## A Theory: Proofs and Complimentary Material

## A. 1 Proof of Lemma 3.1

Lemma. It holds that

$$
\hat{Q}_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)=R_{o}(s)+\gamma^{d} G_{o}(s), \quad \max _{a \neq \pi_{o}(s)} \hat{Q}_{d}^{\pi_{o}}(s, a)=R_{e}(s)+\gamma^{d} G_{e}(s)
$$

with

$$
G_{o}(s) \sim G E V\left(\mu_{o}^{G E V}(s), \sigma_{o}^{G E V}, 0\right), \quad G_{e}(s) \sim G E V\left(\mu_{e}^{G E V}(s), \sigma_{e}^{G E V}, 0\right)
$$

where GEV is the Generalized Extreme Value distribution and

$$
\begin{align*}
\mu_{o}^{G E V}(s) & =\mu_{o}(s)+\sigma_{o} \Phi^{-1}\left(1-\frac{1}{A^{d-1}}\right) \\
\sigma_{o}^{G E V} & =\sigma_{o}\left[\Phi^{-1}\left(1-\frac{1}{e A^{d-1}}\right)-\Phi^{-1}\left(1-\frac{1}{A^{d-1}}\right)\right] \\
\mu_{e}^{G E V}(s) & =\mu_{e}(s)+\sigma_{e} \Phi^{-1}\left(1-\frac{1}{A^{d}-A^{d-1}}\right) \\
\sigma_{e}^{G E V} & =\sigma_{e}\left[\Phi^{-1}\left(1-\frac{1}{e\left(A^{d}-A^{d-1}\right)}\right)-\Phi^{-1}\left(1-\frac{1}{A^{d}-A^{d-1}}\right)\right] . \tag{7}
\end{align*}
$$

The function $\Phi^{-1}$ is the inverse of the CDF of the standard normal distribution.
Proof. For $1 \leq i \leq A^{d-1}$, let $N_{o}^{(i)}$ be independent random variables distributed $N_{o}^{(i)} \sim \mathcal{N}\left(\mu_{o}(s), \sigma_{o}^{2}\right)$. According to (1), we have that

$$
\hat{Q}_{d}^{\pi_{o}}(s, a)=\left[\max _{\left(a_{k}\right)_{k=1}^{d} \in \mathcal{A}}\left[\sum_{t=0}^{d-1} \gamma^{t} r\left(s_{t}, a_{t}\right)\right]+\gamma^{d} \hat{Q}^{\pi_{o}}\left(s_{d}, a_{d}\right)\right]_{s_{0}=s, a_{0}=a}
$$

For the case of $a=\pi_{o}(s)$, using Assumption 2, we can replace the cumulative reward above with $R_{o}(s)$ and be left with

$$
\begin{aligned}
\hat{Q}_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right) & =R_{o}(s)+\left.\gamma^{d} \max _{\left(a_{k}\right)_{k=1}^{d} \in \mathcal{A}} \hat{Q}^{\pi_{o}}\left(s_{d}, a_{d}\right)\right|_{s_{0}=s, a_{0}=\pi_{o}(s)} \\
& =R_{o}(s)+\gamma^{d} \max _{1 \leq i \leq A^{d-1}} N_{o}^{(i)} \\
& =R_{o}(s)+\gamma^{d} G_{o}(s),
\end{aligned}
$$

where the second relation is due to Assumptions 1 together with 2, and the third follows from the maximum of $\left\{N_{o}^{(i)}\right\}_{i=1}^{A^{d-1}}$ being GEV-distributed with the parameters in the statement, derived as in [9]. Similarly, for $a_{0} \neq \pi_{0}(s)$, we will have $A^{d}-A^{d-1}$ iid variables $N_{e}^{(i)} \sim \mathcal{N}\left(\mu_{e}(s), \sigma_{e}^{2}\right)$ and, consequently, $G_{e}(s)$ as given in the statement.

## A. 2 Proof of Lemma 3.2

Lemma. It holds that

$$
\begin{aligned}
\mathbb{E}\left[\hat{Q}_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)\right] & =Q_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)+\gamma^{d} B_{o}\left(\sigma_{o}, A, d\right), \\
\mathbb{E}\left[\max _{a \neq \pi_{o}(s)} \hat{Q}_{d}^{\pi_{o}}(s, a)\right] & =Q_{d}^{\pi_{o}}\left(s, a \neq \pi_{o}(s)\right)+\gamma^{d} B_{e}\left(\sigma_{e}, A, d\right),
\end{aligned}
$$

where the biases $B_{o}$ and $B_{e}$, satisfying $0 \leq B_{o}\left(\sigma_{o}, A, d\right)<B_{e}\left(\sigma_{e}, A, d\right)$, are given by

$$
\begin{align*}
& B_{o}\left(\sigma_{o}, A, d\right)= \begin{cases}0 & \text { if } d=1 \\
\sigma_{o} \Phi^{-1}\left(1-\frac{1}{A^{d-1}}\right)+\gamma_{E M} \sigma_{o}^{G E V} & \text { otherwise }\end{cases} \\
& B_{e}\left(\sigma_{e}, A, d\right)= \begin{cases}0, & \text { if } d=1 \text { and } A=2 \\
\sigma_{e} \Phi^{-1}\left(1-\frac{1}{A^{d}-A^{d-1}}\right)+\gamma_{E M} \sigma_{e}^{G E V}, & \text { otherwise }\end{cases} \tag{8}
\end{align*}
$$

The constant $\gamma_{E M} \approx 0.58$ is the Euler-Mascheroni constant.

Proof. First, obviously, the maximum over a set containing a single random variable has the distribution of that single element. Hence, there is no overestimation bias in the single-element case; i.e., the bias is 0 for the sub-tree of $a=\pi_{o}(s)$ with $d=1$, and the sub-tree of $a \neq \pi_{o}(s)$ with $d=1$ and $A=2$.
Next, let $X$ be s.t. $X \sim \operatorname{GEV}\left(\mu^{\mathrm{GEV}}, \sigma^{\mathrm{GEV}}, 0\right)$. From [9], we have that

$$
\begin{equation*}
\mathbb{E}[X]=\mu^{\mathrm{GEV}}+\gamma_{\mathrm{EM}} \sigma^{\mathrm{GEV}} \tag{9}
\end{equation*}
$$

Applying Lemma 3.1, we have that

$$
\begin{align*}
\mathbb{E}\left[\hat{Q}_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)\right] & =\mathbb{E}\left[R_{o}(s)+\gamma^{d} G_{o}(s)\right] \\
& =R_{o}(s)+\gamma^{d} \mathbb{E}\left[G_{o}(s)\right]  \tag{10}\\
& =R_{o}(s)+\gamma^{d}\left[\mu_{o}(s)+\sigma_{o} \Phi^{-1}\left(1-\frac{1}{A^{d-1}}\right)+\gamma_{\mathrm{EM}} \sigma_{o}^{\mathrm{GEV}}\right]  \tag{11}\\
& =Q_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)+\gamma^{d} B_{o}\left(\sigma_{o}, A, d\right) \tag{12}
\end{align*}
$$

where relation (10) is due to the rewards being deterministic (Assumption2), relation (11) follows from Lemma 3.1 together with (9), and relation (12) is due to the definition of $Q_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)$ in (1), together with Assumptions 1 and 2. The calculation for the expectation $\mathbb{E}\left[\max _{a \neq \pi_{o}(s)} \hat{Q}_{d}^{\pi_{o}}(s, a)\right]$ follows the same steps.
Next, we show that

$$
\begin{equation*}
0 \leq B_{o}\left(\sigma_{o}, A, d\right)<B_{e}\left(\sigma_{e}, A, d\right) \tag{13}
\end{equation*}
$$

For this, let us define

$$
B(n)= \begin{cases}0 & n=1 \\ \gamma_{\mathrm{EM}} \Phi^{-1}\left(1-\frac{1}{e n}\right)+\left(1-\gamma_{\mathrm{EM}}\right) \Phi^{-1}\left(1-\frac{1}{n}\right), & n \neq 1\end{cases}
$$

Notice we now have $B_{o}\left(\sigma_{o}, A, d\right)=\sigma_{o} B\left(A^{d-1}\right)$ and $B_{e}\left(\sigma_{e}, A, d\right)=\sigma_{e} B\left(A^{d}-A^{d-1}\right)$. Since by Assumption 1 we have that $\sigma_{e}>\sigma_{o}>0$, to prove 13) it is sufficient to show that $0 \leq B\left(A^{d-1}\right)<B\left(A^{d}-A^{d-1}\right)$. Since $B(n)$ is composed of two positive monotonically increasing functions, it is a positive monotonically increasing function. So, whenever $A^{d-1}<A^{d}-A^{d-1}$, we also have that $0 \leq B\left(A^{d-1}\right)<B\left(A^{d}-A^{d-1}\right)$ and, consequently, 13). Taking a log, we see it is indeed the case for $A>2$. For $A=2$, we have equality and $B\left(A^{d-1}\right)=B\left(A^{d}-A^{d-1}\right)$. But then, (13) holds again, since $\sigma_{e}>\sigma_{o}>0$.

## A. 3 Proof of Theorem 3.3

Theorem. The relation $\mathbb{E}\left[\hat{Q}_{d}^{B C T S, \pi_{o}}\left(s, \pi_{o}(s)\right)\right]>\mathbb{E}\left[\max _{a \neq \pi_{o}(s)} \hat{Q}_{d}^{B C T S, \pi_{o}}(s, a)\right]$, holds if and only if $Q_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)>\max _{a \neq \pi_{o}(s)} Q_{d}^{\pi_{o}}(s, a)$.

Proof. For the case of $a=\pi_{o}(s)$, we have

$$
\mathbb{E}\left[\hat{Q}_{d}^{\mathrm{BCTS}, \pi_{o}}\left(s, \pi_{o}(s)\right)\right]=\mathbb{E}\left[\hat{Q}_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)\right]=Q_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)+\gamma^{d} B_{o}\left(\sigma_{o}, A, d\right)
$$

where the first relation holds by the definition in (3) and the second relation follows from Lemma 3.2 Using the same arguments, for $a \neq \pi_{o}(s)$,

$$
\begin{aligned}
& \mathbb{E}\left[\max _{a \neq \pi_{o}(s)} \hat{Q}_{d}^{\mathrm{BCTS}, \pi_{o}}(s, a)\right] \\
= & \mathbb{E}\left[\max _{a \neq \pi_{o}(s)} \hat{Q}_{d}^{\pi_{o}}(s, a)\right]-\gamma^{d}\left[B_{e}\left(\sigma_{e}, A, d\right)-B_{o}\left(\sigma_{o}, A, d\right)\right]+\gamma^{d} B_{e}\left(\sigma_{e}, A, d\right) \\
= & \max _{a \neq \pi_{o}(s)} Q_{d}^{\pi_{o}}(s, a)+\gamma^{d} B_{o}\left(\sigma_{o}, A, d\right),
\end{aligned}
$$

and the result immediately follows.

## A. 4 Proof of Lemma 3.4

Lemma. When $A^{d-1} \gg 1$, the correction term in (3) can be approximated with

$$
\begin{equation*}
B_{e}\left(\sigma_{e}, A, d\right)-B_{o}\left(\sigma_{o}, A, d\right) \approx \sqrt{2 \log A}\left(\sigma_{e} \sqrt{d}-\sigma_{o} \sqrt{d-1}\right)-\left(\sigma_{e}-\sigma_{o}\right) / 2 \tag{14}
\end{equation*}
$$

Proof. In the following, we apply the approximation

$$
\begin{equation*}
\Phi^{-1}\left(1-\frac{1}{n}\right) \approx \sqrt{2 \log n}-0.5 \tag{15}
\end{equation*}
$$

from [3], which we empirically show to be highly accurate in Appendix A.7.
By definition (7),

$$
\begin{align*}
\sigma_{e}^{\mathrm{GEV}} & =\sigma_{e}\left[\Phi^{-1}\left(1-\frac{1}{e\left(A^{d}-A^{d-1}\right)}\right)-\Phi^{-1}\left(1-\frac{1}{A^{d}-A^{d-1}}\right)\right] \\
& \approx \sigma_{e}\left[\Phi^{-1}\left(1-\frac{1}{e A^{d}}\right)-\Phi^{-1}\left(1-\frac{1}{A^{d}}\right)\right]  \tag{16}\\
& \approx \sigma_{e}\left[\sqrt{2 \log \left(e A^{d}\right)}-\sqrt{2 \log \left(A^{d}\right)}\right]  \tag{17}\\
& =\sigma_{e} \frac{2 \log \left(e A^{d}\right)-2 \log A^{d}}{\sqrt{\left.2 \log \left(e A^{d}\right)\right)}+\sqrt{2 \log A^{d}}} \\
& \approx \frac{\sigma_{e}}{\sqrt{2 d \log A}} \tag{18}
\end{align*}
$$

where in (16) we applied (15); in relation (17) we use that $A^{d}-A^{d-1} \approx A^{d}$ since $A^{d} \gg 1$; and in relation 18 we approximate $1+\log \left(A^{d}\right) \approx \log \left(A^{d}\right)$, again because $A^{d} \gg 1$.
When $A^{d} \gg 1$, the second case of (8) holds, and thus

$$
\begin{aligned}
B_{e}\left(\sigma_{e}, A, d\right) & =\sigma_{e} \Phi^{-1}\left(1-\frac{1}{A^{d}-A^{d-1}}\right)+\gamma_{\mathrm{EM}} \sigma_{e}^{\mathrm{GEV}} \\
& \approx \sigma_{e}(\sqrt{2 d \log A}-0.5)+\gamma_{\mathrm{EM}} \frac{\sigma_{e}}{\sqrt{2 d \log A}} \\
& \approx \sigma_{e}(\sqrt{2 d \log A}-0.5)
\end{aligned}
$$

where the second relation follows from (15) and (18), while for the third relation we approximate $\sqrt{2 \log \left(A^{d}\right)}\left(1+\frac{\gamma_{\mathrm{EM}}}{2 \log \left(A^{d}\right)}\right) \approx \sqrt{2 \log \left(A^{d}\right)}$ since $A^{d} \gg 1$ (recall that $\gamma_{\mathrm{EM}} \approx 0.58$ ).
Applying the same derivation for $B_{o}$ gives that

$$
\begin{equation*}
B_{o}\left(\sigma_{o}, A, d\right) \approx \sigma_{o}(\sqrt{2(d-1) \log A}-0.5) \tag{19}
\end{equation*}
$$

and the result follows directly.

## A. 5 Proof of Theorem 3.5

To prove Theorem 3.5, we first obtain the following non-approximate result.
Theorem A.1. The policy $\pi_{d}^{B C T S}(s)($ see (4)) chooses a sub-optimal action with probability bounded by:

$$
\begin{equation*}
\operatorname{Pr}\left(\pi_{d}^{B C T S}(s) \notin \arg \max _{a} Q_{d}^{\pi_{o}}(s, a)\right) \leq\left(1+\frac{6\left(Q_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)-\max _{a \neq \pi_{o}(s)} Q_{d}^{\pi_{o}}(s, a)\right)^{2}}{\gamma^{2 d} \pi^{2}\left(\left(\sigma_{o}^{G E V}\right)^{2}+\left(\sigma_{e}^{G E V}\right)^{2}\right)}\right)^{-1} \tag{20}
\end{equation*}
$$

Proof. We first recall Cantelli's inequality [7]: Let $X$ be a real-valued random variable. Then, for some $\lambda>0$,

$$
\begin{equation*}
\operatorname{Pr}(X-\mathbb{E}[X] \geq \lambda) \leq\left(1+\frac{\lambda^{2}}{\operatorname{Var}[X]}\right)^{-1} \tag{21}
\end{equation*}
$$

We choose

$$
X:=\max _{a \neq \pi_{o}(s)} \hat{Q}_{d}^{B C T S, \pi_{o}}(s, a)-\hat{Q}_{d}^{B C T S, \pi_{o}}\left(s, \pi_{o}(s)\right)
$$

We can now split the event of sub-optimal action choice as follows:

$$
\begin{align*}
& \left\{\pi_{d}^{\mathrm{BCTS}}(s) \notin \arg \max _{a} Q_{d}^{\pi_{o}}(s, a)\right\} \\
= & \left\{\pi_{o}(s) \in \arg \max _{a} Q_{d}^{\pi_{o}}(s, a) \cap X>0\right\} \bigcup\left\{\pi_{o}(s) \notin \arg \max _{a} Q_{d}^{\pi_{o}}(s, a) \cap X<0\right\} . \tag{22}
\end{align*}
$$

Note that this division is possible because according to Assumptions 1 and 2, all actions different than $\pi_{o}(s)$ have equal cumulative reward and value in expectation.
Next, we consider the two (deterministic) following possible cases.
Case I: $\pi_{o}(s) \in \arg \max _{a} Q_{d}^{\pi_{o}}(s, a)$. Then, the second event in 22 ) is an empty set and we have that

$$
\begin{equation*}
\left\{\pi_{d}^{\mathrm{BCTS}}(s) \notin \arg \max _{a} Q_{d}^{\pi_{o}}(s, a)\right\}=\{X>0\} \tag{23}
\end{equation*}
$$

Case II: $\quad \pi_{o}(s) \notin \arg \max _{a} Q_{d}^{\pi_{o}}(s, a)$. Then, from symmetry, we will get $\left\{\pi_{d}^{\mathrm{BCTS}}(s) \notin \arg \max _{a} Q_{d}^{\pi_{o}}(s, a)\right\}=\{X<0\}$.
In the rest of the proof, we shall apply Cantelli's inequality to upper bound $P(X>0)$ in Case I, i.e., to bound the sub-optimal action selection event. Afterward, we will explain how Case II yields the same bound as in Case I.
Let us set

$$
\begin{equation*}
\lambda=-\mathbb{E}[X]=Q_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)-\max _{a \neq \pi_{o}(s)} Q_{d}^{\pi_{o}}(s, a) \tag{24}
\end{equation*}
$$

where the second relation follows from the proof of Theorem 3.3 Next, we calculate the variance of $X$ using GEV theory [9]. For $G \sim \operatorname{GEV}(\mu, \sigma, 0)$, we have that $\operatorname{Var}[G]=\sigma^{2} \frac{\pi^{2}}{6}$. Then,

$$
\begin{align*}
\operatorname{Var}[X] & =\operatorname{Var}\left[\max _{a \neq \pi(s)} \hat{Q}_{d}^{B C T S, \pi_{o}}(s, a)-\hat{Q}_{d}^{B C T S, \pi_{o}}\left(s, \pi_{o}(s)\right)\right] \\
& =\operatorname{Var}\left[\max _{a \neq \pi(s)} \hat{Q}_{d}^{B C T S, \pi_{o}}(s, a)\right]+\operatorname{Var}\left[\hat{Q}_{d}^{B C T S, \pi_{o}}\left(s, \pi_{o}(s)\right)\right] \\
& =\operatorname{Var}\left[\max _{a \neq \pi(s)} \hat{Q}_{d}^{\pi_{o}}(s, a)\right]+\operatorname{Var}\left[\hat{Q}_{d}^{\pi_{o}}\left(s, \pi_{o}(s)\right)\right] \\
& =\frac{\gamma^{2 d}\left(\sigma_{e}^{\mathrm{GEV}}\right)^{2} \pi^{2}}{6}+\frac{\gamma^{2 d}\left(\sigma_{o}^{\mathrm{GEV}}\right)^{2} \pi^{2}}{6}, \tag{25}
\end{align*}
$$

where the third relation is because $\operatorname{Var}\left[\hat{Q}_{d}^{B C T S, \pi_{o}}(s, a)\right]=\operatorname{Var}\left[\hat{Q}_{d}^{\pi_{o}}(s, a)\right]$ following (3); the last relation follows from $\hat{Q}_{d}^{\pi_{o}}$ having a GEV distribution as given in Lemma 3.1.
Plugging (23), (24), and (25) into (21) gives the desired result for Case I. Finally, for Case II, we define $Y=-X$ and repeat exactly the same process to upper bound $P(Y>0)$. Since $(\mathbb{E}[Y])^{2}=(\mathbb{E}[X])^{2}$, and $\operatorname{Var}[Y]=\operatorname{Var}[X]$, we obtain exactly the same bound as in Case I. This concludes the proof.

Theorem 3.5 now follows from Theorem A.1 after plugging the approximation from the proof of Lemma 3.4 in 20, and upper bounding the resulting expressions:

$$
\sigma_{e}^{\mathrm{GEv}} \approx \frac{\sigma_{e}}{\sqrt{2 d \log A}} \leq \frac{\sigma_{e}}{\sqrt{d \log A}}, \quad \sigma_{o}^{\mathrm{GEv}} \approx \frac{\sigma_{o}}{\sqrt{2(d-1) \log A}} \leq \frac{\sigma_{o}}{\sqrt{d \log A}}
$$

## A. 6 Proof of Proposition 3.6

Proposition. Let $\widehat{v a r}_{n}[X]$ be the variance estimator based on $n$ samples of $X$. Then,

$$
\widehat{\operatorname{var}}_{n=2}\left[\hat{Q}^{\pi_{o}}(s, a)\right]=\left(\hat{Q}_{1}^{\pi_{o}}(s, a)-\hat{Q}_{0}^{\pi_{o}}(s, a)\right)^{2} / 2=\delta^{2}(s, a) / 2
$$

where $\delta(s, a)$ is the Bellman error.
Proof. We shall employ the known unbiased variance estimator:

$$
\begin{aligned}
& \widehat{\operatorname{var}}_{n=2}\left[\hat{Q}^{\pi_{o}}(s, a)\right] \\
&= \frac{1}{n-1} \sum_{d=0}^{n}\left(\hat{Q}_{d}^{\pi_{o}}(s, a)-\frac{1}{n} \sum_{d=0}^{n} \hat{Q}_{d}^{\pi_{o}}(s, a)\right)^{2} \\
&\stackrel{(\mathrm{n}=}{=} 2) \\
&=\left.\hat{Q}_{0}^{\pi_{o}}(s, a)-\frac{\hat{Q}_{0}^{\pi_{o}}(s, a)+\hat{Q}_{1}^{\pi_{o}}(s, a)}{2}\right)^{2}+\left(\hat{Q}_{0}^{\pi_{o}}(s, a)-\hat{Q}_{1}^{\pi_{o}}(s, a)\right)^{2}+\frac{1}{4}\left(\hat{Q}_{0}^{\pi_{o}}(s, a)-\hat{Q}_{0}^{\pi_{o}}(s, a)\right)^{2} \\
&= \frac{\delta^{2}(s, a)}{2}
\end{aligned}
$$

## A. 7 Approximate bounds bias

To show the validity of our approximation, we numerically evaluate the two sides of (5], i.e. the LHS

$$
B_{e}\left(\sigma_{e}, A, d\right)-B_{o}\left(\sigma_{o}, A, d\right)
$$

as given in (8), vs. the RHS

$$
\sqrt{2 \log A}\left(\sigma_{e} \sqrt{d}-\sigma_{o} \sqrt{d-1}\right)-\left(\sigma_{e}-\sigma_{o}\right) / 2
$$

To compute the values, we arbitrarily choose $A=3, \sigma_{o}=1, \sigma_{e}=4$, and increase $d$. We plot the results in Fig. 6 As seen, the approximation is highly accurate for all values of $A^{d}$.


Figure 6: Quality of approximation in Lemma 3.4

## B Experiments complementary material

## B. 1 Inference timing

We measure the average complete TS time per action selection. We provide the results together with the respective scores per game and depth in Fig. 7 The scores are obtained via BCTS with a BatchBFS implementation, as reported in Section 5.1 As depth increases, the number of explored nodes
in the tree grows exponentially, with runtime increasing accordingly. The plots depict the tradeoff between improved scores and the corresponding price in terms of runtime per action selection.


Figure 7: Score vs. inference time of each TS operation as a function of TS depth.

## B. 2 Training

In our training experiments we use the same training hyper-parameters as the original DQN paper [32].

## B. 3 Ablation study: Training with propagated value from the tree nodes

Here, we present an ablation study for the correction to the Bellman update proposed in [14] in the case of a TS policy. This correction modifies the training target: Instead of bootstrapping the value from the transition sampled from the replay buffer, we use the cumulative reward and value computed during the TS. In Fig. 8, we present training plots for Atari Space-Invaders for DQN with TS of depths 2,3, and 4. Note that for depth 1, the correction is vacuous since it coincides with the classic Bellman update. As seen, for Space-Invaders, the correction improves convergence in all tested depths.


Figure 8: Propagated value from the tree nodes: Ablation study for Space-Invaders convergence. Episodic training cumulative reward of DQN with TS based on 5 seeds. We compare the standard update method with the update based on the propagated value from the tree nodes, as proposed in [14].

Lastly, we summarize the results for all tested games in Table 2 The table reveals that the correction often improves training, though not always. Therefore, we treat it as a hyper-parameter that we sweep over.

## C Training: Wall-clock time

We provide in Fig. 9 convergence plots of DQN with TS for all tested Atari games, for depths 0 to 4. To showcase the time efficiency of using a tree-based policy, we give the scores with respect to

Table 2: Ablation study: Propagated value (PV) from the tree nodes: Ablation study for scores of all tested games. All agents of all depths were trained using DQN with TS, with similar train time that amounts to 30 million frames of DQN. The reported scores were obtained as in [22]: Average of 200 testing episodes, from the agent snapshot that obtained the highest score during training.

| Game | $d=2$ |  | $d=3$ |  | $d=4$ |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  | PV=False | PV=True | PV=False | PV=True | PV=False | PV=True |  |  | 4,794 | 4,693 | 17,929 | 15,434 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| Asteroids | 2,613 | 2,472 | 4,73 |  |  |
| Breakout | 568 | 581 | 420 | 417 | 620 |
| MsPacman | 2,514 | 2,923 | 3,498 | 3,046 | 2,748 |
| SpaceInvaders | 1,314 | 1,602 | 1,204 | 2,132 | 1,452 |
| VideoPinball | 214,168 | 244,052 | 442,347 | 366,670 | 345,742 |

the wall-clock time. We run each training experiment for two days, which amount to roughly 30 million frames for depth 0 and 1 million frames for depth 4 . For each run, we display the average score together with std using 5 seeds.

Fig. 9 reveals the trade-off between deeper TS inference (action selection) time and the improvement thanks to the deeper TS. For regular DQN (depth=0), each inference is the fastest, but generally leads to lower scores than deeper trees. On the other hand, the deepest tree is not necessarily the most time-efficient. In most cases here, there is a sweet-spot in depth 2 or 3 that gives the best score for the same training time.


Figure 9: Training convergence plots for tree search with DQN of depths up to 4.

## D Hardware

We run our training experiments on a 10 core $\operatorname{Intel}(\mathrm{R}) \mathrm{CPU}$ i9-10900X @ 3.70GHz equipped with NVIDIA Quadro GV100 32GB.

