## Appendices

A Preliminary Remarks ..... 15
B Proofs of the Theoretical Results ..... 16
B. 1 Proofs of Lemmas 1, 2, and 3 ..... 16
B. 2 Proofs of Theorems 1 and 2 ..... 18
C Proofs of the Results about Optimal Baselines ..... 21
C. 1 Proof of Theorem 3 ..... 21
C. 2 Remarks about the surrogate optimal baseline ..... 22
C. 3 Proof of Theorem 4 ..... 23
D Pytorch Implementations of the Optimal Baseline ..... 25
E Computations for the Numerical Toy Example ..... 26
F Detailed Hyper-parameter Settings for Experiments ..... 29

## A Preliminary Remarks

Remark 1. The multi-agent state-action value function obeys the bounds

$$
\left|Q_{\boldsymbol{\theta}}^{i_{1}, \ldots, i_{k}}\left(s, \boldsymbol{a}^{\left(i_{1}, \ldots, i_{k}\right)}\right)\right| \leq \frac{\beta}{1-\gamma}, \text { for all } s \in \mathcal{S}, \boldsymbol{a}^{\left(i_{1}, \ldots, i_{k}\right)} \in \mathcal{A}^{\left(i_{1}, \ldots, i_{k}\right)}
$$

Proof. It suffices to prove that, for all $t$, the total reward satisfies $\left|R_{t}\right| \leq \frac{\beta}{1-\gamma}$, as the value functions are expectations of it. We have

$$
\left|R_{t}\right|=\left|\sum_{k=0}^{\infty} \gamma^{k} \mathbf{r}_{t+k}\right| \leq \sum_{k=0}^{\infty}\left|\gamma^{k} \mathbf{r}_{t+k}\right| \leq \sum_{k=0}^{\infty} \gamma^{k} \beta=\frac{\beta}{1-\gamma}
$$

Remark 2. The multi-agent advantage function is bounded.
Proof. We have

$$
\begin{aligned}
& \left|A_{\boldsymbol{\theta}}^{i_{1}, \ldots, i_{k}}\left(s, \boldsymbol{a}^{\left(j_{1}, \ldots, j_{m}\right)}, \boldsymbol{a}^{\left(i_{1}, \ldots, i_{k}\right)}\right)\right| \\
& =\left|Q_{\boldsymbol{\theta}}^{j_{1}, \ldots, j_{m}, i_{1}, \ldots, i_{k}}\left(s, \boldsymbol{a}^{\left(j_{1}, \ldots, j_{m}, i_{1}, \ldots, i_{k}\right)}\right)-Q_{\boldsymbol{\theta}}^{j_{1}, \ldots, j_{m}}\left(s, \boldsymbol{a}^{\left(j_{1}, \ldots, j_{m}\right.}\right)\right| \\
& \leq\left|Q_{\boldsymbol{\theta}}^{j_{1}, \ldots, j_{m}, i_{1}, \ldots, i_{k}}\left(s, \boldsymbol{a}^{\left(j_{1}, \ldots, j_{m}, i_{1}, \ldots, i_{k}\right)}\right)\right|+\left|Q_{\boldsymbol{\theta}}^{j_{1}, \ldots, j_{m}}\left(s, \boldsymbol{a}^{\left(j_{1}, \ldots, j_{m}\right.}\right)\right| \leq \frac{2 \beta}{1-\gamma}
\end{aligned}
$$

Remark 3. Baselines in MARL have the following property

$$
\mathbb{E}_{\mathbf{s} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[b\left(\mathrm{~s}, \mathbf{a}^{-i}\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\right]=\mathbf{0}
$$

Proof. We have

$$
\mathbb{E}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[b\left(\mathrm{~s}, \mathbf{a}^{-i}\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathbf{s}\right)\right]=\mathbb{E}_{\mathbf{s} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a}^{-i} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}^{-i}}\left[b\left(\mathrm{~s}, \mathbf{a}^{-i}\right) \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\right]\right],
$$

which means that it suffices to prove that for any $s \in \mathcal{S}$

$$
\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right]=\mathbf{0} .
$$

We prove it for continuous $\mathcal{A}^{i}$. The discrete case is analogous.

$$
\begin{aligned}
& \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right]=\int_{\mathcal{A}^{i}} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i} \mid s\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right) d a^{i} \\
& =\int_{\mathcal{A}^{i}} \nabla_{\theta^{i}} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i} \mid s\right) d a^{i}=\nabla_{\theta^{i}} \int_{\mathcal{A}^{i}} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i} \mid s\right) d a^{i}=\nabla_{\theta^{i}}(1)=\mathbf{0}
\end{aligned}
$$

## B Proofs of the Theoretical Results

## B. 1 Proofs of Lemmas 1, 2, and 3

In this subsection, we prove the lemmas stated in the paper. We realise that their application to other, very complex, proofs is not always immediately clear. To compensate for that, we provide the stronger versions of the lemmas; we give a detailed proof of the strong version of Lemma 1 which is supposed to demonstrate the equivalence of the normal and strong versions, and prove the normal versions of Lemmas $2 \& 3$, and state their stronger versions as remarks to the proofs.
Lemma 1 (Multi-agent advantage decomposition). For any state $s \in \mathcal{S}$, the following equation holds for any subset of $m$ agents and any permutation of their labels,

$$
A_{\boldsymbol{\theta}}^{1, \ldots, m}\left(s, \boldsymbol{a}^{(1, \ldots, m)}\right)=\sum_{i=1}^{m} A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{(1, \ldots, i-1)}, a^{i}\right) .
$$

Proof. We prove a slightly stronger, but perhaps less telling, version of the lemma, which is

$$
\begin{equation*}
A_{\boldsymbol{\theta}}^{k+1, \ldots, m}\left(s, \boldsymbol{a}^{(1, \ldots, k)}, \boldsymbol{a}^{(k+1, \ldots, m)}\right)=\sum_{i=k+1}^{m} A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{(1, \ldots, i-1)}, a^{i}\right) \tag{13}
\end{equation*}
$$

The original form of the lemma will follow from the above by taking $k=0$.
By the definition of the multi-agent advantage, we have

$$
\begin{aligned}
& A_{\boldsymbol{\theta}}^{k+1, \ldots, m}\left(s, \boldsymbol{a}^{(1, \ldots, k)}, \boldsymbol{a}^{(k+1, \ldots, m)}\right) \\
& =Q_{\boldsymbol{\theta}}^{1, \ldots, k, k+1, \ldots, m}\left(s, \boldsymbol{a}^{(1, \ldots, k, k+1, \ldots, m)}\right)-Q_{\boldsymbol{\theta}}^{1, \ldots, k}\left(s, \boldsymbol{a}^{(1, \ldots, k)}\right)
\end{aligned}
$$

which can be written as a telescoping sum

$$
\begin{aligned}
& Q_{\boldsymbol{\theta}}^{1, \ldots, k, k+1, \ldots, m}\left(s, \boldsymbol{a}^{(1, \ldots, k, k+1, \ldots, m)}\right)-Q_{\boldsymbol{\theta}}^{1, \ldots, k}\left(s, \boldsymbol{a}^{(1, \ldots, k)}\right) \\
& =\sum_{i=k+1}^{m}\left[Q_{\boldsymbol{\theta}}^{(1, \ldots, i)}\left(s, \boldsymbol{a}^{(1, \ldots, i)}\right)-Q_{\boldsymbol{\theta}}^{(1, \ldots, i-1)}\left(s, \boldsymbol{a}^{(1, \ldots, i-1)}\right)\right] \\
& =\sum_{i=k+1}^{m} A_{\theta}^{i}\left(s, \boldsymbol{a}^{(1, \ldots, i-1)}, a^{i}\right)
\end{aligned}
$$

Lemma 2. For any state $s \in \mathcal{S}$, we have

$$
\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[A_{\boldsymbol{\theta}}(s, \mathbf{a})\right]=\sum_{i=1}^{n} \mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}}\left[\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathrm{a}^{i}\right)\right]\right]
$$

Proof. The trick of this proof is to develop a relation on the variance of multi-agent advantage which is recursive over the number of agents. We have

$$
\begin{aligned}
& \operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[A_{\boldsymbol{\theta}}(s, \mathbf{a})\right]=\operatorname{Var}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{v} \sim \pi_{\boldsymbol{\theta}}^{v}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathrm{a}^{(1, \ldots, n)}\right)\right] \\
& =\mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{n} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, n)}\right)^{2}\right] \\
& =\mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{n-1} \sim \pi_{\boldsymbol{\theta}}^{n-1}}\left[\mathbb{E}_{\mathbf{a}^{n} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, n)}\right)^{2}\right]\right. \\
& \left.\quad-\mathbb{E}_{\mathbf{a}^{n} \sim \pi_{\theta}^{n}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, n)}\right)\right]^{2}+\mathbb{E}_{\mathbf{a}^{n} \sim \pi_{\theta}^{n}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, n)}\right)\right]^{2}\right] \\
& = \\
& \quad+\mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{n-1} \sim \pi_{\boldsymbol{\theta}}^{n-1}}\left[\operatorname{Var}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, n)}\right)\right]\right] \\
& \\
& \quad \mathbb{E}_{\mathbf{a}^{1}, \ldots, \mathbf{a}^{n-1} \sim \pi_{\boldsymbol{\theta}}^{n-1}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n-1}\left(s, \mathbf{a}^{(1, \ldots, n-1)}\right)^{2}\right]
\end{aligned}
$$

which, by the stronger version of Lemma 1, given by Equation 13, applied to the first term, equals

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{n-1} \sim \pi_{\boldsymbol{\theta}}^{n-1}}\left[\operatorname{Var}_{\mathrm{a}^{n} \sim \pi_{\theta}^{n}}\left[A_{\boldsymbol{\theta}}^{n}\left(s, \mathbf{a}^{(1, \ldots, n-1)}, \mathrm{a}^{n}\right)\right]\right] \\
& \quad+\mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{n-1} \sim \pi_{\theta}^{n-1}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n-1}\left(s, \mathbf{a}^{(1, \ldots, n-1)}\right)^{2}\right]
\end{aligned}
$$

Hence, we have a recursive relation

$$
\begin{aligned}
& \operatorname{Var}_{\mathrm{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{v} \sim \pi_{\boldsymbol{\theta}}^{v}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, n)}\right)\right] \\
& =\mathbb{E}_{\mathrm{a}^{1} \sim \pi_{\theta}^{1}, \ldots, \mathrm{a}^{n-1} \sim \pi_{\theta}^{n-1}}\left[\operatorname{Var}_{\mathrm{a}^{n} \sim \pi_{\theta}^{n}}\left[A_{\boldsymbol{\theta}}^{n}\left(s, \mathbf{a}^{(1, \ldots, n-1)}, \mathrm{a}^{n}\right)\right]\right] \\
& \quad+\operatorname{Var}_{\mathrm{a}^{1} \sim \pi_{\theta}^{1}, \ldots, \mathrm{a}^{n-1} \sim \pi_{\theta}^{n-1}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n-1}\left(s, \mathbf{a}^{(1, \ldots, n-1)}\right)\right]
\end{aligned}
$$

from which we can obtain

$$
\begin{aligned}
& \operatorname{Var}_{\mathrm{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{v} \sim \pi_{\boldsymbol{\theta}}^{v}}\left[A_{\boldsymbol{\theta}}^{1, \ldots, n}\left(s, \mathrm{a}^{(1, \ldots, n)}\right)\right] \\
& =\sum_{i=1}^{n} \mathbb{E}_{\mathrm{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}}\left[\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathrm{a}^{i}\right)\right]\right]
\end{aligned}
$$

Remark 4. Lemma 2 has a stronger version, coming as a corollary to the above proof; that is

$$
\begin{align*}
& \operatorname{Var}_{\mathrm{a}^{k+1} \sim \pi_{\theta}^{k+1}, \ldots, \mathrm{a}^{n} \sim \pi_{\theta}^{n}}\left[A^{k+1, \ldots, n}\left(s, \boldsymbol{a}^{(1, \ldots, k)}, \mathbf{a}^{(k+1, \ldots, n)}\right)\right] \\
& =\sum_{i=k+1}^{n} \mathbb{E}_{\mathrm{a}^{k+1} \sim \pi_{\theta}^{k+1}, \ldots, \mathrm{a}^{i-1} \sim \pi_{\theta}^{i-1}}\left[\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\theta}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{1, \ldots, k}, \mathbf{a}^{k+1, \ldots, i-1}, \mathrm{a}^{i}\right)\right]\right] . \tag{14}
\end{align*}
$$

We think of it as a corollary to the proof of the lemma, as the fixed joint action $\boldsymbol{a}^{1, \ldots, k}$ has the same algebraic properites, throghout the proof, as state $s$.
Lemma 3. For any state $s \in \mathcal{S}$, we have

$$
\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[A_{\boldsymbol{\theta}}(s, \mathbf{a})\right] \leq \sum_{i=1}^{n} \operatorname{Var}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}, \mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)\right]
$$

Proof. By Lemma 2, we have

$$
\begin{equation*}
\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[A_{\boldsymbol{\theta}}(s, \mathbf{a})\right]=\sum_{i=1}^{n} \mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathbf{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}}\left[\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathbf{a}^{i}\right)\right]\right] \tag{15}
\end{equation*}
$$

Take an arbitrary $i$. We have

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}}\left[\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathrm{a}^{i}\right)\right]\right] \\
& =\mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}}\left[\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathrm{a}^{i}\right)^{2}\right]\right] \\
& =\mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}}\left[\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbb{E}_{\mathrm{a}^{i+1} \sim \pi_{\boldsymbol{\theta}}^{i+1}, \ldots, \mathbf{a}^{n} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[A_{\boldsymbol{\theta}}^{i, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathbf{a}^{(i, \ldots, n)}\right)\right]^{2}\right]\right] \\
& \leq \mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathrm{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}}\left[\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbb{E}_{\mathbf{a}^{i+1} \sim \pi_{\boldsymbol{\theta}}^{i+1}, \ldots, \mathrm{a}^{n} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[A_{\boldsymbol{\theta}}^{i, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathbf{a}^{(i, \ldots, n)}\right)^{2}\right]\right]\right] \\
& =\mathbb{E}_{\mathbf{a}^{1} \sim \pi_{\boldsymbol{\theta}}^{1}, \ldots, \mathbf{a}^{i-1} \sim \pi_{\boldsymbol{\theta}}^{i-1}, \mathrm{a}^{i+1} \sim \pi_{\boldsymbol{\theta}}^{i+1}, \ldots, \mathbf{a}^{n} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i, \ldots, n}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathbf{a}^{(i, \ldots, n)}\right)^{2}\right]\right]
\end{aligned}
$$

The above can be equivalently, but more tellingly, rewritten after permuting (cyclic shift) the labels of agents, in the following way

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i+1, \ldots, n, i}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathbf{a}^{(i+1, \ldots, n, i)}\right)^{2}\right]\right] \\
& =\mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i+1, \ldots, n, i}\left(s, \mathbf{a}^{(1, \ldots, i-1)}, \mathbf{a}^{(i+1, \ldots, n, i)}\right)\right]\right]
\end{aligned}
$$

which, by the strong version of Lemma 1 , equals

$$
\mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{-i}, \mathbf{a}^{i}\right)\right]\right]
$$

which can be further simplified by

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)\right]\right]=\mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right]\right] \\
& =\mathbb{E}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right]=\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \mathbf{a}^{-i}, \mathbf{a}^{i}\right)\right]
\end{aligned}
$$

which, combined with Equation 15, finishes the proof.
Remark 5. Again, subsuming a joint action $a^{(1, \ldots, k)}$ into state in the above proof, we can have a stronger version of Lemma 3,

$$
\begin{align*}
& \operatorname{Var}_{\mathrm{a}^{k+1} \sim \pi_{\boldsymbol{\theta}}^{k+1}, \ldots, \mathrm{a}^{n} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[A_{\boldsymbol{\theta}}^{k+1, \ldots, n}\left(s, \boldsymbol{a}^{(1, \ldots, k)}, \mathbf{a}^{(k+1, \ldots, n)}\right)\right] \\
& \leq \sum_{i=k+1}^{n} \operatorname{Var}_{\mathrm{a}^{k+1} \sim \pi_{\theta}^{k+1}, \ldots, \mathrm{a}^{n} \sim \pi_{\boldsymbol{\theta}}^{n}}\left[A^{i}\left(s, \boldsymbol{a}^{(k+1, \ldots, i-1, i+1, \ldots, n)}, \mathrm{a}^{i}\right)\right] \tag{16}
\end{align*}
$$

## B. 2 Proofs of Theorems 1 and 2

Let us recall the two assumptions that we make in the paper.
Assumption 1. The state space $\mathcal{S}$, and every agent $i$ 's action space $\mathcal{A}^{i}$ is either discrete and finite, or continuous and compact.
Assumption 2. For all $i \in \mathcal{N}, s \in \mathcal{S}$, $a^{i} \in \mathcal{A}^{i}$, the map $\theta^{i} \mapsto \pi_{\boldsymbol{\theta}}^{i}\left(a^{i} \mid s\right)$ is continuously differentiable.
These assumptions assure that the supremum $\sup _{s, a^{i}}\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(a^{i} \mid s\right)\right\|$ exists for every agent $i$. We notice that the supremum $\sup _{s, \boldsymbol{a}^{-i}, a^{i}}\left|A^{i}\left(s, \boldsymbol{a}^{-i}, a^{i}\right)\right|$ exists regardless of assumptions, as by Remark 2, the multi-agent advantage is bounded from both sides.
Theorem 1. The CTDE and DT estimators of MAPG satisfy

$$
\operatorname{Var}_{\mathrm{s}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathbf{g}_{C}^{i}\right]-\operatorname{Var}_{\mathrm{s}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{D}^{i}\right] \leq \frac{B_{i}^{2}}{1-\gamma^{2}} \sum_{j \neq i} \epsilon_{j}^{2} \leq(n-1) \frac{\left(\epsilon B_{i}\right)^{2}}{1-\gamma^{2}}
$$

where $B_{i}=\sup _{s, \boldsymbol{a}}\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\right\|, \epsilon_{i}=\sup _{s, \boldsymbol{a}^{-i}, a^{i}}\left|A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{-i}, a^{i}\right)\right|$, and $\epsilon=\max _{i} \epsilon_{i}$.
Proof. It suffices to prove the first inequality, as the second one is a trivial upper bound. Let's consider an arbitrary time step $t \geq 0$. Let

$$
\begin{aligned}
& \mathbf{g}_{\mathrm{C}, t}^{i}=\hat{Q}\left(\mathrm{~s}_{t}, \mathbf{a}_{t}\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{~s}_{t}, \mathrm{a}_{t}^{i}\right) \\
& \mathbf{g}_{\mathrm{D}, t}^{i}=\hat{Q}^{i}\left(\mathrm{~s}_{t}, \mathrm{a}_{t}^{i}\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{~s}_{t}, \mathrm{a}_{t}^{i}\right)
\end{aligned}
$$

be the contributions to the centralised and decentralised gradient estimators coming from sampling $\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a}_{t} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}$. Note that

$$
\mathbf{g}_{\mathrm{C}}^{i}=\sum_{t=0}^{\infty} \gamma^{t} \mathbf{g}_{\mathrm{C}, t}^{i} \quad \text { and } \quad \mathbf{g}_{\mathrm{D}}^{i}=\sum_{t=0}^{\infty} \gamma^{t} \mathbf{g}_{\mathrm{D}, t}^{i}
$$

Moreover, let $\mathrm{g}_{\mathrm{C}, t, j}^{i}$ and $\mathrm{g}_{\mathrm{D}, t, j}^{i}$ be the $j^{t h}$ components of $\mathbf{g}_{\mathrm{C}, t}^{i}$ and $\mathbf{g}_{\mathrm{D}, t}^{i}$, respectively. Using the law of total variance, we have

$$
\begin{align*}
& \operatorname{Var}_{\mathrm{s} \sim d_{\theta}^{t}, \mathbf{a} \sim \pi_{\theta}}\left[\mathrm{g}_{\mathrm{C}, t, j}^{i}\right]-\operatorname{Var}_{\mathrm{s} \sim d_{\theta}^{t}, \mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right] \\
& =\left(\operatorname{Var}_{\mathrm{s} \sim d_{\theta}^{t}}\left[\mathbb{E}_{\mathbf{a} \sim \pi_{\theta}}\left[\mathrm{g}_{C, t, j}^{i}\right]\right]+\mathbb{E}_{\mathrm{s} \sim d_{\theta}^{t}}\left[\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t, j}^{i}\right]\right]\right) \\
& \quad-\left(\operatorname{Var}_{\mathrm{s} \sim d_{\theta}^{t}}\left[\mathbb{E}_{\mathbf{a} \sim \pi_{\theta}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right]\right]+\mathbb{E}_{\mathrm{s} \sim d_{\theta}^{t}}\left[\operatorname{Var}_{\mathbf{a} \sim \pi_{\theta}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right]\right]\right) \tag{17}
\end{align*}
$$

Noting that $\mathbf{g}_{\mathrm{C}}^{i}$ and $\mathbf{g}_{\mathrm{D}}^{i}$ have the same expectation over $\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}$, the above simplifies to

$$
\begin{align*}
& \mathbb{E}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}}\left[\operatorname{Var}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t, j}^{i}\right]\right]-\mathbb{E}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}}\left[\operatorname{Var}_{\mathbf{a} \sim \boldsymbol{p} \boldsymbol{i}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right]\right] \\
& =\mathbb{E}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}}\left[\operatorname{Var}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t, j}^{i}\right]-\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right]\right] \tag{18}
\end{align*}
$$

Let's fix a state $s$. Using (again) the fact that the expectations of the two gradients are the same, we have

$$
\begin{aligned}
& \operatorname{Var}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t, j}^{i}\right]-\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right] \\
& =\left(\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\left(\mathrm{g}_{\mathrm{C}, t, j}^{i}\right)^{2}\right]-\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t, j}^{i}\right]^{2}\right)-\left(\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\left(\mathrm{g}_{\mathrm{D}, t, j}^{i}\right)^{2}\right]-\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right]^{2}\right) \\
& =\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\left(\mathrm{g}_{\mathrm{C}, t, j}^{i}\right)^{2}\right]-\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\left(\mathrm{g}_{\mathrm{D}, t, j}^{i}\right)^{2}\right] \\
& =\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\left(\mathrm{g}_{\mathrm{C}, t, j}^{i}\right)^{2}-\left(\mathrm{g}_{\mathrm{D}, t, j}^{i}\right)^{2}\right] \\
& =\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta^{i}} \hat{Q}(s, \mathbf{a})\right)^{2}-\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta^{i}} \hat{Q}^{i}\left(s, \mathbf{a}^{i}\right)\right)^{2}\right] \\
& =\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta^{i}}\right)^{2} \mathbb{E}_{\mathbf{a}^{-i} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}^{-i}}\left[\hat{Q}\left(s, \mathbf{a}^{i}, \mathbf{a}^{-i}\right)^{2}-\hat{Q}^{i}\left(s, \mathbf{a}^{i}\right)^{2}\right]\right] .
\end{aligned}
$$

The inner expectation is the variance of $\hat{Q}\left(s, \mathrm{a}^{i}, \mathbf{a}^{-i}\right)$, given $\mathrm{a}^{i}$. We rewrite it as

$$
\begin{aligned}
& =\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta^{i}}\right)^{2} \mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\left(\hat{Q}\left(s, \mathrm{a}^{i}, \mathbf{a}^{-i}\right)-\hat{Q}^{i}\left(s, \mathbf{a}^{i}\right)\right)^{2}\right]\right] \\
& =\mathbb{E}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta^{i}}\right)^{2}\left(\hat{Q}(s, \mathbf{a})-\hat{Q}^{i}\left(s, \mathrm{a}^{i}\right)\right)^{2}\right]
\end{aligned}
$$

Now, recalling that the variance of the total gradient is the sum of variances of the gradient components, we have

$$
\begin{aligned}
& \operatorname{Var}_{\mathbf{a} \sim \pi_{\theta}}\left[\mathbf{g}_{\mathrm{C}, t}^{i}\right]-\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{D}, t}^{i}\right]=\mathbb{E}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right\|^{2}\left(\hat{Q}(s, \mathbf{a})-\hat{Q}^{i}\left(s, \mathbf{a}^{i}\right)\right)^{2}\right] \\
& \leq B_{i}^{2} \mathbb{E}_{\mathbf{a} \sim \pi_{\theta}}\left[\left(\hat{Q}(s, \mathbf{a})-\hat{Q}^{i}\left(s, \mathbf{a}^{i}\right)\right)^{2}\right]=B_{i}^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\left(\hat{Q}(s, \mathbf{a})-\hat{Q}^{i}\left(s, \mathbf{a}^{i}\right)\right)^{2}\right]\right] \\
& =B_{i}^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[\mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\left(\hat{Q}\left(s, \mathbf{a}^{i}, \mathbf{a}^{-i}\right)-\hat{Q}^{i}\left(s, \mathbf{a}^{i}\right)\right)^{2}\right]\right] \\
& =B_{i}^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\hat{A}^{-i}\left(s, \mathbf{a}^{i}, \mathbf{a}^{-i}\right)^{2}\right]\right]=B_{i}^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\operatorname{Var}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\hat{A}^{-i}\left(s, \mathbf{a}^{i}, \mathbf{a}^{-i}\right)\right]\right]
\end{aligned}
$$

which by the strong version of Lemma 3, given in Equation 16, can be upper-bounded by

$$
B_{i}^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[\sum_{j \neq i} \operatorname{Var}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\hat{A}^{j}\left(s, \mathbf{a}^{-j}, \mathbf{a}^{j}\right)\right]\right]=B_{i}^{2} \sum_{j \neq i} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[\operatorname{Var}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\hat{A}^{j}\left(s, \mathbf{a}^{-j}, \mathbf{a}^{j}\right)\right]\right]
$$

Notice that, for any $j \neq i$, we have

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[\operatorname{Var}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\hat{A}^{j}\left(s, \mathbf{a}^{-j}, \mathbf{a}^{j}\right)\right]\right] \\
& =\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[\mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}}\left[\hat{A}^{j}\left(s, \mathbf{a}^{-j}, \mathbf{a}^{j}\right)^{2}\right]\right] \\
& =\mathbb{E}_{\mathbf{a} \sim \pi_{\theta}}\left[\hat{A}^{j}\left(s, \mathbf{a}^{-j}, \mathrm{a}^{j}\right)^{2}\right] \leq \epsilon_{j}^{2}
\end{aligned}
$$

This gives

$$
\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t}^{i}\right]-\operatorname{Var}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t}^{i}\right] \leq B_{i}^{2} \sum_{j \neq i} \epsilon_{j}^{2}
$$

and combining it with Equations 17 and 18 for entire gradient vectors, we get

$$
\begin{equation*}
\operatorname{Var}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{C}, t}^{i}\right]-\operatorname{Var}_{\mathrm{s} \sim d_{\theta}^{t}, \mathbf{a} \sim \pi_{\theta}}\left[\mathbf{g}_{\mathrm{D}, t}^{i}\right] \leq B_{i}^{2} \sum_{j \neq i} \epsilon_{j}^{2} \tag{19}
\end{equation*}
$$

Noting that

$$
\begin{aligned}
& \operatorname{Var}_{\mathrm{s}_{0}: \infty \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\cdot}^{i}\right]=\operatorname{Var}_{\mathrm{s}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \pi_{\boldsymbol{\theta}}}\left[\sum_{t=0}^{\infty} \gamma^{t} \mathbf{g}_{\cdot, t}^{i}\right] \\
& =\sum_{t=0}^{\infty} \operatorname{Var}_{\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\gamma^{t} \mathbf{g}_{\cdot, t}^{i}\right]=\sum_{t=0}^{\infty} \gamma^{2 t} \operatorname{Var}_{\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\cdot, t}^{i}\right]
\end{aligned}
$$

Combining this series expansion with the estimate from Equation 19, we finally obtain

$$
\begin{aligned}
& \operatorname{Var}_{\mathrm{s}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{C}}^{i}\right]-\operatorname{Var}_{\mathrm{s}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{D}}^{i}\right] \\
& \quad \leq \sum_{j=0}^{\infty} \gamma^{2 t}\left(B_{i}^{2} \sum_{j \neq i} \epsilon_{j}^{2}\right) \leq \frac{B_{i}^{2}}{1-\gamma^{2}} \sum_{j \neq i} \epsilon_{j}^{2}
\end{aligned}
$$

Theorem 2. The COMA and DT estimators of MAPG satisfy

$$
\operatorname{Var}_{\mathrm{s}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{C O M A}^{i}\right]-\operatorname{Var}_{\mathrm{s}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathbf{g}_{D}^{i}\right] \leq \frac{\left(\epsilon_{i} B_{i}\right)^{2}}{1-\gamma^{2}}
$$

Proof. Just like in the proof of Theorem 1, we start with the difference

$$
\operatorname{Var}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}, \mathrm{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{COMA}, t, j}^{i}\right]-\operatorname{Var}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}, \mathrm{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right]
$$

which we transform to an analogue of Equation 18:

$$
\mathbb{E}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}}\left[\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{COMA}, t, j}^{i}\right]-\operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{D}, t, j}^{i}\right]\right]
$$

which is trivially upper-bounded by

$$
\mathbb{E}_{\mathrm{s} \sim d_{\boldsymbol{\theta}}^{t}}\left[\operatorname{Var}_{\mathrm{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{COMA}, t, j}^{i}\right]\right]
$$

Now, let us fix a state $s$. We have

$$
\begin{align*}
& \operatorname{Var}_{\mathbf{a} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{COMA}, t, j}^{i}\right]=\operatorname{Var}_{\mathbf{a}^{-i} \sim \pi_{\theta}^{-i}, \mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}} A^{i}\left(s, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)\right] \\
& \leq \mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}, \mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2} A^{i}\left(s, \mathbf{a}^{-i}, \mathbf{a}^{i}\right)^{2}\right] \\
& \leq \epsilon_{i}^{2} \mathbb{E}_{\mathbf{a}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}, \mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right] \tag{20}
\end{align*}
$$

which summing over all components of $\theta^{i}$ gives

$$
\operatorname{Var}_{\mathbf{a} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{COMA}, t}^{i}\right] \leq\left(\epsilon_{i} B_{i}\right)^{2}
$$

Now, applying the reasoning from Equation 19 until the end of the proof of Theorem 1, we arrive at the result

$$
\operatorname{Var}_{\mathrm{S}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{COMA}}^{i}\right]-\operatorname{Var}_{\mathrm{S}_{0: \infty} \sim d_{\boldsymbol{\theta}}^{0: \infty}, \mathbf{a}_{0: \infty} \sim \boldsymbol{\pi}_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{D}}^{i}\right] \leq \frac{\left(\epsilon_{i} B_{i}\right)^{2}}{1-\gamma^{2}}
$$

## C Proofs of the Results about Optimal Baselines

In this section of the Appendix we prove the results about optimal baselines, which are those that minimise the CTDE MAPG estimator's variance. We rely on the following variance decomposition

$$
\begin{aligned}
& \operatorname{Var}_{\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a}_{t} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{C}, t}^{i}(b)\right]=\operatorname{Var}_{\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}}\left[\mathbb{E}_{\mathbf{a}_{t} \sim \pi_{\boldsymbol{\theta}}}\left[\mathbf{g}_{\mathrm{C}, t}^{i}(b)\right]\right]+\mathbb{E}_{\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}}\left[\operatorname{Var}_{\mathbf{a}_{t} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t}^{i}(b)\right]\right] \\
& =\operatorname{Var}_{\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}}\left[\mathbb{E}_{\mathbf{a}_{t} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t}^{i}(b)\right]\right]+\mathbb{E}_{\mathrm{s}_{t} \sim d_{\boldsymbol{\theta}}^{t}}\left[\operatorname{Var}_{\mathbf{a}_{t}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\mathbb{E}_{\mathbf{a}_{t}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathrm{~g}_{\mathrm{C}, t}^{i}(b)\right]\right]+\mathbb{E}_{\mathbf{a}_{t}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\operatorname{Var}_{\mathbf{a}_{t}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathrm{~g}_{\mathrm{C}, t}^{i}(b)\right]\right]\right] \\
& =\underbrace{\operatorname{Var}_{\mathrm{s}_{t} \sim d_{\theta}^{t}}\left[\mathbb{E}_{\mathbf{a}_{t} \sim \pi_{\boldsymbol{\theta}}}\left[\mathrm{g}_{\mathrm{C}, t}^{i}(b)\right]\right]}+\underbrace{\mathbb{E}_{\mathrm{s}_{t} \sim d_{\theta}^{t}}\left[\operatorname{Var}_{\mathbf{a}_{t}^{-i} \sim \pi_{\theta}^{-i}}\left[\mathbb{E}_{\mathbf{a}_{t}^{i} \sim \pi_{\theta}^{i}}\left[\mathrm{~g}_{\mathrm{C}, t}^{i}(b)\right]\right]\right]}+\underbrace{}_{\underbrace{}_{\mathrm{s}_{\boldsymbol{t}} \sim d_{\boldsymbol{\theta}}^{t}, \mathbf{a}_{t}^{-i} \sim \pi_{\boldsymbol{\theta}}^{-i}}\left[\operatorname{Var}_{\mathbf{a}_{a_{t}} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathrm{~g}_{\mathrm{C}, t}^{i}(b)\right]\right]} .
\end{aligned}
$$

Variance from state
Variance from other agents' actions
Variance from agent $i$ 's action

This decomposition reveals that baselines impact the variance via the local variance $\operatorname{Var}_{\mathrm{a}_{t}^{i} \sim \pi_{\theta}^{i}}\left[\mathbf{g}_{\mathrm{C}, t}^{i}(b)\right]$. We rely on this fact in the proofs below.

## C. 1 Proof of Theorem 3

Theorem 3 (Optimal baseline for MAPG). The optimal baseline ( $O B$ ) for the MAPG estimator is

$$
\begin{equation*}
b^{\text {optimal }}\left(\mathbf{s}, \mathbf{a}^{-i}\right)=\frac{\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\mathrm{~s}, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathbf{s}\right)\right\|^{2}\right]}{\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathbf{a}^{i} \mid \mathbf{s}\right)\right\|^{2}\right]} \tag{7}
\end{equation*}
$$

Proof. From the decomposition of the estimator's variance, we know that minimisation of the variance is equivalent to minimisation of the local variance

$$
\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}\left(s, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)-b\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right]
$$

For a baseline $b$, we have

$$
\begin{align*}
& \left.\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\right)\right)\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)\right] \\
& \left.=\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\right)\right)^{2}\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right] \\
& \left.-\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\right)\right)\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)\right]^{2} \\
& \left.=\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\right)\right)^{2}\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right]  \tag{22}\\
& \quad-\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)\right]^{2}
\end{align*}
$$

as $b$ is a baseline. So in order to minimise variance, we shall minimise the term 22.

$$
\begin{aligned}
& \left.\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\right)\right)^{2}\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right] \\
& \left.=\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(b^{2}-2 b \hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)+\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right)\right)\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right] \\
& =b^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right]-2 b \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right] \\
& \\
& +\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right]
\end{aligned}
$$

which is a quadratic in $b$. The last term of the quadratic does not depend on $b$, and so it can be treated as a constant. Recalling that the variance of the whole gradient vector $\mathbf{g}^{i}(b)$ is the sum of variances of its components $\mathrm{g}_{j}^{i}(b)$, we obtain it by summing over $j$

$$
\begin{align*}
& \operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right] \\
& =\sum_{j}\left(b^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right]\right. \\
& \left.-2 b \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\left(\frac{\partial \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)}{\partial \theta_{j}^{i}}\right)^{2}\right]+\text { const }\right) \\
& =b^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right\|^{2}\right] \\
& \quad-2 b \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right\|^{2}\right]+\text { const } \tag{23}
\end{align*}
$$

As the leading coefficient is positive, the quadratic achieves the minimum at

$$
b^{\text {optimal }}=\frac{\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\mathrm{~s}, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)\right\|^{2}\right]}{\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\right\|^{2}\right]}
$$

## C. 2 Remarks about the surrogate optimal baseline

In the paper, we discussed the impracticality of the above baseline. To handle this, we noticed that the policy $\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid s\right)$, at state $s$, is determined by the output layer, $\psi_{\boldsymbol{\theta}}^{i}(s)$, of an actor neural network. With this representation, in order to handle the impracticality of the above optimal baseline, we considered a minimisation objective, the surrogate local variance, given by

$$
\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}(s)\right)\left(\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\left(s, \boldsymbol{a}^{-i}\right)\right)\right]
$$

As a corollarly to the proof, the surrogate version of the optimal baseline (which we refer to as OB) was proposed, and it is given by

$$
b^{*}\left(\mathrm{~s}, \mathbf{a}^{-i}\right)=\frac{\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\mathrm{~s}, \mathbf{a}^{-i}, \mathrm{a}^{i}\right)\left\|\nabla_{\psi_{\boldsymbol{\theta}}^{i}(\mathrm{~s})} \log \pi^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}(\mathrm{~s})\right)\right\|^{2}\right]}{\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left\|\nabla_{\psi_{\boldsymbol{\theta}}^{i}(\mathrm{~s})} \log \pi^{i}\left(\mathrm{a} \mid \psi_{\boldsymbol{\theta}}^{i}(\mathrm{~s})\right)\right\|^{2}\right]} .
$$

Remark 6. The $x_{\psi_{\theta}^{i}}^{i}$ measure, for which $b^{*}\left(s, \boldsymbol{a}^{-i}\right)=\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathbf{a}^{i}\right)\right]$, is generally given by

$$
\begin{equation*}
x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)=\frac{\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\right\|^{2}}{\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left\|\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\right\|^{2}\right]} \tag{24}
\end{equation*}
$$

Let us introduce the definition of the softmax function, which is the subject of the next definition. For a vector $\boldsymbol{z} \in \mathbb{R}^{d}$, we have $\operatorname{softmax}(\boldsymbol{z})=\left(\frac{e^{z_{1}}}{\eta}, \ldots, \frac{e^{z_{d}}}{\eta}\right)$, where $\eta=\sum_{j=1}^{d} e^{z_{j}}$. We write $\operatorname{softmax}\left(\psi_{\boldsymbol{\theta}}^{i}(s)\right)\left(a^{i}\right)=\frac{\exp \left(\psi_{\boldsymbol{\theta}}^{i}(s)\left(a^{i}\right)\right)}{\sum_{\tilde{a}^{i}} \exp \left(\psi_{\boldsymbol{\theta}}^{i}(s)\left(\tilde{a}^{i}\right)\right)}$.
Remark 7. When the action space is discrete, and the actor's policy is $\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)=$ $\operatorname{softmax}\left(\psi_{\boldsymbol{\theta}}^{i}(s)\right)\left(\mathrm{a}^{i}\right)$, then the $x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}$ measure is given by

$$
x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)=\frac{\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\left(1+\left\|\pi_{\boldsymbol{\theta}}^{i}(\mathrm{~s})\right\|^{2}-2 \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \mathrm{s}\right)\right)}{1-\left\|\pi_{\boldsymbol{\theta}}^{i}(\mathrm{~s})\right\|^{2}}
$$

Proof. As we do not vary states s and parameters $\boldsymbol{\theta}$ in this proof, let us drop them from the notation for $\pi_{\boldsymbol{\theta}}^{i}$, and $\psi_{\boldsymbol{\theta}}^{i}(\mathrm{~s})$, hence writing $\pi^{i}\left(\mathrm{a}^{i}\right)=\operatorname{softmax}\left(\psi^{i}\right)\left(\mathrm{a}^{i}\right)$. Let us compute the partial derivatives:

$$
\begin{aligned}
& \frac{\partial \log \pi^{i}\left(a^{i}\right)}{\partial \psi^{i}\left(\tilde{a}^{i}\right)}=\frac{\partial \log \operatorname{softmax}\left(\psi^{i}\right)\left(a^{i}\right)}{\partial \psi^{i}\left(\tilde{a}^{i}\right)}=\frac{\partial}{\partial \psi^{i}\left(\tilde{a}^{i}\right)}\left[\log \frac{\exp \left(\psi^{i}\left(a^{i}\right)\right)}{\sum_{\hat{a}^{i}} \exp \left(\psi^{i}\left(\hat{a}^{i}\right)\right)}\right] \\
& =\frac{\partial}{\partial \psi^{i}\left(\tilde{a}^{i}\right)}\left[\psi^{i}\left(a^{i}\right)-\log \sum_{\hat{a}^{i}} \exp \left(\psi^{i}\left(\hat{a}^{i}\right)\right)\right] \\
& =\mathbf{I}\left(a^{i}=\tilde{a}^{i}\right)-\frac{\exp \left(\psi^{i}\left(\tilde{a}^{i}\right)\right)}{\sum_{\hat{a}^{i}} \exp \left(\psi^{i}\left(\hat{a}^{i}\right)\right)}=\mathbf{I}\left(a^{i}=\tilde{a}^{i}\right)-\pi^{i}\left(\tilde{a}^{i}\right)
\end{aligned}
$$

where $\mathbf{I}$ is the indicator function, taking value 1 if the stetement input to it is true, and 0 otherwise. Taking $e_{k}$ to be the standard normal vector with 1 in $k^{\text {th }}$ entry, we have the gradient

$$
\begin{equation*}
\nabla_{\psi^{i}} \log \pi^{i}\left(a^{i}\right)=\boldsymbol{e}_{a^{i}}-\pi^{i} \tag{25}
\end{equation*}
$$

which has the squared norm

$$
\begin{aligned}
& \left\|\nabla_{\psi^{i}} \log \pi^{i}\left(a^{i}\right)\right\|^{2}=\left\|e_{a^{i}}-\pi^{i}\right\|^{2}=\left(1-\pi^{i}\left(a^{i}\right)\right)^{2}+\sum_{\tilde{a}^{i} \neq a^{i}}\left(-\pi^{i}\left(\tilde{a}^{i}\right)\right)^{2} \\
& =1+\sum_{\tilde{a}^{i}}\left(-\pi^{i}\left(\tilde{a}^{i}\right)\right)^{2}-2 \pi^{i}\left(a^{i}\right)=1+\left\|\pi^{i}\right\|^{2}-2 \pi^{i}\left(a^{a}\right) .
\end{aligned}
$$

The expected value of this norm is

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{a}^{i} \sim \pi^{i}}\left[1+\left\|\pi^{i}\right\|^{2}-2 \pi^{i}\left(a^{i}\right)\right]=1+\left\|\pi^{i}\right\|^{2}-\mathbb{E}_{\mathbf{a}^{i} \sim \pi^{i}}\left[2 \pi^{i}\left(a^{i}\right)\right] \\
& =1+\left\|\pi^{i}\right\|^{2}-2 \sum_{\tilde{a}^{i}}\left(\pi^{i}\left(a^{i}\right)\right)^{2}=1-\left\|\pi^{i}\right\|^{2}
\end{aligned}
$$

which combined with Equation 24 finishes the proof.

## C. 3 Proof of Theorem 4

Theorem 4. The excess surrogate local variance for baseline $b$ satisfies

$$
\Delta \operatorname{Var}(b)=\left(b-b^{*}\left(s, \boldsymbol{a}^{-i}\right)\right)^{2} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left\|\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}(s)\right)\right\|^{2}\right]
$$

In particular, the excess variance of the vanilla MAPG and COMA estimators satisfy

$$
\begin{aligned}
& \Delta \operatorname{Var}_{M A P G} \leq D_{i}^{2}\left(\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{-i}, \mathbf{a}^{i}\right)\right]+Q_{\boldsymbol{\theta}}^{-i}\left(s, \boldsymbol{a}^{-i}\right)^{2}\right) \leq D_{i}^{2}\left(\epsilon_{i}^{2}+\left[\frac{\beta}{1-\gamma}\right]^{2}\right) \\
& \Delta \operatorname{Var}_{C O M A} \leq D_{i}^{2} \operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{-i}, \mathbf{a}^{i}\right)\right] \leq\left(\epsilon_{i} D_{i}\right)^{2}
\end{aligned}
$$

where $D_{i}=\sup _{a^{i}}| | \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(a^{i} \mid \psi_{\boldsymbol{\theta}}^{i}(s)\right) \|$, and $\epsilon_{i}=\sup _{s, a^{-i}, a^{i}}\left|A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{-i}, a^{i}\right)\right|$.

Proof. The first part of the theorem (the formula for excess variance) follows from Equation 23. For the rest of the statements, it suffices to show the first of each inequalities, as the later ones follow directly from the fact that $\left|Q_{\boldsymbol{\theta}}(s, \boldsymbol{a})\right| \leq \frac{\beta}{1-\gamma}, \operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]=\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[A_{\boldsymbol{\theta}}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right]$, and the definition of $\epsilon_{i}$. Let us first derive the bounds for $\Delta \operatorname{Var}_{\text {MAPG }}$. Let us, for short-hand, define

$$
c_{\boldsymbol{\theta}}^{i}:=\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left\|\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}\right)\right\|^{2}\right]
$$

We have

$$
\begin{aligned}
& \Delta \operatorname{Var}_{\mathrm{MAPG}}=\Delta \operatorname{Var}(0)=c_{\boldsymbol{\theta}}^{i} b^{*}\left(s, \boldsymbol{a}^{-i}\right)^{2}=c_{\boldsymbol{\theta}}^{i} \mathbb{E}_{\mathrm{a}^{i} \sim x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]^{2} \\
& \leq c_{\boldsymbol{\theta}}^{i} \mathbb{E}_{\mathrm{a}^{i} \sim x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right]=c_{\boldsymbol{\theta}}^{i} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\frac{\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left\|\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}\right)\right\|^{2}}{c_{\boldsymbol{\theta}}^{i}}\right] \\
& =\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left\|\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}(s)\right)\right\|^{2}\right] \\
& \quad \leq \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2} D_{i}^{2}\right] \\
& =D_{i}^{2}\left(\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right]-\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]^{2}+\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]^{2}\right) \\
& =D_{i}^{2}\left(\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]+\hat{Q}^{-i}\left(s, \boldsymbol{a}^{-i}\right)^{2}\right) \\
& =D_{i}^{2}\left(\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{A}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]+\hat{Q}^{-i}\left(s, \boldsymbol{a}^{-i}\right)^{2}\right)
\end{aligned}
$$

which finishes the proof for MAPG. For COMA, we have

$$
\begin{aligned}
& \Delta \operatorname{Var}_{\mathrm{COMA}}=\Delta \operatorname{Var}\left(\hat{Q}^{-i}\left(s, \boldsymbol{a}^{-i}\right)\right)=c_{\boldsymbol{\theta}}^{i}\left(b^{*}\left(s, \boldsymbol{a}^{-i}\right)-\hat{Q}^{-i}\left(s, \boldsymbol{a}^{-i}\right)\right)^{2} \\
& =c_{\boldsymbol{\theta}}^{i}\left(\mathbb{E}_{\mathbf{a}^{i} \sim x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]-\hat{Q}^{-i}\left(s, \boldsymbol{a}^{-i}\right)\right)^{2} \\
& =c_{\boldsymbol{\theta}}^{i} \mathbb{E}_{\mathbf{a}^{i} \sim x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}}\left[\hat{Q}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-\hat{Q}^{-i}\left(s, \boldsymbol{a}^{-i}\right)\right]^{2} \\
& =c_{\boldsymbol{\theta}}^{i} \mathbb{E}_{\mathbf{a}^{i} \sim x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}}\left[\hat{A}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]^{2} \\
& \leq c_{\boldsymbol{\theta}}^{i} \mathbb{E}_{\mathbf{a}^{i} \sim x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}}\left[\hat{A}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right] \\
& \left.=c_{\boldsymbol{\theta}}^{i} \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{A}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2} \xrightarrow[c_{\boldsymbol{\theta}}]{ } \xrightarrow[\nabla_{\boldsymbol{\theta}}^{i} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}(s)\right) \|^{2}]\right]{c_{\boldsymbol{\theta}}^{i}}\right] \\
& =\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{A}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left\|\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i} \mid \psi_{\boldsymbol{\theta}}^{i}(s)\right)\right\|^{2}\right] \\
& \leq \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[D_{i}^{2} \hat{A}^{i}\left(s, \boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\right] \leq\left(\epsilon_{i} D_{i}\right)^{2}
\end{aligned}
$$

which finishes the proof.

## D Pytorch Implementations of the Optimal Baseline

First, we import necessary packages, which are PyTorch [20], and its nn.functional sub-package. These are standard Deep Learning packages used in RL [1].
import torch, torch.nn.functional as $F$
We then implement a simple method that normalises a row vector, so that its (non-negative) entries sum up to 1 , making the vector a probability distribution.

```
# x: batch of row vectors to normalise to probability mass
normalize = lambda x: F.normalize(x, p=1, dim=-1)
```

The discrete OB is an exact dot product between the measure $x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}$, and available values of $\hat{Q}$.

```
# q: Q values of actions of agent i
# pi: policy of agent i
def optimal_baseline(q, pi):
    M = torch.norm(pi, dim=-1, keepdim = True) ** 2 + 1
    xweight = normalize( (M - 2 * pi) * pi )
    return (xweight * q).sum(-1)
```

In the continuos case, the measure $x_{\psi_{\theta}^{i}}^{i}$ and $Q$-values can only be sampled at finitely many points.

```
# a: sampled actions of agent i
# q: Q values of the sampled actions
# mu, std: parameters of the Gaussian policy of agent i
def optimal_baseline(a, q, mu, std):
    mu_term = torch.norm((a - mu)/std**2, dim=-1)
    std_term = torch.norm(((a - mu)**2 - std**2)/std**3, dim=-1)
    xweight = normalize(mu_term**2 + std_term**2)
    return (xweight * q).sum(-1)
```

We can incorporate it into our MAPG algorith by simply replacing the values of advantage with the values of X , in the buffer. Below, we present a discrete example

```
# compute the policy and sample an action from it
a, pi = actor(obs)
q = critic(obs)
#compute OB
ob = optimal_baseline(q, pi)
# use OB to construct the loss
q = q.gather(-1, a)
pi = pi.gather(-1, a)
X = q-ob
loss = -(X * torch.log(pi)).mean()
```

and a continous one

```
# normal sampling step, where log_pi is the log probability of a
a, log_pi = actor(obs, deterministic=False)
q = critic(obs, a)
# resample m (e.g., m=1000) actions for the observation
obs_m = obs.unsqueeze(0).repeat (m, 1)
a_m, mu_m, std_m = actor(obs, deterministic=False)
# approximate OB
q_m = critic(obs, a_m)
ob = optimal_baseline(a_m, q_m, mu_m, std_m)
# use OB to construct the loss
X = q-ob
loss = -(X * log_pi).mean()
```


## E Computations for the Numerical Toy Example

Here we prove that the quantities in table are filled properly.

| $\mathrm{a}^{i}$ | $\psi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)$ | $\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)$ | $x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}\left(\mathrm{a}^{i}\right)$ | $\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$ | $\hat{A}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$ | $\hat{X}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$ | Method | Variance |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | $\log 8$ | 0.8 | 0.14 | 2 | -9.7 | -41.71 | MAPG | $\mathbf{1 3 2 1}$ |
| 2 | 0 | 0.1 | 0.43 | 1 | -10.7 | -42.71 | COMA | $\mathbf{1 0 1 5}$ |
| 3 | 0 | 0.1 | 0.43 | 100 | 88.3 | 56.29 | OB | $\mathbf{6 7 3}$ |

Proof. In this proof, for convenience, the multiplication and exponentiation of vectors is elementwise. Firstly, we trivially obtain the column $\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)$, by taking softmax over $\psi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)$. This allows us to compute the counterfactual baseline of COMA, which is

$$
\begin{aligned}
& \hat{Q}^{-i}\left(\boldsymbol{a}^{-i}\right)=\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)\right]=\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{Q}\left(\boldsymbol{a}^{-i}, a^{i}\right) \\
& =0.8 \times 2+0.1 \times 1+0.1 \times 100=1.6+0.1+10=11.7
\end{aligned}
$$

By subtracting this value from the column $\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$, we obtain the column $\hat{A}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$.
Let us now compute the column of $x_{\psi_{\theta}^{i}}^{i}$. For this, we use Remark 7. We have

$$
\left\|\pi_{\boldsymbol{\theta}}^{i}\right\|^{2}=0.8^{2}+0.1^{2}+0.1^{2}=0.66
$$

and $1+\left\|\pi_{\boldsymbol{\theta}}^{i}\right\|^{2}-2 \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)=1.66-2 \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)$, which is 0.06 for $\mathrm{a}^{i}=1$, and 1.46 when $\mathrm{a}^{i}=2,3$. For $\mathrm{a}^{i}=1$, we have that

$$
\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\left(1+\left\|\pi_{\boldsymbol{\theta}}^{i}\right\|^{2}-2 \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right)=0.8 \times 0.06=0.048
$$

and for $\mathrm{a}^{i}=2,3$, we have

$$
\pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\left(1+\left\|\pi_{\boldsymbol{\theta}}^{i}\right\|^{2}-2 \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right)=0.1 \times 1.46=0.146
$$

We obtain the column $x_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)$ by normalising the vector $(0.048,0.146,0.146)$. Now, we can compute OB , which is the dot product of the columns $x_{\psi_{\boldsymbol{\theta}}^{i}}^{i}\left(\mathrm{a}^{i}\right)$ and $\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$

$$
b^{*}\left(\boldsymbol{a}^{-i}\right)=0.14 \times 2+0.43 \times 1+0.43 \times 100=0.28+0.43+43=43.71
$$

We obtain the column $\hat{X}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$ after subtracting $b^{*}\left(\boldsymbol{a}^{-i}\right)$ from the column $\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)$.
Now, we can compute and compare the variances of vanilla MAPG, COMA, and OB. The surrogate local variance of an MAPG estimator $\mathbf{g}^{i}(b)$ is

$$
\begin{aligned}
& \operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbf{g}^{i}(b)\right]=\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\left(\boldsymbol{a}^{-i}\right)\right) \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\theta}^{i}\left(a^{i}\right)\right] \\
& =\operatorname{sum}\left(\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\left[\hat{Q}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\left(\boldsymbol{a}^{-i}\right)\right] \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\theta}^{i}\left(a^{i}\right)\right)^{2}\right]-\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\hat{Q}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\left(\boldsymbol{a}^{-i}\right)\right) \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\theta}^{i}\left(a^{i}\right)\right]^{2}\right) \\
& =\operatorname{sum}\left(\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\left(\left[\hat{Q}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)-b\left(\boldsymbol{a}^{-i}\right)\right] \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\theta}^{i}\left(a^{i}\right)\right)^{2}\right]\right)-\operatorname{sum}\left(\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right) \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\theta}^{i}\left(a^{i}\right)\right]^{2}\right)
\end{aligned}
$$

where "sum" is taken element-wise. The last equality follows be linearity of element-wise summing, and the fact that $b$ is a baseline. We compute the variance of vanilla MAPG ( $\mathbf{g}_{\mathrm{MAPG}}^{i}$ ), COMA ( $\mathbf{g}_{\mathrm{COMA}}^{i}$ ),
and OB $\left(\mathbf{g}_{X}^{i}\right)$. Let us derive the first moment, which is the same for all methods

$$
\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right) \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right]=\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{Q}\left(\boldsymbol{a}^{-i}, a^{i}\right) \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\theta}^{i}\left(a^{i}\right)
$$

recalling Equation 25

$$
\begin{aligned}
& =\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{Q}\left(\boldsymbol{a}^{-i}, a^{i}\right)\left(\boldsymbol{e}_{a^{i}}-\pi_{\boldsymbol{\theta}}^{i}\right) \\
& =0.8 \times 2 \times\left[\begin{array}{c}
0.2 \\
-0.1 \\
-0.1
\end{array}\right]+0.1 \times 1 \times\left[\begin{array}{c}
-0.8 \\
0.9 \\
-0.1
\end{array}\right]+0.1 \times 100 \times\left[\begin{array}{c}
-0.8 \\
-0.1 \\
0.9
\end{array}\right] \\
& =\left[\begin{array}{c}
0.32 \\
-0.16 \\
-0.16
\end{array}\right]+\left[\begin{array}{c}
-0.08 \\
0.09 \\
-0.01
\end{array}\right]+\left[\begin{array}{c}
-8 \\
-1 \\
9
\end{array}\right]=\left[\begin{array}{c}
-7.76 \\
-1.07 \\
8.83
\end{array}\right]
\end{aligned}
$$

Now, let's compute the second moment for each of the methods, starting from vanilla MAPG

$$
\begin{aligned}
& \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right)^{2}\right] \\
& =\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{Q}\left(\boldsymbol{a}^{-i}, a^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right)\right)^{2} \\
& =\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{Q}\left(\boldsymbol{a}^{-i}, a^{i}\right)^{2}\left(\boldsymbol{e}_{a^{i}}-\pi_{\boldsymbol{\theta}}^{i}\right)^{2} \\
& =0.8 \times 2^{2} \times\left[\begin{array}{c}
0.2 \\
-0.1 \\
-0.1
\end{array}\right]+0.1 \times 1^{2} \times\left[\begin{array}{c}
-0.8 \\
0.9 \\
-0.1
\end{array}\right]^{2}+0.1 \times 100^{2} \times\left[\begin{array}{c}
-0.8 \\
-0.1 \\
0.9
\end{array}\right]^{2} \\
& =0.8 \times 4 \times\left[\begin{array}{l}
0.04 \\
0.01 \\
0.01
\end{array}\right]+0.1 \times\left[\begin{array}{c}
0.64 \\
0.81 \\
0.01
\end{array}\right]+0.1 \times 10000 \times\left[\begin{array}{c}
0.64 \\
0.01 \\
0.81
\end{array}\right] \\
& =\left[\begin{array}{c}
0.128 \\
0.032 \\
0.032
\end{array}\right]+\left[\begin{array}{l}
0.064 \\
0.081 \\
0.001
\end{array}\right]+\left[\begin{array}{c}
640 \\
10 \\
810
\end{array}\right]=\left[\begin{array}{c}
640.192 \\
10.113 \\
810.033
\end{array}\right]
\end{aligned}
$$

We have

$$
\begin{aligned}
& \operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbf{g}_{\mathrm{MAPG}}^{i}\right] \\
& =\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right)\right)^{2}\right] \\
& -\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{Q}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right) \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right]^{2} \\
& =\left[\begin{array}{c}
640.192 \\
10.113 \\
810.033
\end{array}\right]-\left[\begin{array}{c}
-7.76 \\
-1.07 \\
8.83
\end{array}\right]^{2}=\left[\begin{array}{c}
640.192 \\
10.113 \\
810.033
\end{array}\right]-\left[\begin{array}{c}
60.2176 \\
1.1449 \\
77.9689
\end{array}\right]=\left[\begin{array}{c}
579.9744 \\
8.968 \\
732.064
\end{array}\right]
\end{aligned}
$$

So the variance of vanilla MAPG in this case is

$$
\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\theta}^{i}}\left[\mathbf{g}_{\mathrm{MAPG}}^{i}\right]=1321.007
$$

Let's now deal with COMA

$$
\begin{aligned}
& \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{A}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right)^{2}\right] \\
& =\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{A}^{i}\left(\boldsymbol{a}^{-i}, a^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right)\right)^{2} \\
& =\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{A}^{i}\left(\boldsymbol{a}^{-i}, a^{i}\right)^{2}\left(\boldsymbol{e}_{a^{i}}-\pi_{\boldsymbol{\theta}}^{i}\right)^{2} \\
& =0.8 \times(-9.7)^{2} \times\left[\begin{array}{c}
0.2 \\
-0.1 \\
-0.1
\end{array}\right]^{2}+0.1 \times(-10.7)^{2} \times\left[\begin{array}{c}
-0.8 \\
0.9 \\
-0.1
\end{array}\right]^{2}+0.1 \times 88.3^{2} \times\left[\begin{array}{c}
-0.8 \\
-0.1 \\
0.9
\end{array}\right]^{2} \\
& =0.8 \times 94.09 \times\left[\begin{array}{l}
0.04 \\
0.01 \\
0.01
\end{array}\right]+0.1 \times 114.49 \times\left[\begin{array}{l}
0.64 \\
0.81 \\
0.01
\end{array}\right]+0.1 \times 7796.89 \times\left[\begin{array}{c}
0.64 \\
0.01 \\
0.81
\end{array}\right] \\
& =\left[\begin{array}{l}
3.011 \\
0.753 \\
0.753
\end{array}\right]+\left[\begin{array}{l}
2.327 \\
9.274 \\
0.114
\end{array}\right]+\left[\begin{array}{c}
499.001 \\
7.797 \\
631.548
\end{array}\right]=\left[\begin{array}{c}
504.339 \\
17.824 \\
632.415
\end{array}\right]
\end{aligned}
$$

We have

$$
\begin{aligned}
& \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbf{g}_{\mathrm{COMA}}^{i}\right] \\
= & \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{A}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left(\nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right)^{2}\right]-\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{A}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right) \nabla_{\theta^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right]^{2} \\
& =\left[\begin{array}{c}
504.339 \\
17.824 \\
632.415
\end{array}\right]-\left[\begin{array}{c}
-7.76 \\
-1.07 \\
8.83
\end{array}\right]^{2}=\left[\begin{array}{c}
504.339 \\
17.824 \\
632.415
\end{array}\right]-\left[\begin{array}{c}
60.2176 \\
1.1449 \\
77.9689
\end{array}\right]=\left[\begin{array}{c}
444.1214 \\
16.6791 \\
554.4461
\end{array}\right]
\end{aligned}
$$

and we have

$$
\operatorname{Var}_{\mathrm{a}^{i} \sim \pi_{\theta}^{i}}\left[\mathbf{g}_{\mathrm{COMA}}^{i}\right]=1015.2466
$$

Lastly, we figure out OB

$$
\begin{aligned}
& \mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{X}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right)^{2}\right] \\
& =\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{X}^{i}\left(\boldsymbol{a}^{-i}, a^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right)\right)^{2} \\
& =\sum_{a^{i}=1}^{3} \pi_{\boldsymbol{\theta}}^{i}\left(a^{i}\right) \hat{X}^{i}\left(\boldsymbol{a}^{-i}, a^{i}\right)^{2}\left(\boldsymbol{e}_{a^{i}}-\pi_{\boldsymbol{\theta}}^{i}\right)^{2} \\
& =0.8 \times(-41.71)^{2} \times\left[\begin{array}{c}
0.2 \\
-0.1 \\
-0.1
\end{array}\right]^{2}+0.1 \times(-42.71)^{2} \times\left[\begin{array}{c}
-0.8 \\
0.9 \\
-0.1
\end{array}\right]^{2}+0.1 \times 56.29^{2} \times\left[\begin{array}{c}
-0.8 \\
-0.1 \\
0.9
\end{array}\right]^{2} \\
& =0.8 \times 1739.724 \times\left[\begin{array}{c}
0.04 \\
0.01 \\
0.01
\end{array}\right]+0.1 \times 1824.144 \times\left[\begin{array}{c}
0.64 \\
0.81 \\
0.01
\end{array}\right]+0.1 \times 3168.564 \times\left[\begin{array}{c}
0.64 \\
0.01 \\
0.81
\end{array}\right] \\
& =\left[\begin{array}{c}
55.6712 \\
13.92 \\
13.92
\end{array}\right]+\left[\begin{array}{c}
116.7452 \\
147.756 \\
1.824
\end{array}\right]+\left[\begin{array}{c}
202.788 \\
3.169 \\
256.654
\end{array}\right]=\left[\begin{array}{c}
375.2044 \\
164.845 \\
272.398
\end{array}\right]
\end{aligned}
$$

We have

$$
\begin{aligned}
& \mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\mathbf{g}_{X}^{i}\right] \\
& =\mathbb{E}_{\mathbf{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{X}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right)^{2}\left(\nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\boldsymbol{\theta}}^{i}\left(\mathrm{a}^{i}\right)\right)^{2}\right]-\mathbb{E}_{\mathrm{a}^{i} \sim \pi_{\boldsymbol{\theta}}^{i}}\left[\hat{X}^{i}\left(\boldsymbol{a}^{-i}, \mathrm{a}^{i}\right) \nabla_{\psi_{\boldsymbol{\theta}}^{i}} \log \pi_{\theta}^{i}\left(\mathrm{a}^{i}\right)\right]^{2} \\
& =\left[\begin{array}{c}
375.2044 \\
164.845 \\
272.398
\end{array}\right]-\left[\begin{array}{c}
-7.76 \\
-1.07 \\
8.83
\end{array}\right]^{2}=\left[\begin{array}{c}
375.2044 \\
164.845 \\
272.398
\end{array}\right]-\left[\begin{array}{c}
60.2176 \\
1.1449 \\
77.9689
\end{array}\right]=\left[\begin{array}{c}
314.987 \\
163.7 \\
194.429
\end{array}\right]
\end{aligned}
$$

and we have

$$
\operatorname{Var}_{\mathbf{a}^{i} \sim \pi_{\theta}^{i}}\left[\mathbf{g}_{X}^{i}\right]=673.116
$$

## F Detailed Hyper-parameter Settings for Experiments

In this section, we include the details of our experiments. Their implementations can be found in the following codebase:

$$
\begin{gathered}
\text { https://github.com/morning9393/ } \\
\text { Optimal-Baseline-for-Multi-agent-Policy-Gradients. }
\end{gathered}
$$

In COMA experiments, we use the official implementation in their codebase [7]. The only difference between COMA with and without OB is the baseline introduced, that is, the OB or the counterfactual baseline of COMA [7].

Hyper-parameters used for COMA in the SMAC domain.

| Hyper-parameters | 3 m | 8 m | 2 s 3 z |
| :---: | :---: | :---: | :---: |
| actor lr | $5 \mathrm{e}-3$ | $1 \mathrm{e}-2$ | $1 \mathrm{e}-2$ |
| critic lr | $5 \mathrm{e}-4$ | $5 \mathrm{e}-4$ | $5 \mathrm{e}-4$ |
| gamma | 0.99 | 0.99 | 0.99 |
| epsilon start | 0.5 | 0.5 | 0.5 |
| epsilon finish | 0.01 | 0.01 | 0.01 |
| epsilon anneal time | 50000 | 50000 | 50000 |
| batch size | 8 | 8 | 8 |
| buffer size | 8 | 8 | 8 |
| target update interval | 200 | 200 | 200 |
| optimizer | RMSProp | RMSProp | RMSProp |
| optim alpha | 0.99 | 0.99 | 0.99 |
| optim eps | $1 \mathrm{e}-5$ | $1 \mathrm{e}-5$ | $1 \mathrm{e}-5$ |
| grad norm clip | 10 | 10 | 10 |
| actor network | rnn | rnn | rnn |
| rnn hidden dim | 64 | 64 | 64 |
| critic hidden layer | 1 | 1 | 1 |
| critic hiddem dim | 128 | 128 | 128 |
| activation | ReLU | ReLU | ReLU |
| eval episodes | 32 | 32 | 32 |

As for Multi-agent PPO, based on the official implementation [45], the original V-based critic is replaced by Q-based critic for OB calculation. Simultaneously, we have not used V-based tricks like the GAE estimator, when either using OB or state value as baselines, for fair comparisons.

We provide the pseudocode of our implementation of Multi-agent PPO with OB. We highlight the novel components of it (those unpresent, for example, in [45]) in colour.

```
Algorithm 1 Multi-agent PPO with Q-critic and OB
    Initialize \(\theta\) and \(\phi\), the parameters for actor \(\pi\) and critic \(Q\)
    episode \(_{\text {max }} \leftarrow\) step \(_{\max } /\) batch_size
    while episode \(\leq\) episode \(_{\max }\) do
        Set data buffer \(\boldsymbol{D}=\{ \}\)
        Get initial states \(s_{0}\) and observations \(\boldsymbol{o}_{\mathbf{0}}\)
        for \(t=0\) to batch_size do
            for all agents \(i\) do
                if discrete action space then
                    \(a_{t}^{i}, p_{\pi, t}^{i} \leftarrow \pi\left(o_{t}^{i} ; \theta\right) / /\) where \(p_{\pi, t}^{i}\) is the probability distribution of available actions
            else if continuous action space then
                \(a_{t}^{i}, p_{a, t}^{i} \leftarrow \pi\left(o_{t}^{i} ; \theta\right) / /\) where \(p_{a, t}^{i}\) is the probability density of action \(a_{t}^{i}\)
            end if
            \(q_{t}^{i} \leftarrow Q\left(s_{t}, i, a_{t}^{i} ; \phi\right)\)
            end for
            \(s_{t+1}, o_{t+1}^{n}, r_{t} \leftarrow\) execute \(\left\{a_{t}^{1} \ldots a_{t}^{n}\right\}\)
            if discrete action space then
                Append \(\left[s_{t}, \boldsymbol{o}_{t}, \boldsymbol{a}_{\boldsymbol{t}}, r_{t}, s_{t+1}, \boldsymbol{o}_{\boldsymbol{t + 1}}, \boldsymbol{q}_{\boldsymbol{t}}, \boldsymbol{p}_{\boldsymbol{\pi}, t}\right]\) to \(\boldsymbol{D}\)
            else if continuous action space then
                Append \(\left[s_{t}, \boldsymbol{o}_{\boldsymbol{t}}, \boldsymbol{a}_{\boldsymbol{t}}, r_{t}, s_{t+1}, \boldsymbol{o}_{t+1}, \boldsymbol{q}_{t}, \boldsymbol{p}_{a, t}\right]\) to \(\boldsymbol{D}\)
            end if
        end for
        // from now all agents are processed in parallel in \(D\)
        if discrete action space then
            \(o b \leftarrow\) optimal_baseline \(\left(\boldsymbol{q}, \boldsymbol{p}_{\pi}\right) / /\) use data from \(D\)
        else if continuous action space then
            Resample \(\boldsymbol{a}_{\boldsymbol{t}, \mathbf{1} \ldots m}, \boldsymbol{q}_{t, \mathbf{1} \ldots m} \sim \boldsymbol{\mu}_{\boldsymbol{t}}, \boldsymbol{\sigma}_{\boldsymbol{t}}\) for each \(s_{t}, \boldsymbol{o}_{\boldsymbol{t}}\)
            \(o b \leftarrow\) optimal_baseline \((\boldsymbol{a}, \boldsymbol{q}, \boldsymbol{\mu}, \boldsymbol{\sigma}) / /\) use resampled data
        end if
        \(X \leftarrow q-o b\)
        \(\operatorname{Loss}(\theta) \leftarrow-\operatorname{mean}\left(\boldsymbol{X} \cdot \log \boldsymbol{p}_{a}\right)\)
        Update \(\theta\) with Adam/RMSProp to minimise \(\operatorname{Loss}(\theta)\)
    end while
The critic parameter \(\phi\) is trained with TD-learning [34].
```

Hyper-parameters used for Multi-agent PPO in the SMAC domain.

| Hyper-parameters | 3 s vs $5 \mathrm{z} / 5 \mathrm{~m}$ vs $6 \mathrm{~m} / 6 \mathrm{~h}$ vs $8 \mathrm{z} / 27 \mathrm{~m}$ vs 30 m |
| :---: | :---: |
| actor lr | $1 \mathrm{e}-3$ |
| critic lr | $5 \mathrm{e}-4$ |
| gamma | 0.99 |
| batch size | 3200 |
| num mini batch | 1 |
| ppo epoch | 10 |
| ppo clip param | 0.2 |
| entropy coef | 0.01 |
| optimizer | Adam |
| opti eps | $1 \mathrm{e}-5$ |
| max grad norm | 10 |
| actor network | mlp |
| hidden layper | 1 |
| hidden layer dim | 64 |
| activation | ReLU |
| gain | 0.01 |
| eval episodes | 32 |
| use huber loss | True |
| rollout threads | 32 |
| episode length | 100 |

Hyper-parameters used for Multi-agent PPO in the Multi-Agent MuJoCo domain.

| Hyper-parameters | Hopper(3x1) | Swimmer(2x1) | HalfCheetah(6x1) | Walker(2x3) |
| :---: | :---: | :---: | :---: | :---: |
| actor lr | $5 \mathrm{e}-6$ | $5 \mathrm{e}-5$ | $5 \mathrm{e}-6$ | $1 \mathrm{e}-5$ |
| critic lr | $5 \mathrm{e}-3$ | $5 \mathrm{e}-3$ | $5 \mathrm{e}-3$ | $5 \mathrm{e}-3$ |
| lr decay | 1 | 1 | 0.99 | 1 |
| episode limit | 1000 | 1000 | 1000 | 1000 |
| std x coef | 1 | 10 | 5 | 5 |
| std y coef | 0.5 | 0.45 | 0.5 | 0.5 |
| ob n actions | 1000 | 1000 | 1000 | 1000 |
| gamma | 0.99 | 0.99 | 0.99 | 0.99 |
| batch size | 4000 | 4000 | 4000 | 4000 |
| num mini batch | 40 | 40 | 40 | 40 |
| ppo epoch | 5 | 5 | 5 | 5 |
| ppo clip param | 0.2 | 0.2 | 0.2 | 0.2 |
| entropy coef | 0.001 | 0.001 | 0.001 | 0.001 |
| optimizer | RMSProp | RMSProp | RMSProp | RMSProp |
| momentum | 0.9 | 0.9 | 0.9 | 0.9 |
| opti eps | $1 \mathrm{e}-5$ | $1 \mathrm{e}-5$ | $1 \mathrm{e}-5$ | $1 \mathrm{e}-5$ |
| max grad norm | 0.5 | 0.5 | 0.5 | 0.5 |
| actor network | mlp | mlp | mlp | mlp |
| hidden layper | 2 | 2 | 2 | 2 |
| hidden layer dim | 32 | 32 | 32 | 32 |
| activation | ReLU | ReLU | ReLU | ReLU |
| gain | 0.01 | 0.01 | 0.01 | 0.01 |
| eval episodes | 10 | 10 | 10 | 10 |
| use huber loss | True | True | True | True |
| rollout threads | 4 | 4 | 4 | 4000 |
| episode length | 1000 | 1000 | 1000 | 100 |

For QMIX and COMIX baseline algorithms, we use implementation from their official codebases and keep the performance consistent with the results reported in their original papers [21, 24]. MADDPG is provided along with COMIX, which is derived from its official implementation as well [15].

Hyper-parameters used for QMIX baseline in the SMAC domain.

| Hyper-parameters | 3 s vs $5 \mathrm{z} / 5 \mathrm{~m}$ vs $6 \mathrm{~m} / 6 \mathrm{~h}$ vs $8 \mathrm{z} / 27 \mathrm{~m}$ vs 30 m |
| :---: | :---: |
| critic lr | $5 \mathrm{e}-4$ |
| gamma | 0.99 |
| epsilon start | 1 |
| epsilon finish | 0.05 |
| epsilon anneal time | 50000 |
| batch size | 32 |
| buffer size | 5000 |
| target update interval | 200 |
| double q | True |
| optimizer | RMSProp |
| optim alpha | 0.99 |
| optim eps | $1 \mathrm{e}-5$ |
| grad norm clip | 10 |
| mixing embed dim | 32 |
| hypernet layers | 2 |
| hypernet embed | 64 |
| critic hidden layer | 1 |
| critic hiddem dim | 128 |
| activation | ReLU |
| eval episodes | 32 |

Hyper-parameters used for COMIX baseline in the Multi-Agent MuJoCo domain.

| Hyper-parameters | Hopper(3x1) / Swimmer(2x1) / HalfCheetah(6x1) / Walker(2x3) |
| :---: | :---: |
| critic lr | 0.001 |
| gamma | 0.99 |
| episode limit | 1000 |
| exploration mode | Gaussian |
| start steps | 10000 |
| act noise | 0.1 |
| batch size | 100 |
| buffer size | 1 e 6 |
| soft target update | True |
| target update tau | 0.001 |
| optimizer | Adam |
| optim eps | 0.01 |
| grad norm clip | 0.5 |
| mixing embed dim | 64 |
| hypernet layers | 2 |
| hypernet embed | 64 |
| critic hidden layer | 2 |
| critic hiddem dim | $[400,300]$ |
| activation | ReLU |
| eval episodes | 10 |

Hyper-parameters used for MADDPG baseline in the Multi-Agent MuJoCo domain.

| Hyper-parameters | Hopper(3x1) / Swimmer(2x1) / HalfCheetah(6x1) / Walker(2x3) |
| :---: | :---: |
| actor lr | 0.001 |
| critic lr | 0.001 |
| gamma | 0.99 |
| episode limit | 1000 |
| exploration mode | Gaussian |
| start steps | 10000 |
| act noise | 0.1 |
| batch size | 100 |
| buffer size | 1 e 6 |
| soft target update | True |
| target update tau | 0.001 |
| optimizer | Adam |
| optim eps | 0.01 |
| grad norm clip | 0.5 |
| mixing embed dim | 64 |
| hypernet layers | 2 |
| hypernet embed | 64 |
| actor network | mlp |
| hidden layer | 2 |
| hiddem dim | $[400,300]$ |
| activation | ReLU |
| eval episodes | 10 |

