

1 Appendix

2 A Proof Sketch

3 To better clarify our theoretical results, we provide a proof sketch here. Firstly, we decompose the
 4 performance difference bound under the model-based setting into three terms (Theorem 1). Sec-
 5 ondly, by means of using Return Bound (Theorem 2), we can bound these three terms individually
 6 (Theorem 3). Then, we can do some transformation to get Unified Model Shift and Model Bias
 7 Bound (Theorem 4), which bounds the model shift term and the model bias term in total variation
 8 form. However, due to the intractable property of Δ , we further explore the upper bound of $|\Delta|$
 9 (Theorem 5), finding that Δ can be ignored. Finally, by the Integral Probability Metrics (Lemma 3)
 10 and the property of the Wasserstein distance, we derive the target which bounds the model shift term
 11 and the model bias term in the Wasserstein distance form.

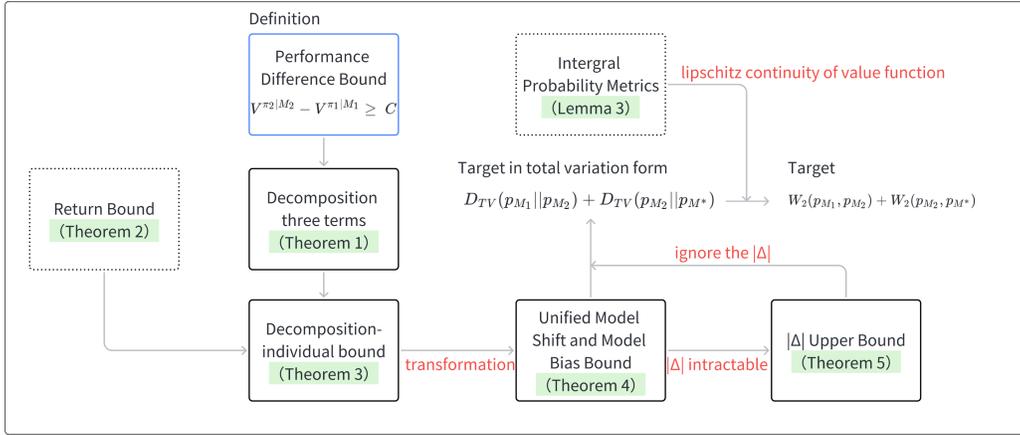


Figure 1: Theoretical sketch of USB-PO.

12 B Useful Lemmas

13 In this section, we provide some proof to support our theoretical analysis.

14 **Lemma 1** (Total variation Distance). *Consider a measurable space (Ω, Σ) and probability measures*
 15 *P and Q are defined on (Ω, Σ) . The total variation distance between P and Q is defined as:*

$$D_{TV}(P||Q) = \sup_{A \in \Sigma} |P(A) - Q(A)| \quad (1)$$

16 *Eq.(1) can be equivalently written as:*

$$D_{TV}(P||Q) = \frac{1}{2} \sum_{\omega \in \Omega} |P(\{\omega\}) - Q(\{\omega\})| \quad (2)$$

17 *Proof:* The proof of this lemma can be found in [7]. □

18 **Lemma 2** (Total Variation Distance of Joint Distributions). *Given two distributions $p(x, y) =$*
 19 *$p(x)p(y|x)$ and $q(x, y) = q(x)q(y|x)$, the total variation distance between them can be bounded*
 20 *as:*

$$D_{TV}(p(x, y)||q(x, y)) \leq D_{TV}(p(x)||q(x)) + \max_x D_{TV}(p(y|x)||q(y|x)) \quad (3)$$

21 *Proof:*

$$\begin{aligned}
D_{TV}(p(x, y)||q(x, y)) &= \frac{1}{2} \sum_{x, y} |p(x, y) - q(x, y)| \\
&= \frac{1}{2} \sum_{x, y} |p(x)p(y|x) - q(x)q(y|x)| \\
&= \frac{1}{2} \sum_{x, y} |p(x)p(y|x) - p(x)q(y|x) + p(x)q(y|x) - q(x)q(y|x)| \\
&\leq \frac{1}{2} \sum_{x, y} p(x)|p(y|x) - q(y|x)| + |p(x) - q(x)|q(y|x) \\
&= \frac{1}{2} \sum_{x, y} p(x)|p(y|x) - q(y|x)| + \frac{1}{2} \sum_x |p(x) - q(x)| \\
&= \mathbb{E}_{x \sim p(x)} [D_{TV}(p(y|x)||q(y|x))] + D_{TV}(p(x)||q(x)) \\
&\leq D_{TV}(p(x)||q(x)) + \max_x D_{TV}(p(y|x)||q(y|x))
\end{aligned} \tag{4}$$

22

□

23 **Lemma 3** (Integral Probability Metrics). *Consider a measurable space (\mathcal{X}, Σ) . The integral probability metric associated with a class \mathcal{F} of real-valued functions on \mathcal{X} is defined as*

$$d_{\mathcal{F}}(P, Q) = \sup_{f \in \mathcal{F}} |\mathbb{E}_{X \sim P}[f(X)] - \mathbb{E}_{Y \sim Q}[f(Y)]| \tag{5}$$

25 where P and Q are probability measures on \mathcal{X} . We demonstrate the following special cases:

26 (a) If $\mathcal{F} = \{f : \|f\|_{\infty} \leq c\}$, then we have

$$d_{\mathcal{F}}(P, Q) = cD_{TV}(P||Q) \tag{6}$$

27 (b) If \mathcal{F} is the set of L -Lipschitz function with a norm $\|\cdot\|$, then we have

$$d_{\mathcal{F}}(P, Q) = LW_1(P, Q) \tag{7}$$

28 In our paper, to distinguish the dynamic transition function, we choose \mathcal{F} to be the class covering V_M^{π} . Since the value function can converge to $\frac{r_{max}}{1-\gamma}$, it only needs to satisfy the L_v -Lipschitz continuity and thus we can get $\frac{r_{max}}{1-\gamma} D_{TV}(p_M||p_{M'}) = L_v W_1(p_M, p_{M'})$ for any arbitrary model M, M' .

32 C Missing Proof

33 **Theorem 1** (Performance Difference Bound Decomposition). *Let $M_i \in \mathcal{M}$ be the evaluated model and $\pi_i \in \Pi$ be the policy derived from the model. The performance difference bound can be decomposed into three terms,*

$$V^{\pi_2|M_2} - V^{\pi_1|M_1} = (V^{\pi_2|M_2} - V_{M_2}^{\pi_2}) - (V^{\pi_1|M_1} - V_{M_1}^{\pi_1}) + (V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1}) \tag{8}$$

36 *Proof:* We introduce two additional terms $V_{M_1}^{\pi_1}$ and $V_{M_2}^{\pi_2}$ that allow the performance difference bound objective to be divided into three operators based on the return bounds, which can be reformulated separately.

$$\begin{aligned}
V^{\pi_2|M_2} - V^{\pi_1|M_1} &= V^{\pi_2|M_2} - V^{\pi_1|M_1} + (V_{M_1}^{\pi_1} - V_{M_1}^{\pi_1}) + (V_{M_2}^{\pi_2} - V_{M_2}^{\pi_2}) \\
&= (V^{\pi_2|M_2} - V_{M_2}^{\pi_2}) - (V^{\pi_1|M_1} - V_{M_1}^{\pi_1}) + (V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1})
\end{aligned} \tag{9}$$

39

□

40 **Theorem 2** (Return Bound). *Let R_{max} denote the bound of the reward function, ϵ_{π} denote $\max_s D_{TV}(\pi_1||\pi_2)$ and $\epsilon_{M_1}^{M_2}$ denote $\mathbb{E}_{(s, a) \sim d_{M_1}^{\pi_1}} [D_{TV}(p_{M_1}||p_{M_2})]$. For two arbitrary policies $\pi_1, \pi_2 \in \Pi$, the expected return under two arbitrary models $M_1, M_2 \in \mathcal{M}$ can be bounded as,*

$$V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1} \geq -2R_{max} \left(\frac{\epsilon_{\pi}}{(1-\gamma)^2} + \frac{\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{M_2} \right) \tag{10}$$

43 *Proof:* We give the thorough proof referring to Lemma B.4 in MBPO [4] as follows.

$$\begin{aligned}
V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1} &= \sum_{t=0}^{\infty} \gamma^t \sum_{s,a} (p_{t,M_2}^{\pi_2}(s,a) - p_{t,M_1}^{\pi_1}(s,a)) r(s,a) \\
&\geq -R_{max} \sum_{t=0}^{\infty} \gamma^t \sum_{s,a} |p_{t,M_2}^{\pi_2}(s,a) - p_{t,M_1}^{\pi_1}(s,a)| \\
&= -2R_{max} \sum_{t=0}^{\infty} \gamma^t D_{TV}(p_{t,M_1}^{\pi_1}(s,a) || p_{t,M_2}^{\pi_2}(s,a))
\end{aligned} \tag{11}$$

44 According to the Lemma 2, we have:

$$\begin{aligned}
D_{TV}(p_{t,M_1}^{\pi_1}(s,a) || p_{t,M_2}^{\pi_2}(s,a)) &\leq D_{TV}(p_{t,M_1}^{\pi_1}(s) || p_{t,M_2}^{\pi_2}(s)) + \max_s D_{TV}(\pi_1(\cdot|s) || \pi_2(\cdot|s)) \\
&= D_{TV}(p_{t,M_1}^{\pi_1}(s) || p_{t,M_2}^{\pi_2}(s)) + \epsilon_{\pi}
\end{aligned} \tag{12}$$

45 Further we expand the first term:

$$\begin{aligned}
&D_{TV}(p_{t,M_1}^{\pi_1}(s) || p_{t,M_2}^{\pi_2}(s)) \\
&= \frac{1}{2} \sum_s |p_{t,M_1}^{\pi_1}(s) - p_{t,M_2}^{\pi_2}(s)| \\
&= \frac{1}{2} \sum_s \left| \sum_{s'} p_{M_1}^{\pi_1}(s|s') p_{t-1,M_1}^{\pi_1}(s') - p_{M_2}^{\pi_2}(s|s') p_{t-1,M_2}^{\pi_2}(s') \right| \\
&\leq \frac{1}{2} \sum_s \sum_{s'} |p_{M_1}^{\pi_1}(s|s') p_{t-1,M_1}^{\pi_1}(s') - p_{M_2}^{\pi_2}(s|s') p_{t-1,M_2}^{\pi_2}(s')| \\
&\leq \frac{1}{2} \sum_{s,s'} p_{t-1,M_1}^{\pi_1}(s') |p_{M_1}^{\pi_1}(s|s') - p_{M_2}^{\pi_2}(s|s')| + p_{M_2}^{\pi_2}(s|s') |p_{t-1,M_1}^{\pi_1}(s') - p_{t-1,M_2}^{\pi_2}(s')| \\
&= \frac{1}{2} \mathbb{E}_{s' \sim p_{t-1,M_1}^{\pi_1}(s')} \left[\sum_s |p_{M_1}^{\pi_1}(s|s') - p_{M_2}^{\pi_2}(s|s')| \right] + D_{TV}(p_{t-1,M_1}^{\pi_1}(s') || p_{t-1,M_2}^{\pi_2}(s')) \\
&= \frac{1}{2} \sum_{t'=1}^t \mathbb{E}_{s' \sim p_{t'-1,M_1}^{\pi_1}(s')} \left[\sum_s |p_{M_1}^{\pi_1}(s|s') - p_{M_2}^{\pi_2}(s|s')| \right] \\
&= \frac{1}{2} \sum_{t'=1}^t \mathbb{E}_{s' \sim p_{t'-1,M_1}^{\pi_1}(s')} \left[\sum_s \left| \sum_a p_{M_1}^{\pi_1}(s,a|s') - p_{M_2}^{\pi_2}(s,a|s') \right| \right] \\
&\leq \frac{1}{2} \sum_{t'=1}^t \mathbb{E}_{s' \sim p_{t'-1,M_1}^{\pi_1}(s')} \left[\sum_{s,a} |p_{M_1}^{\pi_1}(s,a|s') - p_{M_2}^{\pi_2}(s,a|s')| \right] \\
&= \sum_{t'=1}^t \mathbb{E}_{s' \sim p_{t'-1,M_1}^{\pi_1}(s')} D_{TV}(p_{M_1}^{\pi_1}(s,a|s') || p_{M_2}^{\pi_2}(s,a|s')) \\
&\leq \sum_{t'=1}^t \mathbb{E}_{s' \sim p_{t'-1,M_1}^{\pi_1}(s')} [\epsilon_{\pi} + \mathbb{E}_{a \sim \pi_1} [D_{TV}(p_{M_1}(s|s',a) || p_{M_2}(s|s',a))]] \\
&= t\epsilon_{\pi} + \sum_{t'=1}^t \mathbb{E}_{s',a \sim p_{t'-1,M_1}^{\pi_1}(s',a)} D_{TV}(p_{M_1}(s|s',a) || p_{M_2}(s|s',a))
\end{aligned} \tag{13}$$

46 Then move the result of Eq.(13) to Eq.(12), we can get:

$$D_{TV}(p_{t,M_1}^{\pi_1}(s,a) || p_{t,M_2}^{\pi_2}(s,a)) \leq (t+1)\epsilon_{\pi} + \sum_{t'=1}^t \mathbb{E}_{s',a \sim p_{t'-1,M_1}^{\pi_1}(s',a)} D_{TV}(p_{M_1}(s|s',a) || p_{M_2}(s|s',a)) \tag{14}$$

47 Next, we move the result of Eq.(14) to Eq.(11), we can get:

$$\begin{aligned}
& V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1} \\
& \geq -2R_{max} \sum_{t=0}^{\infty} \gamma^t ((t+1)\epsilon_\pi + \sum_{t'=1}^t \mathbb{E}_{s', a \sim p_{t'-1, M_1}^{\pi_1}(s', a)} D_{TV}(p_{M_1}(s|s', a) || p_{M_2}(s|s', a))) \\
& = -2R_{max} \left(\frac{\epsilon_\pi}{(1-\gamma)^2} + \frac{1}{1-\gamma} \sum_{t=1}^{\infty} \gamma^t \mathbb{E}_{s', a \sim p_{t-1, M_1}^{\pi_1}(s', a)} D_{TV}(p_{M_1}(s|s', a) || p_{M_2}(s|s', a)) \right)
\end{aligned} \tag{15}$$

48 Here, we first simplify the second term of the Eq.(15)

$$\begin{aligned}
& \frac{1}{1-\gamma} \sum_{t=1}^{\infty} \gamma^t \mathbb{E}_{s', a \sim p_{t-1, M_1}^{\pi_1}(s', a)} D_{TV}(p_{M_1}(s|s', a) || p_{M_2}(s|s', a)) \\
& = \frac{\gamma}{1-\gamma} \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{s', a \sim p_{t, M_1}^{\pi_1}(s', a)} D_{TV}(p_{M_1}(s|s', a) || p_{M_2}(s|s', a)) \\
& = \frac{\gamma}{(1-\gamma)^2} (1-\gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{s', a \sim p_{t, M_1}^{\pi_1}(s', a)} D_{TV}(p_{M_1}(s|s', a) || p_{M_2}(s|s', a)) \\
& = \frac{\gamma}{(1-\gamma)^2} \mathbb{E}_{s', a \sim d_{M_1}^{\pi_1}(s', a)} D_{TV}(p_{M_1}(s|s', a) || p_{M_2}(s|s', a)) \\
& = \frac{\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{M_2}
\end{aligned} \tag{16}$$

49 Then we bring this result back to Eq.(15) and the proof is complete.

$$V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1} \geq -2R_{max} \left(\frac{\epsilon_\pi}{(1-\gamma)^2} + \frac{\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{M_2} \right) \tag{17}$$

50

51 **Theorem 3** (Decomposition TVD Bound). *Let $\epsilon_{M_i}^{\pi_i}$ denote $\mathbb{E}_{(s,a) \sim d_{M_i}^{\pi_i}} [D_{TV}(p_{M_i} || p_{M^*})]$. Let $M_i \in$*
52 *\mathcal{M} be the evaluated model and $\pi_i \in \Pi$ be the policy derived from the model. The decomposition*
53 *terms can be bounded as,*

$$V^{\pi_2|M_2} - V^{\pi_1|M_1} \geq \frac{2R_{max}\gamma}{(1-\gamma)^2} (\epsilon_{M_1}^{\pi_1} - \epsilon_{M_2}^{\pi_2} - \epsilon_{M_1}^{M_2}) - \frac{2R_{max}\epsilon_\pi}{(1-\gamma)^2} \tag{18}$$

54 *Proof:* According to CMLO [5] and Eq.(10), the term $V^{\pi_1|M_1} - V_{M_1}^{\pi_1}$ can be approximated as
55 $-\frac{2R_{max}\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{\pi_1}$, thus we only need to bound the remaining two terms.

56 For the term $V^{\pi_2|M_2} - V_{M_2}^{\pi_2}$, we use Eq.(10) to bound it.

$$\begin{aligned}
V^{\pi_2|M_2} - V_{M_2}^{\pi_2} & \geq -2R_{max} \left(\frac{\max_s D_{TV}(\pi_2 || \pi_2)}{(1-\gamma)^2} + \frac{\gamma}{(1-\gamma)^2} \epsilon_{M_2}^{\pi_2} \right) \\
& = -\frac{2R_{max}\gamma}{(1-\gamma)^2} \epsilon_{M_2}^{\pi_2}
\end{aligned} \tag{19}$$

57 Similarly, for the term $V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1}$, we can get:

$$V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1} \geq -2R_{max} \left(\frac{\epsilon_\pi}{(1-\gamma)^2} + \frac{\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{M_2} \right) \tag{20}$$

58 We now combine these three bounds together and complete the proof.

$$\begin{aligned}
V^{\pi_2|M_2} - V^{\pi_1|M_1} & = (V^{\pi_2|M_2} - V_{M_2}^{\pi_2}) - (V^{\pi_1|M_1} - V_{M_1}^{\pi_1}) + (V_{M_2}^{\pi_2} - V_{M_1}^{\pi_1}) \\
& \geq -\frac{2R_{max}\gamma}{(1-\gamma)^2} \epsilon_{M_2}^{\pi_2} + \frac{2R_{max}\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{\pi_1} - 2R_{max} \left(\frac{\epsilon_\pi}{(1-\gamma)^2} + \frac{\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{M_2} \right) \\
& = \frac{2R_{max}\gamma}{(1-\gamma)^2} (\epsilon_{M_1}^{\pi_1} - \epsilon_{M_2}^{\pi_2} - \epsilon_{M_1}^{M_2}) - \frac{2R_{max}\epsilon_\pi}{(1-\gamma)^2}
\end{aligned} \tag{21}$$

59

60 **Theorem 4** (Unified Model Shift and Model Bias Bound). Let κ denote the constant $\frac{2R_{max}}{(1-\gamma)^2}$ and
61 Δ denotes $\mathbb{E}_{(s,a)\sim d_{M_1}^{\pi_1}} [D_{TV}(p_{M_2}||p_{M^*})] - \mathbb{E}_{(s,a)\sim d_{M_2}^{\pi_2}} [D_{TV}(p_{M_2}||p_{M^*})]$. Let $M_i \in \mathcal{M}$ be the
62 evaluated model and $\pi_i \in \Pi$ be the policy derived from the model. The unified model shift and
63 model bias bound can be derived as,

$$\begin{aligned} & V^{\pi_2|M_2} - V^{\pi_1|M_1} \\ & \geq \kappa(\gamma \mathbb{E}_{(s,a)\sim d_{M_1}^{\pi_1}} [D_{TV}(p_{M_1}||p_{M^*}) - D_{TV}(p_{M_1}||p_{M_2}) - D_{TV}(p_{M_2}||p_{M^*})] + \Delta) - \epsilon_\pi \end{aligned} \quad (22)$$

64 *Proof:* Based on Eq.(18), we add a new term $\kappa \mathbb{E}_{(s,a)\sim d_{M_1}^{\pi_1}} [D_{TV}(p_{M_2}||p_{M^*})]$ to reformulate the
65 optimization objective. \square

66 **Theorem 5** ($|\Delta|$ Upper Bound). Let $M_i \in \mathcal{M}$ be the evaluated model and $\pi_i \in \Pi$ be the policy
67 derived from the model. The term Δ can be upper bounded as:

$$|\Delta| \leq \frac{2\gamma}{1-\gamma} \mathbb{E}_{(s,a)\sim d_{M_1}^{\pi_1}} [D_{TV}(p_{M_1}||p_{M_2}) \max_{s,a} D_{TV}(p_{M_2}||p_{M^*})] + \frac{2\epsilon_\pi}{1-\gamma} \max_{s,a} D_{TV}(p_{M_2}||p_{M^*}) \quad (23)$$

68 *Proof:* First, we combine these two terms.

$$\begin{aligned} |\Delta| &= |\mathbb{E}_{(s,a)\sim d_{M_1}^{\pi_1}} [D_{TV}(p_{M_2}||p_{M^*})] - \mathbb{E}_{(s,a)\sim d_{M_2}^{\pi_2}} [D_{TV}(p_{M_2}||p_{M^*})]| \\ &= (1-\gamma) \left| \sum_{t=0}^{\infty} \gamma^t \sum_{s,a} (p_{t,M_1}^{\pi_1}(s,a) - p_{t,M_2}^{\pi_2}(s,a)) D_{TV}(p_{M_2}(s'|s,a)||p_{M^*}(s'|s,a)) \right| \\ &\leq (1-\gamma) \max_{s,a} D_{TV}(p_{M_2}(s'|s,a)||p_{M^*}(s'|s,a)) \sum_{t=0}^{\infty} \gamma^t \sum_{s,a} |p_{t,M_1}^{\pi_1}(s,a) - p_{t,M_2}^{\pi_2}(s,a)| \\ &= 2(1-\gamma) \max_{s,a} D_{TV}(p_{M_2}||p_{M^*}) \sum_{t=0}^{\infty} \gamma^t D_{TV}(p_{t,M_1}^{\pi_1}(s,a)||p_{t,M_2}^{\pi_2}(s,a)) \end{aligned} \quad (24)$$

69 Recalling that we get the result of the sum equation above in Eq.(14), and then we have:

$$\begin{aligned} |\Delta| &\leq 2(1-\gamma) \max_{s,a} D_{TV}(p_{M_2}||p_{M^*}) \left(\frac{\epsilon_\pi}{(1-\gamma)^2} + \frac{\gamma}{(1-\gamma)^2} \epsilon_{M_1}^{M_2} \right) \\ &= \frac{2\gamma}{1-\gamma} \mathbb{E}_{s,a\sim d_{M_1}^{\pi_1}} [D_{TV}(p_{M_1}||p_{M_2}) \max_{s,a} D_{TV}(p_{M_2}||p_{M^*})] + \frac{2\epsilon_\pi}{1-\gamma} \max_{s,a} D_{TV}(p_{M_2}||p_{M^*}) \end{aligned} \quad (25)$$

70 \square

71 D Experimental Details

72 D.1 Environment Setup

73 We evaluate the algorithm over a series of MuJoCo [11] continuous control benchmark tasks. To
74 ensure fairness, we use the standard 1000-step version of all the environments. The details of the
environment setup are from OpenAI Gym [1], as shown in Table 1.

Table 1: The general outline of the MuJoCo environment.

Environment-Version	State Dim	Action Dim	Termination
Ant-v2	27	8	obs[0]<0.2 or obs[0] > 1.0
HalfCheetah-v2	17	6	-
Hopper-v2	11	3	obs[1] \geq 0.2 or obs[0] \leq 0.7
Humanoid-v2	45	17	obs[0] < 1.0 or obs[0] > 2.0
InvertedPendulum-v2	4	1	obs[1] > 0.2 or obs[1] < -0.2
Walker2d-v2	17	6	obs[0] \geq 2.0 or obs[0] \leq 0.8 or obs[1] \geq 1.0 or obs[1] \leq -1.0

75

76 **D.2 Baseline implementation**

77 **MFRL Baselines.** We use two state-of-the-art model-free algorithms, i.e. SAC [3] and PPO [9],
 78 to do baseline comparison. To demonstrate the final performance and sampling efficiency of our
 79 method, we train SAC for 3M steps, which is much more than MBRL algorithms. The hyperparam-
 80 eters are consistent with the author’s settings.

81 **MBRL Baselines.** We use several state-of-the-art model-based algorithms to do baseline compar-
 82 ison, covering CMLO [5], MBPO [4], SLBO [8] and STEVE [2]. The implementation of CMLO is
 83 based on the opensource repo published by the author and all of the hyperparameters are set accord-
 84 ing to the paper [5]. Our algorithm USB-PO is implemented based on the opensource repo published
 85 by Janner who is the author of MBPO.

86 We present the final performance on six continuous benchmark tasks in Table 2. The results demon-
 87 strate that our algorithm achieves competitive performance compared to both MBRL and MFRL
 88 baselines over these tasks. Each result in the table shows the average and standard deviation on the
 89 maximum average returns among different random seeds and we choose 250K for HalfCheetah-v2,
 90 300K for Walker2d-v2, 300K for Humanoid-v2, 250K for Ant-v2, 15K for Inverted-Pendulum-v2,
 91 120K for Hopper-v2.

Table 2: The final performance on six continuous benchmark tasks.

		HalfCheetah	Humanoid	Walker2d
MBRL	STEVE	12406.29±458.08	4318.32±853.60	1109.23±1163.74
	SLBO	1915.47±1398.73	459.46±34.27	3107.93±1887.09
	MBPO	12765.67±594.54	5546.77±221.72	4582.06±67.44
	CMLO	10143.55±193.82	5577.01±219.89	4807.60±99.89
	USB-PO	15105.91±177.75	5973.75±110.99	5691.62±162.57
MFRL(@3M steps)	SAC	15012	6207	5879
		Ant	InvertedPendulum	Hopper
MBRL	STEVE	779.72±45.67	778.54±265.51	1131.61±623.52
	SLBO	707.79±218.80	793.24±334.90	898.68±233.21
	MBPO	4926.10±818.38	1000.00±0.00	3436.00±120.72
	CMLO	5123.71±783.97	1000.00±0.00	3495.41±71.02
	USB-PO	6340.84±119.06	1000.00±0.00	3694.22±46.19
MFRL(@3M steps)	SAC	5934	1000	3610

92 **D.3 Hyperparameters**

93 Our algorithm USB-PO is based on MBPO [4] and is implemented according to the opensource repo
 94 published by the MBPO author. Except for the learning rate in phase 2 of our USB-PO algorithm,
 95 the hyperparameters are completely identical to the MBPO settings for all environments. In all
 96 benchmark tasks, we set this learning rate to 1e-4.

97 **D.4 Computing Infrastructure**

98 In Table 3, we list our computing infrastructure and the computational time for training USB-PO
 99 on these six continuous benchmark tasks. Note that the time we report is the cost for 4 random
 100 seeds simultaneously on one graphics card. For Humanoid, only two random seeds can be run
 101 simultaneously because of the limitation of graphics memory.

Table 3: Computing infrastructure and the computational time for each benchmark task compared to MBPO, where the time unit d denotes day and h denotes hour.

	HalfCheetah	Humanoid	Walker2d	Ant	InvertedPendulum	Hopper
CPU	AMD EPYC 7B12 64-Core Processor					
GPU	NVIDIA 2080Ti					
MBPO times	2.46d	1.64d	1.75d	2.88d	3.43h	17.81h
USB-PO times	2.29d	1.51d	1.65d	2.91d	3.42h	18.28h

102 E Comparison with Prior Works

103 In this section, we compare USB-PO with prior theoretical works to emphasize our contribution, as
 104 a complementary to the main paper. First, we give a summary and then show the details as follows.
 105 MBPO-Style does not consider model shift and CMLO-Style rely on a fixed threshold to constrain
 106 model shift. Our algorithm, USB-PO, adaptively adjusts the model updates in a unified manner
 107 (unify model shift and model bias) to get the performance improvement guarantee.

108 **MBPO-Style [4, 10, 6, 12].** They use the return discrepancy bound $V^{\pi|M} \geq V_M^\pi - C(\epsilon_m, \epsilon_\pi)$ to
 109 improve the lower bound on the performance under the real environment, i.e. as long as improving
 110 V_M^π by more than $C(\epsilon_m, \epsilon_\pi)$ can guarantee improvement on $V^{\pi|M}$. Obviously, This scheme is
 111 guaranteed under a fixed model and it does not consider the change in model dynamic during updates
 112 nor the performance variation concerning model shift. Even worse, if the model has some excessive
 113 updates, it is impractical to find a feasible solution to meet the improvement guarantee.

114 **CMLO-Style [5].** They use the performance difference bound under the model-based setting
 115 $V^{\pi_2|M_2} - V^{\pi_1|M_1} \geq C$ to directly consider model shift and model bias. However, they finally
 116 derive a constrained lower-bound optimization problem and use a fixed threshold to constrain model
 117 shift, i.e. $\sup_{s \in \mathcal{S}, a \in \mathcal{A}} D_{TV}(P_{M_1}(\cdot|s, a) || P_{M_2}(\cdot|s, a)) \leq \sigma_{M_1, M_2}$ and determine when to update the
 118 model accordingly. Notably, we find that this fixed threshold plays a key role in the whole algorithm
 119 and needs to be carefully adjusted for each environment. If this threshold is set too low, the model
 120 bias of the following iteration will be large, which impairs the subsequent optimization process. If
 121 this threshold is set too high, the performance improvement can no longer be guaranteed. Addition-
 122 ally, using a fixed threshold during the whole training process makes the algorithm problematic to
 123 adjust adaptively.

124 **USB-PO (Ours).** Following CMLO-Style [5], we also use the performance difference bound un-
 125 der the model-based setting to directly consider model shift and model bias. Compared to relying on
 126 a fixed threshold to constrain model shift, we use a transformation to unify model shift and model
 127 bias into one formulation without the constraint (Theorem 4). Due to the intractable property of
 128 Δ , we further explore the upper bound of $|\Delta|$, finding that Δ can be ignored with respect to model
 129 shift and model bias alone (Theorem 5). Finally, the optimization objective we get can be used
 130 to fine-tune M_2 in a unified manner to adaptively adjust the model updates to get a performance
 131 improvement guarantee. Notably, our algorithm can use the same learning rate of Phase 2 and our
 132 algorithm is robust to this learning rate. To the best of our knowledge, this is the first method that
 133 unifies model shift and model bias and adaptively fine-tunes the model updates during the training
 134 process.

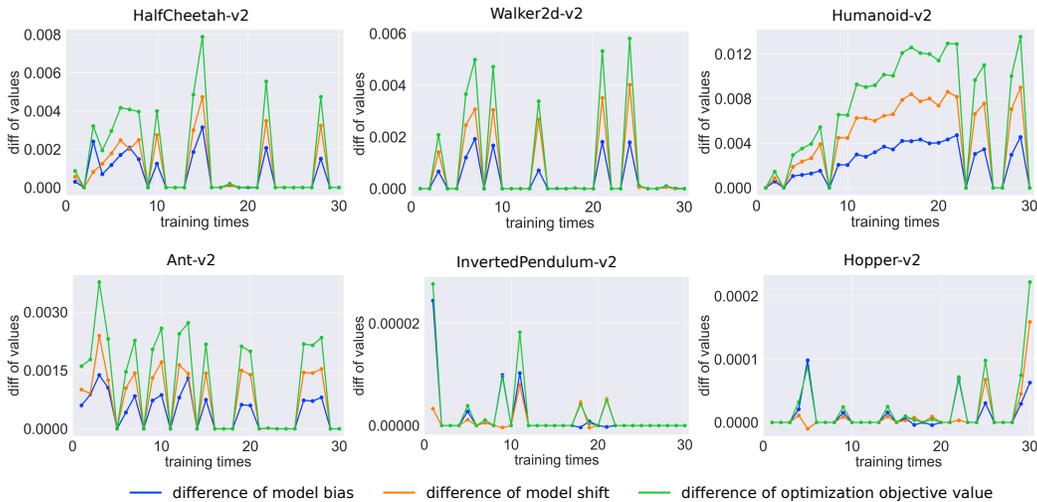


Figure 2: We choose a specific random seed to show the details of the first 30 training times on all benchmark tasks, covering the difference of optimization objective value, the model shift term and the model bias term before and after the fine-tuning process.

135 F Additional Experiment

136 F.1 Working Mechanism Extension

137 To illustrate that the ability of USB-PO to reduce both model shift and model bias potentially is not
138 a coincidence that exists only in the Walker2d environment, we add experimental results in other
139 MuJoCo [11] environments. As shown in Figure 2, when the fine-tuning actually operates, the
140 difference of the model shift term and the model bias term among all of the benchmark tasks are
141 generally both positive, further validating our superiority.

142 F.2 Ablation Study Extension

143 Here, we show the results of the ablation study on all of the MuJoCo benchmark tasks.

144 As shown in Figure 3, only optimizing the model shift term results in a drop in sample efficiency
145 while only optimizing the model bias term leads to performance deterioration. Only fine-tuning the
146 model updates in a unified manner can achieve excellent performance.

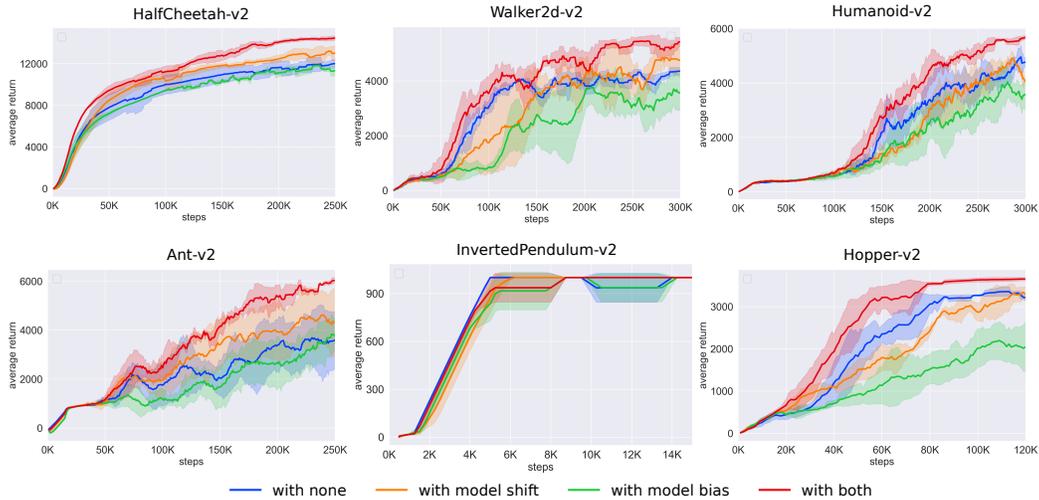


Figure 3: **Optimization Objective Variants** on all MuJoCo benchmark tasks.

147 F.3 Ensemble numbers

148 To illustrate the justification of the ensemble numbers we set, we conduct the ablation experiment
149 on more ensemble numbers containing 3, 5, and 7. As Figure 4 shows, as the number of ensemble
150 models goes up, the performance will be higher and more stable, but it will cost more time. To
151 maintain the balance between performance and time, we finally set the value of this parameter to 7,
152 which is recommended by the MBPO [4] original repo.

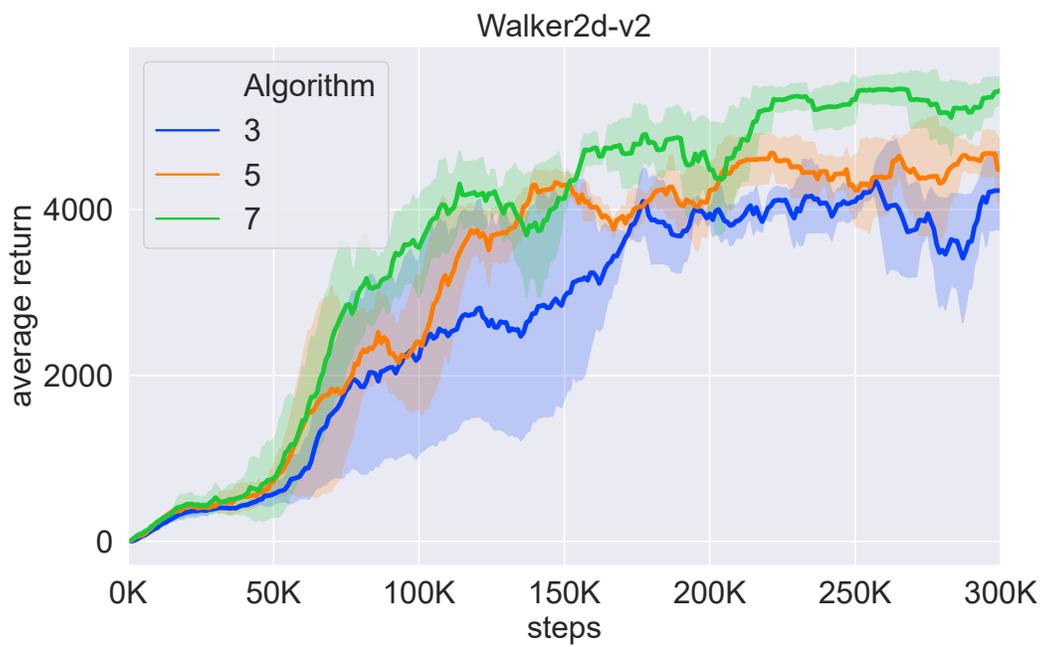


Figure 4: Ablation experimental results of the ensemble model numbers.