

A Proof of Theorem 3.1

Proof. Let us assume that $V_{n-1}(b)$ is a convex function, then we can show that $V_n(b)$ is also convex as follows. First, the value function in Eq. 1 can be expressed by parts as,

$$\begin{aligned} V_n(b) &= \max_a [V_n^a(b)], \\ V_n^a(b) &= \sum_o V_n^{a,o}(b), \\ V_n^{a,o}(b) &= \frac{\rho(b, a)}{|\Omega|} + \psi^{a,o}(b), \text{ and} \\ \psi^{a,o}(b) &= \gamma Pr(o|a, b) V_{n-1}(b^{a,o}). \end{aligned}$$

Also, using the belief update, we can express $b^{a,o} = \frac{P^{a,o} \cdot b}{\|P^{a,o} \cdot b\|_1}$ with $P_{s,s'}^{a,o} = O(s', a, o)T(s, a, s')$. Therefore,

$$\psi^{a,o}(b) = \gamma \|P^{a,o} \cdot b\|_1 V_{n-1}\left(\frac{P^{a,o} \cdot b}{\|P^{a,o} \cdot b\|_1}\right) = \gamma \kappa(P^{a,o} \cdot b)$$

with $\kappa(w) = \|w\|_1 V_{n-1}\left(\frac{w}{\|w\|_1}\right)$.

Here, $\kappa(w)$ is a convex function as it uses the perspective and linear-fractional convexity preserving operations (see App. A.1 for a stand-alone proof). Then, $\psi^{a,o}$ is also convex since convexity is preserved under affine maps. Consequently, $V_n^{a,o}$, V_n^a and V_n are convex because ρ and $\psi^{a,o}$ are convex.⁶ Considering this last result is the inductive step, then $V_n(b)$ is convex over b , because V_0 is convex by definition. \square

A.1 Stand-alone Proof of κ Convexity

Lemma A.1. *Let $w \in \mathbb{R}^n$ and $f : \mathbb{R}^n \mapsto \mathbb{R}$ a convex function. If $\kappa(w) = \|w\|_1 f(\frac{w}{\|w\|_1})$, then $\kappa(w)$ is also a convex function.*

Proof. As stated before, one can use the perspective and linear-fractional convexity preserving operations to directly prove that κ is convex. However, this can also be proved by only using the convexity of f :

$$\begin{aligned} &\kappa(\alpha x + (1 - \alpha)y) \\ &= \|\alpha x + (1 - \alpha)y\|_1 f\left(\frac{\alpha x + (1 - \alpha)y}{\|\alpha x + (1 - \alpha)y\|_1}\right) \\ &= \|\alpha x + (1 - \alpha)y\|_1 f\left(\frac{\alpha \|x\|_1}{\|\alpha x + (1 - \alpha)y\|_1} \cdot \frac{x}{\|x\|_1} + \frac{(1 - \alpha)\|y\|_1}{\|\alpha x + (1 - \alpha)y\|_1} \cdot \frac{y}{\|y\|_1}\right) \\ &\leq \|\alpha x + (1 - \alpha)y\|_1 \left(\frac{\alpha \|x\|_1}{\|\alpha x + (1 - \alpha)y\|_1} f\left(\frac{x}{\|x\|_1}\right) + \frac{(1 - \alpha)\|y\|_1}{\|\alpha x + (1 - \alpha)y\|_1} f\left(\frac{y}{\|y\|_1}\right)\right) \\ &= \alpha \|x\|_1 f\left(\frac{x}{\|x\|_1}\right) + (1 - \alpha)\|y\|_1 f\left(\frac{y}{\|y\|_1}\right) \\ &= \alpha \kappa(x) + (1 - \alpha)\kappa(y). \end{aligned}$$

\square

⁶Convex functions are closed under the sum and the max operators.

B Proofs of Section 4

B.1 Proof of Lemma 4.1

Proof. The largest minimum distance between the boundaries of both simplices (see Fig. 1) is given by the distance between their closest corners, i.e., $\varepsilon' = |(1 - (\mathcal{N} - 1)\varepsilon) - 1| + (\mathcal{N} - 1)|\varepsilon| = 2(\mathcal{N} - 1)\varepsilon$. This is the worst case scenario for $\|b' - b''\|_1$. Then, using the triangular inequality: $\|b - b''\|_1 \leq \|b - b'\|_1 + \|b' - b''\|_1$, and picking the highest possible values for both distances, the bound $\|b - b''\|_1 \leq 2(\mathcal{N} - 1)\varepsilon + \delta_B$ holds. \square

B.2 Proof of Lemma 4.2

Before proving Lemma 4.2, let us first give an equivalent result in the 1-dimensional case.

Lemma B.1. *Let $x_a, x_b \in \mathbb{R}$ ($x_a < x_b$), and $\eta \in (0, x_b - x_a)$ be three scalars. Let also f be a α -Hölder (with constant K_α), bounded and convex function from $[x_a, x_b]$ to \mathbb{R} , f being differentiable everywhere on (x_a, x_b) . Then f is $K_\alpha \eta^\alpha$ -Lipschitz on $[x_a + \eta, x_b - \eta]$.*

Proof. With any $x \in [x_a + \eta, x_b]$ we have (see Fig. 2 for an illustration):

$$\begin{aligned}
 f'(x) &\geq f'(x_a + \eta) && \text{(By convexity of } f\text{)} \\
 &\geq \frac{f(x_a + \eta) - f(x_a)}{\eta} && \text{(For the same reason)} \\
 &\geq -\frac{|f(x_a + \eta) - f(x_a)|}{\eta} \\
 &\geq -K_\alpha \eta^{\alpha-1} && \text{(Because } f \text{ is } \alpha\text{-Hölder).}
 \end{aligned}$$




Figure 2: Illustration for the proof of Lemma B.1

Similarly, for any $x \in [x_a, x_b - \eta]$ we have:

$$\begin{aligned}
 f'(x) &\leq f'(x_b - \eta) \\
 &\leq \frac{f(x_b) - f(x_b - \eta)}{\eta} \\
 &\leq \frac{|f(x_b) - f(x_b - \eta)|}{\eta} \\
 &\leq K_\alpha \eta^{\alpha-1}.
 \end{aligned}$$

Thus, for any $x \in [x_a + \eta, x_b - \eta]$, $|f'(x)| \leq K_\alpha \eta^{\alpha-1}$, so that f is $K_\alpha \eta^\alpha$ -Lipschitz on $[x_a + \eta, x_b - \eta]$. \square

We can now show how the above property extends to an n -simplex (using any norm).

Proof of Lemma 4.2. Let b be a point in Δ_η and let \mathbf{u} be a unit vector parallel to the hyperplane containing Δ . As shown on Fig. 3, the line going through b and directed by \mathbf{u} intersects Δ on the closed segment $S_{\mathbf{u}} = [b + x_a \mathbf{u}, b + x_b \mathbf{u}]$ ($x_a < 0 < x_b$). We can thus define a function $g_{\mathbf{u}}$ from $[x_a, x_b]$ to \mathbb{R} such that $g_{\mathbf{u}} : x \mapsto f(b + x \mathbf{u})$. $g_{\mathbf{u}}$ is then a α -Hölder (with same constant), bounded and convex function from $[x_a, x_b]$ to \mathbb{R} , $g_{\mathbf{u}}$ being differentiable everywhere on (x_a, x_b) .




Figure 3: Illustration for the proof of Lemma 4.2

Let us note that the shortest distance (in any norm) between Δ_η and $bd(\Delta)$ (the boundary of Δ) is lower bounded by η .⁷ The intersection of Δ_η with $S_{\mathbf{u}}$ is thus a segment $S'_{\mathbf{u}} \subseteq [b + (x_a + \eta) \mathbf{u}, b + (x_b - \eta) \mathbf{u}]$. We can therefore apply Lemma B.1 to $g_{\mathbf{u}}$ in $[x_a + \eta, x_b - \eta]$, what tells us that the derivative of $g_{\mathbf{u}}$ at point 0 (and therefore the directional derivative of f at b along \mathbf{u}) is bounded (in absolute value) by $K_\alpha \eta^{\alpha-1}$. This property holding for any \mathbf{u} , we have:

$$\forall b \in \Delta_\eta, \|\nabla f(b)\| \leq K_\alpha \eta^{\alpha-1}.$$

□

The above lemmas naturally extend to the case where f is piecewise differentiable. Please note that the norm of the gradient of Lemma 4.2 is defined as norm-1 for the uses of this paper, but any p -norm can be used as stated in the proof.

B.3 Proof of Theorem 4.3

Proof. Let us pick $b = \operatorname{argmax}_{x \in \Delta} \epsilon_B(x)$, the point where the approximation ω_B presents the worst error. This value can be bounded as follows:

$$\begin{aligned} \epsilon_B(b) &\leq \rho(b) - \omega_{b^*}(b) && \text{(By Eq. 5 and Convexity)} \\ &\leq \rho(b) - \omega_{b''}(b) && (b'' \in B \text{ makes a worse error)} \\ &= \rho(b) - \rho(b'') + (b'' - b) \cdot \nabla \rho(b'') && \text{(By definition of } \omega) \\ &\leq |\rho(b) - \rho(b'')| + |(b'' - b) \cdot \nabla \rho(b'')| && \text{(By triangular inequality)} \\ &\leq K_\alpha \|b - b''\|_1^\alpha + |(b'' - b) \cdot \nabla \rho(b'')| && \text{(By } \alpha\text{-Hölder condition)} \\ &\leq K_\alpha \|b - b''\|_1^\alpha + \|\nabla \rho(b'')\|_\infty \|b'' - b\|_1 && \text{(By Hölder inequality)} \\ &\leq K_\alpha \|b - b''\|_1^\alpha + \|\nabla \rho(b'')\|_1 \|b'' - b\|_1 && \text{(By norm equivalence)} \\ &\leq K_\alpha \|b - b''\|_1^\alpha + K_\alpha \eta^{\alpha-1} \|b'' - b\|_1 && \text{(By Lemma 4.2)} \\ &\leq K_\alpha (2(\mathcal{N} - 1)\varepsilon + \delta_B)^\alpha + K_\alpha \eta^{\alpha-1} (2(\mathcal{N} - 1)\varepsilon + \delta_B) && \text{(By Lemma 4.1)} \\ &= K_\alpha ((2(\mathcal{N} - 1)\varepsilon + \delta_B)^\alpha + (\varepsilon - \delta_B)^{\alpha-1} (2(\mathcal{N} - 1)\varepsilon + \delta_B)) && \text{(By definition of } \eta) \end{aligned}$$

⁷In p -norm, the distance between Δ_η and $bd(\Delta)$ is the distance between the points $(\eta, \frac{1-\eta}{n-1}, \dots, \frac{1-\eta}{n-1})$ and $(0, \frac{1}{n-1}, \dots, \frac{1}{n-1})$, i.e., $d(\Delta_\eta, bd(\Delta)) = (\eta^p + (n-1) \frac{\eta^p}{(n-1)^p})^{1/p} = \eta(1 + \frac{1}{(n-1)^{p-1}})^{1/p} \geq \eta$.

This last result is a generic bound that depends on the choice of ε , with $\varepsilon \in (\delta_B, \frac{1}{\mathcal{N}}]$. If we define ε as a linear function of δ_B , $\varepsilon = \lambda \delta_B$, then the generic bound can be written as,

$$\epsilon_B(b) \leq K_\alpha \left[(2(\mathcal{N} - 1)\lambda + 1)^\alpha + (\lambda - 1)^{\alpha-1} (2(\mathcal{N} - 1)\lambda + 1) \right] \delta_B^\alpha = C \delta_B^\alpha \quad (6)$$

with $\lambda \in (1, \frac{1}{\delta_B \mathcal{N}}]$. □