## A Implementation Details

## A. 1 Hyperparameter Ranges

RL Hyperparameters: Below are the continuous hyperparameters optimized as well as the fixed parameters, which were all taken from [57]. For the categorical variables, we use the data augmentations used in [57, ,37, 35]: crop, grayscale, cutout, cutout-color, flip, rotate, random convolution and color-jitter.

Table 3: DrAC Learned Hyperparameters

| Parameter | Value |
| :--- | :---: |
| PPO Clip $(\epsilon)$ | $[0.01,0.5]$ |
| Learning Rate | $\left[10^{-5}, 10^{-3}\right]$ |
| Entropy Coeff | $[0,0.2]$ |
| Regularization Parameter $\left(\alpha_{r}\right)$ | $[0.01,0.5]$ |

Table 4: DrAC Fixed Hyperparameters

| Parameter | Value |
| :--- | :---: |
| $\gamma$ | 0.999 |
| $\lambda$ | 0.95 |
| \# timesteps per rollout | 256 |
| \# epochs per rollout | 3 |
| \# minibatches per epoch | 8 |
| optimizer | Adam |

PB2 Hyperparameters: We use the same UCB hyperparameters as in [52], which come from [7]. Concretely, we set $c 1=0.2, c 2=0.4$. We believe performance may be increased by selecting these parameters more carefully. However, it is promising to see that this fine tuning is not crucial for performance.

## A. 2 Environment Details

We used five games from the Procgen environment [13]. We chose these games as there appeared to be a greater degree of variance in the optimal data augmentation in [57]. All protocols for the environment are from [57], as we followed the author's open source implementation for agent training ${ }^{4}$, and only varied the hyperparameters and data augmentation chosen.


Figure 6: Example levels from the seven games considered from [13

[^0]
## A. 3 Infrastructure Details

Each trial was run a single GPU, taking around three days to train all four agents. This was due to the nature of our existing compute infrastructure and could likely be made significantly faster with parallelization. The existing code is thus not optimized to work in a distributed fashion, however it will be open sourced in a similar structure to the open source PB2-Rand implementation (part of the Ray library [45]). Since we followed the open source version of PB2-Rand, it is a trivial extension to include PB2-Mult and PB2-Mix, using new kernels provided both in the supplementary material (as part of the research code) or the synthetic function notebook (see the link in Section 6.

## B Additional Results

Here we show the test learning curves in Fig 7 where we see particularly strong improvement in BigFish and CaveFlyer. Notably, PBT performs poorly in many of the tasks.


Figure 7: Test Performance: Learning curves for seven Procgen games. Plots show the mean $\pm 1$ sem for the test performance of the best training agents in the population. Results are averaged over 5 seeds.

Table 5: Train performance for seven Procgen games. $\dagger$ indicates training was conducted for a single agent for 25 M timesteps. Hyperparameters were initialized at optimized values, meaning the effective number of timesteps is much higher. Final performance is the average of 100 trials and results were taken from [57]. $\ddagger$ indicates training was conducted by a population of four agents for 25 M timesteps each, with several hyperparameters initialized at random. Each agent is evaluated for 100 trials on train and test levels and we present the mean train performance of the agent with the best training performance.

| Environment | PPO $\dagger$ | UCB-DrAC $\dagger$ | PBT $\ddagger$ | PB2-Rand $\ddagger$ | PB2-Mult $\ddagger$ | PB2-Mix $\ddagger$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| BigFish | $8.9 \pm 1.5$ | $13.2 \pm 2.2$ | $12.0 \pm 3.3$ | $20.8 \pm 1.8$ | $18.2 \pm 2.0$ | $17.1 \pm 2.9$ |
| CaveFlyer | $6.8 \pm 0.6$ | $5.7 \pm 0.6$ | $7.2 \pm 0.4$ | $7.3 \pm 1.7$ | $7.0 \pm 1.3$ | $7.5 \pm 0.9$ |
| CoinRun | $9.3 \pm 0.3$ | $9.5 \pm 0.3$ | $8.2 \pm 0.8$ | $10.0 \pm 0$ | $9.9 \pm 0.2$ | $9.6 \pm 0.2$ |
| FruitBot | $29.1 \pm 1.1$ | $29.5 \pm 1.2$ | $27.4 \pm 0.2$ | $32.2 \pm 0.7$ | $30.5 \pm 1.4$ | $30.9 \pm 1.2$ |
| Jumper | $8.3 \pm 0.4$ | $8.1 \pm 0.7$ | $8.3 \pm 0.2$ | $9.0 \pm 0.2$ | $8.9 \pm 0.1$ | $9.2 \pm 0.5$ |
| Leaper | $5.5 \pm 0.4$ | $5.3 \pm 0.5$ | $4.1 \pm 0.5$ | $6.9 \pm 2.0$ | $5.6 \pm 2.1$ | $7.1 \pm 2.3$ |
| StarPilot | $29.8 \pm 2.3$ | $35.3 \pm 2.2$ | $33.6 \pm 7.4$ | $44.1 \pm 2.7$ | $40.3 \pm 1.3$ | $41.8 \pm 2.7$ |



Figure 8: Train Performance: Learning curves for all seven Procgen games. Plots show the mean $\pm 1$ sem for the best training performance within the population. Results are averaged over 5 seeds.

In Table 5]we show training performance for the agents shown in the main paper. Interestingly, we see strong training performance for PB2-Rand, indicating a larger generalization gap. We also show training learning curves for all algorithms in Fig. 8, where we once again report the best agent at each timestep.

Is there dependence between data augmentation and hyperparameters in Procgen? Next we consider the dependence between the hyperparameters. For each environment, we fit a linear regression model predicting the change in reward for 500 k steps of training (i.e. one step of PB2), with the independent variable being one of the continuous hyperparameters. Colors correspond to the setting where we condition on an individual category (by creating a subset), while black corresponds to using the entire dataset. We use the data from all trials and show the results in In Fig. 9


Figure 9: t -statistic for the independent variable (one continuous hyperparameter, in the title of each plot) in predicting the change in reward for the given environment. Blue columns represent the data being from one. The dotted line indicates $p<0.05$.

As we see, in some cases the category does not have any impact on the relationship, for example the entropy coefficient for CoinRun: an increase in entropy appears to be negatively related with training performance across all augmentations. However, for some settings such as the clip parameter in BigFish, we see large swings in the $t$-stat depending on category. We even see some settings where the aggregate relationship (grey) is very small, but conditioning on the category makes it worthwhile to tune the hyperparameter, for example the learning rate for CaveFlyer. In this case it may be challenging for PB2-Rand to achieve gains from tuning the learning rate. We note this analysis is purely illustrative and has many assumptions, in particular 1) we use a linear univariate model 2) we ignore the time-varying component.

Learned Schedules The next pages contain learned schedules from all agents trained by our algorithms and baselines. In each case we show all agents for all seeds, with blue corresponding to a config that contributed to the final best (training) agent. The most notable observation across all environments is the stark contrast between PBT and PB2-based methods. PBT clearly increments parameters, and gradually shifts during training, while PB2 rapidly explores the boundaries. We note in some cases that PB2-Mult never re-explores the middle regions, which we hypothesize is due to the lack of data when creating separate models for each of the eight categories.


Figure 10: Learned schedules for BigFish by algorithm, for all seeds. Each point corresponds to one agent.


Figure 11: Learned schedules for CaveFlyer by algorithm, for all seeds. Each point corresponds to one agent.


Figure 12: Learned schedules for CoinRun by algorithm, for all seeds. Each point corresponds to one agent.


Figure 13: Learned schedules for FruitBot by algorithm, for all seeds. Each point corresponds to one agent.


Figure 14: Learned schedules for Jumper by algorithm, for all seeds. Each point corresponds to one agent.


Figure 15: Learned schedules for Leaper by algorithm, for all seeds. Each point corresponds to one agent.


Figure 16: Learned schedules for StarPilot by algorithm, for all seeds. Each point corresponds to one agent.

## C Additional Background

## C. 1 DrAC

In recent times, there has been increased interest in generalization of RL agents, after multiple works showed many RL agents are simply overfitting to a single deterministic training environment [51, 14]. One approach to produce more generalizable agents is data augmentation, which has recently shown impressive results [82, 35, 37, 38]. To formalize the data augmentation step, we follow [35] and define an optimality-invariant state transformation $f: \mathcal{S} \times \mathcal{H} \rightarrow \mathcal{S}$ as a mapping that preserves both the policy $\pi$ and value function $V$, i.e. $V(s)=V(f(s, \nu))$ and $\pi(a \mid s)=\pi(a \mid f(s, \nu)), \forall s \in \mathcal{S}, \nu \in \mathcal{H}$, where $\nu$ are parameters of $f($.$) drawn from the set of possible parameters \mathcal{H}$. In this work we focus on the formulation from [57] who propose additional loss terms for regularizing the policy and value function:

$$
\begin{gather*}
G_{\pi}=\operatorname{KL}\left[\pi_{\theta}(a \mid f(s, \nu]) \mid \pi(a \mid s)\right]  \tag{6}\\
G_{V}=\left(V_{\theta}(f(s, \nu))-V(s)\right)^{2} \tag{7}
\end{gather*}
$$

Combining the PPO objective with $G_{\pi}$ and $G_{V}$ produces the data-regularized actor critic or DrAC objective as follows:

$$
\begin{equation*}
\mathcal{L}_{\mathrm{DrAC}}(\theta)=\mathcal{L}_{\mathrm{PPO}}(\theta)-\alpha_{r}\left(G_{\pi}+G_{V}\right) \tag{8}
\end{equation*}
$$

where $\alpha_{r}$ is the weight of the regularization term, another hyperparameter to consider. The results in [57] show that learning which data augmentation function $f$ to use (from a fixed set) and grid searching over the $\alpha_{r}$ parameter can achieve new state-of-the-art results in challenging Procgen environments.

Table 6: Notation used in the theoretical analysis.

| Variable | Domain | Meaning |
| :---: | :---: | :---: |
| C | $\mathcal{N}$ | number of categories (number of arms) |
| $T$ | $\mathcal{N}$ | maximum number of bandit update (the number of $t_{\text {ready }}$ in PB2) |
| V | $\mathcal{N}$ | number of time segments, how many times the function has been shifted |
| $B$ | $\mathcal{N}$ | batch size (number of parallel agents) |
| $S_{t}$ | list | a list of $B$ selected categories $\left[c_{t, 1}, c_{t, 2}, \ldots, c_{t, B}\right]$ |
| $e$ | $\mathcal{R}^{+}$ | this is Euler's number 2.71828 |
| $\left[p_{t}^{1}, \ldots p_{t}^{C}\right]$ | list | probability vector at iteration $t$ and category $c=1 \ldots . C$ |
| $w_{c}$ | $\mathcal{R}$ | weight for categorical $c$ (or arm $c$ ) |
| $W_{t}=\sum_{c=1}^{C} w_{c}$ | $\mathcal{R}^{+}$ | sum of the weight vector at iteration $t$ |
| $c_{t, b} \in\{1 \ldots C\}$ | $\mathcal{N}$ | a selection at iteration $t$ at agent $b$ |
| $c_{t}=\left[c_{t, 1}, \ldots, c_{t, B}\right]$ | list | a batch of $B$ categorical selections at iteration $t$ |
| $c_{t}^{*} \in\{1 \ldots C\}$ | list | an optimal selection at iteration $t$ |
| $A_{v}^{*}$ | list | a list of $B$ elements taking the (same) optimal selection $c_{t}^{*}$ at iteration/segment $v$ |
| $g_{t}(c)$ | [0, 1] | a gain (reward) occurred at iteration $t$ by pulling an $\operatorname{arm} c$ |
| $\hat{g}_{t}(c)=\frac{g_{t}(c)}{p_{t}^{c}+\gamma} \mathbb{I}\left(c=h_{t}\right)$ | $\mathcal{R}^{+}$ | a normalized gain |
| $\alpha=\frac{1}{T}, \eta=2 \gamma=\sqrt{\frac{2 \ln C}{C T}}$ | $\mathcal{R}^{+}$ | hyperparameters, set by Theorem 1 |
| $\gamma=\sqrt{\frac{\ln C}{2 C T}} \in[0,1]$ | $\mathcal{R}^{+}$ | is the exploration parameter |
| $G_{T}(c)=\sum_{t=1}^{T} g_{t}(c)$ | $\mathcal{R}^{+}$ | a total gain if we select an $\operatorname{arm} c$ entirely |

## D Theoretical Results

We derive the theoretical proofs presented in the main paper.

## D. 1 Time-varying EXP3 Multiple play (Tv.EXP3.M)

We present a new algorithm for parallel (or multiple play) multi-armed bandits in the time-varying setting with adversarial feedback. Particularly, we extend the multiple play EXP3 algorithm (or EXP3.M) [78], to the time-varying setting where the unknown reward distribution of each arm can change arbitrarily, but the total number of change points is no more than $V$ times. We refer to Table 6 for the notations used in the proofs.
We summarize the TV.EXP3.M in Algorithm 5] - this is complementary to the brief Algorithm 2 in the main paper. In Algorithm [5] we maintain a set of probability vectors (step 9) for each arm which will be specified and weighted in the optimal way derived by the theory. Then, at each round we select a batch of $B$ arms for parallel evaluations (step 10), observe the reward (step 11) and update the model (step 12,13,14). The normalization steps 3-7 are to prevent from being biased toward a single arm which performs overwhelmingly well, thus overly exploiting.
Theorem 3. (Theorem 1 in the main paper) Let $T>0, C>0$, set $\alpha=\frac{1}{T}$ and $\gamma=$ $\min \left\{1, \sqrt{\frac{C \ln (C / B)}{(e-1) B T}}\right\}$, we assume the reward distributions to change at arbitrary time instants, but the total number of change points is no more than $V$ times. The expected regret gained by TV.EXP3.M in a batch satisfies the following sublinear regret bound

$$
\mathbb{E}\left[R_{T B}\right] \leq[1+e+V] \sqrt{(e-1) \frac{C T}{B} \ln \frac{C T}{B}}
$$

Proof. We refer to Table 6 for notations. We follow the proof technique presented in [3] and [78] to derive the regret bound. Let $W_{t}=\sum_{c=1}^{C} \omega_{t}^{c}$ and using step 9 in Algorithm[5, we have $\frac{\frac{p_{t}^{c}}{B}-\frac{\gamma}{C}}{1-\gamma}=\frac{\omega_{t}^{c}}{W_{t}}$.

This equation will be used below.

$$
\begin{align*}
& \frac{W_{t+1}}{W_{t}}=\frac{\sum_{c=1}^{C} \omega_{t+1}^{c}}{W_{t}}=\frac{\sum_{c \notin S_{t}(0)} \omega_{t+1}^{c}}{W_{t}}+\frac{\sum_{c \in S_{t}(0)} \omega_{t+1}^{c}}{W_{t}}+\frac{\sum_{c=1}^{C} \frac{e \alpha W_{t}}{C}}{W_{t}} \\
&=\sum_{c \notin S_{t}(0)} \frac{\omega_{t}^{c}}{W_{t}} \times \exp \left(\frac{B \gamma}{C} \hat{g}_{t}(c)\right)+\frac{\sum_{c \in S_{t}(0)} \omega_{t}^{c}}{W_{t}}+e \alpha \\
& \leq e \alpha+\sum_{c \notin S_{t}(0)} \frac{\omega_{t}^{c}}{W_{t}}\left[1+\frac{B \gamma}{C} \hat{g}_{t}(c)+(e-2)\left(\frac{B \gamma}{C}\right)^{2} \hat{g}_{t}^{2}(c)\right]+\frac{\sum_{c \in S_{t}(0)} \omega_{t+1}^{c}}{W_{t}} \text { by } e^{x} \leq 1+x+(e-2) x^{2} \\
&=e \alpha+\sum_{c \notin S_{t}(0)} \frac{\omega_{t}^{c}}{W_{t}}+\sum_{c \notin S_{t}(0)} \frac{\omega_{t+1}^{c}}{W_{t}}+\sum_{c \notin S_{t}(0)} \frac{\frac{p_{t}^{c}}{B}-\frac{\gamma}{C}}{1-\gamma}\left[\frac{B \gamma}{C} \hat{g}_{t}(c)+(e-2)\left(\frac{B \gamma}{C}\right)^{2} \hat{g}_{t}^{2}(c)\right] \\
& \leq e \alpha+1+\frac{B \gamma}{C(1-\gamma)} \sum_{c \notin S_{t}(0)}\left(\frac{p_{t}^{c}}{B}-\frac{\gamma}{C}\right) \hat{g}_{t}(c)+\frac{(e-2)}{1-\gamma}\left(\frac{B \gamma}{C}\right)^{2} \sum_{c \notin S_{t}(0)}\left(\frac{p_{t}^{c}}{B}-\frac{\gamma}{C}\right) \hat{g}_{t}^{2}(c) \\
&=e \alpha+1+\frac{\gamma}{C(1-\gamma)} \sum_{c \notin S_{t}(0)} p_{t}^{c} \hat{g}_{t}(c)-\frac{B \gamma^{2}}{C^{2}(1-\gamma)} \sum_{c \notin S_{t}(0)} \hat{g}_{t}(c) \\
& \frac{(e-2) \gamma^{2} B}{(1-\gamma) C^{2}} \sum_{c \notin S_{t}(0)}^{p_{t}^{c} \hat{g}_{t}^{2}(c)-\frac{(e-2) \gamma^{3} B^{2}}{(1-\gamma) C^{2}} \sum_{c \notin S_{t}(0)} \hat{g}_{t}^{2}(c)} \\
& \leq e \alpha+1+\frac{\gamma}{C(1-\gamma)} \sum_{c \notin S_{t}(0)} p_{t}^{c} \hat{g}_{t}(c)+\frac{(e-2) \gamma^{2} B}{(1-\gamma) C^{2}} \sum_{c \notin S_{t}(0)} p_{t}^{c} \hat{g}_{t}^{2}(c) \\
& \leq e \alpha+1+\frac{\gamma}{C(1-\gamma)} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c)+\frac{(e-2) \gamma^{2} B}{(1-\gamma) C^{2}} \sum_{c \notin S_{t}(0)} \hat{g}_{t}(c) \text { by } x^{2} \leq x, \forall x \in[0,1] \\
& \frac{\gamma}{W_{t+1}}  \tag{9}\\
& W_{t} \leq e \alpha+1+\frac{(e-2) \gamma^{2} B}{C(1-\gamma)} \sum_{c \in S_{t}-S_{t}(0)} \hat{g}_{t}(c) .
\end{align*}
$$

To demonstrate the time-varying property in our bandit problem, we assume the reward distributions change at arbitrary time instants, but the total number of change points is no more than $V$ times. We split the sequence of total decisions $T$ into $V$ segments such that within each segment we have the time-invariant reward function.
We can write $V$ segments as $\left[T_{1}, \ldots T_{2}\right),\left[T_{2}, \ldots, T_{3}\right),\left[T_{V}, \ldots, T_{V+1}\right)$ where $T_{v}$ indicates the starting index of the $v$-th segment, $T_{v+1}$ is the ending index of the $v$-th segment which is also the starting index of the $v+1$-th segment - using the same notation in EXP3.S [3]. Similarly, we can write the optimal sequence $\underbrace{\left[c_{T_{1}}^{*}, \ldots c_{T_{2}}^{*}\right)}_{=c_{T_{1}}^{*}}, \underbrace{\left[c_{T_{2}}^{*}, \ldots, c_{T_{3}}^{*}\right)}_{=c_{T_{2}}^{*}}, \underbrace{\left[c_{T_{V}}^{*}, \ldots, c_{T_{V+1}}^{*}\right]}_{=c_{T_{V}}^{*}}$ where the reward function does not change in each segment, thus takes the same optimal choice $c_{T_{v}}^{*}$.
We consider an arbitrary segment $v$ and denote the length $\Delta_{v}=T_{v+1}-T_{v}$. Furthermore, let define the cumulative gain (or reward) achieved by using TV.EXP3.M strategy within a segment $v$ that the indices are ranging from $T_{v}$ to $T_{v+1}-1$

$$
G_{T V . E X P 3 . M}(v)=\sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}} g_{t}\left(c_{t}=c\right) .
$$

Taking $\ln$ of Eqn. (97, we get

$$
\begin{array}{rlr}
\ln \frac{W_{t+1}}{W_{t}} & \leq \ln \left(e \alpha+1+\frac{\gamma / C}{1-\gamma} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c)+\frac{(e-2)\left(\frac{\gamma}{C}\right)^{2} B}{1-\gamma} \sum_{c=1}^{C} \hat{g}_{t}(c)\right) \\
& \leq e \alpha+\frac{\gamma / C}{1-\gamma} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c)+\frac{(e-2)\left(\frac{\gamma}{C}\right)^{2} B}{1-\gamma} \sum_{c=1}^{C} \hat{g}_{t}(c) \quad \text { by } 1+a \leq e^{a} .
\end{array}
$$

Summing over all indices $t=T_{v}, \ldots, T_{v+1}-1$ within a segment $v$-th:

$$
\begin{aligned}
\ln W_{T_{v+1}}-\ln W_{T_{v}} & \leq \Delta_{v} e \alpha+\frac{\gamma / C}{1-\gamma} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c)+\frac{(e-2)\left(\frac{\gamma}{C}\right)^{2} B}{1-\gamma} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} \hat{g}_{t}(c) \\
& \leq \Delta_{v} e \alpha+\frac{\gamma / C}{1-\gamma} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c)+\frac{(e-2)\left(\frac{\gamma}{C}\right)^{2} B}{1-\gamma} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} \hat{g}_{t}(c) .
\end{aligned}
$$

Let $j=c_{T_{v}}^{*}=\ldots=c_{T_{v+1}-1}^{*}$ be the optimal choice in the sequence $v$, we have $\omega_{j}^{\left(T_{v+1}\right)}=$ $\omega_{j}^{\left(T_{v+1}-1\right)} \times \exp \left(\frac{B \gamma}{C} \hat{g}_{t}(j)\right)+\frac{e \alpha}{C} W_{T_{v+1}-1}, \forall j \notin S_{T_{v+1}-1}(0)$ and $\omega_{j}^{\left(T_{v+1}\right)}=\omega_{j}^{\left(T_{v+1}-1\right)}, \forall j \in$ $S_{T_{v+1}-1}(0)$

$$
\begin{aligned}
\omega_{j}^{\left(T_{v+1}\right)} & =\omega_{j}^{\left(T_{v+1}-1\right)} \times \exp \left(\frac{B \gamma}{C} \hat{g}_{t}(j)\right)+\frac{e \alpha}{C} W_{T_{v+1}-1} \\
& \geq \omega_{j}^{\left(T_{v+1}-1\right)} \times \exp \left(\frac{B \gamma}{C} \hat{g}_{t}(j)\right) \\
& \geq \omega_{j}^{\left(T_{v+1}-2\right)} \times \exp \left(\frac{B \gamma}{C} \hat{g}_{t}(j)\right) \exp \left(\frac{B \gamma}{C} \hat{g}_{\left(T_{v+1}-2\right)}(j)\right) \\
& \geq \omega_{j}^{\left(T_{v}\right)} \times \exp \left(\frac{B \gamma}{C} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j)\right) \\
& \geq \frac{e \alpha}{C} W_{T_{v}} \times \exp \left(\frac{B \gamma}{C} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j)\right) \\
& \geq \frac{\alpha}{C} W_{T_{v}} \times \exp \left(\frac{B \gamma}{C} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j)\right)
\end{aligned}
$$

We now consider the lower bound by taking the optimal (batch) set $A_{v}^{*} \subset[C]$ for $B$ elements with the maximum total of reward within the segment $v: \sum_{c \in A_{v}^{*}} \sum_{t=1}^{T} g_{t}(c)$. Since we have the fact that $\sum_{j \in A_{v}^{*}} \omega_{j}^{\left(T_{v+1}\right)} \geq B\left(\prod_{j \in A_{v}^{*}} \omega_{j}^{\left(T_{v+1}\right)}\right)^{1 / B}$ by Cauchy Swatchz inequality, we continue the above formula
$\ln \left(W_{T_{v+1}-1}\right)-\ln W_{T_{v}} \geq \ln B+\frac{1}{B} \sum_{j \in A_{v}^{*}} \ln \omega_{j}^{\left(T_{v+1}\right)}-\ln W_{T_{v}}$

$$
\begin{aligned}
& =\ln B+\frac{1}{B} \sum_{j \in A_{v}^{*}} \ln \frac{\alpha}{C} W_{T_{v}}+\frac{1}{B} \sum_{j \in A_{v}^{*}} \frac{B \gamma}{C} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j)-\ln W_{T_{v}} \\
& =\ln B+\frac{1}{B} \sum_{j \in A_{v}^{*}} \ln \frac{\alpha}{C}+\sum_{j \in A_{v}^{*}} \frac{\gamma}{C} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j) \\
& =\ln \frac{B \alpha}{C}+\sum_{j \in A_{v}^{*}} \frac{\gamma}{C} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j) .
\end{aligned}
$$

Combining both the lower bound and upper bound, we get

$$
\begin{aligned}
& \ln \frac{B \alpha}{C}+\sum_{j \in A_{v}^{*}} \frac{\gamma}{C} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j) \leq \Delta_{v} e \alpha+\frac{\gamma / C}{1-\gamma} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c) \\
& +\frac{(e-2)\left(\frac{\gamma}{C}\right)^{2} B}{1-\gamma} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} \hat{g}_{t}(c) .
\end{aligned}
$$

Summing over all segments $v=1 \ldots . V$

$$
\begin{align*}
V \ln \frac{B \alpha}{C}+\sum_{v=1}^{V} \sum_{j \in A_{v}^{*}} \frac{\gamma}{C} \sum_{t=T_{v} \mid t: j \neq S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j) \leq & T e \alpha+\frac{\gamma / C}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c) \\
& +\frac{(e-2)\left(\frac{\gamma}{C}\right)^{2} B}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} \hat{g}_{t}(c) \\
\frac{V C}{\gamma} \ln \frac{B \alpha}{C}+\sum_{v=1}^{V} \sum_{j \in A_{v}^{*}} \sum_{t=T_{v} \mid t: j \notin S_{t}(0)}^{T_{v+1-1}} \hat{g}_{t}(j) \leq & \frac{C}{\gamma} T e \alpha+\frac{1}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c) \\
& +\frac{(e-2)\left(\frac{\gamma B}{C}\right)}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} \hat{g}_{t}(c) \\
\sum_{v=1}^{V} \sum_{t=T_{v} \mid j \in S_{t}(0)}^{T_{v+1}-1} \sum_{j \in A_{v}^{*}} g_{t}(j)+\sum_{v=1}^{V} \sum_{j \in A_{v}^{*}} \sum_{t=T_{v} \mid j \neq S_{t}(0)}^{T_{v+1}-1} \hat{g}_{t}(j) \leq & \frac{C}{\gamma} T e \alpha+\frac{1}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}(0)} g_{t}(c) \\
& +\frac{1}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}^{-1}} \sum_{c \in S_{t}-S_{t}(0)} g_{t}(c) \\
& +\frac{(e-2)\left(\frac{\gamma B}{C}\right)}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} \hat{g}_{t}(c) \\
& -\frac{V C}{\gamma} \ln \frac{B \alpha}{C} . \tag{10}
\end{align*}
$$

Let us denote the optimal gain over all iterations and all parallel agents $G_{T B}^{*}=$ $\sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{j \in A_{v}^{*}} g_{t}(j)$. We also denote the cumulative gain achieved by using TV.EXP3.M algorithms $G_{\mathrm{TV.EXP} 3 . \mathrm{M}}=\sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}} g_{t}(c)$. Then, we continue Eqn. 10 as

$$
\begin{equation*}
G_{T B}^{*} \leq \frac{C}{\gamma} T e \alpha+\frac{1}{1-\gamma} G_{\mathrm{TV.EXP} 3 . \mathrm{M}}+\frac{(e-2)\left(\frac{\gamma B}{C}\right)}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} \hat{g}_{t}(c)-\frac{V C}{\gamma} \ln \frac{B \alpha}{C} \tag{11}
\end{equation*}
$$

where the number of element in a batch $\left|S_{t}\right|=\left|A_{j}^{*}\right|=B$. Let take expectation both sides of Eqn. (11), we have $\mathbb{E}\left[\hat{g}_{t}(c) \mid S_{1}, S_{2}, \ldots, S_{i-1}\right]=g_{t}(c)$ from the fact that DepRound [21] selects action $c$ with probability $p_{c}(t)$, we obtain

$$
G_{T B}^{*} \leq \frac{1}{(1-\gamma)} \mathbb{E}\left[G_{\mathrm{TV} . \mathrm{ExP} 3 . \mathrm{M}}\right]+\frac{(e-2)\left(\frac{\gamma B}{C}\right)}{1-\gamma} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} g_{t}(c)+\frac{C}{\gamma} T e \alpha-\frac{V C}{\gamma} \ln \frac{B \alpha}{C}
$$

We finally have

$$
\begin{align*}
(1-\gamma) G_{T B}^{*} & \leq \mathbb{E}\left[G_{\mathrm{TV.ExP} 3 . \mathrm{M}}\right]+(e-2)\left(\frac{\gamma B}{C}\right) \frac{C}{B} G_{T B}^{*}+\frac{(1-\gamma)}{\gamma} C T e \alpha-\frac{(1-\gamma) V C}{\gamma} \ln \frac{B \alpha}{C}  \tag{12}\\
G_{T B}^{*}-\mathbb{E}\left[G_{\mathrm{TV.ExP} 3 . \mathrm{M}}\right] & \leq(e-1) \gamma G_{T B}^{*}+\frac{(1-\gamma)}{\gamma} C T e \alpha+\frac{(1-\gamma) V C}{\gamma} \ln \frac{C}{B \alpha} \\
& \leq(e-1) \gamma T B+\frac{(1-\gamma)}{\gamma} C T e \alpha+\frac{V C}{\gamma} \ln \frac{C}{B \alpha}-V C \ln \frac{C}{B \alpha} \quad \text { by } G_{T B}^{*} \leq T B \\
& \leq(e-1) \gamma T B+\frac{1}{\gamma} C T e \alpha+\frac{V C}{\gamma} \ln \frac{C}{B \alpha} . \tag{13}
\end{align*}
$$

In Eqn. 12, we use the fact that $\sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c=1}^{C} g_{t}(c) \leq \frac{C}{B} G_{T B}^{*}=$ $\frac{C}{B} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}^{-1}} \sum_{c \in A^{*}} g_{t}(c)$.

```
Algorithm 5 TV.EXP3.M algorithm
    Input: \(\gamma=\sqrt{\frac{C \ln (C / B)}{(e-1) B T}}, \alpha=\frac{1}{T}, C\) \#categorical choice, \(T\) \#max iteration, \(B\) \#multiple play
    Init \(\omega_{c}=1, \forall c=1 \ldots C\) and denote \(\eta=\left(\frac{1}{B}-\frac{\gamma}{C}\right) \frac{1}{1-\gamma}\)
    for \(t=1\) to \(T\) do
        if \(\arg \max _{c \in[C]} \omega_{c} \geq \eta \sum_{c=1}^{C} \omega(c)\) then
            \(\nu\) s.t. \(\frac{\nu}{\eta}=\sum_{\omega_{t}(c) \geq \nu} \nu+\sum_{\omega_{t}(c)<\omega_{t}(c)} \omega_{t}(c)\)
            Set \(S_{0}=\left(c: \omega_{t}(c) \geq \nu\right)\) and \(\omega_{t}\left(S_{0}\right)=\nu\)
        else
            Set \(S_{0}=\emptyset\)
        end if
        Compute \(p_{t}^{c}=B\left((1-\gamma) \frac{\omega_{c}}{\sum_{c=1}^{C} \omega_{c}}+\frac{\gamma}{C}\right), \forall c\)
            A batch of \(B\) categorical choices \(S_{t}=\left[c_{t}^{1}, c_{t}^{2}, \ldots, c_{t}^{B}\right]=\operatorname{DepRound}\left(B,\left[p_{t}^{1} p_{t}^{2} \ldots p_{t}^{C}\right]\right)\)
            Observe the reward \(g_{t}(c)=f\left(c_{t}=c\right)\) for \(c \in S_{t}\)
            \(\hat{g}_{t}(c)=\frac{g_{t}(c)}{p_{t}^{c}}, \forall c \in S_{t}\) and \(\hat{g}_{t}(c)=0\) otherwise
            \(\forall c \notin S_{0}\) : update \(\omega_{c}=\omega_{c} \times \exp \left(B \gamma \hat{g}_{t}(c) / C\right)+\frac{e \alpha}{C} \sum_{i=1}^{C} \omega_{c}\),
            \(\forall c=S_{0}\) : update \(\omega_{c}=\omega_{c}+\frac{e \alpha}{C} \sum_{i=1}^{C} \omega_{c}\)
    end for
```

Set $\gamma=\min \left\{1, \sqrt{\frac{C \ln (C / B)}{(e-1) B T}}\right\}$ and $\alpha=\frac{1}{T}$. Then, we have the cumulative regret over all iterations and considering the best element in a batch

$$
\begin{align*}
& \mathbb{E}\left[R_{T B}\right]=\mathbb{E}[\sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \overbrace{\forall c \leq C}^{\text {optimal arms }} g_{t}(c)]-\mathbb{E}[\sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \overbrace{\max _{\forall c \in S_{t}} g_{t}(c)}^{\text {best arm in a batch } S_{t}}]  \tag{14}\\
& \leq \frac{1}{B} G_{T B}^{*}-\frac{1}{B} \mathbb{E}\left[G_{\text {TV.EXP } 3 . \mathrm{M}}\right]  \tag{15}\\
& \leq \frac{1}{B} \sqrt{\frac{C(e-1) \ln (C / B)}{B T}} T+\frac{1}{B} \frac{\sqrt{C(e-1) B T}}{\sqrt{\ln (C / B)}} e+\frac{1}{B} \frac{V C}{\sqrt{\frac{C \ln (C / B)}{(e-1) B T}}} \ln \frac{C T}{B}  \tag{16}\\
& \leq[1+e+V] \sqrt{(e-1) \frac{C T}{B} \ln \frac{C T}{B}} \tag{17}
\end{align*}
$$

where we obtain Eqn. (15) because the best gain should be greater than the average gain of a batch $\sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \max _{\forall c \in S_{t}} g_{t}(c) \geq \frac{1}{B} \sum_{v=1}^{V} \sum_{t=T_{v}}^{T_{v+1}-1} \sum_{c \in S_{t}} g_{t}(c)$.
Given the number of changing points in our reward function is bounded $V \ll \sqrt{T}$ and the number of category $C$ is a constant, our regret bound achieves sublinear regret rate with the number of iterations $T$, i.e., $\lim _{T \rightarrow \infty} \frac{\mathbb{E}\left[R_{T B}\right]}{T}=0$.

In the above derivation, $T$ refers to the number of bandit updates which is equivalent to the number of $t_{\text {ready }}$ in PB2 setting [52].

## D. 2 Theoretical Derivations for PB2-Mult

We adapt Lemma 1 in [52] to handle categorical variables. We first restate some notations used in [52]. Let $F_{t}^{c_{t}}\left(x_{t}\right)$ be an objective function under a given set of continuous hyperparameters $x_{t}$ and categorical variable $c_{t}$ at timestep $t$. An example of $F_{t}^{c_{t}}\left(x_{t}\right)$ could be the reward for a deep RL agent. When training for a total of $T$ steps, our goal is to maximize the final performance $F_{T}^{c_{T}}\left(x_{T}\right)$. We formulate this problem as optimizing the time-varying black-box reward function $f_{t}$, over $\mathcal{D}$.

Lemma 4. Maximizing the final performance $F_{T}$ of a model with respect to a given continuous hyperparameter schedule $\left\{x_{t}\right\}_{t=1}^{T}$ and a categorical hyperparameter schedule $\left\{c_{t}\right\}_{t=1}^{T}$ is equivalent to maximizing the time-varying black-box function $f_{t}\left(x_{t}\right)$ and minimizing the corresponding cumulative regret $r_{t}\left(x_{t}\right)$,

$$
\begin{equation*}
\max F_{T}^{c_{T}}\left(x_{T}\right)=\max \sum_{t=1}^{T} f_{c_{t}}\left(x_{t}\right)=\min \sum_{t=1}^{T} r_{t}\left(x_{t}\right) \tag{18}
\end{equation*}
$$

Proof. At each $t_{\text {ready }}$, we select a categorical variable $c_{t} \in\{1, \ldots, C\}$ and a continuous variable $x_{t} \in \mathcal{D}^{c_{t}}$. We would emphasize that the continuous variables are conditioned on the choice of $c_{t}$. We observe and record noisy observations, $y_{t}=f_{c_{t}}\left(x_{t}\right)+\epsilon_{t}$, where $\epsilon_{t} \sim \mathcal{N}\left(0, \sigma^{2} \mathbf{I}\right)$ for some fixed $\sigma^{2}$. The function $f_{t}$ represents the change in $F_{t}$ after training for $t_{\text {ready }}$ steps, i.e. $f_{c_{t}}\left(x_{t}\right)=F_{t}^{c_{t}}\left(x_{t}\right)-F_{t-t_{\text {ready }}}^{c_{t-t_{\text {ready }}}}\left(x_{t-t_{\text {ready }}}\right)$. We define the best choice at each timestep as $x_{t}^{*}=$ $\arg \max _{x_{t} \in \mathcal{D}^{c} t} f_{c_{t}}\left(x_{t}\right)$, and $c_{t}^{*}=\arg \max _{c_{t}, x_{t}} f_{c_{t}}\left(x_{t}\right)$. The intermediate regret of each decision is defined as $r_{t}=f_{c_{t}^{*}}\left(x_{t}^{*}\right)-f_{c_{t}}\left(x_{t}\right)$ where $f_{c_{t}^{*}}\left(x_{t}^{*}\right)$ is an unknown constant.
We have a reward at the starting iteration $F_{1}^{c_{1}}\left(x_{1}\right)$ as a constant that allows us to write the objective function as:

$$
\begin{equation*}
F_{T}^{c_{T}}\left(x_{T}\right)-F_{1}^{c_{1}}\left(x_{1}\right)=\underbrace{F_{T}^{c_{T}}\left(x_{T}\right)-F_{T-1}^{c_{T-1}}\left(x_{T-1}\right)}_{f_{c_{T-1}}\left(x_{T-1}\right)}+\ldots+\underbrace{F_{3}^{c_{3}}\left(x_{3}\right)-F_{2}^{c_{2}}\left(x_{2}\right)}_{f_{c_{2}}\left(x_{2}\right)}+\underbrace{F_{2}^{c_{2}}\left(x_{2}\right)-F_{1}^{c_{1}}\left(x_{1}\right)}_{f_{c_{1}}\left(x_{1}\right)} \tag{19}
\end{equation*}
$$

Therefore, maximizing the left of Eqn. (19) is equivalent to minimizing the cummulative regret as follows:

$$
\max \left[F_{T}^{c_{T}}\left(x_{T}\right)-F_{1}^{c_{1}}\left(x_{1}\right)\right]=\max \sum_{t=1}^{T} F_{t}^{c_{t}}\left(x_{t}\right)-F_{t-1}^{c_{t-1}}\left(x_{t-1}\right)=\max \sum_{t=1}^{T-1} f_{c_{t}}\left(x_{t}\right)=\min \sum_{t=1}^{T-1} r_{t}\left(x_{t}\right)
$$

where we define $f_{c_{t}}\left(x_{t}\right)=F_{t}^{c_{t}}\left(x_{t}\right)-F_{t-1}^{c_{t-1}}\left(x_{t-1}\right)$, the regret $r_{t}=f_{c_{t}^{*}}\left(x_{t}^{*}\right)-f_{c_{t}}\left(x_{t}\right)$.
We then restate Theorem 2 in [52] which will be used in deriving the convergence guarantee for PB2-Mult.
Theorem 5. (Theorem 2 in $[52])$ Let the domain $\mathcal{D} \subset[0, r]^{d}$ be compact and convex where $d$ is the dimension and suppose that the kernel is such that $f_{t} \sim G P(0, k)$ is almost surely continuously differentiable and satisfies Lipschitz assumptions $\forall L_{t} \geq 0, t \leq \mathcal{T}, \forall j \leq d, p\left(\sup \left|\frac{\partial f_{t}(\mathbf{x})}{\partial \mathbf{x}^{(j)}}\right| \geq L_{t}\right) \leq$ $a e^{-\left(L_{t} / b\right)^{2}}$ for some $a, b$. Pick $\delta \in(0,1)$, set $\beta_{T}=2 \log \frac{\pi^{2} T^{2}}{2 \delta}+2 d \log r d b T^{2} \sqrt{\log \frac{d a \pi^{2} T^{2}}{2 \delta}}$ and define $C_{1}=32 / \log \left(1+\sigma_{f}^{2}\right)$, the PB2 algorithm satisfies the following regret bound after $T$ time steps over $B$ parallel agents with probability at least $1-\delta$ :

$$
R_{T B}=\sum_{t=1}^{T} f_{t}\left(\mathbf{x}_{t}^{*}\right)-\max _{b \leq B} f_{t}\left(\mathbf{x}_{t, b}\right) \leq \sqrt{C_{1} T \beta_{T}\left(\frac{T}{\tilde{N} B}+1\right)\left(\gamma_{\tilde{N} B}+[\tilde{N} B]^{3} \omega\right)}+2
$$

the bound holds for any block length $\tilde{N} \in\{1, \ldots, T\}$ and $B \ll T$.
Next, we are going to derive the main theorem of the paper.
Theorem 6. (Theorem 2 in the main paper) Let the domain for continuous variables $\mathcal{D} \subset[0, r]^{d}$ be compact and convex where $d$ is the dimension and suppose that the kernel is such that $f_{t} \sim G P(0, k)$ is almost surely continuously differentiable and satisfies Lipschitz assumptions $\forall L_{t} \geq 0, t \leq \mathcal{T}, \forall j \leq$ $d, p\left(\sup \left|\frac{\partial f_{t}(\mathbf{x})}{\partial \mathbf{x}^{(j)}}\right| \geq L_{t}\right) \leq a e^{-\left(L_{t} / b\right)^{2}}$ for some $a, b$.
Assume the reward distributions to change at arbitrary time instants, but the total number of change points is no more than $V$. Set $\alpha=\frac{1}{T}, \gamma=\min \left\{1, \sqrt{\frac{C \ln (C / B)}{(e-1) B T}}\right\}, \beta_{T}=2 \log \frac{\pi^{2} T^{2}}{2 \delta}+$ $2 d \log r d b T^{2} \sqrt{\log \frac{d a \pi^{2} T^{2}}{2 \delta}}$, pick $\delta \in(0,1)$ and define $C_{1}=32 / \log \left(1+\sigma_{f}^{2}\right)$, the PB2-Mult algorithm satisfies the following regret bound after $T$ time steps over $B$ parallel agents with probability
at least $1-\delta$ :

$$
\mathbb{E}\left[R_{T B}\right] \leq[1+e+V] \sqrt{(e-1) \frac{C T}{B} \ln \frac{C T}{B}}+\sqrt{C_{1} T_{c_{t}^{*}} \beta_{T_{c_{t}^{*}}}\left(\frac{T_{c_{t}^{*}}}{\tilde{N} B}+1\right)\left(\gamma_{\tilde{N} B}+[\tilde{N} B]^{3} \omega\right)}+2 .
$$

The bound holds for any $\tilde{N} \in\left\{1, \ldots, T_{c_{t}^{*}}\right\}$ and $V \ll \sqrt{T}$.
Proof. We expand the cumulative regret and optimize it using time-varying GP bandit optimization [7, 52]

$$
\begin{aligned}
R_{T B} & =\sum_{t=1}^{T} f^{*}-\max _{b=1 \ldots B} \sum_{t=1}^{T} f_{c_{t, b}}\left(x_{t, b}\right) \\
& =\max _{b=1, \ldots, B} \sum_{t=1}^{T} f_{c_{t, b}^{*}}\left(x_{t, b}\right)-\max _{b=1, \ldots, B} \sum_{t=1}^{T} f_{c_{t, b}}\left(x_{t, b}\right)+\sum_{t=1}^{T} f^{*}-\max _{b=1, \ldots, B} \sum_{t=1}^{T} f_{c_{t, b}^{*}}\left(x_{t, b}\right)
\end{aligned}
$$

where $b \in\{1, \ldots, B\}$ is an agent's index being trained in parallel, $f_{c_{t}^{*}}\left(x_{t}\right)=\max _{\forall c_{t} \in\{1, \ldots, C\}} f_{c_{t}}\left(x_{t}\right)$ and $c_{t}^{*}=\arg \max _{\forall c \in\{1, \ldots C\}} f_{c}\left(x_{t}\right)$ is the optimal categorical choice at iteration $t$.
We bound the two terms separately as follows. The first term is bounded by Theorem 11. We assume the process of generating the arm $c_{t}$ 's reward $f_{c_{t}}():.=f_{c_{t}}\left(x_{t}\right)$ is by the "adversary" that TV.EXP3.M does not have the direct control on the selection of $x_{t}$. Particularly, $x_{t}$ will be chosen by TV-GP-BUCB (as part of PB2) in Eqn. (3). We take the expectation of the first term to have

$$
\begin{aligned}
\mathbb{E}[\max _{b=1 \ldots B} \sum_{t=1}^{T} \underbrace{f_{c_{t, b}^{*}}(.)}_{\text {pull optimal arm }}]-\mathbb{E}[\max _{b=1 \ldots B} \sum_{t=1}^{T} \underbrace{f_{c_{t, b}}(.)}_{\text {pull arm } c_{t}}] & =\mathbb{E}[\sum_{t=1}^{T} \underbrace{f_{c_{t}^{*}}(.)}_{\text {pull optimal arm }}]-\mathbb{E}[\max _{b=1 \ldots B} \sum_{t=1}^{T} \underbrace{f_{c_{t, b}}(.)}_{\text {pull arm } c_{t}}] \\
& =\mathbb{E}\left[\tilde{R}_{T B}\right] \\
& \leq[1+e+V] \sqrt{(e-1) \frac{C T}{B} \ln \frac{C T}{B}}
\end{aligned}
$$

where $R_{T B}$ is cumulative regret of the TV.EXP3.M, defined in Theorem 1
Assuming the best arm (the best categorical choice) $c_{t}^{*}$ can be identified by TV.EXP3.M in Theorem 1. The second term is the regret bound presented in Theorem 2 of the PB2 [52]:

$$
\begin{aligned}
\sum_{t=1}^{T} f_{t}^{*}-\max _{b=1 \ldots B} \sum_{t=1}^{T} f_{c_{t, b}^{*}}\left(x_{t, b}\right)= & \sum_{t=1}^{T} f_{t}^{*}-\max _{b=1 \ldots B} \sum_{t=1}^{T} f_{c_{t}^{*}}\left(x_{t, b}\right) \\
& \leq \underbrace{\sqrt{C_{1} T_{c_{t}^{*}} \beta_{T_{c_{t}^{*}}}\left(\frac{T_{c_{t}^{*}}}{\tilde{N} B}+1\right)\left(\gamma_{\tilde{N} B}+[\tilde{N} B]^{3} \omega\right)+2}}_{\mathcal{O}(P B 2)}
\end{aligned}
$$

where $T_{c_{t}^{*}}$ denotes the number of times the optimal category $c_{t}^{*}$ (or optimal arm) is selected. We follow [7, [52] to denote a block length $\tilde{N} \in\left\{1, \ldots . T_{c_{t}^{*}}\right\}$ in which the function does not change significantly. In the above equation, $\omega \in[0,1]$ is a time-varying hyperparameter which is estimated by maximizing the GP log marginal likelihood, $B$ is a batch size or the number of population agents, $\gamma_{\tilde{N} B}$ is the maximum information gain [74] defined within a block $\tilde{N}$ over $B$ parallel agents.
Note that the theoretical result for PB2 comes with the additional smoothness assumption on the kernel $k$ that holds for some $(a, b)$ and $\forall L_{t} \geq 0$. The kernel satisfies for all dimensions $j=1, \ldots, d$

$$
\begin{equation*}
\forall L_{t} \geq 0, t \leq \mathcal{T}, p\left(\sup \left|\frac{\partial f_{t}(\mathbf{x})}{\partial \mathbf{x}^{(j)}}\right| \geq L_{t}\right) \leq a e^{-\left(L_{t} / b\right)^{2}} \tag{20}
\end{equation*}
$$

We combine the two terms to obtain the final regret bound of the PB2-Mult algorithm:

$$
\mathbb{E}\left[R_{T B}\right] \leq[1+e+V] \sqrt{(e-1) \frac{C T}{B} \ln \frac{C T}{B}}+\sqrt{C_{1} T_{c_{t}^{*}} \beta_{T_{c_{t}^{*}}}\left(\frac{T_{c_{t}^{*}}}{\tilde{N} B}+1\right)\left(\gamma_{\tilde{N} B}+[\tilde{N} B]^{3} \omega\right)}+2 .
$$

The final regret bound is a summation of two sub-linear terms. Therefore, the regret bound grows sublinearly with the number of iterations $T$, i.e., $\lim _{T \rightarrow \infty} \frac{\mathbb{E}\left[R_{T B}\right]}{T}=0$. This is under a common assumption [7, 52] that the time-varying function is correlated, i.e., we have $\omega \rightarrow 0$, thus $\tilde{N} \rightarrow T$. We also note that when $T \rightarrow \infty$, then $T_{c_{T}^{*}} \rightarrow \infty$ due to Theorem 1

## D. 3 Gradients

We optimize the GP hyperparameters by maximizing the log marginal likelihood [58]. We fit the GP hyperparameters by maximizing their posterior probability (MAP), $p\left(\sigma_{x}, \sigma_{t} \mid \mathbf{X}, \mathbf{t}, \mathbf{y}\right) \propto$ $p\left(\sigma_{x}, \sigma_{t}, \mathbf{X}, \mathbf{t}, \mathbf{y}\right)$, which, thanks to the Gaussian likelihood, is available in closed form

$$
\begin{equation*}
\mathcal{L}:=\ln p(\mathbf{y}, \mathbf{X}, \mathbf{t}, \theta)=\frac{1}{2} \mathbf{y}^{T}\left(\boldsymbol{K}+\sigma_{y}^{2} \mathbf{I}_{N}\right)^{-1} \mathbf{y}-\frac{1}{2} \ln \left|\boldsymbol{K}+\sigma_{y}^{2} \mathbf{I}_{N}\right|+\ln p_{\text {hyp }}(\theta)+\mathrm{const} \tag{21}
\end{equation*}
$$

where $\mathbf{I}_{N}$ is the identity matrix in dimension $N$ (the number of points in the training set), and $p_{\text {hyp }}(\theta)$ is the prior over hyperparameters. We optimize Eqn. (21) with a gradient-based optimizer, providing the analytical gradient to the algorithm.

Our time-varying CoCaBO kernel is defined as

$$
k_{z}\left(z, z^{\prime}\right)=(1-\lambda)\left(k_{x t}+k_{c t}\right)+\lambda k_{x t} k_{c t}
$$

where $k_{x t}=k_{\text {Continuous }}\left(x, x^{\prime}\right) \times k_{\text {time } 1}\left(t, t^{\prime}\right), k_{c t}=k_{\text {Categorical }}\left(c, c^{\prime}\right) \times k_{\text {time } 2}\left(t, t^{\prime}\right), k_{\text {Continuous }}\left(x, x^{\prime}\right)=$ $\sigma_{1} \times \exp \left(-\frac{\left\|x-x^{\prime}\right\|^{2}}{l}\right), k_{\text {Categorical }}\left(c, c^{\prime}\right)=\frac{\sigma_{2}}{C} \sum \mathbb{I}\left(c, c^{\prime}\right), k_{\text {timel }}\left(t, t^{\prime}\right)=\left(1-\epsilon_{1}\right)^{\frac{\left|t-t^{\prime}\right|}{2}}$ and $k_{\text {time } 2}\left(t, t^{\prime}\right)=\left(1-\epsilon_{2}\right)^{\frac{\left|t-t^{\prime}\right|}{2}}$.

Our hyperparameters include $\theta=\left\{\epsilon_{1}, \epsilon_{2}, l, \sigma_{1}, \sigma_{2}, \lambda\right\}$. We need to compute the gradients of $\frac{\partial \mathcal{L}}{\partial \epsilon_{1}}, \frac{\partial \mathcal{L}}{\partial \epsilon_{2}}, \frac{\partial \mathcal{L}}{\partial l}, \frac{\partial \mathcal{L}}{\partial \sigma_{1}}, \frac{\partial \mathcal{L}}{\partial \sigma_{2}}, \frac{\partial \mathcal{L}}{\partial \lambda}$ as follows:

- The gradient of $\frac{\partial \mathcal{L}}{\partial \epsilon_{1}}$

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \epsilon_{1}} & =\frac{\partial \mathcal{L}}{\partial k_{z}} \times \frac{\partial k_{z}}{\partial \epsilon_{1}} \\
\frac{\partial k_{z}}{\partial \epsilon_{1}} & =(1-\lambda) k_{\text {Continuous }}\left(x, x^{\prime}\right) \frac{\partial k_{\text {time1 }}}{\partial \epsilon_{1}}+\lambda k_{c t} k_{\text {Continuous }}\left(x, x^{\prime}\right) \frac{\partial k_{\text {time } 1}}{\partial \epsilon_{1}} \\
\frac{\partial k_{\text {time1 }}}{\partial \epsilon_{1}} & =-\frac{\left|t-t^{\prime}\right|}{2}\left(1-\epsilon_{1}\right)^{\frac{\left|t-t^{\prime}\right|}{2}-1}
\end{aligned}
$$

- The gradient of $\frac{\partial \mathcal{L}}{\partial \epsilon_{2}}$

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \epsilon_{2}} & =\frac{\partial \mathcal{L}}{\partial k_{z}} \times \frac{\partial k_{z}}{\partial \epsilon_{2}} \\
\frac{\partial k_{z}}{\partial \epsilon_{2}} & =(1-\lambda) k_{\text {Categorical }}\left(c, c^{\prime}\right) \frac{\partial k_{\text {time } 2}}{\partial \epsilon_{2}}+\lambda k_{x t} k_{\text {Categorical }}\left(c, c^{\prime}\right) \frac{\partial k_{\text {time } 2}}{\partial \epsilon_{2}} \\
\frac{\partial k_{\text {time } 2}}{\partial \epsilon_{2}} & =-\frac{\left|t-t^{\prime}\right|}{2}\left(1-\epsilon_{2}\right)^{\frac{\left|t-t^{\prime}\right|}{2}-1}
\end{aligned}
$$

- The gradient of $\frac{\partial \mathcal{L}}{\partial l}$

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial l} & =\frac{\partial \mathcal{L}}{\partial k_{z}} \times \frac{\partial k_{z}}{\partial l} \\
\frac{\partial k_{z}}{\partial l} & =(1-\lambda) k_{\text {time1 }}\left(t, t^{\prime}\right) \frac{\partial k_{\text {Continuous }}}{\partial l}+\lambda k_{c t} k_{\text {timel }}\left(t, t^{\prime}\right) \frac{\partial k_{\text {Continuous }}}{\partial l} \\
\frac{\partial k_{\text {Continuous }}}{\partial l} & =\frac{\left\|x-x^{\prime}\right\|^{2}}{l^{2}} k_{\text {Continuous }}\left(x, x^{\prime}\right)
\end{aligned}
$$

- The gradient of $\frac{\partial \mathcal{L}}{\partial \sigma_{1}}$

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \sigma_{1}} & =\frac{\partial \mathcal{L}}{\partial k_{z}} \times \frac{\partial k_{z}}{\partial \sigma_{1}} \\
\frac{\partial k_{z}}{\partial \sigma_{1}} & =(1-\lambda) k_{\text {timel }}\left(t, t^{\prime}\right) \frac{\partial k_{\text {Continuous }}}{\partial \sigma_{1}}+\lambda k_{c t} k_{\text {timel }}\left(t, t^{\prime}\right) \frac{\partial k_{\text {Continuous }}}{\partial \sigma_{1}} \\
\frac{\partial k_{\text {Continuous }}}{\partial \sigma_{1}} & =k_{\text {Continuous }}\left(x, x^{\prime}\right)
\end{aligned}
$$

- The gradient of $\frac{\partial \mathcal{L}}{\partial \sigma_{2}}$

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \sigma_{2}} & =\frac{\partial \mathcal{L}}{\partial k_{z}} \times \frac{\partial k_{z}}{\partial \sigma_{2}} \\
\frac{\partial k_{z}}{\partial \sigma_{2}} & =(1-\lambda) k_{\text {time2 }}\left(t, t^{\prime}\right) \frac{\partial k_{\text {Categorical }}}{\partial \sigma_{2}}+\lambda k_{x t} k_{\text {time2 }}\left(t, t^{\prime}\right) \frac{\partial k_{\text {Categorical }}}{\partial \sigma_{2}} \\
\frac{\partial k_{\text {Categorical }}}{\partial \sigma_{2}} & =k_{\text {Categorical }}\left(c, c^{\prime}\right)
\end{aligned}
$$

- The gradient of $\frac{\partial \mathcal{L}}{\partial \lambda}$

$$
\begin{aligned}
& \frac{\partial \mathcal{L}}{\partial \lambda}=\frac{\partial \mathcal{L}}{\partial k_{z}} \times \frac{\partial k_{z}}{\partial \lambda} \\
& \frac{\partial k_{z}}{\partial \lambda}=-\left(k_{c t}+k_{x t}\right)+k_{c t} k_{x t}
\end{aligned}
$$


[^0]:    ${ }^{4}$ see: https://github.com/rraileanu/auto-drac

