mathrm{PathPseudoSub}_{\Delta}$ such that no $(s, a)$ is in $\mathrm{PseudoSub}_{\Delta}$ for all $a \in \mathcal{A}$. We will show that states in PathPseudoSub ${ }_{\Delta}$ will not be explored by UCBVI-Shaped after a few iterations. Intuitively, this happens because the support of the state visitation distribution of the optimal policy does not contain any state in PathPseudoSub ${ }_{\Delta}, \forall \Delta>0$ (or equivalently, $\operatorname{Support}\left(d^{\pi^{\star}}(s, a)\right) \cap$ PathPseudoSub $\left._{\Delta}=\emptyset, \forall \Delta>0\right)$. Hence, once the UCBVI-shaped identifies some sub-optimal state action pairs, the algorithm will not visit these state-action pairs again.

For any state action pair $(s, a)$, define its set of neighboring states Neighbor $(s, a)$ as the set of states with nonzero probability in $\mathbb{P}^{\star}(\cdot \mid s, a)$. By definition of PathPseudoSub ${ }_{\Delta}$, for all $(s, a) \in(\mathcal{S} \times \mathcal{A}) \backslash\left(\left(\right.\right.$ PathPseudoSub $\left._{\Delta} \times \mathcal{A}\right) \cup$ PseudoSub $\left._{\Delta}\right)$, we have:

$$
\begin{equation*}
\text { Neighbor }(s, a) \subseteq \mathcal{S} \backslash \text { PathPseudoSub }_{\Delta} .^{\text {. }} \tag{4}
\end{equation*}
$$

In other words, the neighborhood of any state action pair $(s, a)$ whose state is not in PathPseudoSub ${ }_{\Delta}$ and such that $(s, a)$ is not $\Delta-$ pseudo-suboptimal, are not in PathPseudoSub ${ }_{\Delta}$. We also introduce the notion of "boundary pseudosuboptimal" state action pairs to capture the set of state action pairs that are $\Delta$-suboptimal but whose states are not in PathPseudoSub ${ }_{\Delta}$.

BoundaryPseudoSub $_{\Delta}=\left\{(s, a) \in\right.$ PseudoSub $_{\Delta}$ and $s \notin$ PathPseudoSub $\left._{\Delta}\right\}$.
Naturally, one of these states has to be traversed by any trajectory that contains any state in PathPseudoSub ${ }_{\Delta}$.
Although we only use $\widetilde{Q}^{u}$ in the definition of PseudoSub $_{\Delta}$, PathPseudoSub ${ }_{\Delta}$, and BoundaryPseudoSub $_{\Delta}$, the size of these sets is modulated by the scale of $\beta$ and the width of the intervals $\left[\widetilde{Q}^{l}(s, a), \widetilde{Q}^{u}(s, a)\right]$. We will show that (in the notation of Theorem 3.1) $\mathcal{S}_{\text {pruned }} \approx$ $\mathcal{S} \backslash$ PathPseudoSub ${ }_{\Delta}$. With this notation, the formal version of our main results is stated as follows.
Theorem 5.2. With probability at least $1-6 \delta$, the regret of UCBVI-Shaped is upper bounded by

$$
\begin{aligned}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)= & \mathcal{O}\left(\operatorname { m i n } _ { \Delta } \left(H \beta \widetilde{V}^{\max } \sqrt{\mid \mathcal{S} \backslash \text { PathPseudoSub } \Delta| | \mathcal{A} \left\lvert\, T \ln \frac{\widetilde{V}^{\max }|\mathcal{S}||\mathcal{A}| T}{\delta}\right.}+\right.\right. \\
& \left.\left.\beta^{2}\left(\widetilde{V}^{\max }\right)^{2} H^{1 / 2} \mid \text { BoundaryPseudoSub }\left._{\Delta}\right|^{1 / 2} \ln \frac{\widetilde{V}^{\max }|\mathcal{S}||\mathcal{A}| T}{\delta} \times \min (A(\Delta), B(\Delta))\right)\right)
\end{aligned}
$$

for all $T \in \mathbb{N}$. Where $A(\Delta)=\frac{|\mathcal{S}|^{1 / 2}|\mathcal{A}|^{1 / 2}}{\Delta}$ and $B(\Delta)=\frac{\beta \widetilde{V}^{\max } H^{1 / 2} \mid \text { BoundaryPseudoSub }\left._{\Delta}\right|^{1 / 2}}{\Delta^{2}}$.
Theorem 5.2 instantiates the desiderata of Theorem 3.1. Although this regret upper bound cannot be decomposed into a sum of two term as in Theorem 3.1. For any fixed $\Delta$, the regret upper bound has two components, one where $\mathcal{S}_{\text {remain }}$ can be identified with $\mid \mathcal{S} \backslash$ PathPseudoSub ${ }_{\Delta} \mid$ and a second one that scales logarithmically in $T$. In the next section we flesh out the steps in the proof of this result. The full argument can be found in Appendix A. The bonus scaling also allows us to ameliorate the horizon dependence of the upper bound. Instead of obtaining a $H^{2}$ dependence as the bound of theorem 7.1 in [2], the first term depends on $H \beta \widetilde{V}^{\text {max }}$. Notice that the state dependence in the second term may be only be of order $\log (|\mathcal{S}|)$ if there is a $\Delta$ such that $\mid$ BoundaryPseudoSub ${ }_{\Delta}|\ll| \mathcal{S} \mid$ (albeit at the cost of a quadratic dependence on $1 / \Delta^{2}$ ). Moreover notice that $\mid$ BoundaryPseudoSub ${ }_{\Delta} \mid$ could be much smaller than $\left|\mathrm{PathPseudoSub}_{\Delta}\right|$ (the set of states that are reachable only by visiting states in $\mid$ PseudoSub $_{\Delta} \mid$ ) thus showing that in some cases we can guarantee that even in the low order terms, the regret of UCBVI-Shaped that has polynomial dependence on an effective state space size that may be orders of magnitude smaller than the original one. The full version of this bound, with all the low order terms we have omitted for the sake of readability can be found in Appendix B.1, Theorem B. 11 Note that Theorem 5.2 is a strict generalization to the UCBVI regret upper bounds, as setting $\Delta$ to a value that is smaller than the minimum gap recovers the exact result (Theorem 7.6in [2]).

### 5.2 Proof Intuitions and Sketch for Theorem 5.2

Improved Horizon Dependence. An empirical Bernstein bound shows that adding a bonus scaling (up to low order terms) with $\mathcal{O}\left(\sqrt{\widehat{\operatorname{Var}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{h}^{t}(\cdot \mid s, a)}\left(V_{h+1}^{\star}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \||\mathcal{A}| t}{\delta} / N_{h}^{t}(s, a)}\right)$ is sufficient to ensure optimism. Since $V^{\star}$ is not known, this variance term can be substituted by the empirical second moment of $\beta \widetilde{V}_{h+1}$. Finally, the scaling of these terms can be upper bounded by $\beta \widetilde{V}^{\text {max }}$. Without knowledge of $\widetilde{V}$, achieving this scaling would be challenging, since the only proxy for $V_{h+1}^{\star}$ available is $\widehat{V}_{h+1}^{t}$ or $H$ both of which may vastly overestimate it.
State Pruning. The value function clipping mechanism ensures that $\mathbb{E}_{s^{\prime} \sim \widehat{\mathbb{P}}_{h}^{t}(\cdot \mid s, a)}\left[\widehat{V}_{h+1}^{t}\left(s^{\prime}\right)\right] \leq$ $\beta \mathbb{E}_{s^{\prime} \sim \widehat{P}_{h}^{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}\left(s^{\prime}\right)\right]$ and therefore the empirical gap between $\widehat{Q}_{h}^{t}(s, a)$ and $V^{\star}(s)$ is decreasing at a rate of at most $\mathcal{O}\left(\beta \widetilde{V}^{\max } / \sqrt{N_{h}^{t}(s, a)}\right)$. Since optimism ensures that $\widehat{Q}_{h}^{t}\left(s, \pi^{\star}(a)\right) \geq V^{\star}(s)$, once $\widehat{Q}(s, a)<V^{\star}(s)$, action $a$ will not be chosen anymore. Thus, for any $\Delta$, the number of times any state action pair in PseudoSub ${ }_{\Delta}$ may be visited by UCBVI-Shaped is upper bounded by $\beta^{2}\left(\widetilde{V}^{\max }\right)^{2} / \Delta^{2}$. Since any visit to a state action pair in PathPseudoSub ${ }_{\Delta} \times \mathcal{A} \cup$ PseudoSub $_{\Delta}$ requires a visit to a state in BoundaryPseudoSub ${ }_{\Delta}$, which allows us to bound the total number of visits to a state (or a trajectory containing such a state) in PathPseudoSub ${ }_{\Delta} \times \mathcal{A} \cup$ PseudoSub $_{\Delta}$ by $H \mid$ BoundaryPseudoSub $_{\Delta} \mid \beta^{2}\left(\widetilde{V}^{\max }\right)^{2} / \Delta^{2}$.
Finally, we can split (a version of) the regret decomposition in Eqn. 3 into two sums, one over state action pairs in $\mathcal{U}^{\text {bad }}=$ PathPseudoSub $_{\Delta} \times \mathcal{A} \cup$ PseudoSub $_{\Delta}$ (or trajectories intersecting $\mathcal{U}^{\text {bad }}$ ) and a second one over state action pairs in $\mathcal{U}^{\text {good }}=(\mathcal{S} \times \mathcal{A}) \backslash\left(\right.$ PathPseudoSub $_{\Delta} \times \mathcal{A} \cup$ PseudoSub $\left._{\Delta}\right)$ (or trajectories without intersection with $\mathcal{U}^{\text {bad }}$ ). We can then apply the upper bound on the visitation of state action pairs in PathPseudoSub ${ }_{\Delta} \times \mathcal{A} \cup$ PseudoSub $_{\Delta}$ to derive a regret upper bound over these states scaling with $1 / \Delta$. Using the bound on the number of trajectories that intersect PathPseudoSub ${ }_{\Delta} \times \mathcal{A} \cup$ PseudoSub $_{\Delta}$, the trajectory decomposition yields the term scaling with $1 / \Delta^{2}$ (but having only logarithmic state dependence) in Theorem 5.2. Recall that Eqn. 4 implies that the transition operators over state action pairs in $\mathcal{U}^{\text {good }}$ have support only over $\mathcal{S} \backslash$ PathPseudoSub ${ }_{\Delta}$. Our proofs use this fact to prove a polynomial dependence on $\mid \mathcal{S} \backslash$ PathPseudoSub ${ }_{\Delta} \mid$ and not $|\mathcal{S}|$ in Theorem B. 11 (see Appendix B.1), the full version of Theorem 5.2. Detailed proofs are in Appendix A.

## 6 Practical Considerations: Online Model Selection

Now one may notice that UCBVI-Shaped requires knowledge of the scaling $\beta$ is in order to actually perform the value projection. While this can be pre-provided by a user or set conservatively, in this section we discuss how this can be inferred by viewing this as an online model selection problem. In particular, given a set of $N$ different values of $\beta-\left[\beta_{1}, \beta_{2}, \ldots, \beta_{N}\right]$, each of which parameterizes a different setting of the learning algorithm UCBVI-Shaped $(\beta)$, an online model selection algorithm such as Stochastic CORRAL [1,34] or RegretBalancing [33, 15] jointly infers the value of $\beta$ and learns the appropriate value function online. In Appendix B. 4 we show how these techniques can yield meaningful regret guarantees and we provide pseudocode.

## 7 Numerical Simulations

To show the practical relevance of our analysis on reward shaping we perform some numerical simulations on a family of maze environments with tabular state-action representations, as shown in Fig. 3. These environments have deterministic dynamics and reward. The reward is 0 at all states except a goal sink state, which has reward 1 . These simulations are aimed at studying the following questions: (1) Does reward shaping improve sample complexity in these types of maze environments over uninformed UCBVI? (2) What is the relative importance of the bonus reweighting and the $\widetilde{V}$ projection? (3) How does the "suboptimality" of $\widetilde{V}$ impact the resulting sample complexity? (4) Does introducing decayed shaping actually allow for policy convergence? In these experiments, $\widetilde{V}$ is constructed by scaling the optimal value function $V^{\star}$ by per-state scaling factors sampled independently within the range $\beta$.
 Figure 3: Environments used for numerical simulations. (Left) Open Gridworld (Middle) Single corridor, agent starts bottom right has to reach a goal in top left (Right) Double corridor, agent starts in the middle and has to reach a goal on the left, with many irrelevant states on the right hand side

### 7.1 Does reward shaping help direct exploration over optimism under uncertainty?

We conducted numerical simulations on tabular environments to understand how reward shaping via UCBVI-Shaped provides benefits over standard UCBVI. Fig. 4 shows cumulative regret accumulated with different variants of UCBVI-Shaped (with both projection and bonus scaling), UCBVI-Shaped-P (with only projection), UCBVI-Shaped-BS (with only bonus scaling), UCBVI (standard UCBVI without shaping, as described in [8]). This is benchmarked across the three environments described in Fig. 3, with various levels of imperfect shaping applied by varying $\beta=\{1.5,1.9\}$. As seen from Fig. 4, across all environments UCBVI-shaped with projection and bonus scaling performs most favorably, followed typically by UCBVI-Shaped-P, followed by UCBVI-Shaped-BS and UCBVI, suggesting that reward shaping can significantly help with learning efficiency.


Figure 4: Cumulative regret for learning in various environments with varying amounts of shaping, as compared with UCBVI, and ablations UCBVI-Shaped-BS (no projection) and UCBVI-Shaped-P (no bonus scaling).

### 7.2 How does the effectiveness of reward shaping vary across environments?

We next conducted some numerical simulations across environments to understand how the nature of the environment itself affects the sample complexity of learning with shaping via UCBVI-Shaped. As shown in Fig. 6, we see that UCBVI shaped can be very effective in environments with many irrelevant

(a) Single Corridor

(b) Double Corridor

Figure 5: Effect of suboptimality of reward shaping on the performance of UCBVI-Shaped. While $\beta=1.2,1.5$ don't make much of a difference, very large $\beta$ leads to performance degradation sub-optimal paths like the double corridor environment in Fig. 3, but is relatively less effective in environments where all exploration is directed the same way such as the single corridor. Even incorrect but optimistic shaping will provide guidance towards the goal, making UCBVI relatively less dominant in the single corridor environment as compared to the double corridor. This suggests that in environments where


Figure 6: Visualization of how different environments are affected by reward shaping differently. (Left) intermediate visitations (Right) learning progress of UCBVI vs UCBVI-Shaped. The single corridor environment on the top sees much smaller gains for UCBVI-Shaped compared to double corridor environment.
ruling out an entire part of the exploration space is easy from the shaping, we can expect to see larger benefits.

### 7.3 How does the suboptimality of reward shaping affect learning?

We next compared how different levels of suboptimality of the $\beta$ sandwich term in the reward shaping affect cumulative regret across environments. As shown in Fig. 4 (a) and (d), we see that for environments with open paths (like the open gridworld), the shaping degradation has minimal negative effect until it gets very suboptimal. On the other hand, for corridor and double corridor (Fig. 5 (b)), where there are only a few paths to the goal, suboptimal reward shaping along those paths significantly hamper progress.

### 7.4 Is online UCBVI-Shaped able to infer $\beta$ online without prior knowledge?

As described in Section 6, UCBVI-shaped can be freed of the assumption of $\beta$ being known by performing online model selection of $\beta$ and learning values jointly. In particular, we use the Stochastic CORRAL algorithm [34], a variant of the method introduced in [1] to perform online model selection, with the episodic return being the requisite criterion for updating the model selection distribution. As we see in Fig 7, this scheme is able to show comparable results to when the actual $\beta$ is known beforehand, only degrading as the value of $\beta$ increases. This suggests that online UCBVI-shaped can be practical in regimes with moderate levels of value corruption.


Figure 7: Understanding performance of online model selection in UCBVI-Shaped

## 8 Discussion

In this work, we take a step towards formally analyzing the benefits of reward shaping, proving that effective reward shaping can lead to more efficient learning than uninformed exploration strategies. In our analysis, we study an algorithm that incorporates reward shaping into a modified version of UCBVI, using it to modify bonuses and value function projection. Our analysis shows that incorporating shaped rewards allows for pruning significant parts of the state space and sharpening of optimism in a task directed way. This reduces the dependence of the regret bound on the state space size and on the horizon, depending on the quality of the shaping term and parameters of the MDP. This shows how reward shaping can direct exploration and provide significant sample complexity benefits while retaining asymptotic performance. We hope that this work is a step towards moving sample complexity analysis away from being reward agnostic to actually considering reward shaping more formally in analysis.

## References

[1] A. Agarwal, H. Luo, B. Neyshabur, and R. E. Schapire. Corralling a band of bandit algorithms. In Conference on Learning Theory, pages 12-38. PMLR, 2017.
[2] A. Agarwal, N. Jiang, S. M. Kakade, and W. Sun. Reinforcement learning: Theory and algorithms. CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep, 2019.
[3] A. Agarwal, M. Henaff, S. Kakade, and W. Sun. Pc-pg: Policy cover directed exploration for provable policy gradient learning. Advances in Neural Information Processing Systems, 33: 13399-13412, 2020.
[4] S. Agrawal and R. Jia. Posterior sampling for reinforcement learning: worst-case regret bounds. arXiv preprint arXiv:1705.07041, 2017.
[5] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, et al. Learning dexterous in-hand manipulation. The International Journal of Robotics Research, 39(1):3-20, 2020.
[6] P. Auer, T. Jaksch, and R. Ortner. Near-optimal regret bounds for reinforcement learning. Advances in neural information processing systems, 21, 2008.
[7] A. Ayoub, Z. Jia, C. Szepesvari, M. Wang, and L. Yang. Model-based reinforcement learning with value-targeted regression. In International Conference on Machine Learning, pages 463-474. PMLR, 2020.
[8] M. G. Azar, I. Osband, and R. Munos. Minimax regret bounds for reinforcement learning. In International Conference on Machine Learning, pages 263-272. PMLR, 2017.
[9] Y. Bai, T. Xie, N. Jiang, and Y.-X. Wang. Provably efficient q-learning with low switching cost. Advances in Neural Information Processing Systems, 32, 2019.
[10] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos. Unifying count-based exploration and intrinsic motivation. In Advances in Neural Information Processing Systems, pages 1471-1479, 2016.
[11] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Dębiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, et al. Dota 2 with large scale deep reinforcement learning. arXiv preprint arXiv:1912.06680, 2019.
[12] Y. Burda, H. Edwards, A. Storkey, and O. Klimov. Exploration by random network distillation. arXiv preprint arXiv:1810.12894, 2018.
[13] Q. Cai, Z. Yang, C. Jin, and Z. Wang. Provably efficient exploration in policy optimization. In International Conference on Machine Learning, pages 1283-1294. PMLR, 2020.
[14] C. Cheng, A. Kolobov, and A. Swaminathan. Heuristic-guided reinforcement learning. In M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan, editors, Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, pages 13550-13563, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/ 70d31b87bd021441e5e6bf23eb84a306-Abstract.html.
[15] A. Cutkosky, C. Dann, A. Das, C. Gentile, A. Pacchiano, and M. Purohit. Dynamic balancing for model selection in bandits and rl. In International Conference on Machine Learning, pages 2276-2285. PMLR, 2021.
[16] C. Dann, T. V. Marinov, M. Mohri, and J. Zimmert. Beyond value-function gaps: Improved instance-dependent regret bounds for episodic reinforcement learning. Advances in Neural Information Processing Systems, 34:1-12, 2021.
[17] Y. Efroni, N. Merlis, M. Ghavamzadeh, and S. Mannor. Tight regret bounds for model-based reinforcement learning with greedy policies. Advances in Neural Information Processing Systems, 32, 2019.
[18] A. Faust, K. Oslund, O. Ramirez, A. G. Francis, L. Tapia, M. Fiser, and J. Davidson. PRM-RL: long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning. In 2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018, pages 5113-5120. IEEE, 2018. doi: 10.1109/ICRA. 2018.8461096. URL https://doi.org/10.1109/ICRA.2018.8461096.
[19] D. A. Freedman. On tail probabilities for martingales. the Annals of Probability, pages 100-118, 1975.
[20] R. Fruit, M. Pirotta, A. Lazaric, and R. Ortner. Efficient bias-span-constrained explorationexploitation in reinforcement learning. In International Conference on Machine Learning, pages 1578-1586. PMLR, 2018.
[21] N. Golowich and A. Moitra. Can q-learning be improved with advice? In Conference on Learning Theory, pages 4548-4619. PMLR, 2022.
[22] R. Houthooft, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel. Vime: Variational information maximizing exploration. Advances in neural information processing systems, 29, 2016.
[23] T. Jaksch, R. Ortner, and P. Auer. Near-optimal regret bounds for reinforcement learning. Journal of Machine Learning Research, 11:1563-1600, 2010.
[24] C. Jin, Z. Allen-Zhu, S. Bubeck, and M. I. Jordan. Is q-learning provably efficient? Advances in neural information processing systems, 31, 2018.
[25] C. Jin, Z. Yang, Z. Wang, and M. I. Jordan. Provably efficient reinforcement learning with linear function approximation. In Conference on Learning Theory, pages 2137-2143. PMLR, 2020.
[26] S. M. Kakade et al. On the sample complexity of reinforcement learning. PhD thesis, University College London, 2003.
[27] G. Li, L. Shi, Y. Chen, Y. Gu, and Y. Chi. Breaking the sample complexity barrier to regretoptimal model-free reinforcement learning. Advances in Neural Information Processing Systems, 34, 2021.
[28] K. Li, A. Gupta, A. Reddy, V. H. Pong, A. Zhou, J. Yu, and S. Levine. Mural: Meta-learning uncertainty-aware rewards for outcome-driven reinforcement learning. In International Conference on Machine Learning, pages 6346-6356. PMLR, 2021.
[29] A. Maurer and M. Pontil. Empirical bernstein bounds and sample variance penalization. arXiv preprint arXiv:0907.3740, 2009.
[30] P. Ménard, O. D. Domingues, X. Shang, and M. Valko. Ucb momentum q-learning: Correcting the bias without forgetting. In International Conference on Machine Learning, pages 7609-7618. PMLR, 2021.
[31] A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In ICML, volume 99, pages 278-287, 1999.
[32] A. Pacchiano, P. Ball, J. Parker-Holder, K. Choromanski, and S. Roberts. On optimism in model-based reinforcement learning. arXiv preprint arXiv:2006.11911, 2020.
[33] A. Pacchiano, C. Dann, C. Gentile, and P. Bartlett. Regret bound balancing and elimination for model selection in bandits and rl. arXiv preprint arXiv:2012.13045, 2020.
[34] A. Pacchiano, M. Phan, Y. Abbasi Yadkori, A. Rao, J. Zimmert, T. Lattimore, and C. Szepesvari. Model selection in contextual stochastic bandit problems. Advances in Neural Information Processing Systems, 33:10328-10337, 2020.
[35] A. Pacchiano, P. Ball, J. Parker-Holder, K. Choromanski, and S. Roberts. Towards tractable optimism in model-based reinforcement learning. In Uncertainty in Artificial Intelligence, pages 1413-1423. PMLR, 2021.
[36] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by selfsupervised prediction. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pages 16-17, 2017.
[37] R. Raileanu and T. Rocktäschel. Ride: Rewarding impact-driven exploration for procedurallygenerated environments. In International Conference on Learning Representations, 2020. URL https://openreview.net/forum?id=rkg-TJBFPB.
[38] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. nature, 529(7587):484-489, 2016.
[39] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. nature, 550 (7676):354-359, 2017.
[40] B. C. Stadie, S. Levine, and P. Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. arXiv preprint arXiv:1507.00814, 2015.
[41] H. Tang, R. Houthooft, D. Foote, A. Stooke, O. X. Chen, Y. Duan, J. Schulman, F. DeTurck, and P. Abbeel. \# exploration: A study of count-based exploration for deep reinforcement learning. In Advances in neural information processing systems, pages 2753-2762, 2017.
[42] H. Van Seijen, M. Fatemi, J. Romoff, R. Laroche, T. Barnes, and J. Tsang. Hybrid reward architecture for reinforcement learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips. cc/paper/2017/file/1264a061d82a2edae1574b07249800d6-Paper.pdf.
[43] A. S. Vezhnevets, S. Osindero, T. Schaul, N. Heess, M. Jaderberg, D. Silver, and K. Kavukcuoglu. Feudal networks for hierarchical reinforcement learning. In Proceedings of the 34th International Conference on Machine Learning-Volume 70, pages 3540-3549. JMLR. org, 2017.
[44] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. Nature, 575(7782):350-354, 2019.
[45] Y. Wang, R. Wang, S. S. Du, and A. Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. arXiv preprint arXiv:1912.04136, 2019.
[46] K. Yang, L. Yang, and S. Du. Q-learning with logarithmic regret. In International Conference on Artificial Intelligence and Statistics, pages 1576-1584. PMLR, 2021.
[47] L. Yang and M. Wang. Sample-optimal parametric q-learning using linearly additive features. In International Conference on Machine Learning, pages 6995-7004. PMLR, 2019.
[48] L. Yang and M. Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In International Conference on Machine Learning, pages 10746-10756. PMLR, 2020.
[49] Z. Yang, C. Jin, Z. Wang, M. Wang, and M. I. Jordan. On function approximation in reinforcement learning: Optimism in the face of large state spaces. arXiv preprint arXiv:2011.04622, 2020.
[50] D. Ye, G. Chen, W. Zhang, S. Chen, B. Yuan, B. Liu, J. Chen, Z. Liu, F. Qiu, H. Yu, Y. Yin, B. Shi, L. Wang, T. Shi, Q. Fu, W. Yang, L. Huang, and W. Liu. Towards playing full moba games with deep reinforcement learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pages 621-632. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/ paper/2020/file/06d5ae105ea1bea4d800bc96491876e9-Paper.pdf.
[51] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In L. P. Kaelbling, D. Kragic, and K. Sugiura, editors, 3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings, volume 100 of Proceedings of Machine Learning Research, pages 1094-1100. PMLR, 2019. URL http://proceedings. mlr.press/v100/yu20a.html.
[52] A. Zanette and E. Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In International Conference on Machine Learning, pages 7304-7312. PMLR, 2019.
[53] A. Zanette, D. Brandfonbrener, E. Brunskill, M. Pirotta, and A. Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In International Conference on Artificial Intelligence and Statistics, pages 1954-1964. PMLR, 2020.
[54] Y. Zhai, C. Baek, Z. Zhou, J. Jiao, and Y. Ma. Computational benefits of intermediate rewards for goal-reaching policy learning. Journal of Artificial Intelligence Research, 73:847-896, 2022.
[55] Z. Zhang, Y. Zhou, and X. Ji. Almost optimal model-free reinforcement learningvia referenceadvantage decomposition. Advances in Neural Information Processing Systems, 33:1519815207, 2020.
[56] D. Zhou, J. He, and Q. Gu. Provably efficient reinforcement learning for discounted mdps with feature mapping. In International Conference on Machine Learning, pages 12793-12802. PMLR, 2021.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes] See Section 8
(c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 8
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] This work does not actually use human subjects, and is largely done in simulation. But we have included a discussion in Section 8
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [N/A] Math is used as a theory/formalism, but we don't make any provable claims about it.
(b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Appendix A for link to URL and run instructions in the README in the github repo.
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Appendix A.
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] All plots were created with 3 random seeds with std error bars.
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix A
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [Yes] Envs we used are cited in section 7
(b) Did you mention the license of the assets? [Yes] This is in Appendix B
(c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We published the code and included all environments and assets as a part of this
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] We used three open source domains and collected our own data on these domains.
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [No] We did not include full text since we didn't use an exact script, but we summarized the instructions and included images of the environments used.
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A] Only human involvement was data collection with our system.
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A] Human testers were volunteers

## A Full Proofs

Let $f_{h}: \mathcal{S} \rightarrow \mathbb{R}$ be an arbitrary family of horizon indexed functions satisfying $\left\|f_{h}\right\|_{\infty} \leq B$ for some $B>0$. We assume $f_{h}$ is not a random variable as a function of $\mathcal{S}$. The following upper bounds hold, Lemma A.1. Fix $\delta \in(0,1), \forall t \in \mathbb{N}, s \in \mathcal{S}, a \in \mathcal{A}$ and $h \in[H]$, with probability at least $1-\delta$, we have

$$
\begin{aligned}
\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right| & \leq 16 \sqrt{\frac{\operatorname{Var}\left(f_{h+1}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}+\frac{12 B \frac{2|\mathcal{S} \| \cdot \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)} \\
& \leq 16 B \sqrt{\frac{\ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}+\frac{12 B \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}
\end{aligned}
$$

for all $t \in \mathbb{N}$. Similarly with probability at least $1-\delta$, we have

$$
\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right| \leq 16 \sqrt{\frac{\widehat{\operatorname{Var}}\left(f_{h+1}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}+\frac{12 B \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}
$$

for all $t \in \mathbb{N}$. Where $\widehat{\operatorname{Var}}\left(f_{h+1}\left(s^{\prime}\right) \mid s, a\right)=\frac{1}{N_{h}^{k}(s, a)\left(N_{h}^{k}(s, a)-1\right)} \sum_{1 \leq i<j<N_{h}^{k}(s, a)}\left(f_{h+1}\left(s_{h+1}^{i}\right)-f_{h+1}\left(s_{h+1}^{j}\right)\right)^{2}$.
Proof. Consider a fixed tuple $s, a, t, h \in \mathcal{S} \times \mathcal{A} \times \mathbb{N} \times[H]$. By the definition of $\widehat{\mathbb{P}}_{t}$, we have

$$
\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}=\frac{1}{N_{h}^{t}(s, a)} \sum_{i=1}^{t-1} \mathbf{1}\left\{\left(s_{h^{\prime}}^{i}, a_{h^{\prime}}^{i}\right)=(s, a)\right\} f_{h+1}\left(s_{h^{\prime}+1}^{i}\right)
$$

Now denote $\mathcal{H}_{h, i}$ as the entire history from $t=0$ to iteration $t=i$ where in iteration $i$, the history $\mathcal{H}_{h, i}$ includes all interactions from step 0 up to and including time step $h$. Next, $\forall i=0,1, \ldots, t-1$, define the random variables:

$$
X_{i, h^{\prime}}(s, a)=\mathbf{1}\left\{\left(s_{h^{\prime}}^{i}, a_{h^{\prime}}^{i}\right)=(s, a)\right\} f_{h+1}\left(s_{h^{\prime}+1}^{i}\right)-\mathbb{E}\left[\mathbf{1}\left\{\left(s_{h^{\prime}}^{i}, a_{h^{\prime}}^{i}\right)=(s, a)\right\} f_{h+1}\left(s_{h^{\prime}+1}^{i}\right) \mid \mathcal{H}_{h, i}\right] .
$$

Notice that $\left|X_{i, h^{\prime}}\right| \leq B$, if $\mathbf{1}\left\{\left(s_{h^{\prime}}^{i}, a_{h^{\prime}}^{i}\right)=(s, a)\right\}=1$, else $\left|X_{i, h^{\prime}}\right|=0$ so that

$$
\operatorname{Var}\left(X_{t, h}(s, a) \mid H_{h-1, t}\right) \leq \begin{cases}0 & \text { if }\left|X_{i, h^{\prime}}\right|=0 \\ B^{2} & \text { o.w. }\end{cases}
$$

An anytime Bernstein bound (see Lemma C.2) implies,

$$
\begin{aligned}
\left|\sum_{i=1}^{t-1} X_{i, h^{\prime}}(s, a)\right| & \leq 16 \sqrt{N_{h}^{t}(s, a) \operatorname{Var}\left(f_{h+1}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}+12 B \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta} \\
& \leq 16 B \sqrt{N_{h}^{t}(s, a) \ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta}}+12 B \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}
\end{aligned}
$$

With probability at least $1-\delta$ for all $t \in \mathbb{N}$. Thus, we have

$$
\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right| \leq 16 B \sqrt{\frac{\ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}+\frac{12 B \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}
$$

with probability at least $1-\delta$ for all $t \in \mathbb{N}$. The second inequality can be proven following the exact same proof process as described above, but instead making use of the empirical Bernstein bound of C. 3 instead.

The following bound relates the error of $\widehat{\mathbb{P}}_{t}$ and $\mathbb{P}^{\star}$.
Corollary A. 2 (State-action wise model error under $V^{\star}$ (Lemma 7.3 of [2])). Fix $\delta \in(0,1), \forall t \in$ $[1,2, \ldots T], s \in \mathcal{S}, a \in \mathcal{A}, h \in[H]$, consider $\forall V_{h}^{\star}: \mathcal{S} \rightarrow[0, H]$. With probability at least $1-\delta$, we have

$$
\begin{aligned}
& \left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} V_{h+1}^{\star}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} V_{h+1}^{\star}\right| \leq \\
& \\
& \quad \min \left(16 \sqrt{\left.\frac{\widehat{\operatorname{Var}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\mid s, a)}\left(V_{h+1}^{\star}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}+\frac{12 V_{\max }^{\star}}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}, 2 V_{\max }^{\star}\right),}\right.
\end{aligned}
$$

for all $t \in \mathbb{N}$. Where $\widehat{\operatorname{Var}}\left(V_{h+1}^{\star}\left(s^{\prime}\right) \mid s, a\right)=\frac{1}{N_{h}^{k}(s, a)\left(N_{h}^{k}(s, a)-1\right)} \sum_{1 \leq i<j<N_{h}^{k}(s, a)}\left(V_{h+1}^{\star}\left(s_{h+1}^{i}\right)-V_{h+1}^{\star}\left(s_{h+1}^{j}\right)\right)^{2}$ and $V_{\max }^{\star}=\max _{s} V^{\star}(s)$.

Proof. A direct application of Lemma A. 1 using $B=V_{\max }^{\star}$ yields the desired result.
A direct consequence of Lemma A. 2 and Assumption 1 is that,
Corollary A.3. The $V^{\star}$ "variance bonus" is upper bounded by the empirical $\widetilde{V}$ "second moment bonus":

$$
16 \sqrt{\frac{\widehat{\operatorname{Var}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left(V_{h+1}^{\star}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \||\mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}} \leq 16 \beta \cdot \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{2|\mathcal{S}||\mathcal{A}|}{\delta}}{N_{h}^{t}(s, a)}}
$$

and

$$
\frac{12 V_{\max }^{\star}}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta} \leq \frac{12 \beta \widetilde{V}^{\max }}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}
$$

Proof. For any random variable $X \sim \mathbb{P}$,

$$
\operatorname{Var}(X) \leq \mathbb{E}_{X \sim \mathbb{P}}\left[X^{2}\right]
$$

Therefore, for all $s, a$ and any empirical distribution $\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)$,

$$
\widehat{\operatorname{Var}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left(V_{h+1}^{\star}\left(s^{\prime}\right) \mid s, a\right) \leq \widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\left(V_{h+1}^{\star}\left(s^{\prime}\right)\right)^{2} \mid s, a\right] .
$$

Finally, by Assumption 1,

$$
\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\left(V_{h+1}^{\star}\left(s^{\prime}\right)\right)^{2} \mid s, a\right] \leq \beta^{2} \widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] .
$$

The result follows.
Corollary A. 3 and Lemma A. 1 implies that with probability $1-\delta$ for all $t \in \mathbb{N}, s \in \mathcal{S}, a \in \mathcal{A}$ and $h \in[1, \cdots, H]$ the bonuses $b_{h}^{t}$ from Equation 1 satisfy,

$$
\begin{align*}
\left|\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) \cdot V_{h+1}^{\star}\right| & \stackrel{(i)}{\leq} \min \left(16 \sqrt{\frac{\widehat{\operatorname{Var}\left(V_{h+1}^{\star}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}}{N_{h}^{t}(s, a)}}+\frac{12 V_{\max }^{\star} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}, 2 V_{\max }^{\star}\right) \\
& \stackrel{(i i)}{\leq} b_{h}^{t}(s, a) \tag{5}
\end{align*}
$$

Where inequality (i) follows by Corollary A. 2 and (ii) by Corollary A.3.

Define as $\widehat{V}_{h}^{t}$ to the stage $h$ value function of $\pi_{t}$ as computed in the learned model and using bonus augmented rewards (with bonuses defined as in 1) with the data collected up to round $t-1$. We start by showing $\widehat{V}_{h}^{t}$ is always larger than $V_{h}^{\star}(s)$.
We show the following supporting Lemma, a generalization of Lemma A. 1 where the functions $f_{h}: \mathcal{S} \rightarrow \mathbb{R}$ satisfying $\left\|f_{h}\right\|_{\infty}$ are allowed to be random. This lemma is a slight refinement from its corresponding result in the literature, where the dependence on the size of the support of $\left|\mathbb{P}^{\star}(\cdot \mid s, a)\right|$ is upper bounded by $\mid \mathcal{S}$. Having a more refined control on the size of the support of $\mathbb{P}^{\star}(\cdot \mid s, a)$ is what allows our final bounds to only depend on the size of the effective state space.
Lemma A.4. Fix $\delta \in(0,1), \forall t \in \mathbb{N}, s \in \mathcal{S}, a \in \mathcal{A}, h \in[H]$, with probability at least $1-\delta$, we have

$$
\begin{aligned}
& \left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right| \\
\leq & 16 \beta \cdot \sqrt{\frac{\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{8 B|\mathcal{S} \| \mathcal{A}| t^{3}}{\delta}}{N_{h}^{t}(s, a)}} \\
& +\quad \frac{12 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right|}{N_{h}^{t}(s, a)} \ln \frac{8 B|\mathcal{S}||\mathcal{A}| t^{3}}{\delta}+\frac{1}{t^{2}}
\end{aligned}
$$

simultaneously for all $\left\{f_{h}: \mathcal{S} \rightarrow \mathbb{R}\right\}_{h=0}^{H-1}$ such that $f_{h}(s) \in\left[\widetilde{V}_{h}, \beta \widetilde{V}_{h}\right]$ and $\left\|f_{h}\right\|_{\infty} \leq B$. Where $\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)$ corresponds to the size of the support of distribution $\mathbb{P}^{\star}(\cdot \mid s, a)$.

Proof. The same proof template as in Lemma A. 1 plus a union bound over a covering of the space of $f$ functions yields the desired result. Note that for all $s, a \in \mathcal{S} \times \mathcal{A}$,

$$
\widehat{\operatorname{Var}}\left(f_{h+1}\left(s^{\prime}\right) \mid s, a\right) \leq \widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[f_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \leq \beta^{2} \widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right]
$$

Therefore Lemma A. 1 implies that for any fixed $f_{h+1}$, with probability at least $1-\delta^{\prime}$

$$
\begin{align*}
& \left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right|  \tag{6}\\
\leq & 16 \beta \cdot \sqrt{\frac{\widehat{\operatorname{Var}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left(f_{h+1}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta^{\prime}}}{N_{h}^{t}(s, a)}}+\frac{12 B}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta^{\prime}} \\
\leq & 16 \beta \cdot \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\tilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta^{\prime}}}{N_{h}^{t}(s, a)}}+\frac{12 B}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta^{\prime}} \tag{7}
\end{align*}
$$

for all (fixed) $f \in\left\{f_{h}: \mathcal{S} \rightarrow \mathbb{R}\right\}_{h=0}^{H-1}$ and for all $s, a, h \in \mathcal{S} \times \mathcal{A} \times H$ simultaneously. We now apply a standard $\varepsilon$-net covering argument to the inequality we have proven above. Notice that regardless of the number of samples, support $\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)\right) \subseteq \operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)$. Thus, the only "entries" that matter in $f_{h+1}$ are those corresponding to states in $\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)$.
Let's consider an $\varepsilon$-net of $f_{h+1}$ restricted to $\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)$. Since for any $s \in \mathcal{S}$, we assume $f_{h+1} \in\left[\widetilde{V}_{h+1}, \beta \widetilde{V}_{h+1}\right]$, there exists an epsilon net $\mathcal{N}_{\varepsilon}$ under the infinity norm satisfying $\left|\mathcal{N}_{\varepsilon}\right| \leq$ $\left(\frac{2 B}{\varepsilon}\right)^{\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right|}$. For any $f_{h+1}$ we denote by $f_{h+1}^{\varepsilon}$ its closest element in $\mathcal{N}_{\varepsilon}$ (in the infinity norm over support $\left.\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right)$. The following holds,

$$
\begin{array}{r}
\left|\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right|-\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}^{\varepsilon}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}^{\varepsilon}\right|\right| \leq \\
\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}\right|+\left|\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}^{\varepsilon}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}^{\varepsilon}\right| \leq 2 \varepsilon .
\end{array}
$$

And therefore, setting $\delta^{\prime}=\frac{\delta}{\left(\frac{2 B}{\varepsilon}\right)^{\mid \text {support }\left(\mathbb{P}^{\star} \star(\mid s, q)\right) \mid}} \leq \frac{\delta}{\left|\mathcal{N}_{\varepsilon}\right|}$, a union bound over all elements of $\mathcal{N}_{\varepsilon}$ implies that with probability at least $1-\delta$, we have

$$
\begin{aligned}
&\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right| \\
& \leq 16 \beta \cdot \sqrt{\frac{\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\sim}_{t}(\cdot \mid s, a)}\left[\tilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{4 B|\mathcal{S}||\mathcal{A}| t}{\varepsilon \delta}}{N_{h}^{t}(s, a)}} \\
& \quad+\frac{12 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right|}{N_{h}^{t}(s, a)} \ln \frac{4 B|\mathcal{S}||\mathcal{A}| t}{\varepsilon \delta}+2 \varepsilon
\end{aligned}
$$

for all $f_{h+1}$ simultaneously. Setting $\varepsilon=\frac{1}{2 t^{2}}$ we get that with probability at least $1-\delta$,

$$
\begin{aligned}
& \left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} f_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} f_{h+1}\right| \\
& \leq 16 \beta \cdot \sqrt{\frac{\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\tilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{8 B|\mathcal{S} \| \mathcal{A}| t^{3}}{\delta}}{N_{h}^{t}(s, a)}} \\
& \quad+\frac{12 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right|}{N_{h}^{t}(s, a)} \ln \frac{8 B|\mathcal{S}||\mathcal{A}| t^{3}}{\delta}+\frac{1}{t^{2}}
\end{aligned}
$$

for all $f_{h+1}$ simultaneously and for all $t \in \mathbb{N}$. This completes the result.

## A.1 Proof of Lemma 5.1

In this section we show that optimism holds for all state-action pairs. We restate Lemma 5.1 for the reader's convenience.
Lemma 5.1. With probability at least $1-\delta$ we have

$$
\begin{equation*}
\widehat{V}_{0}^{t}\left(s_{0}\right) \geq V_{0}^{\star}\left(s_{0}\right), \forall s_{0} \in \mathcal{S} ; \quad \text { and } \quad \widehat{Q}_{h}^{t}(s, a) \geq Q_{h}^{\star}(s, a), \quad \forall(s, a) \in \mathcal{S} \times \mathcal{A} \tag{2}
\end{equation*}
$$

for all $t, h \in \mathbb{N} \times[H]$, where $\widehat{V}_{h}^{t}$ is computed via Algorithm 2.

Proof. We prove via induction. At the additional time step $H$ we have $\widehat{V}_{H}^{t}(s)=V_{H}^{\star}(s)=0$ for all $s$. Starting at $h+1$, and assuming we have $\widehat{V}_{h+1}^{t}(s) \geq V_{h+1}^{\star}(s)$ for all $s$, we move to $h$ below.
Consider any $s, a \in \mathcal{S} \times \mathcal{A}$. First if $\widehat{Q}_{h}^{t}(s, a)=H$ then we have $\widehat{Q}_{h}^{t}(s, a) \geq Q_{h}^{\star}(s, a)$. The following inequalities hold,

$$
\begin{aligned}
\widehat{Q}_{h}^{t}(s, a)-Q_{h}^{\star}(s, a) & =b_{h}^{t}(s, a)+\widehat{\mathbb{P}}_{t}(\cdot \mid s, a) \cdot \widehat{V}_{h+1}^{t}-\mathbb{P}^{\star}(\cdot \mid s, a) \cdot V_{h+1}^{\star} \\
& \stackrel{(i)}{\geq} b_{h}^{t}(s, a)+\widehat{\mathbb{P}}_{t}(\cdot \mid s, a) \cdot V_{h+1}^{\star}-\mathbb{P}^{\star}(\cdot \mid s, a) \cdot V_{h+1}^{\star} \\
& =b_{h}^{t}(s, a)+\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) \cdot V_{h+1}^{\star} \\
& \stackrel{(i i)}{\geq} b_{h}^{t}(s, a)-16 \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}-\frac{12 \widetilde{V}^{\text {max }}}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta} \\
& \stackrel{(i i i)}{\geq} 0
\end{aligned}
$$

Where $(i)$ holds because of the inductive optimism assumption. Inequality (ii) holds because by Corollary A. 2 we have,

$$
\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) \cdot V_{h+1}^{\star} \leq 16 \sqrt{\frac{\widehat{\operatorname{Var}}\left(V_{h+1}^{\star}\left(s^{\prime}\right) \mid s, a\right) \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}+\frac{12 V_{\max }^{\star} \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}
$$

and (iii) by Corollary A. 3 applied to the definition of $b_{h}^{t}$.

## B Supporting Results for Section 5.1

A consequence of Lemma A.1, and Corollary A.3,
Corollary B.1. Fix $\delta \in(0,1), \forall t \in \mathbb{N}, s \in \mathcal{S}, a \in \mathcal{A}, h \in[H]$, consider $\forall \widetilde{V}_{h}: \mathcal{S} \mapsto[0, H]$ with probability at least $1-\delta$, we have

$$
\left|\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)^{\top} \widetilde{V}_{h+1}-\mathbb{P}^{\star}(\cdot \mid s, a)^{\top} \widetilde{V}_{h+1}\right| \leq 16 \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}+\frac{12 \widetilde{V}^{\max }}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta} .
$$

We will now show the empirical $\widetilde{Q}$ values can be approximately upper bounded by the $\widetilde{Q}^{u}$ values.
Let's define the "tilde bonus" as $\tilde{b}_{h}^{t}(s, a)=16 \beta \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \hat{\mathbb{P}}_{t}(\mid s, a)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{2|\mathcal{S}||\mathcal{A}| t}{\delta}}{N_{h}^{t}(s, a)}}+\frac{12 \beta \widetilde{V}^{\max }}{N_{h}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}$.
Corollary B.2. With probability at least $1-2 \delta$ the empirical $Q$ function of state action pair $(s, a)$ satisfies,

$$
\widehat{Q}_{h}^{t}(s, a) \leq r(s, a)+b_{h}^{t}(s, a)+\tilde{b}_{h}^{t}(s, a)+\beta \mathbb{P}^{\star}(\cdot \mid s, a)^{\top} \widetilde{V}_{h+1}:=\widetilde{Q}^{u}(s, a)+b_{h}^{t}(s, a)+\tilde{b}_{h}^{t}(s, a)
$$

for all $s, a \in \mathcal{S} \times \mathcal{A}$ and for all $h \in[0, \cdots, H-1]$ and all $t \in \mathbb{N}$.
Proof. By definition $\widehat{Q}_{h}^{t}(s, a)=r(s, a)+b_{h}^{t}(s, a)+\widehat{\mathbb{P}}^{t}(\cdot \mid s, a) \widehat{V}_{h+1}^{t}$, since $\widehat{V}_{h+1}^{t}$ is clipped above by $\beta \widetilde{V}_{h+1}$,

$$
\begin{aligned}
\widehat{Q}_{h}^{t}(s, a) & =r(s, a)+b_{h}^{t}(s, a)+\widehat{\mathbb{P}}^{t}(\cdot \mid s, a) \widehat{V}_{h+1}^{t} \\
& \leq r(s, a)+b_{h}^{t}(s, a)+\beta \widehat{\mathbb{P}}^{t}(\cdot \mid s, a) \widetilde{V}_{h+1} \\
& \leq r(s, a)+b_{h}^{t}(s, a)+\tilde{b}_{h}^{t}(s, a)+\beta \mathbb{P}^{\star}(\cdot \mid s, a)^{\top} \widetilde{V}_{h+1} \\
& =\widetilde{Q}^{u}(s, a)+b_{h}^{t}(s, a)+\tilde{b}_{h}^{t}(s, a) .
\end{aligned}
$$

The last inequality is a consequence of Corollary B.1. The result follows.

Corollary B. 2 implies that once $b_{h}^{t}(s, a)+\tilde{b}_{h}^{t}(s, a)<V^{\star}(s)-\widetilde{Q}^{u}(s, a)$ optimism guarantees that from that point on, action $a$ will never again be taken at state $s$. Indeed, since optimism (see Lemma 5.1) $\widehat{Q}_{h}^{t}\left(s, \pi_{\star}(s)\right)$ is at least $V^{\star}(s)$, action $a$ will be dominated by action $\pi_{\star}(s)$ from that point onward. We make this intuition precise in the following Lemma by upper bounding the number of times action $a$ is taken at state $s$ when $(s, a) \in$ PseudoSub $_{\Delta}$.
Lemma B.3. With probability at least $1-2 \delta$ and for all $\Delta \in[0,1]$ the state action pairs $(s, a) \in$ PseudoSub $_{\Delta}$ satisfy the bound,

$$
N_{h(s)}^{t}(s, a) \leq \frac{8192 \beta^{2} \times\left(\widetilde{V}^{\max }\right)^{2} \cdot \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{\Delta^{2}}
$$

For all $t \in \mathbb{N}$. Where $h(s)$ corresponds to horizon index of the state partitions that contains ${ }^{5}$ state $s$.

[^4]Proof. We start by assuming the events of Lemma 5.1 (optimism) and Corollary B. $2\left(\widehat{Q}_{h}^{t}(s, a) \leq\right.$ $\left.\widetilde{Q}^{u}(s, a)+b_{h}^{t}(s, a)+\tilde{b}_{h}^{t}(s, a)\right)$ hold. We do not need any $\Delta$-dependent high probability event to hold for our results to be valid.
Once $b_{h(s)}^{t}(s, a)+\tilde{b}_{h(s)}^{t}(s, a)<\widetilde{V}(s)-\widetilde{Q}^{u}(s, a)$, action $a$ will never be taken again at state $s$. The following bound holds for $b_{h(s)}^{t}(s, a)+\tilde{b}_{h(s)}^{t}(s, a)$,

$$
\begin{aligned}
b_{h(s)}^{t}(s, a)+\tilde{b}_{h(s)}^{t}(s, a) & \leq 32 \beta \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}(\cdot \mid s, a)}\left[\widetilde{V}_{h(s)+1}^{2}\left(s^{\prime}\right) \mid s, a\right] \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h(s)}^{t}(s, a)}}+\frac{24 \beta \max _{s^{\prime}} \widetilde{V}_{h(s)+1}\left(s^{\prime}\right)}{N_{h(s)}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta} \\
& \leq 32 \beta \max _{s} \widetilde{V}_{h(s)+1}(s) \sqrt{\frac{\ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{N_{h(s)}^{t}(s, a)}}+\frac{24 \beta \max _{s^{\prime}} \widetilde{V}_{h(s)+1}\left(s^{\prime}\right)}{N_{h(s)}^{t}(s, a)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}
\end{aligned}
$$

Assume $(s, a) \in$ PseudoSub $_{\Delta}$. When

$$
N_{h(s)}^{t}(s, a)>\frac{8192 \beta^{2} \times \max _{s^{\prime}, h^{\prime}} \widetilde{V}_{h^{\prime}}^{2}\left(s^{\prime}\right) \cdot \ln \frac{2|\mathcal{S} \||\mathcal{A}| t}{\delta}}{\Delta^{2}}
$$

we get

$$
\begin{equation*}
b_{h(s)}^{t}(s, a)+\tilde{b}_{h(s)}^{t}(s, a)<\Delta \leq V^{\star}(s)-\widetilde{Q}^{u}(s, a) . \tag{8}
\end{equation*}
$$

The result follows because as soon as $b_{h(s)}^{t}(s, a)+\tilde{b}_{h(s)}^{t}(s, a)<\Delta$ holds, the optimism guarantees that from that point on action $a$ will never again be taken at state $s$. Indeed, let $\pi^{\star}(s)$ be the optimal action at state $s$,

$$
\widehat{Q}_{h(s)}^{t}\left(s, \pi^{\star}(s)\right) \geq Q^{\star}\left(s, \pi^{\star}(s)\right)=V^{\star}(s) \geq \widetilde{Q}^{u}(s, a)+\Delta
$$

Where the last inequality holds because by definition $(s, a) \in$ PseudoSub $_{\Delta}$.
Plugging in the bound of Corollary B. 2 , and invoking inequality 8 ,

$$
\widehat{Q}_{h(s)}^{t}\left(s, \pi^{\star}(s)\right) \geq \widetilde{Q}^{u}(s, a)+\Delta>\widetilde{Q}^{u}(s, a)+b_{h(s)}^{t}(s, a)+\tilde{b}_{h(s)}^{t}(s, a) \geq \widehat{Q}_{h(s)}^{t}(s, a)
$$

implying action $a$ will not be selected by the greedy policy induced by the empirical $Q$ function $\widehat{Q}_{h(s)}^{t}$. This finalizes the result.

Lemma B. 4 (Maximum Visitation of PseudoSuboptimal Pairs). With probability at least $1-2 \delta$, and for all $\Delta \in[0,1]$ and $T \in \mathbb{N}$ the number of episodes whose sample trajectories contain a state from PathPseudoSub ${ }_{\Delta}$ is upper bounded by,

$$
\begin{aligned}
& \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap\left(\text { PathPseudoSub }_{\Delta} \cup \text { PseudoSub }_{\Delta}\right) \neq \emptyset\right) \\
\leq & \mid \text { BoundaryPseudoSub }_{\Delta} \left\lvert\, \times \frac{8192 \beta^{2} \times\left(\widetilde{V}^{\mathrm{max}}\right)^{2} \cdot \ln \frac{2|\mathcal{S} \||\mathcal{A}| T}{\delta}}{\Delta^{2}}\right.
\end{aligned}
$$

and

$$
\begin{aligned}
& \sum_{(s, a) \in\left(\left\{\text { PathPseudoSub }_{\Delta} \times \mathcal{A}\right\} \cup \text { PseudoSub }_{\Delta}\right)} N_{h(s)}^{T}(s, a) \\
& \leq H \times \mid \text { BoundaryPseudoSub }_{\Delta} \left\lvert\, \times \frac{8192 \beta^{2} \times\left(\widetilde{V}^{\mathrm{max}}\right)^{2} \cdot \ln \frac{2|\mathcal{S} \| \mathcal{A}| T}{\delta}}{\Delta^{2}}\right.
\end{aligned}
$$

Where $h(s)$ corresponds to horizon index of the state partitions that contains state $s$.

Proof. Let $\tau_{t}$ be the trajectory sampled by our algorithm at time $t$. Notice that whenever $\tau_{t} \cap$ PathPseudoSub ${ }_{\Delta} \neq \emptyset$, there must be a state action pair in $\tau_{t}$ (occurring previous to the state in $\tau_{t}$ that lies in PathPseudoSub ${ }_{\Delta}$ ) that is in BoundaryPseudoSub ${ }_{\Delta}$, namely the first state-action pair belonging to $\mathrm{PseudoSub}_{\Delta}$ in $\tau_{t}$. We obtain,

$$
\sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap\left(\text { PathPseudoSub }_{\Delta} \cup \text { PseudoSub }_{\Delta}\right) \neq \emptyset\right) \leq \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \text { BoundaryPseudoSub }_{\Delta} \neq \emptyset\right)
$$

Conditioning on the event from Lemma B. 3 implies that for all $T \in \mathbb{N}$

$$
\begin{aligned}
\sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \text { BoundaryPseudoSub }_{\Delta} \neq \emptyset\right) & \leq \sum_{(s, a) \in \text { BoundaryPseudoSub }_{\Delta}} N_{h(s)}^{T}(s, a) \\
& \leq \mid \text { BoundaryPseudoSub }_{\Delta} \left\lvert\, \times \frac{8192 \beta^{2} \times \max _{s^{\prime}, h^{\prime}} \widetilde{V}_{h^{\prime}}^{2}\left(s^{\prime}\right) \cdot \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{\Delta^{2}} .\right.
\end{aligned}
$$

Thus,

$$
\begin{aligned}
& \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap\left(\text { PathPseudoSub }_{\Delta} \cup \text { PseudoSub }_{\Delta}\right) \neq \emptyset\right) \\
\leq & \mid \text { BoundaryPseudoSub }_{\Delta} \left\lvert\, \times \frac{8192 \beta^{2} \times \max _{s^{\prime}, h^{\prime}} \widetilde{V}_{h^{\prime}}^{2}\left(s^{\prime}\right) \cdot \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{\Delta^{2}} .\right.
\end{aligned}
$$

Finalizing the proof of the first bullet. To prove the second bullet, for any $t \in \mathbb{N}$ the following inequalities hold

$$
\begin{aligned}
& \left.\sum_{(s, a) \in\left(\text { PathPseudoSub }_{\Delta} \times \mathcal{A} \cup\right. \text { PseudoSub }}^{\Delta}\right) \\
\leq & N_{h(s)}^{t}(s, a) \\
\leq & \sum_{i=1}^{t} \mathbf{1}\left(\tau_{t} \cap\left(\text { PathPseudoSub }_{\Delta} \times \mathcal{A} \cup \text { PseudoSub }_{\Delta}\right) \neq \emptyset\right) \\
\leq & H \sum_{i=1}^{t} \mathbf{1}\left(\tau_{t} \cap \text { BoundaryPseudoSub }_{\Delta} \neq \emptyset\right) \\
\leq & H \sum_{(s, a) \in \text { BoundaryPseudoSub }_{\Delta}} N_{h(s)}^{t}(s, a) \\
\leq & H \times \mid \text { BoundaryPseudoSub }_{\Delta} \left\lvert\, \times \frac{8192 \beta^{2} \times \max _{s^{\prime}, h^{\prime}} \widetilde{V}_{h^{\prime}}^{2}\left(s^{\prime}\right) \cdot \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}}{\Delta^{2}} .\right.
\end{aligned}
$$

The result follows.

Let's now consider a refinement to the upper bound of Equation 3. For any ${ }^{6} T \in \mathbb{N}$,

$$
\begin{align*}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq & \sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H} b_{h}^{t}\left(s_{h}, a_{h}\right)+\left(\widehat{\mathbb{P}}_{t}\left(\cdot \mid s_{h}, a_{h}\right)-\mathbb{P}^{\star}\left(\cdot \mid s_{h}, a_{h}\right)\right) \cdot \widehat{V}_{h+1}^{\pi_{t}}\right] \\
\leq & \underbrace{\sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H} b_{h}^{t}\left(s_{h}, a_{h}\right)+\left(\widehat{\mathbb{P}}_{t}\left(\cdot \mid s_{h}, a_{h}\right)-\mathbb{P}^{\star}\left(\cdot \mid s_{h}, a_{h}\right)\right) \cdot\left(\widehat{V}_{h+1}^{\pi_{t}}-V_{h+1}^{\star}\right)\right]}_{\mathrm{I}}+ \\
& \sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H}\left(\widehat{\mathbb{P}}_{t}\left(\cdot \mid s_{h}, a_{h}\right)-\mathbb{P}^{\star}\left(\cdot \mid s_{h}, a_{h}\right)\right) \cdot V_{h+1}^{\star}\right] \tag{9}
\end{align*}
$$

[^5]By Equation 5 it follows that with probability at least $1-\delta$ for all $s, a, h, T \in \mathcal{S} \times \mathcal{A} \times[H] \times \mathbb{N}$,

$$
\sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H}\left(\widehat{\mathbb{P}}_{t}\left(\cdot \mid s_{h}, a_{h}\right)-\mathbb{P}^{\star}\left(\cdot \mid s_{h}, a_{h}\right)\right) \cdot V_{h+1}^{\star}\right] \leq \sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H} b_{h}^{t}\left(s_{h}, a_{h}\right)\right]
$$

We will now focus on bounding I. We'll make use of Lemma 7.8 from [2]

Lemma B. 5 (Lemma 7.8 from [2]). With probability at least $1-\delta$ for all $t \in \mathbb{N}, s \in \mathcal{S}, a \in \mathcal{A}$, we have,

$$
\left|\widehat{\mathbb{P}}_{t}\left(s^{\prime} \mid s, a\right)-\mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right)\right| \leq \sqrt{\frac{2 \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}}+\frac{2 \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}
$$

Where $h(s)$ corresponds to horizon index of the state partitions that contains state $s$.

A slight modification of the proof ${ }^{7}$ of Lemma 7.7 in [2] yields,

Lemma B. 6 (Support Aware version of Lemma 7.7 in [2]). With probability at least $1-\delta$ for all $t \in \mathbb{N}, s \in \mathcal{S}, a \in \mathcal{A}$ and all $f: \mathcal{S} \rightarrow \mathbb{R}$ satisfying $f: \mathcal{S} \rightarrow[0, B]$ we have,

$$
\left|\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) f\right| \leq \frac{\mathbb{E}_{s^{\prime} \sim \mathbb{P}^{\star}(\cdot \mid s, a)}\left[f\left(s^{\prime}\right)\right]}{H}+B \min \left(\frac{3\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}, 1\right)
$$

Where $h(s)$ corresponds to horizon index of the state partitions that contains state $s$.

[^6]Proof. We start by conditioning on the event that Lemma B. 5 holds. Take any function $f: \mathcal{S} \rightarrow[0, B]$. We have,

$$
\begin{aligned}
& \left|\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) f\right| \leq \sum_{s^{\prime} \in \mathcal{S}}\left|\left(\widehat{\mathbb{P}}_{t}\left(s^{\prime} \mid s, a\right)-\mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right)\right)\right| f\left(s^{\prime}\right) \\
& \leq \sum_{s^{\prime} \in \mathcal{S}} \sqrt{\frac{2 \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right) f^{2}\left(s^{\prime}\right)}{N_{h(s)}^{t}(s, a)}}+\sum_{s^{\prime} \in \mathcal{S}} \frac{2 \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right) f\left(s^{\prime}\right)}{N_{h(s)}^{t}(s, a)} \\
& \stackrel{(i)}{\leq} \sum_{s^{\prime} \in \mathcal{S}} \sqrt{\frac{2 \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) \ln \left(\frac{|\mathcal{S} \||\mathcal{A}| t H}{\delta}\right) f^{2}\left(s^{\prime}\right)}{N_{h(s)}^{t}(s, a)}}+\frac{2 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)} \\
& \stackrel{(\text { ii) }}{\leq} \sqrt{\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right|} \sqrt{\frac{\sum_{s^{\prime} \in \mathcal{S}} 2 \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right) f^{2}\left(s^{\prime}\right)}{N_{h(s)}^{t}(s, a)}+} \\
& \frac{2 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)} \\
& =\sqrt{\frac{2\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| B H \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)} \frac{\sum_{s^{\prime} \in \mathcal{S}} \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) f^{2}\left(s^{\prime}\right)}{B H}}+ \\
& \frac{2 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)} \\
& \stackrel{(\text { iiii })}{\leq} \frac{\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| B H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}+\frac{\sum_{s^{\prime} \in \mathcal{S}} \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) f^{2}\left(s^{\prime}\right)}{B H}+ \\
& \frac{2 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)} \\
& \stackrel{(i v)}{\leq} \frac{\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| B H \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}+\frac{\sum_{s^{\prime} \in \mathcal{S}} \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) f\left(s^{\prime}\right)}{H}+ \\
& \frac{2 B\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}
\end{aligned}
$$

Inequality $(i)$ follows because $\|f\|_{\infty} \leq B$. Inequality (ii) is a result of Cauchy-Schwarz, (iii) uses the inequality $a b \leq \frac{a^{2}+b^{2}}{2}$ and $(i v)$ uses the condition $f \in[0, B]$ again. The above display implies,

$$
\left|\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) f\right| \leq \frac{3\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| B H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}+\frac{\sum_{s^{\prime} \in \mathcal{S}} \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) f\left(s^{\prime}\right)}{H}
$$

Finally, since $f(s) \in[0, B]$,

$$
\left|\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) f\right| \leq B
$$

Combining these last two equations we conclude,

$$
\left|\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right) f\right| \leq \min \left(B, \frac{3\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| B H \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h(s)}^{t}(s, a)}\right)+\frac{\sum_{s^{\prime} \in \mathcal{S}} \mathbb{P}^{\star}\left(s^{\prime} \mid s, a\right) f\left(s^{\prime}\right)}{H}
$$

The result follows.

From now on we'll use the notation

$$
\xi_{h}^{t}(s, a):=(\beta-1) \widetilde{V}^{\max } \min \left(\frac{3\left|\operatorname{support}\left(\mathbb{P}^{\star}(\cdot \mid s, a)\right)\right| H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)}{N_{h}^{t}(s, a)}, 1\right)
$$

Recall that by Assumption 3, if we define $h(s)$ as to horizon index of the state partitions that contains state $s$, we have $N_{h}(s, a)=0$ for all $h \neq h(s)$.
The following corollary holds,
Corollary B.7. With probability at least $1-2 \delta$ for all $t \in \mathbb{N}, s \in \mathcal{S}, a \in \mathcal{A}, h \in[H]$ we have,

$$
\left(\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)-\mathbb{P}^{\star}(\cdot \mid s, a)\right)\left(\widehat{V}_{h+1}^{\pi_{t}}-V_{h+1}^{\star}\right) \leq \frac{\mathbb{E}_{s^{\prime} \sim \mathbb{P}^{\star}(\cdot \mid s, a)}\left[\widehat{V}_{h+1}^{\pi_{t}}\left(s^{\prime}\right)-V_{h+1}^{\star}\left(s^{\prime}\right)\right]}{H}+\xi_{h}^{t}(s, a)
$$

Proof. Under the event of Lemma 5.1 (optimism), it follows that $\widehat{V}_{h+1}^{\pi_{t}}\left(s^{\prime}\right) \geq V_{h+1}^{\star}\left(s^{\prime}\right)$ for all $t \in \mathbb{N}, s \in$ $\mathcal{S}$. Moreover, $\widehat{V}_{h+1}^{\pi_{t}}\left(s^{\prime}\right)-V_{h+1}^{\star}\left(s^{\prime}\right) \leq \beta \widetilde{V}_{h+1}\left(s^{\prime}\right)-V_{h+1}^{\star}\left(s^{\prime}\right) \leq(\beta-1) \widetilde{V}_{h+1}\left(s^{\prime}\right)$.
Thus we can set $B=(\beta-1) \widetilde{V}^{\text {max }}$ in Lemma B.6.
We will now restate a modified version of Lemma 7.10 from [2] that provides us a bound for term I in terms of expectations over sums of $b_{h}^{t}$ and $\xi_{h}^{t}$ terms.
Lemma B. 8 (Notation Adapted version of Lemma 7.10 from [2]). For all $T \in \mathbb{N}$ with probability at least $1-3 \delta$,

$$
\mathrm{I} \leq e \sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H} b_{h}^{t}\left(s_{h}, a_{h}\right)+\xi_{h}^{t}\left(s_{h}, a_{h}\right)\right] .
$$

Thus, with probability at least $1-3 \delta$,

$$
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq e \sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H} 2 b_{h}^{t}\left(s_{h}, a_{h}\right)+\xi_{h}^{t}\left(s_{h}, a_{h}\right)\right]
$$

The following Lemma will allow us to change from a sum of expectations over the played policies to a sum over the sample trajectories.
Lemma B.9. Simultaneously for all $\mathcal{U} \subset \mathcal{S} \times \mathcal{A}$ and all $T \in \mathbb{N}$.

$$
\begin{aligned}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq & e \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}\right)\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& e \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \notin \mathcal{U}\right)\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& \mathcal{O}\left(\beta \widetilde{V}^{\max } \sqrt{H T \ln \left(\frac{T}{\delta}\right)}\right)
\end{aligned}
$$

and

$$
\begin{aligned}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq & e \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U}=\emptyset\right) \sum_{h=1}^{H}\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& \mathcal{O}\left(\beta \widetilde{V}^{\max } \sqrt{H T \ln \left(\frac{T}{\delta}\right)}+\beta H \widetilde{V}^{\max } \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U} \neq \emptyset\right)\right)
\end{aligned}
$$

With probability at least $1-4 \delta$.

Proof. By Lemma B. 8, for all $T \in \mathbb{N}$ with probability at least $1-3 \delta$,

$$
\begin{equation*}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq e \sum_{t=1}^{T} \mathbb{E}_{\tau \sim \pi_{t}}\left[\sum_{h=1}^{H} 2 b_{h}^{t}\left(s_{h}, a_{h}\right)+\xi_{h}^{t}\left(s_{h}, a_{h}\right)\right] \tag{10}
\end{equation*}
$$

Since $\left|2 b_{h}^{t}\left(s_{h}, a_{h}\right)+\xi_{h}^{t}\left(s_{h}, a_{h}\right)\right| \leq \mathcal{O}\left(\beta \widetilde{V}^{\text {max }}\right)$, anytime Hoeffding lemma C. 1 and the union bound implies that for all $T \in \mathbb{N}$ with probability at least $1-4 \delta$,

$$
\begin{equation*}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq e \sum_{t=1}^{T} \sum_{h=1}^{H} 2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\mathcal{O}\left(\beta \widetilde{V}^{\max } \sqrt{H T \ln \left(\frac{T}{\delta}\right)}\right) \tag{11}
\end{equation*}
$$

Observe that for any $\mathcal{U} \in \mathcal{S} \times \mathcal{A}$,

$$
\begin{align*}
\sum_{t=1}^{T} \sum_{h=1}^{H} 2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)= & \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U}=\emptyset\right) \sum_{h=1}^{H}\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U} \neq \emptyset\right) \sum_{h=1}^{H}\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right) \\
\leq & \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U}=\emptyset\right) \sum_{h=1}^{H}\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& \mathcal{O}\left(\beta H \widetilde{V}^{\max } \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U} \neq \emptyset\right)\right) . \tag{12}
\end{align*}
$$

The result follows by combining inequalities 10,11 and 12 .

## B. 1 Full Proof of Theorem 5.2

We are ready to prove our main supporting Lemma,

Lemma B.10. With probability at east $1-6 \delta$ for all $\Delta>0$ and $T \in \mathbb{N}$ simultaneously,

$$
\begin{aligned}
& \sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq \beta H \widetilde{V}^{\max } \sqrt{\left|\mathcal{U} \operatorname{good}^{\text {ood }}\right| T \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta}}+ \\
& \mathcal{O}\left(\operatorname { m i n } \left((\beta-1) \widetilde{V}^{\max }|\mathcal{S}|^{2}|\mathcal{A}| H^{2} \ln \left(\frac{|\mathcal{S}||\mathcal{A}| T H}{\delta}\right) \log (T+1)+\right.\right. \\
& \beta H \widetilde{V}^{\max } \sqrt{\left|\mathcal{U}^{\text {good }}\right| T \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta}}+ \\
& \beta\left(\widetilde{V}^{\max }\right)^{2}\left(\ln \frac{|\mathcal{S} \| \mathcal{A}| T}{\delta}\right) \frac{\sqrt{H|\mathcal{S} \| \mathcal{A}| \mid \text { BoundaryPseudoSub }{ }^{\mid} \mid}}{\Delta}+ \\
& \beta \widetilde{V}^{\max }|\mathcal{S}||\mathcal{A}| \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1), \\
& (\beta-1) \widetilde{V}^{\text {max }} \mid \mathcal{S} \backslash \text { PathPseudoSub }{ }_{\Delta}\left|\mathcal{U}^{\text {good }}\right| H^{2} \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| T H}{\delta}\right) \log (T+1)+ \\
& \beta \widetilde{V}^{\max }\left|\mathcal{U}^{\text {good }}\right| \ln \frac{|\mathcal{S} \| \mathcal{A}| T}{\delta} \ln (T+1)+ \\
& \beta H \widetilde{V}^{\max } \sqrt{\left\lvert\, \mathcal{U}^{\text {good } \left\lvert\, T \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta}\right.}+\right.} \\
& \left.\left.\beta^{3} H\left(\widetilde{V}^{\text {max }}\right)^{3} \mid \text { BoundaryPseudoSub }_{\Delta} \left\lvert\, \times \frac{\ln \frac{|\mathcal{S} \| \mathcal{A}| T}{\delta}}{\Delta^{2}}\right.\right)\right)
\end{aligned}
$$

Where $\mathcal{U}^{\text {good }}=(\mathcal{S} \times \mathcal{A}) \backslash\left(\right.$ PathPseudoSub $_{\Delta} \times \mathcal{A} \cup$ PseudoSub $\left._{\Delta}\right)$.

Proof. We will condition on the events from Lemmas B. 4 and B.9, something that happens with probability at least $1-6 \delta$.
We'll start with the decomposition from Lemma B.9. Set $\mathcal{U}^{\text {good }}=(\mathcal{S} \times$ $\mathcal{A}) \backslash\left[\left(\right.\right.$ PathPseudoSub $\left._{\Delta} \times \mathcal{A}\right) \cup$ PseudoSub $\left._{\Delta}\right]$. In what follows we'll use $\mathcal{U}^{\text {good }}$ to denote this set when convenient. We will also use the notation $\mathcal{U}^{\text {bad }}=(\mathcal{S} \times \mathcal{A}) \backslash \mathcal{U}^{\text {good }}$ to denote the complement of $\mathcal{U}_{\text {good }}$.

Recall that as a consequence of Equation 4 for all state action pairs $(s, a) \in \mathcal{U}^{\text {good }}$, we have that $\operatorname{Neighbor}(s, a) \subseteq \mathcal{S} \backslash$ PathPseudoSub ${ }_{\Delta}$.
Therefore, the support of $\mathbb{P}^{\star}(\cdot \mid s, a)$ and $\widehat{\mathbb{P}}_{t}(\cdot \mid s, a)$ is contained in $\mathcal{S} \backslash$ PathPseudoSub ${ }_{\Delta}$ for $(s, a) \in \mathcal{S} \backslash\left(\right.$ PathPseudoSub $_{\Delta} \times \mathcal{A} \cup$ PseudoSub $\left._{\Delta}\right)$. This implies that for all $(s, a) \in(\mathcal{S} \times$ $\mathcal{A}) \backslash\left(\right.$ PathPseudoSub $_{\Delta} \times \mathcal{A} \cup$ PseudoSub $\left._{\Delta}\right)$, we have

$$
\begin{equation*}
\xi_{h}^{t}(s, a) \leq(\beta-1) \widetilde{V}^{\max } \min \left(\frac{3 \mid \mathcal{S} \backslash \text { PathPseudoSub }_{\Delta} \left\lvert\, H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)\right.}{N_{h}^{t}(s, a)}, 1\right) \forall(s, a) \in \mathcal{U}^{\text {good }} \tag{13}
\end{equation*}
$$

Similarly,

$$
\begin{equation*}
\xi_{h}^{t}(s, a) \leq(\beta-1) \widetilde{V}^{\max } \min \left(\frac{3|\mathcal{S}| H \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| t H}{\delta}\right)}{N_{h}^{t}(s, a)}, 1\right) \forall(s, a) \in \mathcal{U}^{\mathrm{bad}} \tag{14}
\end{equation*}
$$

Let $\mathcal{U}^{\text {good }}=(\mathcal{S} \times \mathcal{A}) \backslash\left(\right.$ PathPseudoSub $_{\Delta} \times \mathcal{A} \cup$ PseudoSub $\left._{\Delta}\right)$. Inequality 13 implies,

$$
\begin{align*}
& \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right) \xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)  \tag{15}\\
\leq & \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right)(\beta-1) \widetilde{V}^{\text {max }} \min \left(\frac{3 \mid \mathcal{S} \backslash \text { PathPseudoSub } \Delta \left\lvert\, H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)\right.}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}, 1\right) \\
\leq & 3(\beta-1) \widetilde{V}^{\text {max }} \mid \mathcal{S} \backslash \text { PathPseudoSub }_{\Delta} \left\lvert\, H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| T H}{\delta}\right) \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{\mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right)}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}\right. \\
(\stackrel{i)}{\leq} & 6(\beta-1) \widetilde{V}^{\text {max }} \mid \mathcal{S} \backslash \text { PathPseudoSub }_{\Delta}| | \mathcal{U}^{\text {good }} \left\lvert\, H^{2} \ln \left(\frac{|\mathcal{S}||\mathcal{A}| T H}{\delta}\right) \log (T+1) .\right.
\end{align*}
$$

Where inequality $(i)$ holds because of Lemma C.5. The exact same argument applied to the upper bound of Equation 15 implies,

$$
\begin{align*}
\sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\mathrm{bad}}\right) \xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right) & \leq 6(\beta-1) \widetilde{V}^{\mathrm{max}}|\mathcal{S}|\left|\mathcal{U}^{\mathrm{bad}}\right| H^{2} \ln \left(\frac{|\mathcal{S} \| \mathcal{A}| T H}{\delta}\right) \log (T+1) \\
& \leq 6(\beta-1) \widetilde{V}^{\mathrm{max}}|\mathcal{S}|^{2}|\mathcal{A}| H^{2} \ln \left(\frac{|\mathcal{S}||\mathcal{A}| T H}{\delta}\right) \log (T+1) \tag{16}
\end{align*}
$$

Let's now bound the sum $\sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right) b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)$.

$$
\begin{aligned}
& \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right) b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right) \\
&= \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right) \min \left(16 \beta \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \mathbb{P}_{t}\left(\cdot \mid s_{h}^{t}, a_{h}^{t}\right)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s_{h}^{t}, a_{h}^{t}\right] \ln \frac{2|\mathcal{S}||\mathcal{A}|}{\delta}}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}}+\frac{12 \beta \widetilde{V}^{\max }}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}, 2 \beta \widetilde{V}^{\max }\right) \\
& \leq \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right)\left(16 \beta \sqrt{\frac{\widehat{\mathbb{E}}_{s^{\prime} \sim \widehat{\mathbb{P}}_{t}\left(\cdot \mid s_{h}^{t}, a_{h}^{t}\right)}\left[\widetilde{V}_{h+1}^{2}\left(s^{\prime}\right) \mid s_{h}^{t}, a_{h}^{t}\right] \ln \frac{2 \mid \mathcal{S | | \mathcal { A } |}}{\delta}}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}}+\frac{12 \beta \widetilde{V}^{\max }}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)} \ln \frac{2|\mathcal{S} \| \mathcal{A}| t}{\delta}\right) \\
& \stackrel{(a)}{\leq} 16 \beta \widetilde{V}^{\max } \sqrt{\ln \frac{2|\mathcal{S} \| \mathcal{A}| T}{\delta}} \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{\mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right)}{\sqrt{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}}+12 \beta \widetilde{V}^{\max } \ln \frac{2|\mathcal{S} \| \mathcal{A}| T}{\delta} \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{\mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right)}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)} \\
&(i) \\
& \leq 32 \beta H \widetilde{V}^{\max } \sqrt{\left|\mathcal{U}^{\text {good }}\right| T \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta}}+24 \beta \widetilde{V}^{\max }\left|\mathcal{U}^{\text {good }}\right| \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1) .
\end{aligned}
$$

Where inequality $(i)$ holds because of Lemmas C. 4 and C.5. The exact same proof argument as in the proof of inequality 17 that leads to inequality $(a)$ yields,

$$
\begin{aligned}
& \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\mathrm{bad}}\right) b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right) \\
& \leq 16 \beta \widetilde{V}^{\max } \sqrt{\ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta}} \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{\mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {bad }}\right)}{\sqrt{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}}+12 \beta \widetilde{V}^{\text {max }} \ln \frac{2|\mathcal{S} \||\mathcal{A}| T}{\delta} \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{\mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {bad }}\right)}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)} \\
& \stackrel{(i)}{\leq} 32 \beta \widetilde{V}^{\text {max }} \sqrt{H\left|\mathcal{U}^{\text {bad }}\right| \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta} \sum_{(s, a) \in(\mathcal{S} \times \mathcal{A}) \backslash \mathcal{U} \text { god }} \sum_{h=1}^{H} N_{h}^{T}(s, a)}+24 \beta \widetilde{V}^{\max }\left|\mathcal{U}^{\text {bad }}\right| \ln \frac{2|\mathcal{S} \| \mathcal{A}| T}{\delta} \ln (T+1) \\
& \stackrel{(i i)}{\leq} 32 \times 8192 \beta\left(\widetilde{V}^{\max }\right)^{2}\left(\ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta}\right) \frac{\sqrt{H\left|\mathcal{U}^{\text {bad }}\right| \mid \text { BoundaryPseudoSub }} \Delta}{\Delta}+ \\
& 24 \beta \widetilde{V}^{\text {max }}\left|\mathcal{U}^{\text {bad }}\right| \ln \frac{2|\mathcal{S} \| \mathcal{A}| T}{\delta} \ln (T+1) \\
& =\mathcal{O}\left(\beta\left(\widetilde{V}^{\max }\right)^{2}\left(\ln \frac{2|\mathcal{S} \| \mathcal{A}| T}{\delta}\right) \frac{\sqrt{H|\mathcal{S} \| \mathcal{A}| \mid \text { BoundaryPseudoSub }_{\Delta} \mid}}{\Delta}\right)+ \\
& \mathcal{O}\left(\beta \widetilde{V}^{\max }|\mathcal{S}||\mathcal{A}| \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1)\right),
\end{aligned}
$$

where inequality $(i)$ follows because of Lemmas C. 4 and C.5. And inequality (ii) follows from Lemma B.4.

Combining inequalities $15,16,17$ and 18 with Lemma B. 9 we get,

$$
\begin{aligned}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq & e \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right)\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& e \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {bad }}\right)\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+\mathcal{O}\left(\beta \widetilde{V}^{\max } \sqrt{H T \ln \left(\frac{T}{\delta}\right)}\right) \\
\leq & \mathcal{O}\left((\beta-1) \widetilde{V}^{\max }|\mathcal{S}|^{2}|\mathcal{A}| H^{2} \ln \left(\frac{\mid \mathcal{S | | \mathcal { A } | T H}}{\delta}\right) \log (T+1)\right)+ \\
& \mathcal{O}\left(\beta H \widetilde{V}^{\max } \sqrt{\left|\mathcal{U} \mathcal{U}^{\text {good }}\right| T \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta}}\right)+ \\
& \mathcal{O}\left(\beta\left(\widetilde{V}^{\max }\right)^{2}\left(\ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta}\right) \frac{\sqrt{H|\mathcal{S}||\mathcal{A}| \mid \text { BoundaryPseudoSub } \mid}}{\Delta}\right)+ \\
& \mathcal{O}\left(\beta \widetilde{V}^{\max }|\mathcal{S}||\mathcal{A}| \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1)\right)
\end{aligned}
$$

We will now derive a bound that has only logarithmic dependence on $\mathcal{S}$ at the cost of a quadratic dependence on $\Delta^{2}$.
By Lemma B.9, inequalities 15 and 17 and Lemma B.4.

$$
\begin{aligned}
& \sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq e \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U}^{\text {bad }}=\emptyset\right) \sum_{h=1}^{H}\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& \mathcal{O}\left(\beta \widetilde{V}^{\max } \sqrt{H T \ln \left(\frac{T}{\delta}\right)}+\beta H \widetilde{V}^{\max } \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U} \neq \emptyset\right)\right) \\
& \leq e \sum_{t=1}^{T} \sum_{h=1}^{H} \mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}^{\text {good }}\right)\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+ \\
& \mathcal{O}\left(\beta \widetilde{V}^{\max } \sqrt{H T \ln \left(\frac{T}{\delta}\right)}+\beta H \widetilde{V}^{\max } \sum_{t=1}^{T} \mathbf{1}\left(\tau_{t} \cap \mathcal{U} \neq \emptyset\right)\right) \\
& \stackrel{(i)}{=} \mathcal{O}\left((\beta-1) \widetilde{V}^{\max }|\mathcal{S} \backslash \operatorname{PathPseudoSub} \Delta|\left|\mathcal{U}^{\text {good }}\right| H^{2} \ln \left(\frac{|\mathcal{S}||\mathcal{A}| T H}{\delta}\right) \log (T+1)\right)+ \\
& \mathcal{O}\left(\beta \widetilde{V}^{\max }\left|\mathcal{U}^{\text {good }}\right| \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1)\right)+ \\
& \mathcal{O}\left(\beta H \widetilde{V}^{\max } \sqrt{\mid \mathcal{U} \text { good } \left\lvert\, T \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta}\right.}\right)+ \\
& \mathcal{O}\left(\beta^{3} H\left(\widetilde{V}^{\max }\right)^{3}\left|\operatorname{BoundaryPseudoSub}{ }_{\Delta}\right| \times \frac{\ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta}}{\Delta^{2}}\right)
\end{aligned}
$$

This holds for all $\Delta \in[0,1]$ simultaneously because Lemma B.4 holds simultaneously for all $\Delta \in[0,1]$ and for all $T \in \mathbb{N}$ because all the events we have conditioned refer to properties that hold for all $T \in \mathbb{N}$. This finalizes the proof of the desired result.

Using the bound $\left|\mathcal{U}^{\text {good }}\right| \leq \mid \mathcal{S} \backslash$ PathPseudoSub $\Delta||\mathcal{A}|$, we have the following version of our results
Theorem B. 11 (Full version of Theorem 5.2). The regret of UCBVI - Shaped satisfies,

$$
\begin{aligned}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)= & \mathcal{O}\left(\operatorname { m i n } _ { \Delta } \left(H \beta \widetilde{V}^{\max } \sqrt{|\mathcal{S} \backslash \operatorname{PathPseudoSub} \Delta||\mathcal{A}| T \ln \frac{\widetilde{V}^{\max }|\mathcal{S}||\mathcal{A}| T}{\delta}}+\right.\right. \\
& \left.\left.\beta \widetilde{V}^{\max } \ln \frac{\widetilde{V}^{\max }|\mathcal{S}||\mathcal{A}| T}{\delta} \times \min (\bar{A}(\Delta), \bar{B}(\Delta))\right)\right) .
\end{aligned}
$$

For all $T \in \mathbb{N}$ with probability at least $1-6 \delta$. Where

$$
\begin{aligned}
\bar{A}(\Delta)= & \frac{\beta \widetilde{V}^{\max } H^{1 / 2}|\mathcal{S}|^{1 / 2}|\mathcal{A}|^{1 / 2} \mid \text { BoundaryPseudoSub }\left._{\Delta}\right|^{1 / 2}}{\Delta}+\frac{(\beta-1)}{\beta}|\mathcal{S}|^{2}|\mathcal{A}| H^{2} \log (T+1) \\
& +|\mathcal{S}||\mathcal{A}| \ln (T+1),
\end{aligned}
$$

and

$$
\begin{aligned}
\bar{B}(\Delta)= & \left.\frac{\beta^{2}\left(\widetilde{V}^{\text {max }}\right)^{2} H \mid \text { BoundaryPseudoSub }_{\Delta} \mid}{\Delta^{2}}+\frac{(\beta-1)}{\beta} \right\rvert\, \mathcal{S} \backslash \text { PathPseudoSub }\left._{\Delta}\right|^{2}|\mathcal{A}| H^{2} \log (T+1) \\
& +\mid \mathcal{S} \backslash \text { PathPseudoSub } \Delta||\mathcal{A}| \ln (T+1) .
\end{aligned}
$$

## B. 2 Generalizations

In this section we briefly discuss what can we say in the case the shaping functions $\left\{\widetilde{Q}_{h}\right\}_{h \in[H]}$ are of the form $\widetilde{Q}_{h}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ such that $Q_{h}^{\star}(s, a) \leq \beta \widetilde{Q}_{h}(s, a)$ instead of our state-only assumption. By doing so we recover some of the results of [21] for which we provide simple proofs. We will use the notation $\widetilde{Q}^{\max }=\max _{s, a, h} \widetilde{Q}_{h}(s, a)$.

Let's consider a generic optimistic RL algorithm $\mathbb{A}$ that computes an optimistic $Q$ function estimator $\left\{\widehat{Q}_{h}^{t}\right\}_{h \in[H]}$ at the beginning of time-step $t$ satisfying $\widehat{Q}_{h}^{t}\left(s, \pi_{\star}(s)\right) \geq Q_{\star}\left(s, \pi_{\star}(s)\right)$. Notice that we only consider optimism at the optimal action $\pi_{\star}(s)$. Let's also assume the policy played by $\mathbb{A}$ at time $t$ equals $\pi_{t}(\mid s, a, h)=\arg \max _{a \in \mathcal{A}} \widehat{Q}_{h}^{t}(s, a)$. We will introduce the notion of a least upper bound stability for optimistic algorithms.
Definition 1. We say that an optimistic algorithm $\mathbb{A}$ is stable w.r.t. to clipping if substituting the $\left\{\widetilde{Q}_{h}^{t}\right\}_{h \in[H]}$ values by their clipped versions

$$
\widehat{Q}_{h}^{t, \text { clipped }}(s, a) \leftarrow \min \left(\widehat{Q}_{h}^{t}(s, a), \beta \widetilde{Q}(s, a)\right)
$$

does not affect its regret guarantees.
Clipping is done after the $\widehat{Q}_{h}^{t}(s, a)$ values have been computed.
Optimism is preserved. When the un-clipped $\left\{\widehat{Q}_{h}^{t}\right\}_{h \in[H]}$ are optimistic, then as long as $\beta \widetilde{Q}_{h}\left(s, \pi_{\star}(s)\right) \geq Q_{\star}\left(s, \pi_{\star}(s)\right)$, the maximum among all the clipped values is also optimistic.
$\max _{a \in \mathcal{A}} \widehat{Q}_{h}^{t, \text { clipped }}(s, a)=\max _{a \in \mathcal{A}}\left[\min \left(\widehat{Q}_{h}^{t}(s, a), \beta \widetilde{Q}(s, a)\right)\right] \geq \min \left(\widehat{Q}_{h}^{t}\left(s, \pi_{\star}(s)\right), \beta \widetilde{Q}\left(s, \pi_{\star}(s)\right)\right) \geq Q^{\star}\left(s, \pi_{\star}(s)\right)$
If $\mathbb{A}$ 's policy at time $t$ equals $\pi_{t}(\mid s, h)=\arg \max _{a \in \mathcal{A}} \widehat{Q}_{h}^{t, \text { clipped }}(s, a)$, as long as the un-clipped $\left\{\widehat{Q}_{h}^{t}\right\}_{h \in[H]}$ are optimistic the support of policy $\pi_{t}$ satisfies

$$
\operatorname{Support}\left(\pi_{t}(\cdot \mid s, a, h)\right)=\left\{a^{\prime} \text { s.t. } \beta \widetilde{Q}_{h}\left(s, a^{\prime}\right) \geq Q_{h}^{\star}\left(s, \pi_{\star}(s, h)\right)\right\}
$$

This is simply because for all $a \in \mathcal{A} \backslash\left\{a^{\prime}\right.$ s.t. $\beta \widetilde{Q}_{h}\left(s, a^{\prime}\right) \geq Q_{h}\left(s, \pi_{\star}(s, h)\right\}$, the clipped $\left\{\widehat{Q}_{h}^{t, \text { clipped }}(s, a)\right\}_{h \in[H]}$ satisfy

$$
\widehat{Q}_{h}^{t, \text { clipped }}(s, a)<Q_{h}^{\star}\left(s, \pi_{\star}(s, h)\right)
$$

In other words, for all $t \in \mathbb{N}$ the clipped algorithm will only visit state-action pairs in $\cup_{s \in \mathcal{S}, h \in[H]}\left\{a\right.$ s.t. $\left.\beta \widetilde{Q}_{h}(s, a) \geq Q_{h}^{\star}\left(s, \pi_{\star}(s, h)\right)\right\}$. Alternatively this can be thought of as if the learner were to interact only with a prunned MDP with state space $\mathcal{S}$ and state (and horizon) dependent action sets $\mathcal{A}(s)$ of the form $\mathcal{A}(s)=\left\{a\right.$ s.t. $\left.\beta \widetilde{Q}_{h(s)}(s, a) \geq Q_{h(s)}^{\star}\left(s, \pi_{\star}(s, h(s))\right)\right\}$.
It is easy to see the corresponding version of UCBVI-Shaped satisfies the following regret guarantee that recovers the min-max rates of Theorem 1.1 in [21],
Lemma B.12. The regret of UCBVI-Shaped with $\widetilde{Q}$-shaping satisfies,

$$
\begin{aligned}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq & \mathcal{O}\left(\beta \widetilde{Q}^{\max } \sqrt{\left|\operatorname{Pair}_{\mathrm{eff}}\right| T \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta}}\right)+ \\
& \mathcal{O}\left(\beta \widetilde{Q}^{\mathrm{max}} H|\mathcal{S}|\left|\operatorname{Pair}_{\mathrm{eff}}\right| \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1)\right)
\end{aligned}
$$

For all $T \in \mathbb{N}$ with probability at least $1-\delta$, where ${ }^{8} \operatorname{Pair}_{\mathrm{eff}}=\{(s, a) \in \mathcal{S} \times \mathcal{A}$ s.t. $a \in \mathcal{A}(s)\}$.
Proof. We borrow the bound from Theorem 7.6 in [2]. Observe that as long as optimism holds (Lemma 5.1), the regret decomposition of Lemma B. 9 is satisfied.

$$
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq e \sum_{t=1}^{T} \sum_{h=1}^{H}\left(2 b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)+\xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)\right)+\mathcal{O}\left(\beta \widetilde{Q}^{\max } \sqrt{H T \ln \left(\frac{T}{\delta}\right)}\right)
$$

Notice that as long as optimism holds, for all $s_{h}^{t}$ the action $a_{h}^{t}$ satisfies $a_{h}^{t} \in \mathcal{A}\left(s_{h}^{t}\right)$. Thus,

[^7]\[

$$
\begin{aligned}
\sum_{t=1}^{T} \sum_{h=1}^{H} \xi_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right) & \leq \sum_{t=1}^{T} \sum_{h=1}^{H} \beta \widetilde{Q}^{\max } \min \left(\frac{3|\mathcal{S}| H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| t H}{\delta}\right)}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}, 1\right) \\
& \leq 3 \beta \widetilde{Q}^{\max }|\mathcal{S}| H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| T H}{\delta}\right) \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{1}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)} \\
& (i) \\
& 6 \beta \widetilde{Q}^{\max }|\mathcal{S}| \operatorname{Pair}_{\mathrm{eff}} \cdot H \ln \left(\frac{|\mathcal{S}||\mathcal{A}| T H}{\delta}\right) \log (T+1)
\end{aligned}
$$
\]

Where inequality $(i)$ follows from Lemma C.5. Similarly,

$$
\begin{aligned}
\sum_{t=1}^{T} \sum_{h=1}^{H} b_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right) \leq & 16 \beta \widetilde{Q}^{\max } \sqrt{\ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta}} \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{1}{\sqrt{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}} \\
& +12 \beta \widetilde{Q}^{\max } \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta} \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{1}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)} \\
& \stackrel{(i)}{\leq} 32 \beta \widetilde{Q}^{\max } \sqrt{\mid \text { Pair }_{\text {eff }} \left\lvert\, T \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta}\right.} \\
& +24 \beta \widetilde{Q}^{\text {max }} \mid \text { Pair }_{\text {eff }} \left\lvert\, \ln \frac{2|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1)\right.
\end{aligned}
$$

Where (i) holds by Lemmas C. 4 and C.5. Combining these last two inequalities yields

$$
\begin{aligned}
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right) \leq & \mathcal{O}\left(\beta \widetilde{Q}^{\max } \sqrt{\left|\operatorname{Pair}_{\mathrm{eff}}\right| T \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta}}\right)+ \\
& \mathcal{O}\left(\beta \widetilde{Q}^{\mathrm{max}} H|\mathcal{S}| \mid \text { Pair }_{\mathrm{eff}} \left\lvert\, \ln \frac{|\mathcal{S}||\mathcal{A}| T}{\delta} \ln (T+1)\right.\right)
\end{aligned}
$$

The result follows.

Similarly, we can obtain a result for instance dependent bounds for a clipped version of the StrongEuler algorithm. We consider the exact same clipping mechanism as in the previous discussion. Optimism guarantees (just as we have described above) that $\operatorname{Support}\left(\pi_{t}(\cdot \mid s, a, h)\right)=\left\{a^{\prime}\right.$ s.t. $\beta \widetilde{Q}_{h}\left(s, a^{\prime}\right) \geq$ $\left.Q_{h}^{\star}\left(s, \pi_{\star}(s, h)\right)\right\}$. Therefore the data generated by the clipped StrongEuler algorithm will be produced from the 'prunned' MDP with state space $\mathcal{S}$ and state-horizon dependent action sets $\mathcal{A}(s)$ (i.e. with state-action space Pair ${ }_{\text {eff }}$ ). Strong optimism holds for the estimates $\widehat{Q}_{h}^{t}(s, a)$ for all $(s, a) \in \operatorname{Pair}_{\text {eff }}$ and all time steps $t$ (notice that $\pi_{*}$ is an optimal policy both in the original as well as in the pruned MDP). The exact same proofs as in [16] (see proof of Theorem 3.2) hold in this case with the only modification being to adapt the argument for MDPs where the number of actions is state-dependent thus yielding the following result,
Lemma B.13. The regret of clipped-StrongEuler satisfies with high probability,

$$
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)=\mathcal{O}\left(\sum_{(s, a) \in \operatorname{Pair}_{\mathrm{eff}}} \frac{Q_{h(s)}^{\star}(s, a)}{\operatorname{gap}_{h(s)}(s, a)} \log (T)\right)
$$

where $\operatorname{gap}_{h}(s, a)=V_{h}^{\star}(s)-Q_{h}^{\star}(s, a)$
Stating this result in terms of the return gaps (see Definition 3.1 in [16]) over Pair ${ }_{\text {eff }}$ is possible. Lemma B. 13 recovers the instance dependent results of [21].

## B. 3 Connections with instance dependent rates for reinforcement learning

Although our main results are not instance dependent, our rates can be much better if the gaps over the effective state space are very small and the effective state space is also very small. Moreover, in contrast with instance dependent rates our results depend on an effective state space size that may be much smaller than the original state space. Our main innovation lies in these state pruning results.

## B. 4 Model Selection

Although UCBVI-Shaped requires knowledge of $\beta$, Algorithm 1 can be used to design a parameter free version. As explained in Section 6 we simply require to initialize algorithm 1 with a set of $N$ different $\beta$ parameter guesses $\left[\beta_{1}, \cdots, \beta_{N}\right]$. A good choice for these is an exponential parameter grid of the form $\beta_{1}=1, \cdots, \beta_{N}=2^{N}$.
In the case we are using the Corral algorithm of [34] (and [1]), it is enough to set the learning rates based on a putative regret bound value of $\sqrt{T}$ (see the proof of Theorem 5.3 in [34] where the learning rate can be set to $\eta=\sqrt{\frac{N}{T}}$. The main term in the regret guarantee of Theorem 5.2 scales with $\sqrt{T}$. Thus, for any choice of $\beta_{i}$ if this value of $\beta_{i}$ was valid, the regret rate would scale as $C_{i} \sqrt{T \log (1 / \delta)}$ for some parameter $C_{i} \in\left[1, H \beta_{i} \widetilde{V}^{\max } \sqrt{|\mathcal{S}||\mathcal{A}|}\right]$. If using regret Balancing as an online model selection mechanism (see [33]) we can use an exponential parameter grid of the form $\left\{\beta_{i} \times\left[1,2, \cdots, H \beta_{i} \widetilde{V}^{\max } \sqrt{|\mathcal{S}||\mathcal{A}|}\right\}_{i=1}^{N}\right.$. The number of base algorithms necessary in this case is at most $N \log \left(H \beta_{N} \widetilde{V}^{\max } \sqrt{|\mathcal{S}||\mathcal{A}|}\right)$.
As a consequence of Theorem 5.3 and as a simple corollary of Theorem 5.3 in [34] or Theorem 5.1 in [33] we conclude that,
Lemma B.14. Provided that $\beta_{N} \widetilde{V}_{h}(s) \geq V_{h}^{\star}(s)$ for all $s \in \mathcal{S}$ and $h \in[H]$ (although it could be that $\beta_{i} \widetilde{V}_{h}(s) \geq V_{h}^{\star}(s)$ for all $s \in \mathcal{S}$ and $h \in[H]$ for $\left.i \ll N\right)$, the expected regret of Algorithm 1 satisfies,

$$
\mathbb{E}\left[\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)\right]=\widetilde{\mathcal{O}}\left(C^{2} \sqrt{N T}\right)
$$

when using Stochastic CORRAL as a model selection strategy. Similarly the expected regret of Algorithm 1 satisfies,

$$
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)=\widetilde{\mathcal{O}}\left(C^{2} \sqrt{N \log \left(H \beta_{N} \widetilde{V}^{\max } \sqrt{|\mathcal{S}||\mathcal{A}|}\right) T \log (1 / \delta)}\right)
$$

with probability at least $1-\mathcal{O}(\boldsymbol{\delta})$ when using Regret Balancing as a model selection strategy.
Proof. This result is (almost) an immediate corollary of Theorems 5.3 and 5.1 in [34] and [34] respectively. In order to apply these results to our setting we only need to note that for the optimal value of $\beta$ (among the choices in $\left[\beta_{1}, \cdots, \beta_{N}\right]$ ) there exists a constant $C \in\left[1,2, \cdots, H \beta_{N} \widetilde{V}^{\max } \sqrt{|\mathcal{S}||\mathcal{A}|}\right]$ such that regret rate of UCBVI-Shaped (Theorem 5.2) can be upper bounded as,

$$
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)=\widetilde{\mathcal{O}}(C \sqrt{T \log (1 / \delta)})
$$

with probability at least $1-\mathcal{O}(\boldsymbol{\delta})$ and where $\widetilde{\mathcal{O}}$ may hide polynomial factors in $|\mathcal{S}|$ and $|\mathcal{A}|$.
Applying the in-expectation results of Theorem 5.3 in [34] for the Stochastic CORRAL model selection algorithm yields (setting $\delta \ll \frac{1}{T}$ ),

$$
\mathbb{E}\left[\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)\right]=\widetilde{\mathcal{O}}\left(C^{2} \sqrt{N T \log (1 / \delta)}\right)
$$

Where $\widetilde{\mathcal{O}}$ may hide polynomial factors in $|\mathcal{S}|$ and $|\mathcal{A}|$ and $T$. Similarly, Regret Balancing yields a bound,

$$
\sum_{t=1}^{T} V^{\star}\left(s_{0}\right)-V^{\pi_{t}}\left(s_{0}\right)=\widetilde{\mathcal{O}}\left(C^{2} \sqrt{N \log \left(H \beta_{N} \widetilde{V}^{\max } \sqrt{|\mathcal{S} \| \mathcal{A}|}\right)} T \log (1 / \delta)\right)
$$

with probability at least $1-\mathcal{O}(\boldsymbol{\delta})$.

The pseudocode for this algoritihm is provided below:

```
Algorithm 3 Online UCBVI - Shaped
    Input reward function \(r\) (assumed to be known), set of \(\beta\) values - \(\left[\beta_{1}, \beta_{2}, \ldots, \beta_{N}\right]\)
    Initialize model selection probability \(p(\beta)\) as a uniform distribution over \(\left[\beta_{1}, \beta_{2}, \ldots, \beta_{N}\right]\)
    for \(t=1, \ldots, T\)
            Sample a value of \(\beta \sim p(\beta)\)
            Run UCBVI-Shaped \((\beta)\) with this sampled beta
            Update \(p(\beta)\) using samples from UCBVI-Shaped using an online model selection algorithm.
    End for
```


## C Supporting Technical Lemmas

Lemma C. 1 (Anytime Hoeffding Inequality [33]). Let $\left\{Y_{\ell}\right\}_{\ell=1}^{\infty}$ be a martingale difference sequence such that $Y_{\ell}$ is $Y_{\ell} \in\left[a_{\ell}, b_{\ell}\right]$ almost surely for some constants $a_{\ell}, b_{\ell}$ almost surely for all $\ell=1, \cdots, t$. then

$$
\sum_{\ell=1}^{t} Y_{\ell} \leq 2 \sqrt{\sum_{\ell=1}^{t}\left(b_{\ell}-a_{\ell}\right)^{2} \ln \left(\frac{12 t^{2}}{\widetilde{\delta}}\right)}
$$

with probability at least $1-\widetilde{\delta}$ for all $t \in \mathbb{N}$ simultaneously.
Lemma C. 2 (Freedman Bounded RVs - unsimplified [19, 8]). Suppose $\left\{X_{t}\right\}_{t=1}^{\infty}$ is a martingale difference sequence with $\left|X_{t}\right| \leq b$. Let

$$
\operatorname{Var}_{\ell}\left(X_{\ell}\right)=\operatorname{Var}\left(X_{\ell} \mid X_{1}, \cdots, X_{\ell-1}\right)
$$

Let $V_{t}=\sum_{\ell=1}^{t} \operatorname{Var}_{\ell}\left(X_{\ell}\right)$ be the sum of conditional variances of $X_{t}$. Then we have that for any $\widetilde{\delta} \in(0,1)$ and $t \in \mathbb{N}$

$$
\mathbb{P}\left(\sum_{\ell=1}^{t} X_{\ell}>2 \sqrt{V_{t}} A_{t}+3 b A_{t}^{2}\right) \leq \widetilde{\delta}
$$

Where $A_{t}=\sqrt{2 \ln \ln \left(2\left(\max \left(\frac{V_{t}}{b^{2}}, 1\right)\right)\right)+\ln \frac{6}{\tilde{\delta}}}$ and $h(s)$ corresponds to horizon index of the state partitions that contains state s.
Lemma C. 3 (Empirical Bernstein Anytime, Theorem 4 of [29]). Let $\ell \geq 2$ and $\left\{Z_{t}\right\}_{t=1}^{\ell}$ be i.i.d. random variables with distribution $Z$ satisfying $\left|Z_{t}\right| \leq b$ for all $t \in[\ell]$ Let the sample variance be defined as,

$$
\operatorname{Var}_{\ell}(Z)=\frac{1}{\ell(\ell-1)} \sum_{1 \leq i<j \leq \ell}\left(Z_{i}-Z_{j}\right)^{2}
$$

With probability at least $1-\delta$,

$$
\mathbb{E}[Z]-\frac{1}{\ell} \sum_{i=1}^{\ell} Z_{i} \leq \sqrt{\frac{4 \operatorname{Var}_{\ell}(Z) \ln \frac{4 \ell^{2}}{\delta}}{\ell}}+\frac{7 b \ln \frac{4 \ell^{2}}{\delta}}{3(\ell-1)}
$$

for all $\ell \in[\mathbb{N}]$.

Lemma C. 4 (Lemma 7.5 of [2]). Consider arbitrary $T$ sequence of trajectories $\tau_{t}=\left\{s_{h}^{t}, a_{h}^{t}\right\}_{h=1}^{H}$, for $t=0,1, \ldots, T$. We have

$$
\begin{equation*}
\sum_{t=1}^{T} \sum_{h=1}^{H} \frac{\mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}\right)}{\sqrt{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}} \leq 2 \sqrt{|\mathcal{U}| \sum_{(s, a) \in \mathcal{U}} N_{h(s)}^{T}(s, a)} \leq 2 \sqrt{|\mathcal{U}| T} \tag{19}
\end{equation*}
$$

Where $\mathcal{U} \subseteq \mathcal{S} \times \mathcal{A}$ is an arbitrary set of state action pairs and $h(s)$ corresponds to horizon index of the state partitions that contains state $s$.

Proof. Consider swapping the order of summation

$$
\begin{align*}
& \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{1}{\sqrt{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}}=\sum_{h=1}^{H} \sum_{t=1}^{T} \frac{1}{\sqrt{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}}=\sum_{(s, a) \in \mathcal{U}} \sum_{i=1}^{N_{h(s)}^{T}(s, a)} \frac{1}{\sqrt{i}}  \tag{20}\\
& \leq 2 \sum_{(s, a) \in \mathcal{U}} \sqrt{N_{h(s)}^{T}(s, a)} \leq 2 \sqrt{|\mathcal{U}| \sum_{(s, a) \in \mathcal{U}} N_{h(s)}^{T}(s, a)} \leq 2 \sqrt{|\mathcal{U}| T}
\end{align*}
$$

where the first inequality use the fact that $\sum_{i=1}^{N} 1 / \sqrt{i} \leq 2 \sqrt{N}$ and the second inequality holds due to Cauchy-Schwarz inequality.
Lemma C. 5 (Sum of inverse counts). Consider arbitrary $T$ sequence of trajectories $\tau_{t}=\left\{s_{h}^{t}, a_{h}^{t}\right\}_{h=1}^{H}$, for $t=0,1, \ldots, T$. We have

$$
\begin{equation*}
\sum_{t=1}^{T} \sum_{h=1}^{H} \frac{\mathbf{1}\left(\left(s_{h}^{t}, a_{h}^{t}\right) \in \mathcal{U}\right)}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)} \leq 2|\mathcal{U}| \log (T+1) \tag{21}
\end{equation*}
$$

Where $\mathcal{U} \subseteq \mathcal{S} \times \mathcal{A}$ is an arbitrary set of state action pairs.
Proof. Consider swapping the order of summation

$$
\begin{align*}
& \sum_{t=1}^{T} \sum_{h=1}^{H} \frac{1}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}=\sum_{h=1}^{H} \sum_{t=1}^{T} \frac{1}{N_{h}^{t}\left(s_{h}^{t}, a_{h}^{t}\right)}=\sum_{(s, a) \in \mathcal{U}} \sum_{i=1}^{N_{h(s)}^{T}(s, a)} \frac{1}{i}  \tag{22}\\
& \leq 2 \sum_{(s, a) \in \mathcal{U}} \log \left(N_{h(s)}^{T}(s, a)+1\right) \leq 2|\mathcal{U}| \log (T+1)
\end{align*}
$$

## D Supporting Algorithms

The UCBVI algorithm $[8,2]$ is described below.

```
Algorithm 4 UCBVI
    Input reward function \(r\) (assumed to be known), confidence parameters
    for \(t=1, \ldots, T\)
        Compute \(\widehat{P}_{t}\) using all previous empirical transition data as \(\widehat{\mathbb{P}}_{t}\left(s^{\prime} \mid s, a\right):=\frac{N_{h}^{t}\left(s, a, s^{\prime}\right)}{N_{h}^{t}(s, a)}, \forall h, s, a, s^{\prime}\).
        Compute reward bonus \(b_{h}^{t}(s, a)=2 H \sqrt{\ln (S A H K / \delta) / N_{h}^{t}(s, a)}\)
        Run Value-Iteration on \(\left\{\widehat{\mathbb{P}}_{t}, r+b_{h}^{t}\right\}_{h=0}^{H-1}\).
        Set \(\pi_{t}\) as the returned policy of VI.
    End for
```

UCBVI works by adding a count based bonus to the rewards and then running value iteration over the empirical model based on these bonus augmented rewards. This encourages optimism of the empirical value functions which ensures an appropriate trade-off between exploration and
exploitation. In contrast with previous approaches with provable guarantees for tabular RL problems such as UCRL [6].
As described in [8], the sample complexity of UCBVI is (up to low order terms) $\widetilde{\mathcal{O}}(H \sqrt{|\mathcal{S}||\mathcal{A}|})$ with variance aware bonuses and $\widetilde{\mathcal{O}}\left(H^{3 / 2} \sqrt{|\mathcal{S} \| \mathcal{A}| T}\right)$ when the bonus terms are computed using simple count based scores as above. Our results are based on a modification of the rates derived in [2] that are of order $\widetilde{\mathcal{O}}\left(H^{2}|\mathcal{S}| \sqrt{|\mathcal{A}| T}\right)$. Despite this base rate having a suboptimal dependence in $H$ and $|\mathcal{S}|$ in contrast with the minimax rate of [8] our results still represent an big improvement w.r.t. the un-shaped rates of [8] when the effective state space is small.

## E Experimental Details

We conducted experiments building on the UCBVI framework outlined in [8]. Specifically, we used the bonus 1 formulation from the UCBVI Algorithm 4 in [8], with the bonus specifically being $\frac{C}{\sqrt{N(s, a)}}$ as the base UCBVI implementation. Our implementation is based on the open source implementation by Ian Osband https://github.com/iosband/TabulaRL, and open source code is attached.
The implementations of UCBVI-BS replaces this bonus by $\frac{C . \tilde{V}}{\sqrt{N(s, a)}}$. In our implementation we chose the scaling to be 0.1 with a hyperparameter sweep. The $\widetilde{V}$ was chosen by scaling the optimal value function $V^{\star}$ by factors chosen between $1+\beta$ and $1-\beta$, with beta as $0.2,0.5,0.9$ (as described in the experimental section).

The environments themselves are open gridworlds of size 8 , a corridor of size $10 x 10$ and the double corridor has size $10 \times 20$. The reward is 0 everywhere except the goal location. Every experiment is averaged over 3 random seeds.

## F Intuition for PseudoSub ${ }_{\Delta}$ and PathPseudoSub ${ }_{\Delta}$

Here we provide some intuition for the concepts of $\mathrm{PseudoSub}_{\Delta}$ and PathPseudoSub ${ }_{\Delta}$ in a deterministic chain environment (actions being left and right with deterministic transitions). The environment has a reward of 1 at state 0 and 0 elsewhere. Discount factor is 0.8 here, and the $\widetilde{V}$ has a sandwich factor $\beta=0.3$. The true optimal value function is shown at the top of Fig. 8. The reward shaping is shown in the middle row of Fig. 8. The concept of pseudosub is shown via the arrows at $(s, a)$ pairs that lie in PseudoSub ${ }_{\Delta}$ for this particular choice of $\widetilde{V}$ and $\Delta$. They basically denote states which are suboptimal with $\Delta$ confidence according to the shaping function $\widetilde{V}$. Now given this definition of PseudoSub ${ }_{\Delta}$, the corresponding PathPseudoSub ${ }_{\Delta}$ is shown in the last row. These are states that can be eliminated via reward shaping. As we can see, around half of the states can be eliminated this way, making exploration much more directed.

## G Results for Corrupted $\widetilde{V}$

We evaluate the behavior of UCBVI-Shaped as we use a corrupted version of $\widetilde{V}$ for the shaping in Fig. 9. In particular, we do this by constructing the sandwiched value function $\widetilde{V}$ as described in our experimental evaluation (in this case with $\beta=1.5$ in the corridor environment) and then adding gaussian noise to corrupt the value function. We experiment with corruptions of $0 ., 0.1,0.5,1.0$ variance, zero-centered gaussian noise. The resulting regret is seen in Fig. 9. As expected, the results show that corruption of $\widetilde{V}$ hurts the results as more and more corruption happens, but still performs better than without any shaping at all.

## H Results for RND

We also ran numerical simulations (in Fig. 10) with a neural network based "pseudo-count" method random network distillation (RND) [12]. This trains a neural network against a random pre-initialized


Figure 8: Illustration of the concepts of $\mathrm{PseudoSub}_{\Delta}$ and $\mathrm{PathPseudoSub}_{\Delta}$. Top row shows the optimal value function for this chain MDP (colorbar shows scale). Middle row shows the $\tilde{Q}$ obtained from the shaping and the $(s, a)$ pairs which are in $\mathrm{PseudoSub}_{\Delta}$ (denoted by arrows). From PseudoSub ${ }_{\Delta}$, the bottom row shows states in PathPseudoSub ${ }_{\Delta}$ which can be pruned out via shaping.


Figure 9: Illustration of behavior with corrupted shaping functions $\widetilde{V}$ with varying levels of corruption
target and uses projection error as a "pseudocount". We ran UCBVI-Shaped-BS with the inverse counts typical in UCBVI-shaped replaced with model error. The results in Fig. 10 are considerably slower than with exact counts likely due to noise in counts estimation, although the results are likely to be improved with environment specific tuning.


Figure 10: Illustration of behavior of UCBVI-Shaped-BS with RND based pseudocounts instead of exact counts.

## I Discussion of Applicability of Reward Shaping Assumption

There can't be any way to obtain an improvement to the UCBVI regret bounds by using a reward shaping mechanism unless there is some information about the optimal policy / optimal value function encoded in the shaping function. In other words, there cannot be any free lunch in this setting. Prior knowledge is fundamental for reward shaping to be successful.
In this work we have chosen to make the assumption that the side information available to us $\widetilde{V}$ can be used to sandwich the value of $V^{\star}$. There are many examples where this assumption may hold.

1. Maze like environments. In any goal metrized goal oriented environment (for example a maze) with a stay-in-place action, if the position of the goal state is (roughly) known and the reward function equals -1 at every step from a state to a neighboring state and 0 for the stay-in-place action at the goal state, using the metric distance from the current state to the goal state may be a good estimator of the number of steps needed to reach the goal $\left(-V^{\star}\right)$. The accuracy of this approximation will depend on how off this heuristic is.
2. Sim to Real. We may think of training a robot in a simulation environment and use the learned optimal value function in the simulation as our input $\widetilde{V}$. Having this information at hand should make training in real much easier than from scratch depending on how accurately the $\widetilde{V}$ estimator from the simulation tracks $V^{*}$. Under the assumption that dynamics are lipschitz bounded, and there is bounded state estimation error, dynamics estimation error, the resulting value function also shares the sandwich property of the true value function.
3. Object manipulation. In this setting even if the objective is to manipulate a huge combination of objects, an appropriate combination of euclidean distances of objects can serve as a good proxy for $V^{\star}$. While this may underestimate the distance to the goal especially in the presence of obstacles, in most cases a sandwich term $\beta$ can be found that bounds the true value function (under the assumption that the number of obstacles is bounded). This becomes particularly pronounced as the number of objects increases and there is a combinatorially large state space, all of which cannot possibly be explored.
4. Motion planning: Recent methods have attempted to combine motion planning and RL [18]. One such instantation would involve using an anytime motion planning to generate a reward function to guide RL. An anytime motion planner may start by generating suboptimal paths, and gets better over time. If the motion planning is run only for some time so as to get a beta-sandwiched suboptimal path (and hence reward shaping as a sandwich on $V^{\star}$ ) to guide RL, this would then provide suboptimal shaping that can then be improved with subsequent reinforcement learning as described in our work.
5. First person RPG: For games like first person RPGs, reward shaping terms may look like a sum of euclidean distances to different objects of interest. This will underestimate the true value function will sandwich the optimal value function depending on how many obstacles are in the environment. This environment will also likely benefit from pruning off large parts of the state space.

In all of these situations, it is possible to construct a multiplicative approximation to $V^{\star}$ to be used as $\widetilde{V}$.

## J Limitations

A limitation of our work is that the reward shaping has to be provided before hand rather than learned or inferred, which would be very interesting to explore in future work. Additionally, the shaping bound we have is a point-wise one rather than a correlation based bound in expectation, which can be susceptible to be outlier rewards which have a large sandwich term $\beta$. This should be investigated in future work in more detail. Additionally, it is likely that reward shaping can directly allow for horizon reduction in value computation, which should be considered in future work more directly. The behavior of these types of methods under function approximation also may be quite different, which should be studied more systematically.

(a) Quantitative comparison of effect of (b) Visualization of non decaying agent get-(c) Visualization of decaying agent bonus decay ting stuck (UCBVI-Shaped) reaching goal
Figure 11: Effect of bonus decay on the performance of UCBVI-Shaped vs standard reward shaping under shaping misspecification. As we can see, without decay the agent can stuck in arbitrarily sub-optimal points, whereas with decay the agent easily converges to an optimal solution.

## J. 1 Does decaying suboptimal reward shaping help over standard shaped rewards?

We ran simulations to understand whether the adaptive decay of $\widetilde{V}$ by $\frac{1}{\sqrt{N(s, a)}}$ actually provides tangible benefits, specifically when the shaping is suboptimal. We compared UCBVI-Shaped with a variant where $\widetilde{V}$ is added to the reward linearly, as done in many practical RL implementations. As seen in Fig. 11, with suboptimal shaping, UCBVI-shaped dampens the effect of the suboptimal shaping and succeed, whereas simple addition of shaping leads to the agent becoming trapped. This suggests that not only can reward shaping improve sample efficiency, but incorporating reward shaping via UCBVI-Shaped can mitigate the potential bias from suboptimal shaping.


[^0]:    *Equal Contributions

[^1]:    ${ }^{2}$ Our main result, Theorem 5.2, is slightly more complex than this statement. We have chosen this simplified form to aid the reader to form the right intuition.

[^2]:    ${ }^{3}$ The $\widehat{\mathbb{P}}_{t}$ here is dependent on the horizon $h$, but since we have assumed (Assumption 3) the states $s$ are $h$-indexed, we will use $\widehat{\mathbb{P}}_{t}$ for notation simplicity.

[^3]:    ${ }^{4}$ We add the qualifier pseudo to our name to distinguish between suboptimality as captured by our surrogate $Q$ values and true suboptimality between the true values of $Q_{h}^{\star}$.

[^4]:    ${ }^{5}$ Recall that by Assumption 3 the states are assumed to be $h$-indexed and therefore the state space can be written as $\mathcal{S}=\cup_{h \in[H]} \mathcal{S}_{h}$. This also implies that $N_{h}^{t}(s)=0$ for all $h \neq h(s)$.

[^5]:    ${ }^{6}$ From now on we'll use $T$ to index the final timestep of the regret sequence we are bounding.

[^6]:    ${ }^{7}$ Instead of using $|\mathcal{S}|$ to uniformly bound the support of $\mathbb{P}^{\star}(\cdot \mid s, a)$ we explicitly write the bound in terms of its support

[^7]:    ${ }^{8}$ Recall we have assumed that $\mathcal{S}$ is $h$ indexed so that any state $s \in \mathcal{S}$ can be accessed only during in-episode time-step $h(s)$.

