
Stochastic Shortest Path with Sparse Adversarial Costs

Emmeran Johnson^{1*} Alberto Rumi^{2,3*} Ciara Pike-Burke¹ Patrick Rebeschini⁴

¹ Department of Mathematics, Imperial College London

² Università degli Studi di Milano

³ CENTAI Institute

⁴ Department of Statistics, University of Oxford

Abstract

We study the adversarial Stochastic Shortest Path (SSP) problem with sparse costs under full-information feedback. In the known transition setting, existing bounds based on Online Mirror Descent (OMD) with negative-entropy regularization scale with $\sqrt{\log SA}$, where SA is the size of the state-action space. While we show that this is optimal in the worst-case, this bound fails to capture the benefits of sparsity when only a small number $M \ll SA$ of state-action pairs incur cost. In fact, we also show that the negative-entropy is inherently non-adaptive to sparsity: it provably incurs regret scaling with $\sqrt{\log S}$ on sparse problems. Instead, we propose a family of ℓ_r -norm regularizers ($r \in (1, 2)$) that adapts to the sparsity and achieves regret scaling with $\sqrt{\log M}$ instead of $\sqrt{\log SA}$. We show this is optimal via a matching lower bound, highlighting that M captures the effective dimension of the problem instead of SA . Finally, in the unknown transition setting the benefits of sparsity are limited: we prove that even on sparse problems, the minimax regret for any learner scales polynomially with SA .

1 Introduction

The Stochastic Shortest Path (SSP) problem is a fundamental model in reinforcement learning [1, 32], which describes tasks where an agent interacts with an environment over episodes and must reach a designated goal state within each episode while minimizing accumulated costs. This covers problems such as car navigation while trying to avoid traffic jams, or internet routing. Recently, this classical setting has been extended to the adversarial regime, where costs may vary arbitrarily between episodes [28, 8, 6, 38] and the goal is to obtain theoretical guarantees robust to any cost-generation mechanism. Under *full-information* feedback where the full cost vector is observed after each episode and known transitions, current algorithms achieve regret bounds that scale as $\mathcal{O}(\sqrt{DKT_\star \log SAT_\star})$, where D is the diameter (the smallest expected hitting time of any policy from any state), $T_\star$ is the expected hitting time of the optimal policy, S is the number of states, A is the number of actions, and K is the number of episodes. These bounds are independent of any cost structure and are shown to be minimax optimal up to logarithmic factors in [8].

In SSP, the size of the state-action space SA – which we consider and refer to as the *dimension* of the problem – appears in the minimax regret as $\sqrt{\log SA}$. While in the worst-case this is unimprovable (we show this in Theorem 3.1), many real-world problems have costs with structural properties that may be leveraged for improved regret. A common property often considered in the statistics and machine learning literature [36] is sparsity, which can naturally arise for SSP problems. For instance, in the car navigation example, the number of traffic jams is usually much smaller than the number of roads. Motivated by this, we consider *sparse* SSP problems where M , the maximum number of state-action pairs with non-zero cost in an episode, can be much smaller than SA .

In such scenarios, the regret bounds should capture some dependence on M , reflecting an improvement in performance on easier sparse problems as $M \rightarrow 1$ and recovering the standard bounds on

*Equal contribution, corresponding author: emmeran.johnson17@imperial.ac.uk

worst-case problems as $M \rightarrow SA$. In fact, for the so-called experts setting² ($S = 1$), the minimax regret scales with $\sqrt{KMA^{-1} \log A}$ instead of $\sqrt{K \log A}$ in the worst-case [19], providing a polynomial improvement in the dimension A . Furthermore, this is achieved with Online Mirror Descent (OMD) with negative entropy regularization. However for SSP where the problem dimension also includes the size of the state space S , we show that existing approaches also based on OMD with the same negative entropy regularization [8] fail to exploit sparse costs. We construct a sparse SSP problem where this algorithm suffers $\sqrt{\log S}$ regret (Theorem 3.1), providing no improvement in terms of S compared to the minimax regret for the non-sparse worst-case $M = SA$ problem. The failure of existing SSP methods to exploit sparsity, let alone match the polynomial improvements from the experts setting, leads us to ask the following questions:

Does sparsity improve the minimax-regret in the full-information feedback SSP problem? How much?

We answer the first question positively for the known-transition case by designing a family of regularizers based on ℓ_r -norms for $r \in (1, 2)$ for which we show a $\mathcal{O}(\sqrt{DKT_\star \log MT_\star})$ regret bound that depends logarithmically on the sparsity level M , rather than on the size of the state-action space SA , without requiring the knowledge of M in advance. This family of regularizers interpolates between the negative entropy and the squared Euclidean norm (see Section 4), allowing flexibility for much weaker regularization on sparse points in sparse settings and recovering existing algorithms (and guarantees) in the non-sparse setting.

We show that the above $\sqrt{\log M}$ dependence on M is unimprovable by constructing a matching lower bound. Interestingly, this establishes that the benefit of sparsity in SSP is logarithmic in SA instead of polynomial, as in the simpler experts problem, thus answering the second question. It also highlights that M plays the role of effective dimension, replacing the general dimension SA in controlling the scaling of the minimax regret.

While the benefits of sparsity in SSP are only logarithmic, we emphasize that due to the often combinatorial nature of the state-action space, these improvements can be significant. For instance, in many problems the size of the state-action space grows exponentially in some parameters, while the assigned costs remain linear or even constant. This occurs in many real-world problems (e.g. [31, 30]) in which settings exploiting sparsity can lead to polynomial improvements.

Finally, it is natural to ask whether sparsity may provide similar benefits in the unknown transition setting. However, in Theorem 5.1, we show a lower bound with polynomial dependence on SA in a sparse SSP instance with unknown transitions. This illustrates that in the unknown transition setting the sparsity level M does not play the same role of effective dimension, and that the general dimension SA is crucial in controlling the scaling of the minimax regret polynomially, motivating our focus on the known transitions setting. In particular, this result shows that sparse problems with combinatorial state-action spaces will remain very challenging.

Our results provide a complete characterisation of the benefits of sparsity in removing dimension dependence (i.e. SA) for adversarial SSP problems under full-information feedback.

1.1 Contributions

We highlight our main contributions below:

- We design a family of ℓ_r -norm ($r \in (1, 2)$) regularizers for OMD that allows interpolation between the negative entropy and squared Euclidean norm, adjusting its geometry to the sparsity of the cost functions (see Section 4). The regularizer naturally plugs into the standard OMD analysis.
- We show OMD with the above regularizer achieves sparsity-adaptive regret bounds of order $\mathcal{O}(\sqrt{DKT_\star \log MT_\star})$ (Theorem 4.1). We also give a parameter-free version achieving the same bound (Theorem 4.4) that does not require prior knowledge of the sparsity level M nor the expected hitting time of the optimal policy $T_\star$ (the only unknown parameters).
- We establish a lower bound of order $\Omega(\sqrt{DKT_\star \log M})$ (Theorem 4.6), matching our regret guarantees up to a logarithmic factor of $T_\star$ (already present in prior work [8]) and improving over [8] in the $M = SA$ non-sparse setting by including the $\sqrt{\log SA}$ dependence.

²The experts setting is the single-state full-information feedback A -action online learning problem [5, 14].

- We show that OMD with the negative entropy used in prior work [28, 8] suffers regret at least $\Omega(\sqrt{K \log S})$ even when $M = 3$ (Theorem 3.1). This rules out the negative entropy as a viable regularizer in the sparse setting and provides justification for the use of our regularizer.
- We establish that results independent of SA are not achievable in the unknown transitions setting via a lower bound in the sparse ($M = 1$) setting of order $\Omega(D\sqrt{SAK})$ (Theorem 5.1).

Technical Contributions: Proving these results requires new technical ideas. For the general sparse lower-bound, we derive a result on the expectation of the maximum of asymmetric zero-mean random walks, generalizing the result for the symmetric case from [25]. The negative-entropy-specific lower-bound relies on the careful design of an MDP with skewed initial occupancy measures that highlights both the reasons for the failure of the negative entropy as well as the more general difficulty of the stochastic nature of SSP problems.

1.2 Related works

Regret minimisation for SSP problems under full-information feedback was initiated by a line of work studying stochastic costs [32, 29, 33, 11, 7, 9, 17]. In the adversarial setting, it was first studied by [28] in the known transition case. Their bounds were later improved by [8]. There have since been many extensions: [6] consider the unknown transition setting, [38] establish dynamic regret bounds, [10] consider a policy optimisation approach in the unknown transition setting.

Regret minimisation for SSP problems under bandit feedback where only the costs of the visited state-action pairs in an episode are revealed to the learner has also been studied both in the stochastic [10] and adversarial settings [8, 6, 10]. In the adversarial known transition setting, the minimax regret is of the order $\sqrt{KDT_*SA}$ (ignoring log terms) [8]. It is an interesting future direction to study the sparse SSP problem with bandit feedback and understand if the regret scales with $\sqrt{M}$ instead of $\sqrt{SA}$, in which case M would play the same role of effective dimension as in the setting we consider.

Regret minimisation with sparse costs was studied in the classical online learning setting [19] ($S = 1$). The minimax regret goes from $\mathcal{O}(\sqrt{K \log A})$ to $\mathcal{O}(\sqrt{KMA^{-1} \log A})$ under full-information feedback (experts problem). For rewards instead of costs, it goes from $\mathcal{O}(\sqrt{K \log A})$ to $\mathcal{O}(\sqrt{K \log M})$, which matches the benefits of sparsity we establish for the SSP problem with costs. Note that we restrict our focus to costs since it is unclear how to interpret rewards within the SSP framework. Under bandit feedback, the sparse minimax regret goes from $\tilde{\mathcal{O}}(\sqrt{KA})$ to $\tilde{\mathcal{O}}(\sqrt{KM})$ for both rewards and costs [19, 4]. The above minimax regrets can also be achieved by sparse-agnostic methods [19, 34]. Finally, sparsity was also considered in the case of stochastic losses by [20].

2 Preliminaries

2.1 Problem setting

We consider the *Stochastic Shortest Path* (SSP) problem with adversarial costs. The environment is modeled as a Markov Decision Process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, s_0, g)$ along with a sequence of cost functions $\{c_k\}_{k=1}^K$ chosen by an oblivious adversary over K episodes. $\mathcal{S}$ is the state space with cardinality $S = |\mathcal{S}|$, and $s_0 \in \mathcal{S}$ is the fixed starting state. The goal state g is a special absorbing state not included in $\mathcal{S}$. $\mathcal{A}$ is the action space with cardinality $A = |\mathcal{A}|$ and we assume for simplicity that it is the same in every state. Let $\Gamma = \mathcal{S} \times \mathcal{A}$ denote the set of all state-action pairs. The dynamics in the MDP are given by the *known* transition function P , where $P(s'|s, a)$ specifies the probability of moving to state $s' \in \mathcal{S} \cup \{g\}$ after taking action a in state s .

Each episode begins in state s_0 and proceeds with the learner selecting actions until the goal state g is reached. When the goal state is reached, the current episode ends and a new one begins. At the start of each episode k , the adversary selects a cost function $c_k : \Gamma \rightarrow [0, 1]$, assigning a cost to each state-action pair. We denote the sparsity level as $M = \max_k \sum_{(s,a) \in \Gamma} \mathbb{I}\{c_k(s, a) > 0\}$ the maximum number of non-zero costs in an episode. We work in the full-information setting where the entire function c_k is revealed to the learner at the end of the episode. The objective is to minimize the total cost over all episodes, which requires a balance of minimizing the accumulated costs while ensuring the goal state is reached efficiently.

We use super-scripts to denote the time-step within an episode and sub-scripts to denote the episode: e.g. (s_k^t, a_k^t) refers to the state-action pair at the t -th time-step of the k -th episode. We sometimes omit the sub-script when referring to an arbitrary episode. We now define some important concepts:

- A **stationary policy** π is a mapping such that $\pi(\cdot|s)$ is a probability distribution over the choice of action $a \sim \pi(\cdot|s)$ in state s . A policy is called **proper** if it reaches the goal g in finite time from any initial state in $\mathcal{S}$ with probability one, and **improper** if not. Let Π_p be the set of all stationary proper policies. We assume the existence of at least one proper policy.
- The **expected hitting time** $T^\pi(s)$ is the expected number of steps required to reach g from state s under π . Letting $I_\pi(s)$ be the random number of time-steps used to reach the goal state when executing a policy π in an episode starting from state s , then $T^\pi(s) = \mathbb{E}[I_\pi(s)]$. For any proper policy π , $I_\pi(s)$ and $T^\pi(s)$ are finite for all $s \in \mathcal{S}$.
- The **fast policy** π_f is the deterministic policy that minimizes the worst-state expected hitting time, and the **diameter** D of the MDP is the corresponding expected hitting time:

$$\pi_f = \arg \min_{\pi \in \Pi_p} \max_{s \in \mathcal{S}} T^\pi(s), \quad D = \max_{s \in \mathcal{S}} T^{\pi_f}(s) = \max_{s \in \mathcal{S}} \min_{\pi \in \Pi_p} T^\pi(s).$$

Since the transition function P is known, both the fast policy π_f and the diameter D can be computed offline prior to the learning process. We assume $D \geq 1$.

- The **cost-to-go** function $J_c^\pi : \mathcal{S} \rightarrow [0, \infty)$ is the expected cost suffered during an episode executing policy π and starting from state s , given a cost function c and a proper policy π . It is defined as

$$J_c^\pi(s) = \mathbb{E} \left[\sum_{t=1}^{I_\pi(s)} c(s^t, a^t) \middle| P, \pi, s^1 = s \right],$$

where the expectation is with respect to the randomness in the action sampling and state transitions. We use J_k^π to denote the cost-to-go from the initial state s_0 using the cost function c_k in episode k .

- The **regret** R_K is the primary measure of performance by which the learner is evaluated. It is the difference between the total cost over all episodes of the policies $\pi_1, \dots, \pi_K$ chosen by the learner, and the total cost of the best proper deterministic policy in hindsight, $\pi^* \in \arg \min_{\pi \in \Pi_p} \sum_{k=1}^K J_k^\pi$:

$$R_K = \sum_{k=1}^K \sum_{t=1}^{I_{\pi_k}(s_0)} c_k(s_k^t, a_k^t) - \sum_{k=1}^K J_k^{\pi^*}.$$

- The **occupancy measure** $q_\pi \in \mathbb{R}_{\geq 0}^\Gamma$ of a proper policy π is the expected number of visits to state-action pairs in an episode executing policy π starting from s_0 :

$$q_\pi(s, a) = \mathbb{E} \left[\sum_{i=1}^{I_\pi(s_0)} \mathbb{I}\{s^i = s, a^i = a\} \middle| P, \pi, s^1 = s_0 \right].$$

The marginal $q_\pi(s) = \sum_{a \in \mathcal{A}} q_\pi(s, a)$ gives the expected number of visits to state s . Given a vector $q \in \mathbb{R}_{\geq 0}^\Gamma$, if it corresponds to a valid occupancy measure, the corresponding policy π_q can be recovered via normalization as $\pi_q(a|s) = q(s, a) / \sum_{a'} q(s, a')$ [39, 28].

2.2 SSP as online linear optimisation and online mirror descent

Occupancy measures allow the cost-to-go to be expressed in a linear form:

$$J_k^\pi = \sum_{(s,a) \in \Gamma} q_\pi(s, a) c_k(s, a) = \langle q_\pi, c_k \rangle.$$

If the learner executes a stationary proper policy π_k in episode k , the expected regret can thus be reformulated as an online linear optimisation problem on the space of occupancy measures:

$$\mathbb{E}[R_K] = \sum_{k=1}^K \left\{ J_k^{\pi_k} - J_k^{\pi^*} \right\} = \sum_{k=1}^K \langle q_{\pi_k} - q_{\pi^*}, c_k \rangle.$$

Online linear optimisation is a well studied problem and can be solved using Online Mirror Descent (OMD) (see e.g. [24]). In the SSP framework, OMD is applied on the space of occupancy measures corresponding to proper policies with expected hitting time bounded by some $T > 0$ defined as:

$$\Delta(T) = \left\{ q \in \mathbb{R}^\Gamma : \sum_{(s,a) \in \Gamma} q(s, a) \leq T, \quad \forall s \in \mathcal{S} : \sum_{a \in \mathcal{A}} q(s, a) - \sum_{(s', a') \in \Gamma} P(s | s', a') q(s', a') = \mathbb{I}\{s = s_0\} \right\}.$$

The first constraint ensures the expected hitting time is bounded by T , while the second is a flow constraint ensuring the vector corresponds to the occupancy measure of a policy. The regret bounds of OMD will hold against any fixed comparator policy as long as T is large enough such that $\Delta(T)$ contains the occupancy measure of the optimal policy, i.e. $q_{\pi^*} \in \Delta(T)$ or $T \geq T_*$ where we denote by $T_* = T^{\pi^*}(s_0)$ the expected hitting time of π^* . OMD with a strictly convex differentiable regularizer ψ and step-size η selects occupancy measures computed through the update

$$q_1 = \arg \min_{q \in \Delta(T)} \psi(q), \quad q_{k+1} = \arg \min_{q \in \Delta(T)} \left\{ \eta \cdot \langle q, c_k \rangle + D_\psi(q, q_k) \right\}, \quad (1)$$

where $D_\psi(x, y) = \psi(x) - \psi(y) - \langle \nabla \psi(y), x - y \rangle$ is the Bregman divergence with respect to ψ . This update can be computed efficiently for all the regularizers we will discuss (see Appendix B). As discussed in the previous section, we can easily recover via normalization the corresponding policy π_{q_k} that will be executed by the learner.

If the regularizer satisfies for some $\alpha > 0$, any $q \in \mathbb{R}^\Gamma$ and all $k \geq 1$:

$$\nabla \psi(q) \in [\nabla \psi(q_k), \nabla \psi(q_k) - \eta c_k] \implies \nabla^2 \psi(q) \geq \alpha \nabla^2 \psi(q_k), \quad (2)$$

(this is satisfied by many common regularizers), then a standard result (see e.g. Theorem 6 in [4], Theorem 5.5 in [3]) gives the following general regret bound for OMD:

$$\sum_{k=1}^K \langle q_k - q_{\pi^*}, c_k \rangle \leq \underbrace{\frac{\psi(q_{\pi^*}) - \psi(q_1)}{\eta}}_{\text{Penalty}} + \underbrace{\frac{\eta}{2\alpha} \sum_{k=1}^K \|c_k\|_{\nabla^2 \psi(q_k)^{-1}}^2}_{\text{Stability}} \quad (3)$$

where $\|q\|_A^2 = \sum_{s,a,s',a'} q(s,a)A((s,a),(s',a'))q(s',a')$ for a matrix $A \in \mathbb{R}^{\Gamma \times \Gamma}$. Various regret bounds can be obtained by instantiating the above with different regularizers. In particular, [8] use the negative-entropy to obtain a $\mathcal{O}(\sqrt{DKT_* \log SAT_*})$ regret bound.

3 Failure of negative entropy regularization

In the general non-sparse setting, [8] use the negative entropy to achieve a regret of $\mathcal{O}(\sqrt{DKT_* \log SAT_*})$, which in the non-sparse setting has optimal dependence on SA (as we show later in Theorem 4.4). Despite this success, the negative entropy fails to benefit from sparsity in its dependence on S , as shown by the result below. As we will see in Section 4, this establishes the negative entropy as a sub-optimal choice for sparse SSP problems.

Theorem 3.1. *For any $S \geq 6$, there exists an SSP instance with a fixed horizon of 3, sparsity level $M = 3$, an action space of size $A = 2$ and state space of size S such that the regret of OMD (1) with negative-entropy regularization and any step-size $\eta > 0$ after K episodes is $\mathbb{E}[R_K] = \Omega(\min\{\sqrt{K \log S}, K\})$.*

This result shows that despite the SSP instance being sparse ($M = 3$), the regret of OMD with negative entropy regularization nevertheless scales as $\sqrt{\log S}$, which is the same dependence on S as in the non-sparse setting. For sparse problems, the negative entropy provides no improvement on the regret with respect to S . This highlights that existing approaches and regularizers are inadequate to appropriately exploit sparse problems and motivates considering alternate regularizers specifically designed for the geometry of sparse problems, as we do in the next section.

To better understand the failure of negative entropy regularization in sparse settings, we highlight the main intuition behind the lower bound construction and defer the details of the proof to Appendix A.

Proof intuition: The key idea is to reduce SSP to an experts problem with 2 actions and a heavily skewed initial distribution over the actions. The initial occupancy measure played by OMD in (1) is $q_1 = \arg \min_{q \in \Delta(T)} \psi(q)$. For most regularizers, including the negative entropy, this encourages q_1 to be uniform across the state-action space while maintaining the constraints on the flow and expected hitting time. Since we consider a fixed-horizon MDP, only the flow constraint is relevant.

Consider the SSP problem shown in Figure 1 with $N = S - 2$. Since $s_1, \dots, s_N$ constitute a large majority of the states (especially for large N), $\psi(q)$ is mainly affected by the values of q in these




Figure 1: MDP for the reduction to a skewed experts problem with 2 actions: $\mathcal{S} = \{s_0, s_g, s_1, \dots, s_N\}$ ($N = S - 2$), $\mathcal{A} = \{a_1, a_2\}$. The transitions are given by $p(s_g|s_0, a_1) = 1, p(g|s_g, a) = 1$ for all $a \in \mathcal{A}$, for $i \geq 1$: $p(s_i|s_0, a_2) = 1/N, p(g|s_i, a) = 1$ for all $a \in \mathcal{A}$.

states. In order to minimize $\psi(q)$, q_1 needs to ensure the expected number of visits to these N states is sufficiently high. However since for any q and any $i \geq 1$, $q(s_i) = \frac{1}{N}q(s_0, a_2)$, for $q_1(s_i)$ to be sufficiently large then $q_1(s_0, a_2)$ needs to be much larger (by a factor of N). This results in q_1 being heavily skewed towards a_2 in s_0 . For the negative entropy, this gives specifically $q_1(s_0, a_1) \approx \frac{1}{\sqrt{N}}$.

If the costs in all states but s_0 are set to 0, the problem is sparse ($M = 2$) and reduces to a experts problem with 2 actions where the initial probability for the first action, which in our case is $q_1(s_0, a_1)$, scales as $1/\sqrt{N}$. The regret for OMD with the negative entropy in this setting can be shown to scale for any step-size at least as $\Omega(\sqrt{K \log N}) = \Omega(\sqrt{K \log S})$, providing the dependence from the statement of the theorem. To prove this formally for the SSP reduction, we use the above construction coupled with a non-skewed reduction and careful setting of the costs. We include the details in Appendix A.

Finally, we remark that this failure comes from the negative entropy stretching euclidean distance near the boundary of the space in such a way that two nearby points in terms of euclidean distance can be arbitrarily far in terms of negative entropy. This makes it hard for OMD to recover from the initial occupancy measure $q_1(s_0, a_1) \approx 1/\sqrt{N}$ ($\rightarrow 0$ as N increases) unless the step-size is unreasonably large. This property does not generalize to all regularizers and in fact provides insights for designing a regularizer to appropriately handle sparsity. In particular, the regularizer we consider in the next section does not suffer from the same issue because the stretching of euclidean distance is finite since its gradient does not diverge at the boundary (i.e. as $q(s, a) \rightarrow 0$) unlike the negative entropy.

4 The benefits of sparsity

In this section, we show that it is possible to achieve a regret bound of order $\mathcal{O}(\sqrt{DKT_\star \log(MT_\star)})$, where M is the maximum number of non-zero entries in the cost. This is our main result and together with the lower bound in Theorem 4.6 establishes that the sparsity level M acts as a measure of effective dimension instead of the state-action space size SA for SSP with full-information feedback.

In the previous section, we showed and discussed that the negative entropy, the regularizer used in OMD by existing methods, is inadequate to handle sparse SSP problems. Motivated by this failure, we consider alternate regularizers. However, identifying a suitable regularizer poses two key challenges. Firstly, it must work for SSP and the associated technical complexities compared to other simpler online learning problems. In particular, it needs to match the dependence in terms of the other non-sparsity-related quantities appearing in the regret of the negative entropy (i.e. $D, T_\star, K$). Second, it must explicitly leverage sparsity to improve performance. We propose the following family of regularizers parameterised by $p > 1$:

$$\psi_p(q) = p \cdot \left(-1 + \|q\|_{1+1/p}^{1+1/p} \right) = p \cdot \left(-1 + \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} |q(s, a)|^{1+1/p} \right). \quad (4)$$

As $p \rightarrow \infty$, the regularizer in Equation (4) converges to the negative entropy. On the other hand, as $p \rightarrow 1$, the regularizer converges to the squared Euclidean norm that enforces much weaker regularization on sparse points. Therefore, ψ_p allows smooth interpolation between dense and sparse regimes via the tunable parameter p (see Figure 2 for a comparison). In particular, ψ_p for small p




Figure 2: Bregman Divergence between a deterministic distribution $x = [0, 1]$ and the uniform distribution $y = [1/2, 1/2]$ for our regularizer ψ_p , squared Euclidean norm ψ_E and negative entropy H for increasing values of p .

induces an OMD update that is able to easily move to and away from the boundary of the space, making it robust to the skewed initial occupancy measure on the SSP instance which caused the failure of the negative entropy in Section 3. The parameter p also controls a trade-off between the stability and penalty term in (3), which ultimately will enable the removal of the dependence on SA .

Versions of this family of regularizers can be found in the convex optimization literature [18, 22]. As far as we are aware, its use with OMD is novel. A regularizer involving an r -norm with $r \in (1, 2]$ has been used but the norm is squared rather than to the power of r (see e.g. Section 6.7 in [24]). Our regularizer is also similar in flavor to the Tsallis-entropy in the sense that it converges to the negative entropy in the limit of its parameter.

We note that OMD with the above regularizer can be implemented efficiently for any p : the projection step of OMD over $\Delta(T)$ can be written as a convex optimization problem as in [28], which can be solved efficiently (details in Appendix B).

We can now turn to our main result, which establishes regret bounds that scale with the sparsity level M for OMD with the regularizer in (4) when M is known.

Theorem 4.1. *Consider OMD with ψ_p as regularizer. If $T > e$ is such that $q_{\pi^*} \in \Delta(T)$, $\eta = \sqrt{\frac{pT^{1+1/p}}{KDM^{1/p}}}$, $p = \log(TM)$, then $\mathbb{E}[R_K] \leq \mathcal{O}\left(\sqrt{DKT \log(MT)}\right)$.*

We present below the outline of the proof and include the missing details in Appendix C.

Proof. It can be shown that ψ_p satisfies the condition (2) with $\alpha = 1$, allowing us to use the bound in (3) as a starting point. Using that $q_\star \in \Delta(T)$, we can bound the Penalty term:

$$\psi_p(q_\star) - \psi_p(q_1) = p \left(\|q_\star\|_{1+1/p}^{1+1/p} - \|q_1\|_{1+1/p}^{1+1/p} \right) \leq p \cdot \|q_\star\|_1^{1+1/p} \leq p \cdot T^{1+1/p} \quad (5)$$

It can also be shown that $\nabla^2 \psi_p(q)^{-1} = \text{diag}\left(\frac{p}{p+1} q^{1-1/p}\right)$. Using that $c_k(s, a)^2 \leq c_k(s, a)$, we get

$$\begin{aligned} \|c_k\|_{\nabla^2 \psi_p(q_k)^{-1}}^2 &\leq \frac{p}{p+1} \sum_{s,a} c_k(s, a) q_k(s, a)^{1-1/p} \leq \|c_k\|_1 \sum_{s,a} \frac{c_k(s, a)}{\|c_k\|_1} q_k(s, a)^{1-1/p} \\ &\leq \|c_k\|_1 \left(\sum_{s,a} \frac{c_k(s, a)}{\|c_k\|_1} q_k(s, a) \right)^{1-1/p} \\ &= \|c_k\|_1^{1/p} \langle c_k, q_k \rangle^{1-1/p} \leq M^{1/p} \max\{1, \langle c_k, q_k \rangle\} \leq M^{1/p} (1 + \langle c_k, q_k \rangle), \end{aligned} \quad (6)$$

where the key step (6) uses Jensen's inequality on the concave function $x^{1-1/p}$ ($p > 1$) and probability distribution $c_k/\|c_k\|_1$. Plugging this into the Stability term and combining with (5):

$$\sum_{k=1}^K \langle q_k - q_{\pi^*}, c_k \rangle \leq \frac{pT^{1+1/p}}{\eta} + \frac{\eta M^{1/p} K}{2} + \frac{\eta M^{1/p}}{2} \sum_{k=1}^K \langle q_k, c_k \rangle$$

$$\begin{aligned}
&\Rightarrow \sum_{k=1}^K \langle q_k - q_{\pi^*}, c_k \rangle \leq \frac{1}{1 - \frac{\eta M^{1/p}}{2}} \left(\frac{pT^{1+1/p}}{\eta} + \frac{\eta M^{1/p} K}{2} + \frac{\eta M^{1/p}}{2} \sum_{k=1}^K \langle q_*, c_k \rangle \right) \\
&\Rightarrow \sum_{k=1}^K \langle q_k - q_{\pi^*}, c_k \rangle \leq \frac{2pT^{1+1/p}}{\eta} + 2\eta M^{1/p} DK,
\end{aligned}$$

where the last step uses that $\sum_{k=1}^K \langle q_*, c_k \rangle \leq \sum_{k=1}^K \langle q_{\pi_f}, c_k \rangle \leq K \|q_{\pi_f}\|_1 \|c_k\|_\infty \leq DK$, $D \geq 1$ and $\eta \leq 4M^{-1/p} \iff 1 - \eta M^{1/p}/2 \geq 1/2$. Tuning $\eta = \sqrt{\frac{pT^{1+1/p}}{KDM^{1/p}}}$ (so $\eta \leq 4M^{-1/p}$ for sufficiently large K) and $p = \log(TM)$ gives the result. $\square$

Provided we can suitably select $T \approx T_*$ (see Section 4.1), this result establishes that sparsity does lead to an improvement in the minimax regret. In Section 4.2, we show that the dependence on M and SA is optimal, ruling out polynomial improvements from sparsity such as in the experts setting [19]. This highlights that M acts as the effective dimension of the problem instead of SA . In particular, if the sparsity level M is constant, then we obtain a dimension-independent regret of $\mathcal{O}(\sqrt{DKT \log T})$.

Remark 4.2. Although we express the bound in terms of the sparsity level M , it can be seen that the analysis above holds more generally if M is instead an upper bound on the ℓ_1 norm of the costs: $M = \max_k \|c_k\|_1$. This relaxation allows our result to cover "softly sparse" cost structures and aligns with the notion of first-order bounds commonly studied in the online learning literature [23, 37, 35].

Remark 4.3. The above result does not recover the M/A polynomial improvement in the special case of the expert setting. This can be recovered through a regret bound of a slightly different flavor which includes the hitting time of the uniform policy. We include the details and subtleties in Appendix D but the upshot is that the necessity to reach the goal state in SSP creates a fundamental difference in the benefits of sparsity compared to the expert setting.

4.1 Sparse-agnostic parameter-free upper bound

The procedure in Theorem 4.4 assumes knowledge of the sparsity level M to tune the parameter p of our regularizer and uses knowledge of the expected hitting time of the optimal policy T_* to consider OMD over the space of suitable occupancy measures. We now adapt existing techniques to remove both of these assumptions and derive fully parameter-free guarantees.

For the unknown sparsity level, we use the same approach as in [19]. We divide the K episodes into batches. Within each batch, we independently run OMD tuning the parameter p of our regularizer with the sparsity level observed up to the current batch, as described in Algorithm 1 in Appendix C.2.

For the unknown expected hitting time of the optimal policy T_* , we can exploit the same meta-algorithm technique as in [8], using the sparse agnostic algorithms introduced above as base learners. We run $N \approx \log K$ instances of Algorithm 1 where the j -th instance sets its parameter T as $b(j) \approx 2^j$. Therefore, there exists a good instance j_* such that $b(j_*)$ is close to the unknown T_* . The regret of a scale-invariant meta-algorithm, described for completeness in Algorithm 2 in Appendix C.2, closely matches that of this good instance.

Together, these two techniques yield the following parameter-free regret bound (proof in Appendix C.2):

Theorem 4.4. If $K > \max(T_*, \frac{T_*}{D} \log(T_* M))$ and $T_* > e$, Algorithm 2 guarantees $\mathbb{E}[R_K] \leq \tilde{\mathcal{O}}\left(\sqrt{DKT_* \log(MT_*)} + T_*\right)$, where the notation $\tilde{\mathcal{O}}$ hides double-logarithmic factors.

The leading term matches the regret bound from Theorem 4.1, while the second does not depend on M or K . Therefore, running a procedure that does not assume knowledge of M and T_* comes at no additional cost in terms of the regret bound (up to double-logarithmic factors). We also note that it is common for log-log factors to be ignored in parameter free results with expert-like algorithms [12, 16]. It is also possible to obtain a bound that holds with high-probability since the high-probability analysis given in [8] can easily be adapted to work with our regularizer.

Remark 4.5. The assumption $K \geq \frac{T_*}{D} \log(T_* M)$ or $K \geq \frac{T_*}{D} \log(T_* SA)$ in the non-sparse setting is actually non-restrictive since it is required for the upper-bound to be meaningful:

$$\sqrt{DKT_* \log(MT_*)} \leq T_* K \iff K \geq \frac{T_*}{D} \log(T_* M).$$

In particular, it is likely that there is a gap between the behavior of the minimax regret between the "low-dimensional" setting which we study and a high-dimensional setting where $K \ll \frac{T_\star}{D} \log(T_\star M)$. The high-dimensional problem is yet to be explored, even in the non-adversarial setting and could be an interesting avenue of future research. Indeed, all prior works on SSP have implicitly studied the problem in low-dimension, which comes with an implicit assumption that K is sufficiently large.

4.2 Lower bound

In this section, we provide a general lower bound for sparse SSP problems.

Theorem 4.6. *For any $D, T_\star, K, S, A$ with $T_\star \geq D \geq 3 \log S$, $S(A-1) \geq 400$, $K \geq \frac{800T_\star}{D} \log M$ and $M \geq 101$, there exists an SSP instance with stochastic M -sparse costs, S states and A actions such that its diameter is D , the expected hitting time of the optimal policy is $T_\star$, and the expected regret with respect to the randomness of the losses for any learner after K episodes is $\mathbb{E}[R_K] \geq \Omega(\sqrt{KT_\star D \log M})$.*

For general M (> 100), the lower bound matches the upper bound established in Theorem 4.1 in its dependence on M , characterizing the minimax regret for general sparse problems (up to a $\log T_\star$ term). For $M = SA$, our result gives a $\Omega(\sqrt{KT_\star D \log SA})$ lower bound improving on the $\Omega(\sqrt{KT_\star D})$ lower bound of [8]. In particular, this establishes the optimal dependence on the size of the state-action space SA in the minimax regret for the general non-sparse SSP problem.

Proof intuition: The proof is based on the combination of an SSP instance from [8] and a probabilistic costs construction, which then requires some non-trivial arguments to extend to the sparse SSP problem. We give an overview of the construction and defer the details to Appendix E.

The MDP construction is essentially a reduction to a non-sparse experts problem with $\mathcal{O}(M)$ actions. First, there is a reduction to an experts problem with $\mathcal{O}(SA)$ actions. Then within these, there are $\mathcal{O}(M)$ good actions, while the remaining are bad. The good actions suffer small costs in expectation and can lead directly to the goal-state. The bad actions are zero-cost but all lead to the same unique bad state, where only one action leads to the goal-state and suffers high cost. This allows a big proportion of the actions to be bad while still guaranteeing sparsity and forcing the learner to only consider the $\mathcal{O}(M)$ good actions, completing the reduction to the non-sparse experts problem with $\mathcal{O}(M)$ actions.

However, we cannot directly apply lower bounds for the experts problem because of subtleties in the reduction and the cost-generating mechanism. We use a similar approach to the experts lower bounds by sampling the costs i.i.d. from a Bernoulli distribution, however with a scaled parameter to ensure the reduction above holds. The regret in this stochastic environment can then be expressed as the maximum of asymmetric zero-mean i.i.d. random walks, capturing how much better the optimal policy can be by choosing the best action after the i.i.d. Bernoulli costs have been sampled for all episodes. The result then follows from a technical result on the expectation of this maximum that we derive in Appendix G. We note that the reduction and costs are constructed in such a way that the diameter of the MDP and expected hitting time of the optimal policy are indeed D and $T_\star$.

5 Unknown transition setting

In this section, we consider the setting where the transitions are unknown and show through the following lower bound that the benefits of sparsity are limited.

Theorem 5.1. *For any D, K, S, A with $S \geq 2, A \geq 16$, $D \geq 2$ and $K \geq SA$, there exists an SSP instance with $M = 1$, S states and A actions such that its diameter is D and the expected regret for any learner without knowledge of the transitions after K episodes is $\mathbb{E}[R_K] \geq \Omega(D\sqrt{SAK})$.*

The above result establishes that the minimax regret for the unknown transition setting must scale polynomially with SA , regardless of the sparsity. In particular, this highlights the limited benefits of sparsity in removing the dependence on the state-action space size in the unknown transition setting, which is in stark contrast to the known transition setting.

The proof is based on an SSP instance used by [29] to prove an $\Omega(D\sqrt{SAK})$ lower bound in the unknown transition non-sparse setting. It turns out that this instance can be adapted such that the cost is only non-zero for a single state-action pair, while keeping the regret lower bound unchanged, giving the above result. We include the details in Appendix F.

6 Conclusion, limitations and future-work

In this work, we studied the SSP problem under sparse adversarial costs and full-information feedback. When the transitions are known, we have shown that existing methods fail to appropriately exploit sparsity. Instead, we designed a family of regularizers to use with Online Mirror Descent that allowed us to characterize the sparse minimax regret, establishing the extent of the benefits of sparsity in this setting. When the transitions are unknown, we showed that even the sparse minimax regret scales polynomially in the size of the state-action space, suggesting fundamental limits in such settings.

Our results open up many further directions of research. In particular, we established the benefits of sparsity under known transition as limited to logarithmic, however, there could be structural properties of an MDP that could break this logarithmic limit and achieve polynomial benefits. Moreover, we have limited our focus to the adversarial full-information feedback setting, but the study of sparse SSP problems in other settings, such as partial feedback, stochastic environments, or structured decision problems remains unexplored.

Acknowledgments and Disclosure of Funding

Emmeran Johnson is funded by EPSRC through the Modern Statistics and Statistical Machine Learning (StatML) CDT (grant no. EP/S023151/1). Alberto Rumi was funded by European Lighthouse of AI for Sustainability project (ELIAS). Patrick Rebeschini was funded by UK Research and Innovation (UKRI) under the UK government’s Horizon Europe funding guarantee (grant no. EP/Y028333/1).

We would like to thank the reviewers and meta-reviewers for their time and feedback.

References

- [1] Dimitri P Bertsekas and John N Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- [2] A. V. Boyd. Inequalities for mills’ ratio. *Reports of Statistical Application Research (Union of Japanese Scientists and Engineers)*, 6(1959):44–46, 1959.
- [3] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [4] Sébastien Bubeck, Michael Cohen, and Yuanzhi Li. Sparsity, variance and curvature in multi-armed bandits. *Proceedings of Algorithmic Learning Theory*, 83:111–127, 2018.
- [5] Nicolo Cesa-Bianchi, Yoav Freund, David Haussler, David P Helmbold, Robert E Schapire, and Manfred K Warmuth. How to use expert advice. *Journal of the ACM (JACM)*, 44(3):427–485, 1997.
- [6] Liyu Chen and Haipeng Luo. Finding the stochastic shortest path with low regret: the adversarial cost and unknown transition case. *Proceedings of the 38th International Conference on Machine Learning*, 139:1651–1660, 2021.
- [7] Liyu Chen, Mehdi Jafarnia-Jahromi, Rahul Jain, and Haipeng Luo. Implicit finite-horizon approximation and efficient optimal algorithms for stochastic shortest path. *Advances in Neural Information Processing Systems*, 34:10849–10861, 2021.
- [8] Liyu Chen, Haipeng Luo, and Chen-Yu Wei. Minimax regret for stochastic shortest path with adversarial costs and known transition. *Proceedings of the 34th Conference on Learning Theory*, 134:1180–1215, 2021.
- [9] Liyu Chen, Rahul Jain, and Haipeng Luo. Improved no-regret algorithms for stochastic shortest path with linear MDP. *Proceedings of the 39th International Conference on Machine Learning*, 162:3204–3245, 2022.
- [10] Liyu Chen, Haipeng Luo, and Aviv Rosenberg. Policy optimization for stochastic shortest path. *Proceedings of the 35th Conference on Learning Theory*, 178:982–1046, 2022.
- [11] Alon Cohen, Yonathan Efroni, Yishay Mansour, and Aviv Rosenberg. Minimax regret for stochastic shortest path. *Advances in Neural Information Processing Systems*, 34:28350–28361, 2021.

- [12] Ashok Cutkosky and Francesco Orabona. Black-box reductions for parameter-free online learning in banach spaces. In *Proceedings of the 31st Conference On Learning Theory*, volume 75, pages 1493–1529, 2018.
- [13] Benjamin Doerr. An elementary analysis of the probability that a binomial random variable exceeds its expectation. *Statistics & Probability Letters*, 139:67–74, 2018.
- [14] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [15] Armengol Gasull and Frederic Utzet. Approximating mills ratio. *Journal of Mathematical Analysis and Applications*, 420(2):1832–1853, 2014.
- [16] Andrew Jacobsen and Ashok Cutkosky. Parameter-free mirror descent. In *Proceedings of 35th Conference on Learning Theory*, volume 178, pages 4160–4211, 2022.
- [17] Mehdi Jafarnia-Jahromi, Liyu Chen, Rahul Jain, and Haipeng Luo. Posterior sampling-based online learning for the stochastic shortest path model. *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence*, 216:922–931, 2023.
- [18] Anatoli B. Juditsky and Arkadii S. Nemirovski. First Order Methods for Nonsmooth Convex Large-Scale Optimization, I: General Purpose Methods. In *Optimization for Machine Learning*, pages 1–28. MIT Press, 2010.
- [19] Joon Kwon and Vianney Perchet. Gains and losses are fundamentally different in regret minimization: The sparse case. *Journal of Machine Learning Research*, 17(227):1–32, 2016.
- [20] Joon Kwon, Vianney Perchet, and Claire Vernade. Sparse stochastic bandits. In *Proceedings of the 30th Conference on Learning Theory*, volume 65, pages 1269–1270, 2017.
- [21] Brendan D McKay. On littlewood’s estimate for the binomial distribution. *Advances in Applied Probability*, 21(2):475–478, 1989.
- [22] Yurii Nesterov and Arkadi Nemirovski. On first-order algorithms for l1/nuclear norm minimization. *Acta Numerica*, 22:509–575, 2013.
- [23] Gergely Neu. First-order regret bounds for combinatorial semi-bandits. *Proceedings of the 28th Conference on Learning Theory*, 40:1360–1375, 2015.
- [24] Francesco Orabona. A modern introduction to online learning. *Preprint arXiv:1912.13213*, 2019.
- [25] Francesco Orabona and Dávid Pál. Optimal non-asymptotic lower bound on the minimax regret of learning with expert advice. *Preprint arXiv:1511.02176*, 2015.
- [26] Christos Pelekis and Jan Ramon. A lower bound on the probability that a binomial random variable is exceeding its mean. *Statistics & Probability Letters*, 119:305–309, 2016.
- [27] Herbert Robbins. A remark on stirling’s formula. *The American mathematical monthly*, 62(1): 26–29, 1955.
- [28] Aviv Rosenberg and Yishay Mansour. Stochastic shortest path with adversarially changing costs. *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, 2021.
- [29] Aviv Rosenberg, Alon Cohen, Yishay Mansour, and Haim Kaplan. Near-optimal regret bounds for stochastic shortest path. *Proceedings of the 37th International Conference on Machine Learning*, 119:8210–8219, 2020.
- [30] Rongye Shi, Peter Steenkiste, and Manuela M. Veloso. Sc-m*: A multi-agent path planning algorithm with soft-collision constraint on allocation of common resources. *Applied Sciences*, 9 (19), 2019.
- [31] Pavel Surynek, T. K. Satish Kumar, and Sven Koenig. Multi-agent path finding with capacity constraints. In *Advances in Artificial Intelligence*, 2019.

- [32] Jean Tarbouriech, Evrard Garcelon, Michal Valko, Matteo Pirodda, and Alessandro Lazaric. No-regret exploration in goal-oriented reinforcement learning. *Proceedings of the 37th International Conference on Machine Learning*, 119:9428–9437, 2020.
- [33] Jean Tarbouriech, Runlong Zhou, Simon S Du, Matteo Pirodda, Michal Valko, and Alessandro Lazaric. Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret. *Advances in Neural Information Processing Systems*, 34:6843–6855, 2021.
- [34] Taira Tsuchiya, Shinji Ito, and Junya Honda. Stability-penalty-adaptive follow-the-regularized-leader: Sparsity, game-dependency, and best-of-both-worlds. *Advances in Neural Information Processing Systems*, 36:47406–47437, 2023.
- [35] Andrew J Wagenmaker, Yifang Chen, Max Simchowitz, Simon Du, and Kevin Jamieson. First-order regret in reinforcement learning with linear function approximation: A robust estimation approach. *Proceedings of the 39th International Conference on Machine Learning*, 162:22384–22429, 2022.
- [36] Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.
- [37] Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. *Proceedings of the 31st Conference On Learning Theory*, 75:1263–1291, 2018.
- [38] Peng Zhao, Long-Fei Li, and Zhi-Hua Zhou. Dynamic regret of online Markov decision processes. *Proceedings of the 39th International Conference on Machine Learning*, 162:26865–26894, 2022.
- [39] Alexander Zimin and Gergely Neu. Online learning in episodic markovian decision processes by relative entropy policy search. *Advances in Neural Information Processing Systems*, 26, 2013.

NeurIPS Paper Checklist

1. Claims

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

Answer: [\[Yes\]](#)

Justification: In the contributions section of the Introduction, we have pointers and references to where all the results we claim can be found.

Guidelines:

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

2. Limitations

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

Answer: [\[Yes\]](#)

Justification: We discuss limitations in Section 6.

Guidelines:

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

3. Theory assumptions and proofs

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

Answer: [\[Yes\]](#)

Justification: All of the theoretical results we present are followed by proofs or pointers / references to the appendices where the proofs can be found.

Guidelines:

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

4. Experimental result reproducibility

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

Answer: [\[NA\]](#)

Justification: The paper does not include experiments.

Guidelines:

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

- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

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

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

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

6. Experimental setting/details

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

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

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

7. Experiment statistical significance

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

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

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

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

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

9. Code of ethics

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

Answer: [Yes]

Justification: Our work is purely theoretical and conforms with the code of ethics.

Guidelines:

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

10. **Broader impacts**

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

Answer: [NA]

Justification: Our work is purely theoretical and poses no societal impact.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

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

Answer: [NA]

Justification: Our work is purely theoretical and poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.

- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

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

Answer: [NA]

Justification: The paper does not use existing assets.

Guidelines:

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

13. New assets

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

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

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

14. Crowdsourcing and research with human subjects

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

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

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

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

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Failure of the negative entropy

In this appendix, we prove our lower bound result for OMD with the negative entropy from Section 3. We first restate the result.

Theorem 3.1. *For any $S \geq 6$, there exists an SSP instance with a fixed horizon of 3, sparsity level $M = 3$, an action space of size $A = 2$ and state space of size S such that the regret of OMD (1) with negative-entropy regularization and any step-size $\eta > 0$ after K episodes is $\mathbb{E}[R_K] = \Omega(\min\{\sqrt{K \log S}, K\})$.*

Proof. Fix K even, $S \geq 6$, $A = 2$ and $N = S - 5$. We first describe the SSP instance. Consider the following MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p, s_0, g)$, where $\mathcal{S} = \{s_0, s_0^L, s_0^R, s_1^R, \dots, s_N^R, s_1\}$ and $\mathcal{A} = \{a_1, a_2\}$. The transitions and costs (in each episode k) are defined as:

- s_0 : $p(s_0^L | s_0, a) = p(s_0^R | s_0, a) = 1/2$ and $c_k(s_0, a) = 0$ for all $a \in \mathcal{A}$.
- s_0^L : $p(s_g^L | s_0^L, a) = 1$ for all $a \in \mathcal{A}$ and $c_k(s_0^L, a_1) = \frac{1+(-1)^k}{2}$, $c_k(s_0^L, a_2) = 1/2$.
- s_0^R : $p(s_g^R | s_0^R, a_1) = 1$, $p(s_i^R | s_0^R, a_2) = 1/N$ and $c_k(s_0^R, a_1) = 0$, $c_k(s_0^R, a_2) = 1$.
- s_i^R : $p(g | s_i^R, a) = 1$ and $c_k(s_i^R, a) = 0$ for all $a \in \mathcal{A}$.
- s_g^L : $p(g | s_g^L, a) = 1$ and $c_k(s_g^L, a) = 0$ for all $a \in \mathcal{A}$.
- s_g^R : $p(g | s_g^R, a) = 1$ and $c_k(s_g^R, a) = 0$ for all $a \in \mathcal{A}$.

An illustration is given in Figure 3. This SSP instance has a fixed horizon of 3 in the sense that all policies have a hitting time of exactly 3 (the states s_g^L and s_g^R are added to guarantee this). As a result we have that $T_\star = D = 3$. Also note that there are at most 3 state-action pairs that have non-zero cost, therefore the sparsity level $M = 3$.




Figure 3: Diagram illustrating MDP construction for the proof of Theorem 3.1. When an action is not specified for an edge, then both actions give the same transition and cost. If an edge has a number in black, it is a transition probability; if it does not then the transition is deterministic. The costs are given in red. The formal description of the MDP is given above.

From Appendix B.1 in [28] (we can ignore the optimization over λ because we are in a fixed horizon setting), the update of OMD with negative entropy for any $k \geq 0$ can be computed by solving a convex optimization problem:

$$q_{k+1}(s, a) = q_k(s, a) e^{B_k^{v_{k+1}}(s, a)}, \quad \text{where} \quad B_k^v(s, a) = v(s) - \eta c_k(s, a) - \sum_{s' \in \mathcal{S}} p(s' | s, a) v(s'),$$

$$v_{k+1} = \arg \min_v \mathcal{D}_k(v),$$

$$\mathcal{D}_k(v) = \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_k(s, a) e^{B_k^v(s, a)} - v(s_0),$$

with $q_0(s, a) = 1$ and $c_0(s, a) = 0$. This allows us to compute exactly the points played by the algorithm on the SSP instance described above, and in turn compute the regret, from which the result will follow.

In the following few pages, we compute the occupancy measures played by OMD with the negative entropy on the SSP instance described earlier for all episodes, using the convex optimization problem above. We begin by computing expressions for $B_k^v(s, a)$ in each state:

- $B_k^v(s_0, a) = v(s_0) - \frac{1}{2}v(s_0^L) - \frac{1}{2}v(s_0^R)$ for all $a \in \mathcal{A}$
- $B_k^v(s_0^L, a) = v(s_0^L) - \eta c_k(s_0^L, a) - v(s_g^L)$ for all $a \in \mathcal{A}$
- $B_k^v(s_g^L) = v(s_g^L)$
- $B_k^v(s_0^R, a_1) = v(s_0^R) - v(s_g^R)$
- $B_k^v(s_0^R, a_2) = v(s_0^R) - \eta c_k(s_0^R, a_2) - \frac{1}{N} \sum_{i=1}^N v(s_i^R) = v(s_0^R) - \eta c_k(s_0^R, a_2) - v(s_1^R)$ since by symmetry $v(s_i^R) = v(s_1^R)$ for all $i \geq 1$ and any v solving the convex optimization problem specified in the OMD update.
- $B_k^v(s_i^R, a) = v(s_i^R) = v(s_1^R)$ for all $a \in \mathcal{A}$
- $B_k^v(s_g^R, a) = v(s_g^R)$ for all $a \in \mathcal{A}$

Plugging these into the optimization problem, we obtain (recall the notation $q(s) = \sum_{a \in \mathcal{A}} q(s, a)$):

$$\begin{aligned}
v_{k+1} &= \arg \min_v \mathcal{D}_k(v) = \arg \min_v \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_k(s, a) e^{B_k^v(s, a)} - v(s_0) \\
&= \arg \min_v q_k(s_0, a_1) e^{v(s_0) - 0.5v(s_0^L) - 0.5v(s_0^R)} + q_k(s_0, a_2) e^{v(s_0) - 0.5v(s_0^L) - 0.5v(s_0^R)} \\
&\quad + q_k(s_0^L, a_1) e^{v(s_0^L) - \eta c_k(s_0^L, a_1) - v(s_g^L)} + q_k(s_0^L, a_2) e^{v(s_0^L) - \eta c_k(s_0^L, a_2) - v(s_g^L)} \\
&\quad + q_k(s_g^L, a_1) e^{v(s_g^L)} + q_k(s_g^L, a_2) e^{v(s_g^L)} \\
&\quad + q_k(s_0^R, a_1) e^{v(s_0^R) - v(s_g^R)} + q_k(s_0^R, a_2) e^{v(s_0^R) - \eta c_k(s_0^R, a_2) - v(s_1^R)} \\
&\quad + \sum_{i=1}^N \left\{ q_k(s_i^R, a_1) e^{v(s_1^R)} + q_k(s_i^R, a_2) e^{v(s_1^R)} \right\} \\
&\quad + q_k(s_g^R, a_1) e^{v(s_g^R)} + q_k(s_g^R, a_2) e^{v(s_g^R)} \\
&\quad - v(s_0) \\
&= \arg \min_v q_k(s_0) e^{v(s_0) - 0.5v(s_0^L) - 0.5v(s_0^R)} \\
&\quad + q_k(s_0^L, a_1) e^{v(s_0^L) - \eta c_k(s_0^L, a_1) - v(s_g^L)} + q_k(s_0^L, a_2) e^{v(s_0^L) - \eta c_k(s_0^L, a_2) - v(s_g^L)} \\
&\quad + q_k(s_g^L) e^{v(s_g^L)} \\
&\quad + q_k(s_0^R, a_1) e^{v(s_0^R) - v(s_g^R)} + q_k(s_0^R, a_2) e^{v(s_0^R) - \eta c_k(s_0^R, a_2) - v(s_1^R)} \\
&\quad + \sum_{i=1}^N \left\{ q_k(s_i^R) e^{v(s_1^R)} \right\} \\
&\quad + q_k(s_g^R) e^{v(s_g^R)} \\
&\quad - v(s_0).
\end{aligned}$$

This being a convex optimization problem, it can be solved by differentiating and setting to 0:

$$\begin{aligned}
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0)} &= q_k(s_0) e^{v(s_0) - 0.5v(s_0^L) - 0.5v(s_0^R)} - 1 = 0 \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^L)} &= -0.5 q_k(s_0) e^{v(s_0) - 0.5v(s_0^L) - 0.5v(s_0^R)} + q_k(s_0^L, a_1) e^{v(s_0^L) - \eta c_k(s_0^L, a_1) - v(s_g^L)} + q_k(s_0^L, a_2) e^{v(s_0^L) - \eta c_k(s_0^L, a_2) - v(s_g^L)} \\
&= -0.5 + q_k(s_0^L, a_1) e^{v(s_0^L) - \eta c_k(s_0^L, a_1) - v(s_g^L)} + q_k(s_0^L, a_2) e^{v(s_0^L) - \eta c_k(s_0^L, a_2) - v(s_g^L)} = 0
\end{aligned}$$

$$\begin{aligned}\frac{\partial \mathcal{D}_k(v)}{\partial v(s_g^L)} &= -q_k(s_0^L, a_1)e^{v(s_0^L) - \eta c_k(s_0^L, a_1) - v(s_g^L)} - q_k(s_0^L, a_2)e^{v(s_0^L) - \eta c_k(s_0^L, a_2) - v(s_g^L)} + q_k(s_g^L)e^{v(s_g^L)} \\ &= q_k(s_g^L)e^{v(s_g^L)} - 0.5 = 0\end{aligned}$$

$$\begin{aligned}\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^R)} &= -0.5q_k(s_0)e^{v(s_0) - 0.5v(s_0^L) - 0.5v(s_0^R)} + q_k(s_0^R, a_1)e^{v(s_0^R) - v(s_g^R)} + q_k(s_0^R, a_2)e^{v(s_0^R) - \eta c_k(s_0^R, a_2) - v(s_1^R)} \\ &= -0.5 + q_k(s_0^R, a_1)e^{v(s_0^R) - v(s_g^R)} + q_k(s_0^R, a_2)e^{v(s_0^R) - \eta c_k(s_0^R, a_2) - v(s_1^R)} = 0\end{aligned}$$

$$\frac{\partial \mathcal{D}_k(v)}{\partial v(s_1^R)} = -q_k(s_0^R, a_2)e^{v(s_0^R) - \eta c_k(s_0^R, a_2) - v(s_1^R)} + e^{v(s_1^R)} \sum_{i=1}^N q_k(s_i^R) = 0$$

$$\frac{\partial \mathcal{D}_k(v)}{\partial v(s_g^R)} = -q_k(s_0^R, a_1)e^{v(s_0^R) - v(s_g^R)} + q_k(s_g^R)e^{v(s_g^R)} = 0.$$

Let's look specifically at the case $k = 0$ ($q_0(s, a) = 1, q_0(s) = 2, c_0(s, a) = 0$ for all $s \in \mathcal{S}, a \in \mathcal{A}$). For the left part of the MDP we have:

$$\begin{aligned}\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^L)} = 0 &\implies 2e^{v(s_0^L) - v(s_g^L)} = 0.5 \implies e^{v(s_0^L)} = 0.25e^{v(s_g^L)} \\ \frac{\partial \mathcal{D}_k(v)}{\partial v(s_g^L)} = 0 &\implies e^{v(s_g^L)} = 0.25 \implies e^{v(s_0^L)} = 0.25^2 \\ &\implies q_1(s_0^L, a) = e^{B_0^v(s_0^L, a)} = \frac{0.25^2}{0.25} = 0.25, \text{ for all } a \in \mathcal{A}.\end{aligned}$$

For the right part of the MDP, we have:

$$\begin{aligned}\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^R)} = 0 &\implies e^{v(s_0^R) - v(s_g^R)} + e^{v(s_0^R) - v(s_1^R)} = 0.5 \implies e^{v(s_0^R)} = \frac{0.5}{e^{-v(s_g^R)} + e^{-v(s_1^R)}} \\ \frac{\partial \mathcal{D}_k(v)}{\partial v(s_1^R)} = 0 &\implies e^{v(s_0^R) - v(s_1^R)} = 2Ne^{v(s_1^R)} \implies e^{v(s_1^R)} = \frac{1}{\sqrt{2N}}e^{0.5v(s_0^R)} \\ \frac{\partial \mathcal{D}_k(v)}{\partial v(s_g^R)} = 0 &\implies e^{v(s_0^R) - v(s_g^R)} = 2e^{v(s_g^R)} \implies e^{v(s_g^R)} = \frac{1}{\sqrt{2}}e^{0.5v(s_0^R)} \\ &\implies e^{v(s_0^R)} = \frac{0.5}{\sqrt{2}e^{-0.5v(s_0^R)} + \sqrt{2N}e^{-0.5v(s_0^R)}} \\ &\implies e^{0.5v(s_0^R)} = \frac{0.5}{\sqrt{2} + \sqrt{2N}} \\ &\implies e^{v(s_0^R)} = \frac{0.25}{(\sqrt{2} + \sqrt{2N})^2} \\ &\implies q_1(s_0^R, a_1) = e^{B_0^v(s_0^R, a_1)} = e^{v(s_0^R) - v(s_g^R)} = \sqrt{2}e^{0.5v(s_0^R)} = \frac{0.5}{1 + \sqrt{N}} \\ q_1(s_0^R, a_2) &= e^{B_0^v(s_0^R, a_2)} = e^{v(s_0^R) - v(s_1^R)} = \sqrt{2N}e^{0.5v(s_0^R)} = \frac{0.5\sqrt{N}}{1 + \sqrt{N}} \\ q_1(s_1^R, a) &= e^{B_0^v(s_1^R, a)} = e^{v(s_1^R)} = \frac{1}{\sqrt{2N}}e^{0.5v(s_0^R)} = \frac{0.25}{\sqrt{N}(1 + \sqrt{N})} \\ q_1(s_g^R, a) &= e^{B_0^v(s_g^R, a)} = e^{v(s_g^R)} = \frac{1}{\sqrt{2}}e^{0.5v(s_0^R)} = \frac{0.25}{1 + \sqrt{N}}.\end{aligned}$$

Let's now look at general $k \geq 1$: Since q_k is an occupancy measure, it satisfies the properties of the dynamics of the MDP (see the definition of $\Delta(T)$ in Section 2) and we have that for any $s \in \mathcal{S}$: $\sum_{a \in \mathcal{A}} q_k(s, a) = \sum_{s' \in \mathcal{S}} \sum_{a' \in \mathcal{A}} p(s|s', a')q_k(s', a') + \mathbb{I}\{s = s_0\}$. In particular, this gives

- $s = s_0$: $q_k(s_0) = 1$
- $s = s_0^L$: $q_k(s_0^L) = 0.5q_k(s_0) = 0.5$

- $s = s_g^L: q_k(s_g^L) = q_k(s_0^L) = 0.5$
- $s = s_0^R: q_k(s_0^R) = 0.5q_k(s_0) = 0.5$
- $s = s_g^R: q_k(s_g^R) = q_k(s_0^R, a_1)$
- $s = s_1^R: q_k(s_1^R) = \frac{1}{N}q_k(s_0^R, a_2)$

This leads to the following simplifications in the derivatives of $\mathcal{D}_k(v)$:

$$\begin{aligned}
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0)} &= e^{v(s_0) - 0.5v(s_0^L) - 0.5v(s_0^R)} - 1 = 0 \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^L)} &= -0.5 + q_k(s_0^L, a_1)e^{v(s_0^L) - \eta c_k(s_0^L, a_1) - v(s_g^L)} + (1/2 - q_k(s_0^L, a_1))e^{v(s_0^L) - \eta c_k(s_0^L, a_2) - v(s_g^L)} = 0 \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_g^L)} &= 0.5e^{v(s_g^L)} - 0.5 = 0 \implies e^{v(s_g^L)} = 1 \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^R)} &= -0.5 + q_k(s_0^R, a_1)e^{v(s_0^R) - v(s_g^R)} + (1/2 - q_k(s_0^R, a_1))e^{v(s_0^R) - \eta - v(s_1^R)} = 0 \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_1^R)} &= -q_k(s_0^R, a_2)e^{v(s_0^R) - \eta - v(s_1^R)} + q_k(s_0^R, a_2)e^{v(s_1^R)} = 0 \implies e^{v(s_0^R)} = e^{\eta + 2v(s_1^R)} \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_g^R)} &= -q_k(s_0^R, a_1)e^{v(s_0^R) - v(s_g^R)} + q_k(s_0^R, a_1)e^{v(s_g^R)} = 0 \implies e^{v(s_0^R)} = e^{2v(s_g^R)}.
\end{aligned}$$

Left part of the MDP:

$$\begin{aligned}
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^L)} &= 0 \implies q_k(s_0^L, a_1)e^{v(s_0^L) - \eta \frac{1+(-1)^k}{2}} + (0.5 - q_k(s_0^L, a_1))e^{v(s_0^L) - 0.5\eta} = 0.5 \\
&\implies e^{v(s_0^L)} = \frac{0.5}{q_k(s_0^L, a_1)e^{-\eta \frac{1+(-1)^k}{2}} + (0.5 - q_k(s_0^L, a_1))e^{-0.5\eta}} \\
k+1=2 &\implies e^{v(s_0^L)} = \frac{0.5}{0.25 + 0.25e^{-0.5\eta}} \\
&\implies q_2(s_0^L, a_1) = q_1(s_0^L, a_1)e^{B_1^v(s_0^L, a_1)} = 0.25e^{v(s_0^L) - v(s_g^L)} = \frac{0.5}{1 + e^{-0.5\eta}} \\
k+1=3 &\implies e^{v(s_0^L)} = \frac{0.5}{q_1(s_0^L, a_1)e^{-\eta} + (0.5 - q_1(s_0^L, a_1))e^{-0.5\eta}} \\
&\implies e^{v(s_0^L)} = \frac{0.5}{\frac{0.5}{1+e^{-0.5\eta}}e^{-\eta} + \frac{0.5e^{-0.5\eta}}{1+e^{-0.5\eta}}e^{-0.5\eta}} = 0.25 \frac{e^\eta}{q_2(s_0^L, a_1)} \\
&\implies q_3(s_0^L, a_1) = q_2(s_0^L, a_1)e^{v(s_0^L) - \eta - v(s_g^L)} = 0.25 \\
k \text{ even} &\implies q_k(s_0^L, a_1) = \frac{0.5}{1 + e^{-0.5\eta}} \\
k \text{ odd} &\implies q_k(s_0^L, a_1) = 0.25,
\end{aligned}$$

where the last two lines follow by a straightforward induction. Hence the losses suffered by OMD on the left part of the MDP are:

$$\begin{aligned}
\sum_{k=1}^K \left\{ q_k(s_0^L, a_1)c_k(s_0^L, a_1) + q_k(s_0^L, a_2)c_k(s_0^L, a_2) \right\} &= \sum_{k=1}^K \left\{ q_k(s_0^L, a_1) \cdot \frac{1+(-1)^k}{2} + 0.5 \cdot (0.5 - q_k(s_0^L, a_1)) \right\} \\
&= \sum_{k=1}^K \left\{ q_k(s_0^L, a_1) \cdot \frac{(-1)^k}{2} + 0.25 \right\} \\
&= 0.25K + 0.5 \sum_{t=1}^{K/2} \left\{ q_{2t}(s_0^L, a_1) - q_{2t-1}(s_0^L, a_1) \right\}
\end{aligned}$$

$$\begin{aligned}
&= 0.25K + 0.5 \sum_{t=1}^{K/2} \left\{ \frac{0.5}{1 + e^{-0.5\eta}} - 0.25 \right\} \\
&= 0.25K + 0.5 \frac{K}{2} \frac{0.5 - 0.25 - 0.25e^{-0.5\eta}}{1 + e^{-0.5\eta}} \\
&= 0.25K + \frac{K}{16} \cdot \frac{1 - e^{-0.5\eta}}{1 + e^{-0.5\eta}} \\
&= 0.25K + \frac{K}{16} \cdot \min\left\{\frac{\eta}{5}, \frac{1}{2}\right\}. \tag{7}
\end{aligned}$$

Right part of the MDP:

$$\begin{aligned}
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_0^R)} = 0 &\implies e^{v(s_0^R)} = \frac{0.5}{q_k(s_0^R, a_1)e^{-v(s_g^R)} + (1/2 - q_k(s_0^R, a_1))e^{-\eta - v(s_1^R)}} \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_1^R)} = 0 &\implies v(s_0^R) = \eta + 2v(s_1^R) \\
\frac{\partial \mathcal{D}_k(v)}{\partial v(s_g^R)} = 0 &\implies v(s_0^R) = v(s_g^R) \\
&\implies e^{v(s_0^R)} = \frac{0.5}{q_k(s_0^R, a_1)e^{-0.5v(s_0^R)} + (1/2 - q_k(s_0^R, a_1))e^{-\eta - 0.5v(s_0^R) + 0.5\eta}} \\
&\implies e^{0.5v(s_0^R)} = \frac{0.5}{q_k(s_0^R, a_1) + (1/2 - q_k(s_0^R, a_1))e^{-0.5\eta}} = \frac{0.5}{q_k(s_0^R, a_1) + q_k(s_0^R, a_2)e^{-0.5\eta}} \\
&\implies q_{k+1}(s_0^R, a_1) = q_k(s_0^R, a_1)e^{B_k^v(s_0^R, a_1)} = q_k(s_0^R, a_1)e^{v(s_0^R) - v(s_g^R)} = q_k(s_0^R, a_1)e^{0.5v(s_0^R)} \\
&\implies q_{k+1}(s_0^R, a_1) = 0.5 \frac{q_k(s_0^R, a_1)}{q_k(s_0^R, a_1) + (1/2 - q_k(s_0^R, a_1))e^{-0.5\eta}} \\
&\implies \frac{q_{k+1}(s_0^R, a_1)}{q_{k+1}(s_0^R, a_2)} = \frac{q_{k+1}(s_0^R, a_1)}{(1/2 - q_{k+1}(s_0^R, a_1))} = \frac{q_k(s_0^R, a_1)}{(1/2 - q_k(s_0^R, a_1))e^{-0.5\eta}} = e^{0.5\eta} \frac{q_k(s_0^R, a_1)}{q_k(s_0^R, a_2)} \\
&\implies \frac{q_{k+1}(s_0^R, a_1)}{q_{k+1}(s_0^R, a_2)} = e^{0.5k\eta} \frac{q_1(s_0^R, a_1)}{q_1(s_0^R, a_2)} = e^{0.5k\eta} \frac{0.5}{0.5\sqrt{N}} = \frac{1}{\sqrt{N}} e^{0.5k\eta} \\
&\implies q_{k+1}(s_0^R, a_2) = \frac{0.5}{1 + \frac{1}{\sqrt{N}}e^{0.5k\eta}} = \frac{0.5\sqrt{N}}{\sqrt{N} + e^{0.5\eta k}}.
\end{aligned}$$

This also holds for $k = 0$ (as shown above). Hence, the losses suffered by OMD on the right part of the MDP are

$$\begin{aligned}
\sum_{k=1}^K q_k(s_0^R, a_2)c_k(s_0^R, a_2) &= \sum_{k=1}^K \frac{0.5\sqrt{N}}{\sqrt{N} + e^{0.5k\eta}} \\
&\geq \int_1^{K+1} \frac{0.5\sqrt{N}}{\sqrt{N} + e^{0.5\eta x}} dx = 0.5K - \int_1^{K+1} \frac{0.5e^{0.5\eta x}}{\sqrt{N} + e^{0.5\eta x}} dx \\
&= 0.5K - \left[\frac{1}{\eta} \log(\sqrt{N} + e^{0.5\eta x}) \right]_1^{K+1} \\
&= 0.5K - \frac{1}{\eta} \log(\sqrt{N} + e^{0.5\eta(K+1)}) + \frac{1}{\eta} \log(\sqrt{N} + e^{0.5\eta}) \\
&\geq 0.5K - \frac{1}{\eta} \log(2e^{0.5\eta(K+1)}) + \frac{1}{\eta} \log \sqrt{N} \quad \text{assuming } \sqrt{N} \leq e^{0.5\eta(K+1)} \\
&= 0.5K - 0.5(K+1) - \frac{1}{\eta} \log 2 + \frac{1}{2\eta} \log N \\
&= -0.5 + \frac{1}{2\eta} \log \frac{N}{4}.
\end{aligned}$$

If $\sqrt{N} > e^{0.5\eta(K+1)}$, then we have

$$\begin{aligned}
0.5K - \frac{1}{\eta} \log(\sqrt{N} + e^{0.5\eta(K+1)}) + \frac{1}{\eta} \log(\sqrt{N} + e^{0.5\eta}) &= 0.5K + \frac{1}{\eta} \log\left(\frac{\sqrt{N} + e^{0.5\eta}}{\sqrt{N} + e^{0.5\eta(K+1)}}\right) \\
&\geq 0.5K + \frac{1}{\eta} \log\left(\frac{e^{0.5\eta(K+1)} + e^{0.5\eta}}{2e^{0.5\eta(K+1)}}\right) \\
&\geq 0.5K + \frac{1}{\eta} \log\left(\frac{1 + e^{-0.5\eta K}}{2}\right) \\
&\geq 0.25K,
\end{aligned}$$

using that $\frac{1+e^{-0.5Kx}}{2} \geq e^{-0.25Kx}$ since $\cosh(x) \geq 1$. So we have

$$\sum_{k=1}^K q_k(s_0^R, a_2) c_k(s_0^R, a_2) \geq \min\left\{-0.5 + \frac{1}{2\eta} \log \frac{N}{4}, 0.25K\right\}. \quad (8)$$

Combining the losses from the left part in (7) and from the right part in (8), we have:

$$\begin{aligned}
\sum_{k=1}^K \langle q_k, c_k \rangle &= \sum_{k=1}^K \left\{ q_k(s_0^L, a_1) c_k(s_0^L, a_1) + q_k(s_0^L, a_2) c_k(s_0^L, a_2) \right\} + \sum_{k=1}^K \left\{ q_k(s_0^R, a_2) c_k(s_0^R, a_2) \right\} \\
&\geq 0.25K + \frac{K}{16} \cdot \min\left\{\frac{\eta}{5}, \frac{1}{2}\right\} + \min\left\{-0.5 + \frac{1}{2\eta} \log \frac{N}{4}, 0.25K\right\}.
\end{aligned}$$

Regret lower-bound: consider $q_\star$ defined as follows:

- $q_\star(s_0, a) = 1/2$ for all $a \in \mathcal{A}$
- $q_\star(s_0^L, a_2) = 1/2$, $q_\star(s_0^L, a_1) = 0$
- $q_\star(s_g^L, a) = 1/4$ for all $a \in \mathcal{A}$
- $q_\star(s_0^R, a_1) = 1/2$, $q_\star(s_0^R, a_2) = 0$, $q_\star(s_i^R, a) = 0$
- $q_\star(s_g^R, a) = 1/4$ for all $a \in \mathcal{A}$

It is straightforward to check that $q_\star$ satisfies the flow constraints and is an occupancy measure. We obtain

$$\begin{aligned}
\sum_{k=1}^K \langle q_\star, c_k \rangle &= \sum_{k=1}^K \left\{ q_\star(s_0^L, a_2) \cdot 0.5 \right\} = 0.25K \\
\implies R_K &\geq \sum_{k=1}^K \langle q_k - q_\star, c_k \rangle \geq \frac{K}{16} \cdot \min\left\{\frac{\eta}{5}, \frac{1}{2}\right\} + \min\left\{-0.5 + \frac{1}{2\eta} \log \frac{N}{4}, 0.25K\right\} \\
&\geq \min\left\{\frac{1}{2} \sqrt{\frac{1}{10} K \log \frac{N}{4}}, \frac{K}{32}\right\} - 0.5.
\end{aligned}$$

Recalling that $N = S - 5$, we have $R_K = \Omega(\min\{\sqrt{K \log S}, K\})$ for an MDP where the sparsity level is $M = 3$, concluding the proof. $\square$

B Efficient implementation of OMD using our regularizer

In this section, we describe how the OMD update with our regularizer from Section 4 defined in (4) can be computed efficiently. This closely follows Appendix B.1 of [28], who provide a similar description for the negative entropy.

Recall the regularizer $\psi_p(q) = p \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q(s, a)^{1+1/p} - p$ for $q \in \mathbb{R}_{\geq 0}^\Gamma$. We have

$$\nabla \psi_p(q) = (p+1) \cdot q(s, a)^{1/p}.$$

The Bregman divergence is defined as:

$$\begin{aligned} D_{\psi_p}(q, q') &= \psi_p(q) - \psi_p(q') - \langle \nabla \psi_p(q'), q - q' \rangle \\ &= \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ p \cdot q(s, a)^{1+1/p} - p \cdot q'(s, a)^{1+1/p} \right\} \\ &\quad - (p+1) \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q'(s, a)^{1/p} q(s, a) - q'(s, a)^{1+1/p} \right\} \\ &= \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q'(s, a)^{1+1/p} + q(s, a) \cdot [p \cdot q(s, a)^{1/p} - (p+1) \cdot q'(s, a)^{1/p}] \right\}. \end{aligned}$$

Recall that OMD with the above regularizer computes the occupancy measures as follows - see (1):

$$q_1 = \arg \min_{q \in \Delta(T)} \psi_p(q), \quad q_{k+1} = \arg \min_{q \in \Delta(T)} \left\{ \eta \cdot \langle q, c_k \rangle + D_{\psi_p}(q, q_k) \right\}.$$

As shown in [24] (Theorem 6.15), each of these steps can be split into an unconstrained minimization step, and a projection step. Thus, q_1 can be computed as follows

$$\begin{aligned} q'_1 &= \arg \min_{q \in \mathbb{R}_{\geq 0}^\Gamma} \psi_p(q) \\ q_1 &= \arg \min_{q \in \Delta(T)} D_{\psi_p}(q, q'_1), \end{aligned}$$

where q'_1 has a closed-form solution $q'_1(s, a) = 1$ for every $s \in \mathcal{S}$ and $a \in \mathcal{A}$. Similarly, q_{k+1} is computed as follows for every $k = 1, \dots, K-1$:

$$\begin{aligned} q'_{k+1} &= \arg \min_{q \in \mathbb{R}_{\geq 0}^\Gamma} \left\{ \eta \cdot \langle q, c_k \rangle + D_{\psi_p}(q, q_k) \right\} \\ q_{k+1} &= \arg \min_{q \in \Delta(T)} D_{\psi_p}(q, q'_{k+1}), \end{aligned}$$

where again q'_{k+1} has a closed-form solution $q'_{k+1}(s, a) = \left[q_k(s, a)^{1/p} - \frac{\eta}{p+1} c_k(s, a) \right]_+$ for every $s \in \mathcal{S}$ and $a \in \mathcal{A}$ (follows from straightforwardly differentiating above objective and setting to 0 and accounting for the non-negativity of occupancy measures) - we use notation $a_+ = \max\{0, a\}$.

For the projection step, we start by formulating it as a constrained convex optimization problem:

$$\begin{aligned} \min_{q \in \mathbb{R}^\Gamma} D_{\psi_p}(q, q'_{k+1}) \quad \text{s.t.} \quad & \sum_{a \in \mathcal{A}} q(s, a) - \sum_{s' \in \mathcal{S}} \sum_{a' \in \mathcal{A}} P(s'|s, a) q(s', a') = \mathbb{I}\{s = s_0\} \quad \forall s \in \mathcal{S} \\ & \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q(s, a) \leq T \\ & q(s, a) \geq 0 \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \end{aligned}$$

The problem can be solved by considering the Lagrangian with Lagrange multipliers λ and $\{v(s)\}_{s \in \mathcal{S}}$:

$$\begin{aligned} \mathcal{L}(q, \lambda, v) &= D_{\psi_p}(q, q'_{k+1}) + \lambda \left(\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q(s, a) - T \right) + \sum_s v(s) \left(\sum_{s', a'} P(s'|s, a) q(s', a') + \mathbb{I}\{s = s_0\} - \sum_a q(s, a) \right) \\ &= D_{\psi_p}(q, q'_{k+1}) + \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q(s, a) \left(\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a) v(s') - v(s) \right) + v(s_0) - \lambda T, \end{aligned}$$

Differentiating the Lagrangian with respect to any $q(s, a)$ and setting to 0, we get

$$\begin{aligned}\frac{\partial \mathcal{L}(q, \lambda, v)}{\partial q(s, a)} &= \nabla \psi_p(q)(s, a) - \nabla \psi_p(q'_{k+1})(s, a) + \lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s) \\ &= (p+1)q(s, a)^{1/p} - (p+1)q'_{k+1}(s, a)^{1/p} + \lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s) = 0. \\ \implies q_{k+1}(s, a) &= \left[q'_{k+1}(s, a)^{1/p} - \frac{\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s)}{p+1} \right]_+^p.\end{aligned}$$

This formula is also valid for $k = 0$ by setting $c_0(s, a) = 0$ and $q_0(s, a) = 1$ for every $s \in \mathcal{S}$ and $a \in \mathcal{A}$.

To compute the value of λ and v at the optimum, we write the dual problem $\mathcal{D}(\lambda, v) = \min_q \mathcal{L}(q, \lambda, v)$ by substituting q_{k+1} back into $\mathcal{L}$:

$$\mathcal{D}(\lambda, v) = D_{\psi_p}(q_{k+1}, q'_{k+1}) + \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_{k+1}(s, a) \left(\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s) \right) + v(s_0) - \lambda T.$$

Recall that $q'_{k+1}(s, a) = \left[q_k(s, a)^{1/p} - \frac{\eta}{p+1} c_k(s, a) \right]_+^p$, so (ignoring terms independent of λ, v , e.g. $q'_{k+1}(s, a)$):

$$\begin{aligned}D_{\psi_p}(q_{k+1}, q'_{k+1}) &= \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q'_{k+1}(s, a)^{1+1/p} + q_{k+1}(s, a) \cdot [pq_{k+1}(s, a)^{1/p} - (p+1)q'_{k+1}(s, a)^{1/p}] \right\} \\ &\propto \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q_{k+1}(s, a) \cdot \left[p \left(q'_{k+1}(s, a)^{1/p} - \frac{\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s)}{p+1} \right)_+ - (p+1)q'_{k+1}(s, a)^{1/p} \right] \right\} \\ &= \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q_{k+1}(s, a) \cdot \left[p \left(q'_{k+1}(s, a)^{1/p} - \frac{\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s)}{p+1} \right) - (p+1)q'_{k+1}(s, a)^{1/p} \right] \right\} \quad \text{since if } q_{k+1}(s, a) = 0, \text{ then the whole term is 0} \\ &= \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q_{k+1}(s, a) \cdot \left[-q'_{k+1}(s, a)^{1/p} - \frac{p}{p+1} \left(\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s) \right) \right] \right\} \\ \implies \mathcal{D}(\lambda, v) &\propto \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q_{k+1}(s, a) \cdot \left[-q'_{k+1}(s, a)^{1/p} - \frac{p}{p+1} \left(\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s) \right) + \left(\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s) \right) \right] \right\} + v(s_0) - \lambda T \\ &= \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left\{ q_{k+1}(s, a) \cdot \left[-q'_{k+1}(s, a)^{1/p} + \frac{\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s)}{p+1} \right] \right\} + v(s_0) - \lambda T \\ &= - \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} q_{k+1}(s, a)^{1+1/p} + v(s_0) - \lambda T \\ &= - \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left[q'_{k+1}(s, a)^{1/p} - \frac{\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s)}{p+1} \right]_+^{1+p} + v(s_0) - \lambda T.\end{aligned}$$

Maximizing the dual gives λ and v or equivalently, we can minimize the negation of the dual:

$$\lambda_{k+1}, v_{k+1} = \arg \min_{\lambda \geq 0, v} \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left[q'_{k+1}(s, a)^{1/p} - \frac{\lambda + \sum_{s' \in \mathcal{S}} P(s'|s, a)v(s') - v(s)}{p+1} \right]_+^{1+p} - v(s_0) + \lambda T.$$

This is a convex optimization problem subject only to non-negativity constraints, and can be efficiently solved using iterative methods (e.g. gradient descent).

C The benefits of sparsity - upper bounds

We restate the main theorem proved in Section 4.

Theorem 4.1. *Consider OMD with ψ_p as regularizer. If $T > e$ is such that $q_{\pi^*} \in \Delta(T)$, $\eta = \sqrt{\frac{pT^{1+1/p}}{KDM^{1/p}}}$, $p = \log(TM)$, then $\mathbb{E}[R_K] \leq \mathcal{O}\left(\sqrt{DKT \log(MT)}\right)$.*

We include the missing details from the proof given in Section 4. First, recall from (4) that

$$\begin{aligned} \psi_p(q) &= p \cdot \left(-1 + \|q\|_{1+1/p}^{1+1/p}\right) = p \cdot \left(-1 + \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} |q(s, a)|^{1+1/p}\right) \\ \implies \frac{\partial \psi_p(q)}{\partial q(s, a)} &= (p+1)q(s, a)^{1/p} \\ \implies \frac{\partial^2 \psi_p(q)}{\partial q(s, a)^2} &= \left(1 + \frac{1}{p}\right)q(s, a)^{1/p-1} \\ \implies \nabla \psi_p(q) &= (p+1)q^{1/p}, \quad \nabla^2 \psi_p(q) = \text{diag}\left(\frac{p+1}{p}q^{1/p-1}\right) \end{aligned}$$

We implicitly assumed here that ψ_p is defined on $\mathbb{R}_{\geq 0}^\Gamma$. The missing details are:

- ψ_p satisfies the condition (2) with $\alpha = 1$:

$$\begin{aligned} \nabla \psi_p(q) \in \left[\nabla \psi_p(q_k), \nabla \psi_p(q_k) - \eta c_k\right] &\implies q^{1/p}(s, a) \leq q_k^{1/p}(s, a) \\ &\implies \frac{1}{q(s, a)} \geq \frac{1}{q_k(s, a)} \\ &\implies \frac{1}{q^{1-1/p}(s, a)} \geq \frac{1}{q_k^{1-1/p}(s, a)} \\ &\implies \nabla^2 \psi_p(q) \geq \nabla^2 \psi_p(q_k). \end{aligned}$$

- $\nabla^2 \psi_p(q)^{-1} = \text{diag}\left(\frac{p}{p+1}q^{1-1/p}\right)$: follows directly from the expression for $\nabla^2 \psi_p(q)$ above.

We now turn to the description of the parameter-free algorithm and the proof of its corresponding regret bound (Theorem 4.4).

C.1 Sparse-agnostic bound

For the unknown sparsity level, we use the same approach as in [19], dividing the episode horizon into segments, where each segment will run OMD from scratch with an increasing sparsity level guess. Crucially, there will be at most $O(\log \log M)$ such segments.

Define

- $M = \max_{k \in [K]} \|c_k\|_0$, the true sparsity level across the horizon.
- $B = \lceil \log_2 \log_2 M \rceil$, the maximum number of segments.
- $m(b) = 2^{2^b}$, the assumed sparsity level during the b -th segment (or interval $I(b)$ below). The reason to use a double exponential is that this sparse-agnostic procedure brings an extra B factor to the regret bound: if we use 2^b then $B = O(\log M)$ which harms the regret bound.
- for $1 \leq b \leq B$, $\tau(b) = \min \{1 \leq k \leq K \mid \|c_k\|_0 > m(b)\}$, the first episode in which the sparsity level of the loss vector exceeds $m(b)$. We also define $\tau(0) = 0$ and $\tau(B) = K$.

Using this notation, we can partition the horizon $[K]$ as intervals $(I(b))_{b \in [B]}$ according to the episodes $\tau(b)$ where the thresholds $m(b)$ are first exceeded. For $1 \leq b \leq B$:

$$I(b) = \begin{cases} [\tau(b-1) + 1, \tau(b)] & \text{if } \tau(b-1) < \tau(b) \\ \emptyset & \text{if } \tau(b-1) = \tau(b) \end{cases}$$

Let $b_k = \min \{b \geq 1 \mid \tau(b) \geq k\}$ be the index of the only interval to which episode k belongs.

Now we define the OMD parameters used in interval $I(b)$, in which we essentially use the parameters from Theorem 4.1 assuming the sparsity level is $m(b)$:

Algorithm 1 Sparse-Agnostic Mirror Descent

Input: T, K, D
Initialize: $p \leftarrow \log 2^2 T$, $\eta \leftarrow \sqrt{T^{1+1/p}/(pDK2^{2/p})}$, $b \leftarrow 1$, $q_1 \leftarrow \arg \min_{q \in \Delta(T)} \psi_p(q)$
for $k = 1, \dots, K$ **do**
 Play q_k and Observe c_k
 if $\|c_k\|_0 \leq 2^{2^b}$ **then**
 $q_{k+1} = \arg \min_{q \in \Delta_T} \langle q, c_k \rangle + D_{\psi_p}(q, q_k)$
 else
 $b \leftarrow \lceil \log_2 \log_2 \|c_k\|_0 \rceil$
 $p \leftarrow \log 2^{2^b} T$
 $\eta \leftarrow \sqrt{pT^{1+1/p}/(DK2^{2/p})}$
 $q_{k+1} \leftarrow \arg \min_{q \in \Delta(T)} \psi_p(q)$
 end if
end for

- the parameter of our regularizer is $p(b) = \log(m(b)T)$.
- the step-size is $\eta(b) = \sqrt{\frac{p(b)T^{1+1/p(b)}}{DKm(b)^{1/p(b)}}}$.

Recall that our regularizer with parameter p is given by $\psi_p(q) = p \left(-1 + \|q\|_{1+1/p}^{1+1/p} \right)$. At episode k , we use the parameter $p(b_k)$ defined above, i.e. using the index value b_k of the interval $I(b_k)$ to which episode k belongs to. The OMD update is then defined by:

$$q_k = \nabla \psi_{p(b_k)}^* \left(\eta(b_k) \sum_{k' < k, k' \in I(b_k)} c_{k'} \right), \quad k = 1, \dots, K.$$

The full procedure is given in Algorithm 1. The following lemma shows the cost of being sparse-agnostic is an additive T term and a double-logarithmic factor in the sparsity level M .

Lemma C.1. *Consider running Algorithm 1. If $T > e$ is such that $q_{\pi^*} \in \Delta(T)$, then $\mathbb{E}[R_K] \leq \mathcal{O}(TB + B\sqrt{DKT \log(MT)})$.*

Proof. Fix (interval) $b \in [B]$. On the time interval $I(b)$, we run OMD with regularizer $\psi_{p(b)}$, learning rate $\eta(b)$ and we consider the (expected) interval regret $R(b) = \sum_{k \in I(b)} \langle q_{\pi_k}, c_k \rangle - \min_{\pi \in \Pi_p} \sum_{k \in I(b)} \langle q_{\pi^*}, c_k \rangle$. Crucially we know that up to the last time step of the interval, we have a bound $m(b)$ on the sparsity for all rounds but the last.

Since $J_k^\pi \leq T$ for any k , we just consider the regret on the rounds not including $\tau(b)$ for which we suffer a regret of at most T :

$$R(b) = \sum_{k \in I(b)} \langle q_{\pi_k} - q_{\pi^*}, c_k \rangle \leq T + \sum_{\substack{k \in I(b) \\ k < \tau(b)}} \langle q_{\pi_k} - q_{\pi^*}, c_k \rangle.$$

For the other rounds we follow similar steps as in the proof of Theorem 4.1:

$$\begin{aligned}
\sum_{\substack{k \in I(b) \\ k < \tau(b)}} \langle q_{\pi_k} - q_{\pi^*}, c_k \rangle &\leq \frac{p(b)T^{1+1/p(b)}}{\eta(b)} + \frac{\eta(b)}{2} \sum_{\substack{k \in I(b) \\ k < \tau(b)}} \|c_k\|_1^{1/p(b)} (\langle c_k, q_k \rangle + 1) \\
&\leq \frac{p(b)T^{1+1/p(b)}}{\eta(b)} + \frac{\eta(b)m(b)^{1/p(b)}}{2} \sum_{\substack{k \in I(b) \\ k < \tau(b)}} (\langle c_k, q_k \rangle + 1) \\
\implies \sum_{\substack{k \in I(b) \\ k < \tau(b)}} \langle q_{\pi_k} - q_{\pi^*}, c_k \rangle &\leq \frac{1}{1 - \eta(b)m(b)^{1/p(b)}} \left[\frac{p(b)T^{1+1/p(b)}}{\eta(b)} + \frac{\eta(b)m(b)^{1/p(b)}}{2} \sum_{\substack{k \in I(b) \\ k < \tau(b)}} (\langle c_k, q_{\pi^*} \rangle + 1) \right]
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{2p(b)T^{1+1/p(b)}}{\eta(b)} + \eta(b)m(b)^{1/p(b)} \sum_{\substack{k \in I(b) \\ k < \tau(b)}} (\langle c_k, q_{\pi^*} \rangle + 1) \\
&\leq \frac{2p(b)T^{1+1/p(b)}}{\eta(b)} + \eta(b)m(b)^{1/p(b)} \sum_{k \in I(b)} (\langle c_k, q_{\pi^*} \rangle + 1),
\end{aligned}$$

where we used that $\frac{1}{1-\eta(b)m(b)^{1/p(b)}} \leq 2$ which is the case if $K > 8eT \log(MT)$ using how we defined $\eta(b)$. Using $\eta(b) = \sqrt{\frac{p(b)T^{1+1/p(b)}}{DKm(b)^{1/p(b)}}}$, we have

$$\begin{aligned}
R(b) &\leq T + \sqrt{p(b)T^{1+1/p(b)}DKm(b)^{1/p(b)}} \cdot \left(2 + \frac{\sum_{k \in I(b)} (\langle c_k, q_{\pi^*} \rangle + 1)}{DK}\right) \\
&\leq T + \sqrt{TDKe \log(m(b)T)} \cdot \left(2 + \frac{\sum_{k \in I(b)} (\langle c_k, q_{\pi^*} \rangle + 1)}{DK}\right) \\
&\leq T + \sqrt{2TDKe \log(MT)} \cdot \left(2 + \frac{\sum_{k \in I(b)} (\langle c_k, q_{\pi^*} \rangle + 1)}{DK}\right),
\end{aligned}$$

where we used that $p(b) = \log(m(b)T)$ and that

$$\begin{aligned}
b \leq B \leq 1 + \log_2 \log_2 M &\implies m(b) = 2^{2^b} \leq 2^{2^{1+\log_2 \log_2 M}} = 2^{2 \cdot \log_2 M} = M^2 \\
&\implies \log(m(b)T) = \log(M^2 T) \leq 2 \log(MT).
\end{aligned}$$

Then

$$\begin{aligned}
E[R_K] &\leq \sum_{b=1}^B R(b) \\
&\leq TB + \sqrt{2TDKe \log(MT)} \cdot \left(2B + \sum_{b=1}^B \frac{\sum_{k \in I(b)} (\langle c_k, q_{\pi^*} \rangle + 1)}{DK}\right) \\
&= TB + \sqrt{2TDKe \log(MT)} \cdot \left(2B + \frac{1}{D} + \frac{1}{DK} \sum_{k=1}^K \langle c_k, q_{\pi^*} \rangle\right) \\
&\leq TB + 4B\sqrt{2TDKe \log(MT)},
\end{aligned}$$

where the last step uses that $\sum_{k=1}^K \langle q_{\pi^*}, c_k \rangle \leq \sum_{k=1}^K \langle q_{\pi_f}, c_k \rangle \leq K \|q_{\pi_f}\|_1 \|c_k\|_\infty \leq DK$ and $D \geq 1$. $\square$

C.2 Fully parameter-free bound

We now turn our attention to the unknown hitting time of the optimal policy $T_\star$, where we can exploit the same technique presented in [8].

We run $N \approx \log K$ instances of Algorithm 1 where the j -th instance will set its parameter T as $b(j)$ which is roughly 2^j , so that there always exists an instance $j_\star$ such that $b(j_\star)$ is close to the unknown $T_\star$. Specifically, we run a scale invariant meta algorithm with a correction term as in [8] to obtain the desired bound (details in Algorithm 2).

Theorem 4.4. *If $K > \max(T_\star, \frac{T_\star}{D} \log(T_\star M))$ and $T_\star > e$, Algorithm 2 guarantees $\mathbb{E}[R_K] \leq \tilde{\mathcal{O}}\left(\sqrt{DKT_\star \log(MT_\star)} + T_\star\right)$, where the notation $\tilde{\mathcal{O}}$ hides double-logarithmic factors.*

Proof. We closely follow the steps of the proof of Theorem 2 in [8]. We have

- $b(1) \geq T^{\pi_f}(s_0)$ so for all instances $\Delta(b(j))$ is non-empty and the instance is well-defined.
- Let $j^\star$ be the index of the instance with smallest $b(j^\star)$ that is larger than $T^\star$, i.e. $\frac{b(j^\star)}{2} \leq T_\star \leq b(j^\star)$. This instance exists since $b(N) \geq K > T_\star$.

Algorithm 2 Fully Parameter-Free Online Mirror Descent for Sparse SSPs

Define $j_0 = \lceil \log_2 T^{\pi_f}(s_0) \rceil - 1$, $b(j) = 2^{j_0+j}$, $N = \lceil \log_2 K \rceil - j_0$, $\eta_j = (\sqrt{Db(j)K \log(b(j)M)})^{-1/2}$
Define $\psi_p(p) = \sum_{j=1}^N \frac{1}{\eta_j} p(j) \log p(j)$
Initialize: $p_1 \in \Delta_N$, such that $p_1(j) = \frac{\eta_j}{\eta_1 N}$, $\forall j \neq 1$
Initialize: N instances of Algorithm 1, with j -th instance $T = b(j)$
for $k = 1, \dots, K$ **do**
 Obtain occupancy measures q_k^j for $j \in [N]$
 Sample $j_k \sim p_k$, execute policy induced by $q_k^{j_k}$
 Receive c_k and send it to all instances.
 Compute $\ell_k(j) = \langle q_k^j, c_k \rangle$, $a_k(j) = 4\eta_j \ell_k^2(j)$
 Update $p_{k+1} = \arg \min_{p \in \Delta_N} \langle p, \ell_k + a_k \rangle + D_{\psi_p}(p, p_k)$
end for

We start by decomposing the regret into the regret of the meta algorithm *w.r.t.* finding $j_\star$ and the regret of the $j_\star$ instance *w.r.t.* the best policy:

$$\begin{aligned} \mathbb{E}[R_K] &= \sum_{k=1}^K \sum_{j=1}^N p_k(j) \langle q_k^j, c_k \rangle - \sum_{k=1}^K \langle q_{\pi^\star}, c_k \rangle \\ &= \sum_{k=1}^K \sum_{j=1}^N p_k(j) \langle q_k^j, c_k \rangle - \langle q_k^{j_\star}, c_k \rangle + \sum_{k=1}^K \langle q_k^{j_\star} - q_{\pi^\star}, c_k \rangle \\ &= \underbrace{\sum_{k=1}^K \langle p_k - e_{j_\star}, \ell_k \rangle}_{\text{Meta-Regret}} + \underbrace{\sum_{k=1}^K \langle q_k^{j_\star} - q_{\pi^\star}, c_k \rangle}_{j_\star\text{-Regret}} \end{aligned}$$

where we consider p_k, ℓ_k as N -dimensional vectors and $e_{j_\star}$ as the basis vector with the $j_\star$ coordinate equal to 1.

By Lemma C.1 the $j_\star$ -Regret is bounded by $\mathcal{O}\left(Bb(j_\star) + B\sqrt{Db(j_\star)K \log(b(j_\star)M)}\right) = \mathcal{O}\left(BT_\star + B\sqrt{DT_\star K \log(T_\star M)}\right)$.

This also allows to say that:

$$\begin{aligned} \sum_{k=1}^K \langle q_k^{j_\star}, c_k \rangle &\leq \sum_{k=1}^K \langle q_{\pi^\star}, c_k \rangle + \mathcal{O}\left(BT_\star + B\sqrt{DT_\star K \log(T_\star M)}\right) \\ &\leq \sum_{k=1}^K \langle q_{\pi_f}, c_k \rangle + \mathcal{O}\left(BT_\star + B\sqrt{DT_\star K \log(T_\star M)}\right) \\ &\leq DK + \mathcal{O}\left(BT_\star + B\sqrt{DT_\star K \log(T_\star M)}\right), \end{aligned}$$

which we will make use of just below.

For the meta algorithm regret, we can use Lemma 12 of [8] which guarantees that:

$$\begin{aligned} \mathbb{E}\left[\sum_{k=1}^K \langle p_k - e_{j_\star}, \ell_k \rangle\right] &= \mathcal{O}\left(\frac{2 + \log\left(N\sqrt{\frac{b(j_\star)}{b(1)}}\right)}{\eta_{j_\star}} + 4\eta_{j_\star} b(j_\star) \sum_{k=1}^K \langle q_k^{j_\star}, c_k \rangle\right) \\ &= \mathcal{O}\left(\frac{2 + \log\left(N\sqrt{\frac{b(j_\star)}{b(1)}}\right)}{\eta_{j_\star}} + 4\eta_{j_\star} b(j_\star) \left(DK + BT_\star + B\sqrt{DT_\star K \log(T_\star M)}\right)\right) \end{aligned}$$

$$\begin{aligned}
&= \tilde{\mathcal{O}}\left(\frac{\log T_\star}{\eta_{j_\star}} + \eta_{j_\star} T_\star \left(DK + \sqrt{DT_\star K \log(T_\star M)}\right)\right) \\
&= \tilde{\mathcal{O}}\left(\frac{\log T_\star}{\eta_{j_\star}} + \eta_{j_\star} DT_\star K\right) \\
&= \tilde{\mathcal{O}}\left(\sqrt{DT_\star K \log T_\star}\right),
\end{aligned}$$

where we used

- the notation $\tilde{\mathcal{O}}$ to ignore *double*-logarithmic factors.
- $K \geq \frac{T_\star}{D} \log(T_\star M)$.
- $\eta_{j_\star} = \sqrt{\frac{\log T_\star}{DT_\star K}}$.
- $D \geq 1$.

Combining everything gives the result. □

D State level sparsity

In this section, we consider a different notion of state-level sparsity:

$$M' = \max_k \max_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mathbb{I}\{c_k(s, a) > 0\} \leq A.$$

From Theorem 4 in [38], we have a cumulative loss bound for a version of OMD with negative entropy regularisation:

$$\mathbb{E}[R_K] \leq \tilde{\mathcal{O}}\left(\sqrt{T_\star \sum_{k=1}^K J_k^{\pi_\star}}\right),$$

from which we can exploit the state-level sparsity:

$$\begin{aligned} \sum_{k=1}^K J_k^{\pi_\star} &= \min_{\pi} \sum_{k=1}^K \mathbb{E}\left[\sum_{t=1}^{I_\pi(s_0)} c_k(s^t, a^t) | P, \pi, s^1 = s_0\right] \\ &\leq \sum_{k=1}^K \mathbb{E}\left[\sum_{t=1}^{I_{\pi_u}(s_0)} c_k(s^t, a^t) | P, \pi_u, s^1 = s_0\right] \\ &= \sum_{k=1}^K \mathbb{E}\left[\sum_{t=1}^{I_{\pi_u}(s_0)} \frac{1}{A} \sum_{a \in \mathcal{A}} c_k(s^t, a) | P, \pi_u, s^1 = s_0\right] \\ &\leq \frac{KM'}{A} \mathbb{E}[I_{\pi_u}(s_0) | P, \pi_u, s^1 = s_0] \\ &= \frac{KM'}{A} T^u(s_0), \end{aligned}$$

where π_u is the uniform policy ($\pi_u(a|s) = 1/A$) and T^u is its corresponding hitting-time. This gives the following regret bound:

$$\mathbb{E}[R_K] \leq \tilde{\mathcal{O}}\left(\sqrt{KT_\star \frac{M'}{A} T^u(s_0)}\right).$$

We can actually relate this result back to our original notion of sparsity since we know that $M' \leq M$. If $M \leq A$, then we can non-trivially bound M' by M and achieve a regret of $\tilde{\mathcal{O}}\left(\sqrt{KT_\star \frac{M}{A} T^u(s_0)}\right)$.

This result highlights that it is possible to achieve polynomial improvements from sparsity if we consider state-level sparsity M' or $M < A$. However, this comes at the cost of a $T^u(s_0)$ factor. In the worst case, this additional factor will cancel the polynomial improvement. It could be an interesting avenue of future research to understand specific structural properties of the MDP that may lead to real polynomial improvements.

In the experts setting ($S = 1$), we have $M = M'$ and $T^u(s_0) = 1$ and this bound recovers the $\frac{M}{A}$ improvement of the expert setting. This provides some further insights into the performance of OMD with negative entropy regularisation and that in particular issues arise when there is at least 1 state with non-sparse costs even though most other states may have sparse costs.

E The benefits of sparsity - lower bound under sparsity

In this appendix, we prove our sparse lower bound result from Section 4.2. We first restate the result.

Theorem 4.6. *For any D, T_*, K, S, A with $T_* \geq D \geq 3 \log S$, $S(A-1) \geq 400$, $K \geq \frac{800T_*}{D} \log M$ and $M \geq 101$, there exists an SSP instance with stochastic M -sparse costs, S states and A actions such that its diameter is D , the expected hitting time of the optimal policy is T_* , and the expected regret with respect to the randomness of the losses for any learner after K episodes is $\mathbb{E}[R_K] \geq \Omega(\sqrt{KT_*D \log M})$.*

Proof. Fix $B = \lceil \frac{\log S/2}{\log 2} \rceil - 2$. Fix $N = 2^{B+1} \geq 2^{\frac{\log S/2}{\log 2} - 1} = \frac{S}{4}$. Fix $L = \min\{M-1, N \cdot (A-1)\} \geq \frac{M}{8}$. Fix $D' = D - B - 2$, $T' = T_* - B - 1$, with $T_* \geq D$.

We first describe the SSP instance with stochastic costs. Consider the following MDP $\mathcal{M} = (S, \mathcal{A}, p, s_0, g)$ illustrated in Figure 4 and that we formally define below.




Figure 4: Diagram illustrating MDP construction for the proof of Theorem 4.6. Details are given below.

The first part of the states are represented by a binary tree of depth $B+2$ and allow us to formerly consider the N states at the bottom of the tree that matter, while avoiding an assumption on the existence of a state with $A \approx S$ actions as was done in prior work [8]. Each non-leaf node corresponds to a state with two actions transitioning (deterministically) to the left or right child respectively. The total number of nodes in the tree is

$$\sum_{i=0}^{B+1} 2^i = 2^{B+2} - 1 \leq 2^{\frac{\log S/2}{\log 2} + 1} - 1 = 2 \frac{S}{2} - 1 = S - 1.$$

The total number of leaf nodes is $N = 2^{B+1} \geq \frac{S}{4}$. Denote the set of states corresponding to the leaf nodes by $\mathcal{S}_\ell = \{s_1, \dots, s_N\}$. The root node is s_0 . There is also one additional state denoted by f (recall that the number of states in the tree is $\leq S-1$).

We consider the same action set across each state: $\mathcal{A} = \{a_1, \dots, a_{A-1}, a_f\}$. In the states of the binary tree where we have only described two actions, we can consider the other actions to remain in the same state deterministically with 0 cost.

The transitions and costs are defined as follows:

- For all states and actions in the tree that are not leaves, the transitions are specified above. The costs are all 0.
- For $s_i \in \mathcal{S}_\ell$, and $a_j \in \mathcal{A}$

- if $j = f$, then $p(f|s_i, a_f) = 1$ and $c_k(s_i, a_f) = 0$.
- if $j \in \{1, \dots, A-1\}$ and $j + (A-1) \cdot (i-1) \leq L$, then $p(g|s_i, a_j) = \frac{1}{T'}$, $p(s_i|s_i, a_j) = 1 - \frac{1}{T'}$ and the cost is an independent sample from a Bernoulli distribution at each episode k : $c_k(s_i, a_j) \sim \text{Ber}\left(\frac{D'}{2T'}\right)$.
- if $j \in \{1, \dots, A-1\}$ and $j + (A-1) \cdot (i-1) > L$, then a_j is the same as a_f , i.e. $p(f|s_i, a_j) = 1$ and $c_k(s_i, a_j) = 0$.
- For f ,
 - $p(g|f, a_f) = \frac{1}{D'}$, $p(f|f, a_f) = 1 - \frac{1}{D'}$ and $c_k(f, a_f) = 1$.
 - for all $a_j \in \mathcal{A} \setminus \{a_f\}$, $p(f|f, a_j) = 1$ and $c_k(f, a_j) = 0$.

Denote the above distribution for c_k by $\mathcal{D}$. In each episode there are at most $L+1 \leq M$ non-zero costs, ensuring the condition on sparsity is respected.

For $i \in \{1, \dots, N\}$, let $\mathcal{A}_i$ correspond to the actions in state $s_i \in \mathcal{S}_L$ which can transition directly to g and $\mathcal{A} \setminus \mathcal{A}_i$ corresponds to the actions which deterministically transition to f (e.g. if $(A-1) \cdot i \leq L$, then $\mathcal{A}_i = \{a_1, \dots, a_{A-1}\}$). For any proper policy π independent of the stochastically generated costs in episode k , we have

$$\begin{aligned}
\mathbb{E}_{c_k \sim \mathcal{D}}[J_k^\pi(s_i)] &= \mathbb{E}_{c_k \sim \mathcal{D}}\left[\sum_{a \in \mathcal{A}_i} \pi(a|s_i) \left(c_k(s_i, a) + \left(1 - \frac{1}{T'}\right) J_k^\pi(s_i)\right) + J_k^\pi(f) \sum_{a \notin \mathcal{A}_i} \pi(a|s_i)\right] \\
&= \sum_{a \in \mathcal{A}_i} \pi(a|s_i) \left(\mathbb{E}_{c_k \sim \mathcal{D}}[c_k(s_i, a)] + \left(1 - \frac{1}{T'}\right) \mathbb{E}_{c_k \sim \mathcal{D}}[J_k^\pi(s_i)]\right) + D' \cdot \sum_{a \notin \mathcal{A}_i} \pi(a|s_i) \\
&= \sum_{a \in \mathcal{A}_i} \pi(a|s_i) \left(\frac{D'}{2T'} + \left(1 - \frac{1}{T'}\right) \mathbb{E}_{c_k \sim \mathcal{D}}[J_k^\pi(s_i)]\right) + D' \cdot \left(1 - \sum_{a \in \mathcal{A}_i} \pi(a|s_i)\right) \\
\implies \mathbb{E}_{c_k \sim \mathcal{D}}[J_k^\pi(s_i)] &\left(1 - \left(1 - \frac{1}{T'}\right) \sum_{a \in \mathcal{A}_i} \pi(a|s_i)\right) = \frac{D'}{2T'} \sum_{a \in \mathcal{A}_i} \pi(a|s_i) + D' \cdot \left(1 - \sum_{a \in \mathcal{A}_i} \pi(a|s_i)\right) \\
\implies \mathbb{E}_{c_k \sim \mathcal{D}}[J_k^\pi(s_i)] &= \frac{\frac{D'}{2T'} \sum_{a \in \mathcal{A}_i} \pi(a|s_i) + D' \cdot \left(1 - \sum_{a \in \mathcal{A}_i} \pi(a|s_i)\right)}{1 - \left(1 - \frac{1}{T'}\right) \sum_{a \in \mathcal{A}_i} \pi(a|s_i)} \\
&= \frac{D'}{2} \cdot \frac{\frac{1}{T'} \sum_{a \in \mathcal{A}_i} \pi(a|s_i) + 2\left(1 - \sum_{a \in \mathcal{A}_i} \pi(a|s_i)\right)}{\frac{1}{T'} \sum_{a \in \mathcal{A}_i} \pi(a|s_i) + \left(1 - \sum_{a \in \mathcal{A}_i} \pi(a|s_i)\right)} \\
&\geq \frac{D'}{2}.
\end{aligned}$$

The optimal policy π^* is the policy that takes actions in the binary tree to reach state s_{i^*} and then $\pi^*(a_{j^*}|s_{i^*}) = 1$ for $i^*, j^* = \arg \min_{i, j: j + (A-1) \cdot i \leq L} \sum_{k=1}^K c_k(s_i, a_j)$. We have $J_k^{\pi^*}(s_0) = J_k^{\pi^*}(s_{i^*})$ and for any $k \geq 1$

$$\begin{aligned}
J_k^{\pi^*}(s_{i^*}) &= c_k(s_{i^*}, a_{j^*}) + \left(1 - \frac{1}{T'}\right) J_k^{\pi^*}(s_{i^*}) \\
\implies J_k^{\pi^*}(s_0) &= T' c_k(s_{i^*}, a_{j^*}) \\
\implies \sum_{k=1}^K J_k^{\pi^*}(s_0) &= T' \sum_{k=1}^K c_k(s_{i^*}, a_{j^*}) = T' \min_{i, j: j + (A-1) \cdot i \leq L} \sum_{k=1}^K c_k(s_i, a_j).
\end{aligned}$$

Hence,

$$\mathbb{E}_{c_1, \dots, c_K \sim \mathcal{D}}[R_K] \geq \frac{D'}{2} \cdot K - T' \cdot \mathbb{E}_{c_1, \dots, c_K \sim \mathcal{D}}\left[\min_{i, j: j + (A-1) \cdot i \leq L} \sum_{k=1}^K c_k(s_i, a_j)\right]$$

$$\begin{aligned}
&= T' \cdot \left(\frac{D'}{2T'} \cdot K - \mathbb{E}_{c_1, \dots, c_K \sim \mathcal{D}} \left[\min_{i, j: j + (A-1) \cdot i \leq L} \sum_{k=1}^K c_k(s_i, a_j) \right] \right) \\
&= T' \cdot \mathbb{E}_{c_1, \dots, c_K \sim \mathcal{D}} \left[\max_{i, j: j + (A-1) \cdot i \leq L} \sum_{k=1}^K \left(\frac{D'}{2T'} - c_k(s_i, a_j) \right) \right]
\end{aligned}$$

We now apply Theorem G.1 with $p = 1 - \frac{D'}{2T'} \geq \frac{1}{2}$, $d = L \geq 100$ (since $S(A-1) \geq 400$ and $M \geq 101$) and $n = K \geq \frac{800T_\star}{D} \log M \geq \frac{400T'}{D'} \log M \geq 200 \frac{p}{1-p} \log d$. We obtain:

$$\begin{aligned}
\sup_{c_1, \dots, c_K} \mathbb{E}[R_K] &\geq \mathbb{E}_{c_1, \dots, c_K \sim \mathcal{D}}[R_K] \geq 0.02T' \sqrt{K \left(1 - \frac{D'}{2T'}\right) \cdot \frac{D'}{2T'} \cdot \log L - 1.5T'} \\
&= \Omega\left(\sqrt{KT_\star D \log M}\right),
\end{aligned}$$

since $L \geq M/8$. Note that since $T_\star \geq D$, the hitting-time of the fast-policy is $D' + B + 2 = D$ and the hitting time of the optimal is $T' + B + 1 = T_\star$, as required. This concludes the proof. $\square$

F Lower bound under unknown transitions

Theorem 5.1. *For any D, K, S, A with $S \geq 2, A \geq 16, D \geq 2$ and $K \geq SA$, there exists an SSP instance with $M = 1$, S states and A actions such that its diameter is D and the expected regret for any learner without knowledge of the transitions after K episodes is $\mathbb{E}[R_K] \geq \Omega(D\sqrt{SAK})$.*




Figure 5: base case

Proof. The idea is to inject sparsity into the lower bound construction of [29] and to see if sparsity helps. Imagine the simple SSP in Figure 5, where at state s_0 there are A available actions, all with zero cost, while in the state f there is only one deterministic action with unit cost going back to s_0 . Among them, there exists an action a_* such that the transition probabilities are given by: $P(g \mid s_0, a) = \frac{1}{D} - \epsilon \mathbb{I}(a \neq a_*)$, and consequently, $P(f \mid s_0, a) = 1 - \frac{1}{D} + \epsilon \mathbb{I}(a \neq a_*)$. The cost is therefore only suffered when the selected action transitions to the f state. This will therefore not increase the hitting time of any proper deterministic policy while still inducing the desired sparsity.

Clearly, the optimal policy plays a_* at every time step to reach the goal as fast as possible and therefore $J^{\pi^*}(s_0) = D - 1$.

Now, denote with N_k the number of steps that the learner spends in s_0 in episode k and N_k^* the number of steps that the learner picks action a_* in episode k . Note that N_k is also the total cost that the learner suffers during episode k minus one (since the last transition will not be paid). Thanks to our construction we can still prove Lemma C.1 in [29] as follows:

Lemma F.1.

$$\mathbb{E}[N_k] - 1 - J^{\pi^*} = \epsilon \mathbb{E}[N_k - N_k^*]$$

Proof.

$$\begin{aligned} \mathbb{E}[N_k] &= \sum_{t=1}^{\infty} P[s_t = s_0] \\ &= 1 + \sum_{t=2}^{\infty} P[s_t = s_0] \\ &= 1 + \sum_{t=2}^{\infty} P[s_t = s_0 \mid s_{t+1} = s_0, a_{t-1} = a_*] P[s_{t+1} = s_0, a_{t-1} = a_*] \\ &\quad + \sum_{t=2}^{\infty} P[s_t = s_0 \mid s_{t+1} = s_0, a_{t-1} \neq a_*] P[s_{t+1} = s_0, a_{t-1} \neq a_*] \\ &= 1 + \left(1 - \frac{1}{D}\right) \sum_{t=2}^{\infty} P[s_{t+1} = s_0, a_{t-1} = a_*] + \left(1 - \frac{1-\epsilon}{D}\right) \sum_{t=2}^{\infty} P[s_{t+1} = s_0, a_{t-1} \neq a_*] \end{aligned}$$

Rearranging gives:

$$\mathbb{E}[N_k] - D = \epsilon \mathbb{E}[N_k - N_k^*]$$

Adding and subtracting 1 gives the desired result. $\square$

Hence:

$$\mathbb{E}[R_K] = \sum_{k=1}^K \sum_{i=1}^{I_k} c_k(s_k^i, a_k^i) - \sum_{k=1}^K J^{\pi^*}(s_0) = \mathbb{E}[N_K] - 1 - J^{\pi^*} = \epsilon[N - N^*]$$

where $N = \sum_{k=1}^K N_k$ and $N^* = \sum_{k=1}^K N_k^*$. Since we recovered Lemma C.1 in [29] as the starting block of the proof, following the derivation we can lower bound N in expectation and upper bound the expected value of N^* to retrieve

Lemma F.2 (Theorem C.4 in [29]). *Suppose that $D \geq 2$, $\epsilon \in (0, 1/8)$ and $A > 16$, for the problem described above we have:*

$$\mathbb{E}[R_K] \geq \epsilon K D \left(\frac{1}{8} - 2\epsilon \sqrt{\frac{2K}{A}} \right)$$

Now consider the following MDP. Let $\mathcal{S}$ be the set of states disregarding g and s_0 . The initial state s_0 has only one action which leads uniformly at random into one of the states $s \in \mathcal{S}$, where each one has its own optimal action a_s^* . Then the transition distributions are defined $P(g | a_s^*, s) = 1/D$, $P(s | a_s^*, s) = 1 - 1/D$, and $P(g | a, s) = (1 - \epsilon)/D$, $P(s | a, s) = 1 - (1 - \epsilon)/D$ for any other action $a \in \mathcal{A} \setminus \{a_s^*\}$. Note that for each state, the learner is faced with a simple problem as the one described above. Therefore, we can apply Lemma F.2 for each state separately and lower bound the learner's expected regret the sum of the regrets suffered at each state, which would depend on the number of times each state $s \in \mathcal{S}$ is visited from the initial state. Since reaching each state has uniform probability, there are many states (constant fraction) that are chosen $\Theta(K/S)$ times. Summing the regret bounds and choosing ϵ , gives the desired bound.

Denote by K_s the number of episodes the state $s \in \mathcal{S}$ is visited:

$$\mathbb{E}[R_K] \geq \sum_{s \in \mathcal{S}} \mathbb{E} \left[\epsilon K_s D \left(\frac{1}{8} - 2\epsilon \sqrt{\frac{2K_s}{A}} \right) \right] = \frac{\epsilon K D}{8} - 2\epsilon^2 D \sqrt{\frac{2}{A}} \sum_{s \in \mathcal{S}} \mathbb{E} [K_s^{3/2}]$$

Then:

$$\sum_{s \in \mathcal{S}} \mathbb{E} [K_s^{3/2}] \leq \sum_{s \in \mathcal{S}} \sqrt{\mathbb{E} [K_s]} \sqrt{\mathbb{E} [K_s^2]} = \sum_{s \in \mathcal{S}} \sqrt{\mathbb{E} [K_s]} \sqrt{\mathbb{E} [K_s^2] + \mathbb{V} [K_s]} = \sum_{s \in \mathcal{S}} \sqrt{\frac{K}{S}} \sqrt{\frac{K^2}{S^2} + \frac{K(S-1)}{S^2}} \leq K \sqrt{\frac{2K}{S}}$$

Leading to:

$$\mathbb{E}[R_K] \geq \frac{\epsilon K D}{8} - 2\epsilon^2 D K \sqrt{\frac{2K}{SA}} \geq \frac{1}{1024} D \sqrt{SAK}$$

for $\epsilon = 1/64 \sqrt{SA/K}$ $K \geq SA$, concluding the proof. $\square$

G Lower bound on the maximum of asymmetric zero-mean random walks

We extend the lower bound of [25] to asymmetric zero-mean random-walks. We consider $p \geq 1/2$ because it simplifies the proof below (lower-bounding ψ by 1 and upper-bounding C in proof below) and is what we need in the proof of Section 4.6 in Appendix E (we use $p = 1 - D/2T^*$).

Theorem G.1. Fix $p \in [\frac{1}{2}, 1 - \frac{1}{n}]$. Consider random walks $Z_i^{(n)} = \sum_{t=1}^n X_t^i$, where

$$X_t^i = \begin{cases} -p, & \text{w.p. } 1-p \\ 1-p, & \text{w.p. } p. \end{cases}$$

If $n \geq 200 \frac{p}{1-p} \log d$ (also ensures that $p \leq 1 - \frac{1}{n}$) and $d \geq 100$. Then,

$$\mathbb{E}[\max_{1 \leq i \leq d} Z_i^n] \geq 0.02 \sqrt{np(1-p) \log d} - 1.5.$$

Proof. We follow the same lines as [25] who show a special case of the result for $p = 1/2$. We generalize it to $p > 1/2$.

Consider $Z^{(n)} = \sum_{t=1}^n X_t$, a random-walk of length n , then $B_n = Z^{(n)} + pn \sim B(n, p)$, Binomial distribution with parameters n and p .

G.1 1st part of the proof:

The 1st part of the proof is all about providing a lower bound on $\mathbb{P}(B_n \geq pn + t - 1)$ in (10) for any $t \in [1, np + 1]$.

Lemma G.2 (Generalized version of Lemma 4 of [25], Theorem 2 of [21]). Let n, k be integers satisfying $n \geq 1$ and $pn \leq k \leq n$. Define $x = \frac{k-pn}{\sqrt{p(1-p)n}}$. Then $B_n \sim B(n, p)$ satisfies

$$\mathbb{P}(B_n \geq k) \geq \sqrt{n} \binom{n-1}{k-1} p^{k-1/2} (1-p)^{n-k+1/2} \cdot \frac{1-\Phi(x)}{\phi(x)},$$

where $\phi(x)$ is the PDF of a standard Normal and $\Phi(x)$ is the CDF. The proof can be found in [21].

Denote $D(p, q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ as the KL-divergence between two Bernoullis.

Lemma G.3 (Generalized version of Theorem 5 of [25]). Let n, k be integers satisfying $n \geq 1$, and $np \leq k \leq n$. Define $x = \frac{k-pn}{\sqrt{p(1-p)n}}$. Then $B_n \sim B(n, p)$ satisfies

$$\mathbb{P}(B_n \geq k) \geq \frac{\exp\left(-nD\left(\frac{k}{n}, p\right)\right)}{e^{1/6} \sqrt{2\pi}} \cdot \frac{1-\Phi(x)}{\phi(x)}.$$

Proof. For $k = n$, we verify the statement of the theorem directly. The left hand side is $\mathbb{P}(B_n \geq n) = p^n$. The right hand side is smaller because $\exp\{-nD(1, p)\} = p^n$ and for $x = \sqrt{n \frac{1-p}{p}} > 0$, we have $\frac{1-\Phi(x)}{\phi(x)} \leq \sqrt{2}$ (see e.g. Section 3.3 in [15]).

For $np \leq k < n$, we first bound the binomial coefficient $\binom{n}{k}$. Stirling's formula for the factorial [27] gives for any $n \geq 1$,

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n < n! < e^{1/12} \sqrt{2\pi n} \left(\frac{n}{e}\right)^n.$$

Since $0 < np \leq k \leq n-1$, we can use this approximation for k, n and $n-k$ and obtain

$$\begin{aligned} \binom{n}{k} &= \frac{n!}{k!(n-k)!} \\ &> \frac{n^n e^{-n} \sqrt{2\pi n}}{(e^{1/12} k^k e^{-k} \sqrt{2\pi k}) \cdot (e^{1/12} (n-k)^{n-k} e^{-(n-k)} \sqrt{2\pi(n-k)})} \\ &= \frac{1}{e^{1/6} \sqrt{2\pi}} \left(\frac{n}{k}\right)^k \left(\frac{n}{n-k}\right)^{n-k} \sqrt{\frac{n}{k(n-k)}} \end{aligned}$$

$$= \frac{1}{e^{1/6}\sqrt{2\pi}} \frac{1}{p^k(1-p)^{n-k}} \exp\left\{-nD\left(\frac{k}{n}, p\right)\right\} \sqrt{\frac{n}{k(n-k)}},$$

since

$$\begin{aligned} D\left(\frac{k}{n}, p\right) &= \frac{k}{n} \log\left(\frac{k}{np}\right) + \left(1 - \frac{k}{n}\right) \log\left(\frac{1-k/n}{1-p}\right) = -\frac{k}{n} \log\left(\frac{np}{k}\right) - \frac{n-k}{n} \log\left(\frac{n(1-p)}{n-k}\right) \\ \implies \exp\left\{-nD\left(\frac{k}{n}, p\right)\right\} &= \left(\frac{np}{k}\right)^k \cdot \left(\frac{n(1-p)}{n-k}\right)^{n-k} = p^k(1-p)^{n-k} \left(\frac{n}{k}\right)^k \left(\frac{n}{n-k}\right)^{n-k}. \end{aligned}$$

Since $k \geq 1$, we can write the binomial coefficient as $\binom{n-1}{k-1} = \frac{k}{n} \binom{n}{k}$. By Lemma G.2, we have

$$\begin{aligned} \mathbb{P}(B_n \geq k) &\geq \sqrt{n} \binom{n-1}{k-1} p^{k-1/2} (1-p)^{n-k+1/2} \cdot \frac{1-\Phi(x)}{\phi(x)} \\ &= \sqrt{n} \frac{k}{n} \binom{n}{k} p^{k-1/2} (1-p)^{n-k+1/2} \cdot \frac{1-\Phi(x)}{\phi(x)} \\ &\geq \frac{1}{e^{1/6}\sqrt{2\pi}} \frac{k}{\sqrt{n}} \sqrt{\frac{n}{k(n-k)}} \frac{p^{k-1/2} (1-p)^{n-k+1/2}}{p^k(1-p)^{n-k}} \exp\left\{-nD\left(\frac{k}{n}, p\right)\right\} \cdot \frac{1-\Phi(x)}{\phi(x)} \\ &= \frac{1}{e^{1/6}\sqrt{2\pi}} \sqrt{\frac{k}{n-k}} \cdot \sqrt{\frac{1-p}{p}} \exp\left\{-nD\left(\frac{k}{n}, p\right)\right\} \cdot \frac{1-\Phi(x)}{\phi(x)}. \end{aligned}$$

The result follows from $\sqrt{\frac{k}{n-k}} \geq \sqrt{\frac{np}{n-np}} = \sqrt{\frac{p}{1-p}}$ for $np \leq k \leq n-1$. $\square$

For $k = pn + xn$, the 2nd-order Taylor approximation of $u(x) = D\left(\frac{k}{n}, p\right) = D(p+x, p)$ around 0 is $\frac{x^2}{2p(1-p)}$. We define $\psi : [-p, 1-p] \rightarrow \mathbb{R}$ as the ratio of the divergence and the approximation:

$$\psi(x) = D(p+x, p) \cdot \frac{2p(1-p)}{x^2}.$$

In particular, we have that $1 \leq \psi(x) \leq \frac{p(1-p)}{(x+p)(1-p-x)}$ for $x \in [0, 1-p]$. This can be shown using Taylor's theorem on $u(x)$: for some $z \in [0, x]$,

$$\begin{aligned} D(p+x, p) &= \frac{u^{(2)}(z)}{2} x^2 = \left(\frac{1}{z+p} + \frac{1}{1-p-z}\right) \frac{x^2}{2} = \frac{x^2}{2(z+p)(1-p-z)} \\ \implies \frac{x^2}{2p(1-p)} &\leq D(p+x, p) \leq \frac{x^2}{2(x+p)(1-p-x)}, \end{aligned} \tag{9}$$

since $\frac{1}{(x+p)(1-p-x)}$ is increasing on $[0, 1-p]$.

Let $t \in [1, np+1]$ be a real number. By Lemma G.3 and Lemma 1 in [25] (also Mill's ratio for standard Gaussian [2]), we have

$$\begin{aligned} \mathbb{P}(B_n \geq pn+t-1) &= \mathbb{P}(B_n \geq \lceil pn+t-1 \rceil) \\ &\geq \frac{\exp\left(-nD\left(\frac{\lceil pn+t-1 \rceil}{n}, p\right)\right)}{e^{1/6}\sqrt{2\pi}} \cdot \frac{\pi}{\pi \frac{\lceil pn+t-1 \rceil - np}{\sqrt{p(1-p)n}} + \sqrt{2\pi}} \\ &\geq \frac{\exp\left(-nD\left(\frac{pn+t}{n}, p\right)\right)}{e^{1/6}\sqrt{2\pi}} \cdot \frac{\pi}{\pi \frac{pn+t-np}{\sqrt{p(1-p)n}} + \sqrt{2\pi}} \\ &= \frac{\exp\left(-nD\left(p + \frac{t}{n}, p\right)\right)}{e^{1/6}\sqrt{2\pi}} \cdot \frac{\pi}{\pi \frac{t}{\sqrt{p(1-p)n}} + \sqrt{2\pi}} \\ &= e^{-1/6} \exp\left(-\frac{1}{2p(1-p)} \psi\left(\frac{t}{n}\right) \cdot \frac{t^2}{n}\right) \cdot \frac{1}{\frac{\sqrt{2\pi}t}{\sqrt{np(1-p)}} + 2}. \end{aligned} \tag{10}$$

G.2 2nd part of the proof:

We can now turn to the actual proof of the result. Define the event A equal to the case that at least one of the $Z_i^{(n)}$ is greater or equal to $C\sqrt{np(1-p)\log d} - 1$. We will show this event / threshold controls the expectation of the maximum. First, we define C and provide some upper and lower bounds for it. Denote by $f(d) = \sqrt{2 - \frac{2\log\log d}{\log d}}$, then

$$C = C(d, n) = \frac{1}{\sqrt{\psi\left(\sqrt{\frac{2p(1-p)\log d}{n}}\right)}} \sqrt{2 - \frac{2\log\log d}{\log d}} = \frac{1}{\sqrt{\psi\left(\sqrt{\frac{2p(1-p)\log d}{n}}\right)}} f(d). \quad (11)$$

We bound the two factors separately:

- $z = \sqrt{\frac{2p(1-p)\log d}{n}} \in [0, \frac{1}{10}(1-p)]$ for $n \geq 200\frac{p}{1-p}\log d$ and so
$$1 \leq \psi(z) \leq \frac{p(1-p)}{(z+p)(1-p-z)} \leq \frac{1-p}{(1-p-\frac{1-p}{10})} = \frac{10}{9}. \quad (12)$$

- The function $f(d)$ is as in [25]: decreasing on $(1, e^e]$, increasing on $[e^e, +\infty)$, and $\lim_{d \rightarrow \infty} f(d) = \sqrt{2}$. Therefore for all $d \in [5, \infty)$,

$$1.12 \leq f(e^e) \leq f(d) \leq \max\{f(5), \sqrt{2}\} = \sqrt{2}$$

This gives for $n \geq 200\frac{p}{1-p}\log d$,

$$1 \leq \frac{1.12}{\sqrt{10/9}} \leq C(d, n) \leq \sqrt{2} \quad (13)$$

Since $p \geq 1/2$, if $n \geq 200\frac{p}{1-p}\log d$, then $n > \frac{200}{p(1-p)\log d}$ (if $d \geq 8$) and $n \geq 200\frac{1-p}{p}\log d$. The above implies:

$$1 < C\sqrt{np(1-p)\log d} \leq np \leq np + 1. \quad (14)$$

Finally, we bound the quantity of interest:

$$\begin{aligned} \mathbb{E}\left[\max_{1 \leq i \leq d} Z_i^n\right] &= \mathbb{E}\left[\max_{1 \leq i \leq d} Z_i^n | A\right] \cdot \mathbb{P}(A) + \mathbb{E}\left[\max_{1 \leq i \leq d} Z_i^n | A^C\right] \cdot (1 - \mathbb{P}(A)) \\ &\geq \mathbb{E}\left[\max_{1 \leq i \leq d} Z_i^n | A\right] \cdot \mathbb{P}(A) + \mathbb{E}\left[Z_1^{(n)} | A^C\right] \cdot (1 - \mathbb{P}(A)) \\ &= \mathbb{E}\left[\max_{1 \leq i \leq d} Z_i^n | A\right] \cdot \mathbb{P}(A) + \mathbb{E}\left[Z_1^{(n)} | Z_1^{(n)} < C\sqrt{np(1-p)\log d} - 1\right] \cdot (1 - \mathbb{P}(A)) \\ &\geq \mathbb{E}\left[\max_{1 \leq i \leq d} Z_i^n | A\right] \cdot \mathbb{P}(A) + \mathbb{E}\left[Z_1^{(n)} | Z_1^{(n)} \leq 0\right] \cdot (1 - \mathbb{P}(A)) \quad \text{by (14)} \\ &\geq (C\sqrt{np(1-p)\log d} - 1) \cdot \mathbb{P}(A) + \mathbb{E}\left[Z_1^{(n)} | Z_1^{(n)} \leq 0\right] \cdot (1 - \mathbb{P}(A)). \end{aligned} \quad (15)$$

First, we lower bound $\mathbb{E}\left[Z_1^{(n)} | Z_1^{(n)} \leq 0\right]$. Let $\beta = \frac{1}{1 - \sqrt{\frac{2n}{\pi[n(1-p)](n - [n(1-p)])}}}$. For $n \geq \frac{200}{p(1-p)}\log d$, we have $n \geq \frac{205}{\pi p(1-p)} \geq \frac{200 + \pi p}{\pi(1-p)p}$ and $\beta \leq \frac{10}{9}$.

Then Lemma 2.2 in [26] combined with Lemma 8 in [13] give that for $Y_n \sim B(n, 1-p)$:

$$\mathbb{E}[Y_n | Y_n \geq n(1-p)] < n(1-p) + \beta\sqrt{np(1-p)} < n(1-p) + \frac{10}{9}\sqrt{np(1-p)}.$$

Since $B_n = Z^{(n)} + pn \sim B(n, p)$ can be written as $n - Y_n$, we have:

$$\mathbb{E}\left[Z_1^{(n)} | Z_1^{(n)} \leq 0\right] = \mathbb{E}\left[B_n | B_n \leq np\right] - np$$

$$\begin{aligned}
&= \mathbb{E}[n - Y_n | n - Y_n \leq np] - np \\
&= n - \mathbb{E}[Y_n | Y_n \geq n(1-p)] - np \\
&\geq n - n(1-p) - \frac{10}{9}\sqrt{np(1-p)} - np \\
&= -\frac{10}{9}\sqrt{np(1-p)}. \tag{16}
\end{aligned}$$

Next, we lower-bound $\mathbb{P}(A)$:

$$\begin{aligned}
\mathbb{P}(A) &= 1 - \mathbb{P}(A^C) \\
&= 1 - \left(\mathbb{P}\left[Z_1^{(n)} < C\sqrt{np(1-p)\log d} - 1\right] \right)^d \\
&= 1 - \left(\mathbb{P}\left[B_n < np + C\sqrt{np(1-p)\log d} - 1\right] \right)^d \\
&= 1 - \left(1 - \mathbb{P}\left[B_n \geq np + C\sqrt{np(1-p)\log d} - 1\right] \right)^d \\
&\geq 1 - \exp\left(-d \cdot \mathbb{P}\left[B_n \geq np + C\sqrt{np(1-p)\log d} - 1\right]\right) \quad \text{since } 1 - x \leq e^{-x} \\
&\geq 1 - \exp\left(-d \cdot \frac{e^{-1/6} \exp\left(-\frac{1}{2p(1-p)}\psi\left(\frac{C\sqrt{np(1-p)\log d}}{n}\right) \cdot \frac{C^2 np(1-p)\log d}{n}\right)}{\frac{\sqrt{2\pi}C\sqrt{np(1-p)\log d}}{\sqrt{np(1-p)}} + 2}\right) \quad \text{using (10) and (14)} \\
&= 1 - \exp\left(-d \cdot \frac{e^{-1/6} \exp\left(-\frac{1}{2}\psi\left(C\sqrt{p(1-p)\log d/n}\right) \cdot C^2 \log d\right)}{\sqrt{2\pi}C\sqrt{\log d} + 2}\right) \\
&= 1 - \exp\left(-\frac{e^{-1/6} d^{1-\frac{C^2}{2}} \psi\left(C\sqrt{p(1-p)\log d/n}\right)}{\sqrt{2\pi}C\sqrt{\log d} + 2}\right) \\
&\geq 1 - \exp\left(-\frac{e^{-1/6} d^{1-\frac{C^2}{2}} \psi\left(\sqrt{\frac{2p(1-p)\log d}{n}}\right)}{2\sqrt{\pi \log d} + 2}\right) \quad \text{by (13)}.
\end{aligned}$$

We now use that $d^{1-\frac{C^2}{2}} \psi\left(\sqrt{2p(1-p)\log d/n}\right) = \log d$ by the definition of C in (11). Hence, we obtain:

$$\mathbb{P}(A) \geq 1 - \exp\left(-\frac{e^{-1/6} \log d}{2\sqrt{\pi \log d} + 2}\right) = 1 - g(d), \quad \text{for } g(d) = \exp\left(-\frac{e^{-1/6} \log d}{2\sqrt{\pi \log d} + 2}\right). \tag{17}$$

Putting everything together: we plug (16) and (17) into (15) to get

$$\begin{aligned}
\mathbb{E}\left[\max_{1 \leq i \leq d} Z_i^n\right] &\geq (C\sqrt{np(1-p)\log d} - 1) \cdot (1 - g(d)) - \frac{10}{9}g(d)\sqrt{np(1-p)} \\
&= \frac{f(d) \cdot (1 - g(d))}{\sqrt{\psi\left(\sqrt{\frac{2p(1-p)\log d}{n}}\right)}} (\sqrt{np(1-p)\log d} - 2) - \frac{10}{9}g(d)\sqrt{np(1-p)} \quad \text{using (11)} \\
&\geq \frac{f(d)(1 - g(d))}{\sqrt{10/9}} (\sqrt{np(1-p)\log d} - 2) - \frac{10}{9}g(d)\sqrt{np(1-p)} \quad \text{using (12)} \\
&= \sqrt{np(1-p)\log d} \cdot \left(\frac{f(d)(1 - g(d))}{\sqrt{10/9}} - \frac{10}{9} \cdot \frac{g(d)}{\sqrt{\log(d)}} \right) - 2 \frac{f(d)(1 - g(d))}{\sqrt{10/9}} \\
&\geq 0.02\sqrt{np(1-p)\log d} - 1.5,
\end{aligned}$$

for $d \geq 100$ (we also used that $\sqrt{np(1-p)\log d} > 2$ in the 2nd inequality). This gives the result. $\square$