Fixed-Budget Change Point Identification in Piecewise Constant Bandits

Joseph Lazzaro Imperial College London

Abstract

We study the *piecewise constant bandit* problem where the expected reward is a piecewise constant function with one change point (discontinuity) across the action space [0, 1] and the learner's aim is to locate the change point. Under the assumption of a fixed exploration budget, we provide the first non-asymptotic analysis of policies designed to locate abrupt changes in the mean reward function under bandit feedback. We study the problem under a large and small budget regime, and for both settings establish lower bounds on the error probability and provide algorithms with near matching upper bounds. Interestingly, our results show a separation in the complexity of the two regimes. We then propose a regime adaptive algorithm which is near optimal for both small and large budgets simultaneously. We complement our theoretical analysis with experimental results in simulated environments to support our findings.

1 INTRODUCTION

In many settings, we are interested in sequentially learning to detect a change point/discontinuity in a piecewise constant function. For example, Park et al. (2021, 2023) study the development of materials with physical behaviours which abruptly change under different experimental conditions, such as temperature and pressure. Learning where these changes occur can help predict the quality of production techniques and improve efficiency. However, experiments are expensive and time consuming, so we would like to sequentially choose experimental parameters to learn where the Ciara Pike-Burke Imperial College London

abrupt changes occur as quickly as possible. The same problem also occurs when mapping out the edge of a cliff on the floor of the ocean (Hayashi et al., 2019). Exploring the whole ocean naively in a grid would be extremely expensive, so we would like to develop strategic ways of minimising the number of times and locations where we measure the depth of the ocean floor. Developing methods to minimize the number of samples needed to detect a change point could also help us develop computationally efficient subsampling methods in offline change point analysis, leading to yet more applications in different fields (Lu et al., 2020).

In this paper, we study this problem, which we refer to as the *Piecewise Constant Bandit Problem*. Here, the underlying reward function is a piecewise constant function on [0, 1] and the learner's aim is to sequentially select points to query in order to identify the unknown change point x^* as accurately as possible after a fixed number of samples. We assume that whenever we select a point $x_t \in [0,1]$, we receive a noisy observation of the unknown piecewise constant function at that point. In contrast to other bandit problems on a continuous action space, in the piecewise constant bandit problem, our goal is to detect where the change in mean occurs rather than identify the optimal arm. Moreover, the abrupt change in the mean reward function violates the smoothness conditions of existing bandit methods for infinite action spaces (e.g. Kleinberg et al., 2008; Srinivas et al., 2010; Bubeck et al., 2011) meaning that new techniques need to be developed. The fixed-budget assumption on the number of samples also necessitates the development of new methods. When the budget tends to infinity, some asymptotic methods have been proposed (Hall and Molchanov, 2003; Lan et al., 2009). However in most practical cases, we only have a finite budget so it is essential to develop a non-asymptotic understanding of the complexity of the problem and develop theoretically justified policies for the fixedbudget piecewise constant bandits problem.

We study the piecewise constant bandit problem in environments with exactly one change point across a one dimensional action space and sub-Gaussian noise, see

Proceedings of the 28th International Conference on Artificial Intelligence and Statistics (AISTATS) 2025, Mai Khao, Thailand. PMLR: Volume 258. Copyright 2025 by the author(s).



Figure 1: Example of a piecewise constant mean reward function, f, across action space [0, 1] with change point x^* , change in mean of Δ from μ_1 to μ_2 and 10 arbitrarily chosen noisy observations in red.

Figure 1 for an example. While this setting may appear restrictive, it turns out that significant innovation is required to develop optimal methods for this setting, and we hope that these ideas will inspire solutions to more complex problems. We focus on the non-asymptotic fixed-budget problem where we are given a fixed number of queries and our aim is to return an estimated change point that minimizes the error probability. Finding optimal solutions to the piecewise constant bandit problem is non-trivial. Indeed, it requires distributing a finite number of samples across an infinite action space, comparing them to detect a change in mean, and allocating sufficient samples near the unknown change in order to confidently determine its location. We make the following contributions: (i) We characterise the difficulty of the piecewise constant bandits problem for both large and small budgets by proving lower bounds that show a separation in difficulty of the two regimes. This is in contrast to most of the fixed-budget bandits literature, where sufficiently large budgets are explicitly or implicitly assumed (e.g. Locatelli et al., 2016; Carpentier and Locatelli, 2016; Cheshire et al., 2021). Our proof techniques are novel and lead to improved lower bounds in related problems such as Thresholding Bandits (see Section 4.2). (ii) We adapt two elimination algorithms, based on Sequential Halving and Binary Search (e.g. Karnin et al., 2013; Cheshire et al., 2021) to our setting and prove that these have near matching upper bounds for error probability in both regimes. (iii) We propose a regime adaptive method which is near optimal across both regimes simultaneously. (iv) We complement our theoretical results with experiments in simulated environments.

2 RELATED WORK

Pure Exploration in Finite Action Spaces In fixed-budget Best Arm Identification there is a fixed

budget of samples and the aim is to minimize the probability of failing to identify the (approximate) arm with the highest expected reward (Mannor et al., 2004; Carpentier and Locatelli, 2016; Audibert et al., 2010). This motivates the objective we consider in Section 3. Other relevant pure exploration problems are Noisy Binary Search (Karp and Kleinberg, 2007: Nowak, 2009; Gretta and Price, 2024), Binary Classification with Noise (Locatelli et al., 2016; Castro and Nowak, 2008), and the (Monotonic) Thresholding Bandits problem (Cheshire et al., 2021). In these problems different assumptions are made on the reward function f and noise distribution, however they all aim to locate where f crosses a known threshold θ . This is simpler than the piecewise constant bandit problem since all observations can be compared with the known threshold θ . In our case, we do not have a known threshold θ and so we have to compare observations with each other across an infinite action space to determine where the change in mean occurs, making the problem more challenging. Nonetheless, we are able to extend ideas from Cheshire et al. (2021) to develop nearly minimax optimal algorithms for piecewise constant bandits. In clustering with bandit feedback (Yang et al., 2022; Yavas et al., 2024; Thuot et al., 2024) arms are sorted in to a known number of clusters, each containing arms with the same expected reward. Unlike our setting, they do not assume any structure in the underlying reward function. Moreover, these clustering problems have not been considered in the fixed-budget setting we consider.

Infinite Action Spaces Bandit problems with continuous action spaces assume the reward function is linear (e.g. Abbasi-Yadkori et al., 2011), convex (e.g. Agarwal et al., 2011), Lipschitz (e.g. Bubeck et al., 2011) or smooth enough to be modelled by a Gaussian Process (GP) (e.g. Srinivas et al., 2010). In our setting, the mean reward function has an abrupt change, which violates the smoothness assumptions of these works.

Change Points and Non-Stationary Bandits Both offline and online change point detection for time series have been well studied in statistics (Aminikhanghahi and Cook, 2017). Online change point analysis has been used in non-stationary bandits. Here, the mean reward of the arms evolves over time with some abrupt changes (e.g. Garivier and Moulines, 2011; Deng et al., 2022; Hong et al., 2023). Conversely, in our setting, the reward changes abruptly over the action space, but is stationary across time.

Change Points Across the Action Space The active learning literature has studied sequential meth-

ods for learning entire piecewise continuous functions (e.g. Gramacy and Lee, 2008; Park et al., 2021, 2023). While these papers only provide empirical results, Castro et al. (2005) develop a two-stage algorithm with near minimax optimal expected squared ℓ_2 error (in a problem-independent sense and up to log terms). This measures the accuracy of estimating the entire reward function. In our work, however, we focus on methods for estimating the *locations* of the discontinuities in the reward function. Our PAC methods are minimax optimal in a problem-dependent sense.

Hall and Molchanov (2003) and Lan et al. (2009) consider fixed-budget multi-stage methods to sequentially estimate discontinuities in piecewise smooth functions with exactly one change point. Their first stage uses a portion of the budget to sample evenly across the action space to construct a confidence interval (CI) for the change point. In the following stages this process is repeated within the previous CI (backtracking to an earlier CI if they no longer detect a change). The theoretical guarantees provided only hold as the budget tends to infinity, which clashes with the motivation of fixed-budget problems with an extremely limited number of samples. In certain finite-budget problems, their methods perform suboptimally. In particular, if there are not many samples one side of the change point, the initial CI for the change point can be unreliable and the final estimate performs poorly (i.e. the algorithm is unstable for changes near the boundary - see Section 7). Moreover, their results are worst case and do not provide any insights into how performance changes depending on the problem instance. Thus we provide the first non-asymptotic problem dependent bounds on the performance of algorithms for piecewise constant bandits which apply for most realistic budgets. We also provide near matching problem-dependent lower bounds demonstrating that our methods are near optimal, and correctly adapt to the difficulty of the problem.

Hayashi et al. (2019) consider a class of active change point detection problems, where the aim is to locate changes in reward functions that are piecewise constant, piecewise linear, or contain other types of change points. They propose an anytime meta-algorithm, ACPD, which requires a statistical model for the type of change point and noise distribution to calculate "change scores" for different regions of the space. They run Bayesian Optimisation using these scores at each iteration. No theoretical guarantees for ACPD have been provided. Our simulations in Section 7 demonstrate that, while ACPD can perform well on 'easy' problems, there are settings where our theoretically grounded methods perform significantly better, while also being computationally cheaper than ACPD.

3 PROBLEM SETTING

We consider the piecewise constant bandit problem with a fixed budget T. Here, in each round $t = 1, \ldots, T$, we choose an action $x_t \in \mathcal{A} = [0, 1]$ and observe a reward $y_t \in \mathbb{R}$. We consider a set of environments Vwhere the mean reward function $f : [0, 1] \mapsto \mathbb{R}$ is a piecewise constant function with exactly one change point $x^* \in [0, 1)$. Then, the reward we observe from playing x_t in round t is

$$y_t = f(x_t) + \epsilon_t,$$

where $f(x) = \mu_1 \mathbb{I}\{x \le x^*\} + \mu_2 \mathbb{I}\{x > x^*\}$ is a piecewise constant function with $\mu_1 \ne \mu_2 \in \mathbb{R}$. We assume the random noise, ϵ_t , is i.i.d and σ^2 sub-Gaussian with mean-zero. Importantly, we assume that the values of $\mu_1 \ne \mu_2 \in \mathbb{R}, \sigma^2 \in [0, \infty)$, and $x^* \in [0, 1)$ are unknown to the learner. We subsequently denote the change in mean reward as $\Delta = |\mu_1 - \mu_2|$ and $V(\Delta, \sigma)$ as the set of environments with change in mean reward Δ, σ^2 sub-Gaussian noise, and any change point $x^* \in [0, 1)$.

Our goal is to estimate the change point, x^* , which separates the two reward distributions across [0, 1]. In particular, given some budget T, our objective is to generate an estimate after T rounds, \hat{x}_T , which is close to the true unknown change point with high probability. Finding a change point exactly is impossible in a continuous action space, so we define an *acceptable tolerance* $\eta > 0$ and aim to return an estimate within η of the true change point. Let Π be the set of all policies which return an estimate for the change point \hat{x}_T after budget T. We aim to find a policy $\pi \in \Pi$ such that in environment $v \in V$ with change point x_v^* ,

$$\mathbb{P}_{\pi,v}(|\hat{x}_T - x_v^*| < \eta) > 1 - \delta \tag{1}$$

with δ as small as possible. Here $\mathbb{P}_{\pi,v}$ is the measure induced by the interactions between the policy π and environment v (we henceforth drop the subscript whenever it is clear which policy and environment we are referring to). We assume the learner is given fixed values for T and η , and their goal is to develop a policy that returns an estimate \hat{x}_T which satisfies (1) with δ as small as possible. Equivalently, we want to minimise the probability that we fail to estimate the change point with sufficient accuracy after a given number of observations. Defining the failure event as $F_{T,v,\eta} = \{ |\hat{x}_T - x_v^*| \ge \eta \},$ we can state this objective as minimizing $\mathbb{P}_{\pi,v}(F_{T,v,\eta})$. Note that this objective is similar to those seen in other pure exploration problems such as PAC best arm identification (e.g. Mannor et al., 2004). This objective is also practically relevant as a practitioner will only care about a particular level of precision across an infinite space, hence it is natural to include a pre-specified tolerance η .

4 LARGE BUDGET

We first study the setting where the budget, T, is large enough to efficiently explore the space and accurately locate the change point, while still considering nonasymptotic methods. We consider an extension of the binary search with backtracking algorithm which was used for Monotonic Thresholding Bandits and Noisy Binary Search (Cheshire et al., 2021; Karp and Kleinberg, 2007). Informally, the idea is to split our budget up into phases and in each phase sample the leftmost, mid, and rightmost point in the action space repeatedly. If there is more evidence that a change in mean occurs on the left half than the right half, then we eliminate the right half of the space. Similarly, if there is more evidence that the change occurs on the right half, we eliminate the left half of the space. If, however, at any point we think the remaining region no longer contains the change point, we go back a step and undo our previous elimination (i.e. backtrack).

Naively applying the binary search with backtracking algorithm from Cheshire et al. (2021) to our setting would not work. In Thresholding Bandits, all observation could simply be compared with a known threshold θ . However, in our case we need to compare *unknown* reward distributions across the action space with each other in order to find a change point. In particular, if we want to correctly identify when a change point is *not* in some region $B \subset \mathcal{A}$ with high probability, then it is crucial to have data from *both* sides of the change point. Therefore, we additionally sample outside of region B, at the boundaries of the action space 0 and 1, in each phase. We call our algorithm **Sequential Halving** with Backtracking (SHB). SHB is detailed in Algorithm 1 by setting backtracking=True (a standalone version is in Appendix A).

In more detail, in SHB, we split our budget T into $J = \lceil 6 \log(1/2\eta) \rceil$ phases. In each phase $j \in \{1, ..., J\}$, we define the set of sampling points as $\mathcal{A}^j = \{0, a_1^j, a_2^j, a_3^j, 1\}$ and begin the phase by playing each action in \mathcal{A}^j $t_j = \lfloor \frac{T}{5J} \rfloor$ times. The set of sampling points in \mathcal{A}^j consist of the endpoints and midpoint of the noneliminated region, together with the extreme points 0,1 needed for backtracking. This is illustrated in Figure 2. In phase 1 we have not eliminated anything yet, so we initialise with $a_1^j = 0, a_2^j = 1/2, a_3^j = 1$. In every phase $j = 1, \ldots, J$, we compute the empirical mean of the observations from playing each of the five actions actions t_j times and denote them $\hat{\mu}_{0,t_j}, \hat{\mu}_{a_1^j,t_j}, \hat{\mu}_{a_2^j,t_j}, \hat{\mu}_{1,t_j}, \hat{\mu}_{1,t_j}$, respectively. These estimates help identify which half of the remaining action space $[a_1^j, a_3^j)$ contains the change point.

Firstly, we determine whether there is more evidence for x^* being in the left or right half of this remaining **Algorithm 1** Sequential halving (with backtracking) - SH(B)

1: Input: $\eta \in (0, 1/2)$, budget T, and 2: **Input:** backtracking = True or False 3: $\mathcal{A}^1 = \{0, a_1^1, a_2^1, a_3^1, 1\} \leftarrow \{0, 0, 1/2, 1, 1\}$ 4: if backtracking then $J \leftarrow [6\log(1/2\eta)], \quad t_i \leftarrow \lfloor \frac{T}{5T} \rfloor$ 5:6: **else** $J \leftarrow \lceil \log_2(1/2\eta) \rceil, \quad t_i \leftarrow \lfloor \frac{T}{2T} \rfloor$ 7: 8: end if 9: for phase j in 1, ..., J do Play actions a_1^j, a_2^j, a_3^j , each t_j times 10:if backtracking then 11: Play actions 0, 1, each t_i times 12:end if 13:if backtracking and $E_{P,j}$ from (4) holds then 14: $\mathcal{A}^{j+1} \leftarrow P(\bar{\mathcal{A}}^j)$ 15:▷ Backtrack else if $E_{R,i}$ from (2) holds then 16: $\mathcal{A}^{j+1} \leftarrow R(\mathcal{A}^j)$ \triangleright Zoom in to the right 17:else if $E_{L,j}$ from (3) holds then 18: $\mathcal{A}^{j+1} \leftarrow L(\mathcal{A}^j)$ \triangleright Zoom in to the left 19:20:end if 21: end for 22: **Return:** $\hat{x}_T = a_2^{J+1}$

space in phase j, i.e we determine whether x^* is in $[a_1^j, a_2^j)$ or $[a_2^j, a_3^j)$. In particular if the event

$$E_{R,j} = \left\{ |\hat{\mu}_{a_1^j, t_j} - \hat{\mu}_{a_2^j, t_j}| < |\hat{\mu}_{a_2^j, t_j} - \hat{\mu}_{a_3^j, t_j}| \right\}, \quad (2)$$

holds then we suspect that the change point is in the right half of the remaining action space, namely $x^* \in [a_2^j, a_3^j)$. This is because $E_{R,j}$ occurs when we observe a bigger change in (empirical) mean reward between actions a_2^j and a_3^j , than between a_1^j and a_2^j , suggesting that the change in distribution occurred between a_2^j and a_3^j . (It turns out that this intuition is theoretically justified from change point analysis, see Appendix G.) In such a case, we 'zoom in to the right' by eliminating the left half of the remaining action space and define our actions for the next phase as $\mathcal{A}^{j+1} = R(\mathcal{A}^j)$, for the operator

$$R(\{0, a_1^j, a_2^j, a_3^j, 1\}) = \{0, a_2^j, (a_2^j + a_3^j)/2, a_3^j, 1\},\$$

so $\mathcal{A}^{j+1} = \{0, a_1^{j+1}, a_2^{j+1}, a_3^{j+1}, 1\}$ for $a_1^{j+1} = a_2^j, a_2^{j+1} = (a_2^j + a_3^j)/2, a_3^{j+1} = a_3^j$. See Figure 2 for an illustration.

On the other hand, if the converse is true and event

$$E_{L,j} = \left\{ |\hat{\mu}_{a_1^j, t_j} - \hat{\mu}_{a_2^j, t_j}| \ge |\hat{\mu}_{a_2^j, t_j} - \hat{\mu}_{a_3^j, t_j}| \right\} \quad (3)$$

holds, this suggests that the change point is in the left half of the remaining action space, $x^* \in [a_1^j, a_2^j)$. In



Figure 2: Example illustration of action space $\mathcal{A} = [0,1]$ in (Top) phase j with sampling points $\mathcal{A}^j = \{0, a_1^j, a_2^j, a_3^j, 1\}$ and (Bottom) phase j + 1 with sampling points $\mathcal{A}^{j+1} = \{0, a_1^{j+1}, a_2^{j+1}, a_3^{j+1}, 1\}$, where the shaded regions have been eliminated. In this example, $E_{R,j}$ held and region $[a_1^j, a_2^j)$ was eliminated in phase j.

this case, we 'zoom in to the left' by eliminating the right half of the space, $[a_2^j, a_3^j)$, and update our actions for the next phase as $\mathcal{A}^{j+1} = L(\mathcal{A}^j)$ for the operator

$$L(\{0, a_1^j, a_2^j, a_3^j, 1\}) = \{0, a_1^j, (a_1^j + a_2^j)/2, a_2^j, 1\}.$$

It is, of course, possible that the change point is not within our remaining action space $[a_1^j, a_3^j)$ at all. Hence, before we even consider checking events $E_{L,j}$ and $E_{R,j}$, we consider the event $E_{P,j}$,

$$E_{P,j} = \left\{ Q_1 < \frac{3}{4} \max\left(Q_2, \, Q_3\right) \right\} \tag{4}$$

for
$$Q_{1} = \left| \frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} - \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2} \right|,$$
$$Q_{2} = \left| \frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} - \hat{\mu}_{1,t_{j}} \right|,$$
$$Q_{3} = \left| \hat{\mu}_{0,t_{j}} - \frac{\hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{3} \right|.$$

If $E_{P,j}$ holds then we suspect that $x^* \notin [a_1^j, a_3^j)$ so we believe the change point is elsewhere. Intuitively $E_{P,j}$ occurs when the change in empirical mean is smaller across the region $[a_1^j, a_3^j)$ than elsewhere, see Appendix **G** for more details. If $E_{P,j}$ holds then we 'zoom out' one step by un-doing the previous elimination and update our actions for the next phase as $\mathcal{A}^{j+1} = P(\mathcal{A}^j)$. The operator P returns the sampling points \mathcal{A}^i from which we previously zoomed in to get to \mathcal{A}^j . More formally, we define $P(\mathcal{A}^j)$ as an operator that outputs a set of actions from a previous phase, \mathcal{A}^i such that $i \leq j - 1$ and we zoomed in from \mathcal{A}_i to obtain our current set of actions \mathcal{A}^j . Namely $P(\mathcal{A}^j) = \mathcal{A}^i$ for an $i \in \{1, \ldots, J\}$ such that $\mathcal{A}^j = R(\mathcal{A}^i)$ or $\mathcal{A}^j = L(\mathcal{A}^i)$.

At the end of the J phases, we are left with the set $\mathcal{A}^{J+1} = \{0, a_1^{J+1}, a_2^{J+1}, a_3^{J+1}, 1\}$ representing the leftmost point, midpoint, and right-most point of the remaining action space that has not been eliminated. We estimate the change point, x^* , by taking the midpoint of this remaining region, namely $\hat{x}_T = a_2^{J+1}$.

4.1 Upper Bound

We upper bound the probability that the SHB algorithm fails to estimate the true change point x^* up to the acceptable tolerance η . Namely, we upper bound $\mathbb{P}(F_{T,v,\eta})$, with failure event $F_{T,v,\eta} = \{|\hat{x}_T - x_v^*| \ge \eta\}$.

Theorem 1. Let $\eta < 1/4$. Consider SHB in an environment $v \in V(\Delta, \sigma)$. Then, for $T > 60 \log(1/2\eta)$,

$$\mathbb{P}(F_{T,v,\eta}) \le \exp\left(-\frac{\Delta^2}{600\sigma^2}T + 13\log(1/2\eta)\right).$$

Proof. We sketch the proof here, leaving details for Appendix C. Note that, since SHB is an extension of (Cheshire et al., 2021), proofs for some technical lemmas are similar to existing work. In particular, we first note that we don't need to make the correct decision in eliminating/backtracking in every phase. As long as we make correct decisions in at least 3/4of the phases, then our final estimate for the change point, \hat{x}_T , will be within η of x^* . However, our elimination/backtracking criteria are more involved than (Cheshire et al., 2021) because we are not simply comparing the means with a known threshold, but comparing multiple means with each other to identify which region we believe the change point to be in. Hence we must develop new techniques to show that we make the correct decision in each phase with high probability, regardless of the position of x^* .

Note that $\log(1/2\eta)$ only appears additively in the exponent in Theorem 1. This log dependence on η has been attained in related settings such as Thresholding Bandits or Active Binary Classification (Cheshire et al., 2021; Castro and Nowak, 2008). However, unlike these previous works, we show that this log dependence on η is unimprovable for larger budgets in Section 4.2, and we will show in Section 5 that a different dependence on η appears for small budgets.

4.2 Lower Bound

We provide a lower bound to show that the error probability from our SHB algorithm is minimax optimal up to constants, for **large budgets** $T \ge T_1 := \frac{\sigma^2}{\Delta^2}(1.59 \log(\lfloor \frac{1}{2\eta} \rfloor) - 2 \log(2))$. Note the upper bound in Theorem 1 holds for any sub-Gaussian reward distributions, whereas in this section we assume that the rewards are Gaussian.

Theorem 2. Let $\overline{V} \subset V(\Delta, \sigma)$ be the set of environments with change in mean Δ and Gaussian random noise with variance σ^2 . Then, for $T \geq T_1$,

$$\inf_{\pi \in \Pi} \sup_{v \in \tilde{V}} \mathbb{P}_{v,\pi}(F_{T,v,\eta}) \geq \frac{1}{8} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + \frac{1}{2}\log\left(\frac{1}{2}\left(\left\lfloor\frac{1}{2\eta}\right\rfloor - 1\right)\right)\right).$$

Proof. We sketch the proof here, see Appendix D for full proof. Suppose we have a policy π which has reasonably good failure probability regardless of the environment, namely $\forall v \in \overline{V}$, $\mathbb{P}_{v,\pi}(F_{T,v,\eta}) < 1/2$. We would expect this policy to explore the action space sufficiently well. Then, using a sequence of changeof-measure arguments between three carefully chosen environments, we show such a policy π will always make a mistake with some probability (involving η) in at least one environment. We then extend this to show that for any policy $\pi \in \Pi$, when the budget is sufficiently large we must incur the stated failure probability. \Box

Comparing the upper bound in Theorem 1 and the minimax lower bound in Theorem 2, we see that the terms in the exponent of both the upper and lower bounds match up to constants. Hence for large budgets, $T > T_1$, SHB is minimax optimal up to constants. It is also interesting to note that in both the upper and lower bounds, η only has an additive effect in the exponent, and the η terms do not scale multiplicatively with T. This means that as η becomes very small (i.e. when we need our estimates' accuracy to be very high), this only has an additive effect in the exponent of the probability of failure, and does not affect the *rate* at which the failure probabilities decay with increased T. In the simpler Monotonic Thresholding Bandits Problem (Cheshire et al., 2021; Castro and Nowak, 2008), a similar additive dependence on η ($\eta \approx 1/2K$ in their setting) has appeared in upper bounds on the error probability. However, there is currently no lower bound for Thresholding Bandits to show that this dependence on η is unavoidable. As a consequence of Theorem 2, we can provide tighter lower bounds for Thresholding Bandits with large budgets which formalises the effect of η , and may be of independent interest.

5 SMALL BUDGET

For problems where the budget is small and gives little time for exploration, it is natural to consider omitting the exploratory backtracking actions from SHB to provide a more exploitative algorithm. We show that the resulting **Sequential Halving (SH)** algorithm (Algorithm 1 with backtracking=False) is sufficient to attain near-optimality for small budgets $T < T_1 := \frac{\sigma^2}{\Delta^2} (1.59 \log(\lfloor \frac{1}{2\eta} \rfloor) - 2 \log(2))$. The SH algorithm is written explicitly in Appendix A.

In SH, we split our budget up into $J = \lceil \log_2(1/2\eta) \rceil$ phases. Then in every phase, using only samples from the leftmost, mid, and rightmost points of the remaining action space, we eliminate half of this remaining action space. In particular, when event $E_{R,j}$ holds (defined the same as in (2)) we eliminate the left half of the action space. When $E_{L,j}$ holds (again, same as (3)) we eliminate the right half of the action space. By our choice of the number of phases, the width of the final region will be less than 2η . Hence, we estimate the midpoint of this region to be the change point.

In this SH algorithm we are more exploitative and avoid additionally sampling the actions 0, 1, as needed for backtracking in SHB, which gives better performance for smaller budgets. However, we will see this comes at the cost of worse error probabilities for larger budget problems. Note that SH requires $T \geq 3\lceil \log_2(1/2\eta) \rceil$ in order for there to be at least one sample for each action in every phase. This assumption is reasonable since even in the noiseless setting with known μ_1, μ_2 , the minimum number of samples required to guarantee $|\hat{x}_T - x^*| \leq \eta$ is $T \geq \lceil \log_2(1/2\eta) \rceil$ (Sikorski, 1982).

5.1 Upper Bound

The failure probability of the SH algorithm is bounded in Theorem 3, with proofs in Appendix F.

Theorem 3. Under the SH algorithm in an environment $v \in V(\Delta, \sigma)$, for $T \geq 3\lceil \log_2(1/2\eta) \rceil$, and $F_{T,v,\eta} = \{ |\hat{x}_T - x_v^*| \geq \eta \},$

$$\mathbb{P}(F_{T,v,\eta}) < 2 \left\lceil \log_2\left(\frac{1}{2\eta}\right) \right\rceil \exp\left(\frac{-T\Delta^2}{36\sigma^2 \log_2(1/2\eta)}\right).$$

It is important to note that in Theorem 3 the $\log_2(1/2\eta)$ term is in the denominator of the exponent, multiplying the leading term. This comes from the fact that, in order to achieve our objective of $|\hat{x}_T - x^*| \leq \eta$ with SH, we need to eliminate the correct half of the action space in *every* phase (unlike in SHB where by backtracking we need only make the correct decision in some proportion of the phases - see Appendix C). Therefore for SH, η does affect the rate at which the

failure probability decreases. This is in contrast to Theorem 1 and Theorem 2 where η 's involvement is only additive in the exponent and does not affect the rate. This indicates SH is suboptimal for large budgets, which we discuss further in Appendix B. However, we now show that SH is minimax optimal up to constants in the small budget regime.

5.2 Lower Bound

To understand the influence of η on the difficulty of the small budget problem, we consider covering/Fano arguments similar to Chapter 15 of (Wainwright, 2019). From this we obtain a minimax lower bound shown in Theorem 4, the proof of which is in Appendix E. By comparing the upper bound for SH in Theorem 3 with the lower bound in Theorem 4, we see that SH is near minimax-optimal in the small budget regime, $T < T_1$, up to an additive $\log \log(1/\eta)$ term in the exponent.

Theorem 4. Let $\overline{V} \subset V(\Delta, \sigma)$ be the set of environments with change in mean Δ and Gaussian random noise with variance σ^2 . When $T < T_1$, we have

$$\inf_{\pi \in \Pi} \sup_{v \in \bar{V}} \mathbb{P}_{\pi,v}(F_{T,v,\eta}) \ge \exp\left(-\frac{\Delta^2 T + 2\sigma^2 \log(2)}{\sigma^2 \log(\lfloor \frac{1}{2\eta} \rfloor)}\right).$$

Theorem 4 holds for hold for small budgets, $T < T_1 := \frac{\sigma^2}{\Delta^2}(1.59 \log(\lfloor 1/2\eta \rfloor) - 2 \log(2))$ while Theorem 2 holds for large budgets $T \ge T_1$. These results together give a full characterisation of the difficulty of the fixed-budget piecewise constant bandit problem. Interestingly they show a separation in the achievable error probability (with respect to η) depending on the budget regime. Similar separations in complexity have been observed in active learning (Dasgupta, 2005), although those results hold for fixed confidence problems with a labelling oracle, which we do not have in our setting.

6 ADAPTIVE ALGORITHM

Dependent on what regime we are in, it may be better to use SH or SHB. For example, our theoretical results show that when T is smaller than a threshold of order $\frac{\sigma^2}{\Delta^2} \log(1/\eta)$ our guarantee for SH (Theorem 3) is better than SHB (Theorem 1), suggesting SH is better suited for such smaller budgets. The converse can hold when we consider larger budgets. Note that while these observations come from comparing upper bounds on the performance of SH and SHB, we see experimentally that these conclusions hold (Section 7). Moreover, our lower bounds show that our guarantees are tight up to constants or $\log \log(1/\eta)$ terms.

In practice, we may not know whether SH or SHB is more appropriate for our problem setting, since the budget threshold depends on unknown problem parameters, Δ, σ . While it is reasonable to assume σ is $known^1$ (e.g. by restricting rewards to be bounded, see Audibert et al. (2010)), Δ is still unknown. To deal with this, we propose a regime adaptive method, Adaptive Sequential Halving (SHA), which performs near-optimally regardless of the setting we are in. The main idea is to use some small portion of the budget to identify whether SH or SHB is more appropriate to play for the remaining budget. In particular, we first sample the actions 0, 1 a total of L < T times and use these samples to estimate the change in mean $\hat{\Delta}$. We use Δ to estimate a budget threshold τ . Then if the budget is smaller than τ , we use SH for the remainder of the budget and if T is larger than the threshold τ we use SHB. SHA takes as input a parameter $\gamma > 0$ and uses the estimated threshold

$$\tau = \gamma \frac{\sigma^2}{\hat{\Delta}^2} \log\left(\frac{1}{2\eta}\right). \tag{5}$$

SHA is written explicitly in Appendix A, Algorithm 4.

To run SHA we have to choose appropriate hyperparamaters γ, L . While we could pick γ to match T_1 (from Theorems 2, 4), it turns out there are better choices both in theory and practice. In Theorem 5 we show that there exists a universal choice for γ, L such that the SHA algorithm is near optimal for both small and large budgets simultaneously. Choices for γ, L which perform well in practice are seen in Section 7.

Theorem 5. Let L = BT for some $B \in (2/T, 1-2/T)$, $\gamma \in \left((\sqrt{104/B} + \sqrt{1.59})^2, 1800/(1-B) \right)$, and define $\ell_{\eta} = \log_2(1/2\eta)$. In an environment $v \in V(\Delta, \sigma)$, using SHA and with universal constants c_1, c_2, c_3

$$\mathbb{P}(F_{T,v,\eta}) \leq \begin{cases} 4 \left\lceil \ell_{\eta} \right\rceil \exp\left(\frac{-c_{1}(1-B)\Delta^{2}T}{\sigma^{2}\ell_{\eta}}\right), & T < T_{1} \\ 5 \exp\left(-\frac{c_{2}B(1-B)\Delta^{2}}{\sigma^{2}}T + c_{3}\ell_{\eta}\right), & T \geq T_{1} \end{cases}$$

The proof of Theorem 5 is given in Appendix H. Theorem 5 shows that SHA simultaneously matches the lower bounds in Theorems 2 and 4 for $T \ge T_1$ and $T < T_1$, up to constants and $\log \log(1/\eta)$ terms.

7 EXPERIMENTS

Comparing SH, SHB, and SHA In Figure 3a we demonstrate that SH or SHB can perform better than the other depending on the problem setting. We consider a synthetic environment with $\Delta = 2, \sigma = 8, x^* = 0.7, \eta = 10^{-8}$. Figure 3a shows the error as a function

¹While this assumption is not needed for SH/SHB, we assume σ^2 is known here to isolate the key difficulty of the adaptive problem.



Figure 3: Proportion of final estimates more than η away from x^* against the inputted budget, T, with Gaussian rewards, $\Delta = 2$ and 90% CIs. (a) compares the SH, SHB, and SHA by running each algorithm 1000 times with different budgets. (b,c) both compare SH, SRR, and ACPD. We run SH and SRR 500 times each at different budgets, while the anytime ACPD algorithm is run a total of 500 times for T = 60 and we plot the evolution of ACPD's failure probability.

of inputted budget. We see that SH has a smaller failure probability for smaller budgets. However, as the budget gets larger, SHB performs better and reaches near-zero probability faster. This supports our theoretical observations in Section 6. Similar results hold for other problem instances (see Appendix B). We compare SHA to SH/SHB in the same environment. We chose the hyperparameters L = T/20 and $\gamma = 120$ for SHA as these worked well across a variety of environments (see Appendix B). We see from Figure 3a that SHA can outperform SHB for small budgets and outperform SH for large budgets. Therefore, we observe that SHA performs well for all the budgets simultaneously.

Comparison with existing work We compare our proposed algorithms with existing methods. In Figures 3b and 3c we consider synthetic environments with $\Delta = 2, \sigma = 1, \eta = 0.1$ only varying the location of the change point from $x^* = 0.7$ to $x^* = 0.01$, respectively. We compare our SH algorithm to SRR (Hall and Molchanov, 2003) and ACPD (Hayashi et al., 2019). Due to the significant computational expense² of repeated GP regression in the ACPD algorithm, we focus on smaller budget problems (where simulations are tractable) and consequently compare these algorithms to our small budget algorithm, SH. For ACPD, we use parameters suggested by Hayashi et al. (2019) along with knowledge that rewards are Gaussian. We initialise ACPD with 10 random actions, and use GP-UCB (Srinivas et al., 2010) with Matérn kernel of smoothness 5/2 for the Bayesian Optimisation. For the SRR algorithm we use the proposed parameters in Section 4 of (Hall and Molchanov, 2003). In particular we choose the number of actions in each stage $m = \log^{2+\beta}(T)$ and the CI radius parameter $\lambda = \log^{1+\alpha}(T)$, with $\alpha = 1, \beta = 0.1$. We show that tuning these parameters for particular environments does not lead to significant performance boosts in Appendix B. There we also show that the method by Lan et al. (2009) (which is similar to SRR) also has very similar empirical performance to SRR.

We first compare the performances of these algorithms when the change point is $x^* = 0.7$ in Figure 3b. Here we see that the failure probability of ACPD decreases fastest while SRR and SH are similar to each other, decreasing at a slower rate towards zero. However, when the change point is closer to the boundary $(x^* =$ (0.01) in Figure 3c we see that the performance of SRR and ACPD becomes significantly worse, whereas the performance of SH is relatively unchanged.³ Intuitively this occurs because existing methods are based on an initial widespread exploration of the action space before quickly focusing in on where they believe x^* to be. However when the change is near the boundary or the problem is challenging this initial exploration may have not have a sufficient number of samples from both sides of x^* meaning that it is hard to detect a change in mean reward (i.e. distinguish noise vs signal). Hence, existing methods can miss the change point with high probability. On the other hand our proposed methods start by sampling on the boundaries of the space and slowly zoom in to where we believe the change to be, ensuring that we do not miss the change point (with high probability). This demonstrates that existing methods' performance is dependent on the change point being near the centre of the action space whereas our proposed policies perform well regardless of the location of x^* . Our methods are also simple,

 $^{^2\}mathrm{Computational}$ costs are discussed further in Appendix B.

³Recall that we would not expect the curves for our proposed algorithms to be strictly monotonic due to the floor/ceiling functions, for example in the definition of our SH algorithm (see Algorithm 1, line 7).

computationally inexpensive, and accompanied with tight non-asymptotic optimality guarantees.

8 DISCUSSION

In this paper, we studied the piecewise constant bandit problem and provided the first non-asymptotic problem dependent theoretical analysis of the problem. We developed two algorithms, SH and SHB, which achieve nearly minimax optimal error probabilities under different conditions on the budget. We then combined these two algorithms into a regime adaptive method SHA which is near optimal in both regimes simultaneously. We complemented our theoretical results with simulations and provided a comparison to existing methods.

A natural extension of the piecewise constant bandit problem would be to allow for multiple abrupt changes in the reward function. However, we suspect that this would require significant innovation in order to carry out a large amount of additional (carefully chosen) exploration to locate and distinguish individual change points. Furthermore, if the number of change points were unknown, this would also need to be estimated within the limited budget. Another extension could be to consider an action space which is multi-dimensional or, more broadly, any metric space. We note however, this will be challenging as our objective will no longer be to identify a single change point in the action space, but instead to identify an entire boundary between disjoint subsets of the action space. Hence, significantly more samples would be required to ensure enough data is near the whole boundary to get a good estimate of it and we would likely need some additional assumptions on the shape of the boundary to make any analysis tractable. Finally, it would be interesting to extend our methods to the piecewise smooth setting as in Hall and Molchanov (2003). We expect significant innovation to be required to extend our non-asymptotic results to these settings, although we hope our methods will provide a useful starting point. The findings in this paper already represent a significant advancement in understanding the complexity of choosing samples to learn the location of a change point in unknown, noisy environments.

Acknowledgments

The authors would like to thank Professor Niall Adams for support and insightful discussions about this problem. Joseph Lazzaro was supported by a Roth Scholarship from the Department of Mathematics, Imperial College London.

References

- Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In Advances in Neural Information Processing Systems, 2011.
- A. Agarwal, D. P. Foster, D. J. Hsu, S. M. Kakade, and A. Rakhlin. Stochastic convex optimization with bandit feedback. In *Advances in Neural Information Processing Systems*, 2011.
- S. Aminikhanghahi and D. Cook. A survey of methods for time series change point detection. *Knowledge* and *Information Systems*, 2017.
- J.-Y. Audibert, S. Bubeck, and R. Munos. Best arm identification in multi-armed bandits. In *COLT 2010 The 23rd Conference on Learning Theory*, 2010.
- L. D. Brown, T. T. Cai, and A. DasGupta. Interval Estimation for a Binomial Proportion. *Statistical Science*, 2001.
- S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári. X-armed bandits. *Journal of Machine Learning Re*search, 2011.
- A. Carpentier and A. Locatelli. Tight (lower) bounds for the fixed budget best arm identification bandit problem. In *Annual Conference Computational Learning Theory*, 2016.
- R. M. Castro and R. D. Nowak. Minimax bounds for active learning. *IEEE Transactions on Information Theory*, 2008.
- R. M. Castro, R. M. Willett, and R. D. Nowak. Faster rates in regression via active learning. In Advances in Neural Information Processing Systems, 2005.
- J. Chen and A. K. Gupta. Parametric Statistical Change Point Analysis: With Applications to Genetics, Medicine, and Finance; 2nd ed. Springer, 2012.
- J. Cheshire, P. Menard, and A. Carpentier. Problem dependent view on structured thresholding bandit problems. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.
- S. Dasgupta. Coarse sample complexity bounds for active learning. Advances in Neural Information Processing Systems, 2005.
- Y. Deng, X. Zhou, B. Kim, A. Tewari, A. Gupta, and N. Shroff. Weighted gaussian process bandits for nonstationary environments. In *International Conference* on Artificial Intelligence and Statistics, 2022.
- A. Garivier and E. Moulines. On upper-confidence bound policies for switching bandit problems. In *Algorithmic Learning Theory*, 2011.

- A. Garivier, P. Ménard, and G. Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 2016.
- R. Gramacy and H. Lee. Adaptive design and analysis of supercomputer experiments. *Technometrics*, 2008.
- L. Gretta and E. Price. Sharp noisy binary search with monotonic probabilities. In 51st International Colloquium on Automata, Languages, and Programming (ICALP 2024), 2024.
- P. Hall and I. Molchanov. Sequential methods for design-adaptive estimation of discontinuities in regression curves and surfaces. *The Annals of Statistics*, 2003.
- S. Hayashi, Y. Kawahara, and H. Kashima. Active change-point detection. In Proceedings of The Eleventh Asian Conference on Machine Learning, 2019.
- K. Hong, Y. Li, and A. Tewari. An optimization-based algorithm for non-stationary kernel bandits without prior knowledge. In *International Conference on Artificial Intelligence and Statistics*, 2023.
- Z. Karnin, T. Koren, and O. Somekh. Almost optimal exploration in multi-armed bandits. In *Proceedings* of the 30th International Conference on Machine Learning, 2013.
- R. M. Karp and R. D. Kleinberg. Noisy binary search and its applications. In ACM-SIAM Symposium on Discrete Algorithms, 2007.
- R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, 2008.
- Y. Lan, M. Banerjee, and G. Michailidis. Change-point estimation under adaptive sampling. *The Annals of Statistics*, 2009.
- T. Lattimore and C. Szepesvári. Bandit Algorithms. Cambridge University Press, 2020.
- A. Locatelli, M. Gutzeit, and A. Carpentier. An optimal algorithm for the thresholding bandit problem. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016.
- Z. Lu, M. Banerjee, and G. Michailidis. Intelligent sampling for multiple change-points in exceedingly long time series with rate guarantees. In ArXiv Preprint ArXiv:1710.07420, 2020.
- S. Mannor, J. Tsitsiklis, K. Bennett, and N. Cesabianchi. The sample complexity of exploration in the multi-armed bandit problem. 2004.
- R. Nowak. Noisy generalized binary search. In Proceedings of the 23rd International Conference on Neural Information Processing Systems, 2009.

- C. Park, P. Qiu, J. Carpena-Núñez, R. Rao, M. Susner, and B. Maruyama. Sequential adaptive design for jump regression estimation. In ArXiv Preprint ArXiv:1904.01648, 2021.
- C. Park, R. Waelder, B. Kang, B. Maruyama, S. Hong, and R. Gramacy. Active learning of piecewise gaussian process surrogates. In ArXiv Preprint ArXiv:2301.08789, 2023.
- K. Sikorski. Bisection is optimal. Numerische Mathematik, 40:111–118, 1982.
- N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML 2010 Proceedings*, 27th International Conference on Machine Learning, 2010.
- V. Thuot, A. Carpentier, C. Giraud, and N. Verzelen. Active clustering with bandit feedback. In ArXiv Preprint ArXiv:2406.11485, 2024.
- M. J. Wainwright. High-Dimensional Statistics: A Non-Asymptotic Viewpoint. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- J. Yang, Z. Zhong, and V. Y. F. Tan. Optimal clustering with bandit feedback. *Journal of Machine Learning Research*, 2022.
- R. C. Yavas, Y. Huang, V. Y. F. Tan, and J. Scarlett. A general framework for clustering and distribution matching with bandit feedback. In ArXiv Preprint ArXiv:2409.05072, 2024.

Checklist

- 1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model.
 [Yes] The mathematical setting was outlined clearly in Section 3 and detailed descriptions of the proposed algorithms are in Sections 4, 5, and 6.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm.[Yes] The paper focuses on analysing properties of the proposed algorithms. Theorems 1, 3, and 5 provide our main gaurantees on failure probability.
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes] We include code in the supplementary material, with more details on how to run the code/where to see requirements discussed in Appendix B.

- 2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes] All theoretical results presented in the paper (namely Theorems 1,2,3,4, and 5) are accompanied by assumptions/settings either in the statement of the result, in the adjacent text, or in the problem setting Section 3.
 - (b) Complete proofs of all theoretical results. [Yes] These can all be found in the supplementary material. Specific appendices are pointed to throughout the main paper.
 - (c) Clear explanations of any assumptions. [Yes] Any assumptions are either made in the statement of the theoretical claims or in the surrounding text.
- 3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL).[Yes] This is provided in the supplementary material, with more details in Appendix B.
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes] Parameters are discussed in Section 7 for the main simulations, with further discussion in Appendix B.
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes] We formally define measures in Section 3 and detail the construction of simulated results in Section 7, with more details on construction of confidence intervals in Appendix B.
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes] This is detailed in Appendix B.
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]

- (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
- 5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

 $\mathbf{2}$

 $\mathbf{4}$

 $\mathbf{7}$

APPENDIX

Contents	
1 INTRODUCTION	1
2 RELATED WORK	2
3 PROBLEM SETTING	3
4 LARGE BUDGET	4
4.1 Upper Bound	5
4.2 Lower Bound	6
5 SMALL BUDGET	6
5.1 Upper Bound	6
5.2 Lower Bound	7
6 ADAPTIVE ALGORITHM	7
7 EXPERIMENTS	7
8 DISCUSSION	9
A EXPLICIT ALGORITHMS	14
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS	14 15
 A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values 	14 15 15
 A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms 	14151516
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime	 14 15 16 17
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η	 14 15 16 17 18
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals	 14 15 16 17 18 19
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals B.6 Code	 14 15 16 17 18 19 19
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals B.6 Code	 14 15 16 17 18 19 19 20
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals B.6 Code Code C.1 Proof for Theorem 1	 14 15 16 17 18 19 19 20 20
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals B.6 Code Composition of FOR LARGE BUDGET UPPER BOUND C.1 Proof for Theorem 1 C.2 Proof of Lemma 1	 14 15 16 17 18 19 19 20 20 20 20
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals B.6 Code C PROOFS FOR LARGE BUDGET UPPER BOUND C.1 Proof for Theorem 1 C.2 Proof of Lemma 1 C.3 Proof of Lemma 2	 14 15 16 17 18 19 19 20 20 20 20 22
A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals B.6 Code C1 Proof for LARGE BUDGET UPPER BOUND C.1 Proof for Theorem 1 C.2 Proof of Lemma 1 C.3 Proof of Lemma 2 C.4 Proof of Lemma 3	 14 15 16 17 18 19 19 20 20 20 20 20 22 25
 A EXPLICIT ALGORITHMS B ADDITIONAL EXPERIMENTS B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values B.2 Comparison of All Algorithms B.3 Computational Complexity and Runtime B.4 SH vs ACPD with Smaller η B.5 Confidence intervals B.6 Code C PROFS FOR LARGE BUDGET UPPER BOUND C.1 Proof for Theorem 1 C.2 Proof of Lemma 1 C.3 Proof of Lemma 2 C.4 Proof of Lemma 3 	 14 15 16 17 18 19 19 20 20 20 20 20 20 22 25 28

	D.3 Proof of Theorem 2	30
E	PROOFS FOR SMALL BUDGET LOWER BOUND	32
	E.1 Proof of Theorem 4	32
F	PROOFS FOR SMALL BUDGET UPPER BOUND	33
	F.1 Proof for Theorem 3	33
G	CHANGE POINTS AND INTUITION FOR ELIMINATION CRITERIA	34
	G.1 Estimation in Offline Change Point Analysis	34
	G.2 Justification of Elimination Criteria	34
н	PROOFS FOR ADAPTIVE ALGORITHM	36
	H.1 Proof of Theorem 5	36

EXPLICIT ALGORITHMS Α

We explicitly write the SH, SHB and SHA algorithms below.

Algorithm 2 Sequential halving with backtracking (SHB)

1: Input: $\eta \in (0, 1/2)$, and budget T 2: $\mathcal{A}^1 = \{0, a_1^1, a_2^1, a_3^1, 1\} \leftarrow \{0, 0, 1/2, 1, 1\}$ 3: $J \leftarrow \left\lceil 6 \log(1/2\eta) \right\rceil$ 4: $t_j \leftarrow \left\lfloor \frac{T}{5J} \right\rfloor$ 5: for phase j in 1, ..., J do Play each action in $\mathcal{A}^j = \{0, a_1^j, a_2^j, a_3^j, 1\}, t_j$ times 6: 7:if $E_{P,i}$ from (4) holds then $\mathcal{A}^{j+1} \leftarrow \mathcal{P}(\mathcal{A}^j)$ 8: else if $E_{R,j}$ from (2) holds then 9: $\mathcal{A}^{j+1} \leftarrow R(\mathcal{A}^j)$ \triangleright Zoom in to the right 10:else if $E_{L,j}$ from (3) holds then $\mathcal{A}^{j+1} \leftarrow L(\mathcal{A}^j)$ 11: 12: \triangleright Zoom in to the left end if 13:14: **end for** 15: **Return:** $\hat{x}_T = a_2^{J+1}$

Algorithm 3 Sequential halving without backtracking (SH)

Input: $\eta \in (0, 1/2)$, and budget T $\{a_1^1, a_2^1, a_3^1\} \leftarrow \{0, 1/2, 1\}$ $J \leftarrow \left\lceil \log_2(1/2\eta) \right\rceil$ $t_j \leftarrow \left\lfloor \frac{T}{3J} \right\rfloor$ for phase j in 1, ..., J do Play each action a_1^j, a_2^j, a_3^j, t_j times if $E_{R,j}$ from (2) holds then $\{a_1^{j+1}, a_2^{j+1}, a_3^{j+1}\} \leftarrow \{a_2^j, (a_2^j + a_3^j)/2, a_3^j\}$ else if $E_{L,j}$ from (3) holds then $\{a_1^{j+1}, a_2^{j+1}, a_3^{j+1}\} \leftarrow \{a_1^j, (a_1^j + a_2^j)/2, a_2^j\}$ end if end for Return $\hat{x}_T = a_2^J$

 \triangleright Zoom in to the right

▷ Backtrack

 \triangleright Zoom in to the left

Algorithm 4 Adaptive Sequential Halving (SHA)	
1: Input: $\gamma > 0, L > 2$	▷ Hyperparameters
2: Input: $\eta \in (0, 1/2), \sigma > 0$ and budget T	
3: Play actions 0,1 each $L/2$ times	\triangleright Initial phase to estimate threshold
4: Calculate $\hat{\Delta}_L = \hat{\mu}_{1,L} - \hat{\mu}_{0,L} $	
5: Calculate $\tau = \gamma \frac{\sigma^2}{\Delta_t^2} \log\left(\frac{1}{2\eta}\right)$	\triangleright Threshold τ from (5)
6: if $T \ge \tau$ then	
7: $\hat{x}_{T-L}^{SHB} \leftarrow SHB(T-L,\eta)$	\triangleright Play SHB for remainder of budget
8: Return: $\hat{x}_T^{SHA} = \hat{x}_{T-L}^{SHB}$	
9: end if	
10: if $T < \tau$ then	
11: $\hat{x}_{T-L}^{SH} \leftarrow SH(T-L,\eta)$	▷ Play SH for remainder of budget
12: Return: $\hat{x}_T^{SHA} = \hat{x}_{T-L}^{SH}$	
13: end if	

B ADDITIONAL EXPERIMENTS

B.1 Further Empirical Comparisons of SH, SHB and SHA with Other η Values

In Figure 3a from Section 7, we demonstrate that when η is very small ($\eta = 10^{-8}$ in that case) the SH algorithm will outperform SHB for smaller budgets, whereas the SHB algorithm will reach near zero failure probabilities faster than SH. This matches our discussion of our theoretical results in Section 6. This also matches our intuition. Namely, in order for SH to return a good estimate, it must eliminate the correct half of the action space in *all* of the $\log_2(1/2\eta)$ phases. This occurs with very small probability when η is small (i.e. the number of phases is large) and therefore we can benefit from using SHB which explores more and only requires correct decisions in some portion of the phases for a good estimate. In Figure 3a we also illustrate the desired effect of SHA, which performs well for both large and small budgets simultaneously.

In Figure 4 we now consider two other settings to compare SH, SHB, and SHA where our simulations provide similar conclusions. In particular, in Figures 4a and 4b we consider settings with $\eta = 10^{-7}$ and 10^{-6} respectively. In both cases, we once again see that for smaller budgets SH outperforms SHB, while SHB is able to reach near-zero failure probabilities faster than SH in larger budgets. Hence, SH is still more appropriate to use for settings with smaller budgets whereas SHB is more appropriate to use for settings with larger budgets. Additionally we show that SHA, with the hyperparameters L = T/20 and $\gamma = 120$ performs well in both of these additional settings in Figures 4a and 4b. In particular, for both of these new settings, SHA outperforms SHB for smaller budgets and outperforms SH for larger budgets. We conclude SHA performs well for all budgets simultaneously. We note that the hyperparameters L = T/20 and $\gamma = 120$ are the same as those used for Figure 3a, suggesting that this is a good choice of hyperparameters for SHA that work well in a variey of settings. (This will be further supported by Figure 5.)



Figure 4: Proportion of final estimates more than η away from x^* against the inputted budget, T, with Gaussian rewards, $\Delta = 2$ and 90% CIs. SH, SHB, and SHA were each run 1000 times with different budgets.

In Figure 5, we also compare SH, SHB, and SHA in settings with larger choices of η . In particular we consider $\eta = 10^{-4}$ and $\eta = 10^{-2}$ in Figures 5a and 5b, respectively. From these figures, we firstly note that our adaptive algorithm SHA, with the hyperparameters L = T/20 and $\gamma = 120$, still performs well regardless of the budget. Secondly, from Figures 5a and 5b (as well as 4a and 4b), we observe the interesting phenomenon that the advantages of using SHB over SH become less significant as η becomes larger. This matches our intuition. In particular, for SH to return a good estimate, it must make the correct decision in all $\log_2(1/2\eta)$ phases. This can occur with high probability when η is large (i.e. the number of phases is small). Therefore the exploitative nature of SH in settings with larger η can outweigh the benefits from the additional exploration in SHB (which can allow

mistakes in some phases). These observations also match our theoretical results. In particular for larger values of η , the the $\frac{T\Delta^2}{36\sigma^2 \log_2(1/2\eta)}$ term from our guarantee for SH (Theorem 3) can be similar to or larger than the $\frac{\Delta^2}{600\sigma^2}T$ term from our guarantee for SHB (Theorem 1). In which case our guarantee for SH would be similar to or better than SHB.



Figure 5: Proportion of final estimates more than η away from x^* against the inputted budget, T, with Gaussian rewards, $\Delta = 2$ and 90% CIs. SH, SHB, and SHA were each run 1000 times with different budgets.



B.2 Comparison of All Algorithms

Figure 6: Proportion of final estimates more than η away from x^* against the inputted budget, T, with Gaussian rewards, $\Delta = 2$ and 90% CIs. SH, SHB, SHA, SRR, and SR were each run 500 times with different budgets, while the anytime ACPD algorithm is run a total of 500 times for T = 60 and we plot the evolution of ACPD's failure probability.

We include the additional algorithms of SHB, SHA and SR (Lan et al., 2009) to to the setting of Figures 3b and 3c in Figures 6a and 6b here. Recall from Section 7 that, due to the very computationally expensive existing methods (e.g. ACPD), we consider settings with very limited budgets. Our small budget algorithm SH is much better suited to these very limited budgets and hence this was the main algorithm we compared with existing

work. Furthermore, in this setting SHB (and similarly SHA) requires a budget $T \ge 50$ in order for there to be at least one action for each sampling point in each phase. We plot the performance of SHB/SHA for such budgets in Figures 6a and 6b. More thorough simulations comparing SH, SHB, and SHA can be seen in Appendix B.1.

SRR Algorithm and Tuning In the Sequential Refinement with Reassessment Algorithm (SRR), Hall and Molchanov (2003) propose a multi-stage method where they first spend half of the budget uniformly exploring the space and they then spend every subsequent stage sampling m actions evenly across a confidence interval for the change point constructed in the previous stage. The width of this confidence interval is influenced by parameter λ . In each stage they also have a reassessment criteria in which they test for the presence of a change point during the current stage (within the current CI for the change point) and if it is not significant with level $1 - \epsilon$ confidence, they "reassess" and return to the previous confidence interval constructed. In Section 4 of Hall and Molchanov (2003), the authors propose setting $m = \log^{2+\beta}(T)$ and $\lambda = \log^{1+\alpha}(T)$ with $0 < \alpha < 1 + \beta$ with $\beta > 0$. Since we would like for there to be at least two stages and to satisfy these conditions, for experiments in Figures 3b and 3c (similarly 6a, 6b) we set $\beta = 0.1, \alpha = 1$.

One might be concerned that the negative performance by SRR compared to SH when the change point is near the boundary in Figure 3c, might be due to poorly chosen parameters α, β . In practice, we will not know which parameters perform best in which environment. However, we can show that tuning these parameters does not significantly improve the performance in these settings. To do so, we first note that β should be set to be at most 0.4 since otherwise this would mean for $T \leq 60$ that there is only one stage. With budget T = 60 for 10 values of β (between 0 and 0.4), and at each value for β we run 5 values for α (between 0 and $1 + \beta$) 500 times. Across these 50 different α, β pairs chosen, we select the pair with minimal failure probability. We did this tuning for both environments studied in Figures 3b and 3c individually. For each we then plot the performance of SRR with these tuned parameters against the SH algorithm in Figures 7a, 7b.

SR Algorithm We refer to the algorithm proposed by Lan et al. (2009) as Sequential Refinement (SR). This is because the algorithm itself is extremely similar to SRR, except the policy never reassess once it has "zoomed in" to a particular region of the action space. (Furthermore, they do not spend half of the budget initially exploring the space.) In particular, the SR algorithm begins by splitting the budget into L stages. Then, in each stage SR plays actions evenly across a confidence interval for the change point (constructed using samples in the previous stage). For their multi-stage method Lan et al. (2009) generally recommend splitting the budget evenly into L stages such that in each phase SR plays around 30 - 50 actions. Hence, for these very limited budget settings with $T \leq 60$ seen in Figures 3b and 3c, we choose the smallest number of stages L = 2. For the construction of the confidence intervals we use the proposed form shown in equation (10) of Lan et al. (2009), which requires knowledge of the signal to noise ratio Δ/σ . Note that we would not have access to this in practice. We plot the performance of SR in settings described in Figures 6a and 6b. Here we see that the the failure probability when running SR is very similar to SRR regardless of the position of the change point ($x^* = 0.7$ or $x^* = 0.01$). Furthermore, SH still significantly outperforms SR when the change is near the boundary $x^* = 0.01$.

B.3 Computational Complexity and Runtime

As mentioned in Section 7, while there are some settings in which ACPD can outperform our proposed algorithms, it comes at a significant additional computational cost due to the repeated GP regression used in ACPD. Consequently the computational complexity of ACPD, given a budget of T, is at least $O(T^3)$. Whereas, the computational complexity of our proposed methods is O(T). This is reflected by the runtimes from our simulation studies in Figures 3b and 3c, displayed in Table 1 where we see that the time taken by ACPD in both experiments is significantly larger than SH.

Table 1: Run times for algorithms SH, SRR, and ACPD required to run the simulation studies for Figures 3b and 3c. Simulations were run on an i7 CPU with 8GB RAM.

Setting	SH Algorithm	SRR Algorithm	ACPD Algorithm
Figure 3b	7.8s	155.9s	5004.2s
Figure 3c	8.2s	132.3s	5053.5s



Figure 7: Proportion of final estimates more than η away from x^* against the inputted budget, T, with Gaussian rewards, $\Delta = 2$ and 90% CIs. SH, and tuned SRR were each run 500 times with different budgets. In (a) the tuned constants for SRR tuned were $\alpha \approx 0.28, \beta \approx 0.41$. In (b) the tuned constants were $\alpha \approx 1.29, \beta \approx 0.29$.

B.4 SH vs ACPD with Smaller η



Figure 8: Proportion of final estimates more than $\eta = 0.009$ away from $x^* = 0.01$ against the inputted budget, T, with Gaussian rewards with $\sigma = 1$, $\Delta = 2$ and 90% CIs. SH, and ACPD were each run 500 times with different budgets.

In Figure 3c, we illustrate that there are settings where the proposed SH algorithm significantly outperforms ACPD. In particular, we considered an environment with change point $x^* = 0.01$ and acceptable error $\eta = 0.1$, meaning that estimates for the change point which were placed at the boundary of the space ($\hat{x}_T = 0$) are classified as successes. However, if we restrict the region of acceptable estimates for the change point by instead considering $\eta = 0.009$, the region of successful estimates of the change point is strictly within (0, 1) and a similar comparison between SH and ACPD still holds through the same intuition. This is displayed in Figure 8 where we see that SH still significantly outperforms ACPD since the failure probability decreases at a much faster rate as the budget increases.

B.5 Confidence intervals

Confidence intervals in all plots for the failure probabilities of different algorithms were calculated using a simple Gaussian approximation, as seen in equation (1) of (Brown et al., 2001). This was done using the set of all results at each budget for each algorithm.

B.6 Code

Code can be found at the following GitHub repository: https://github.com/JosephLazzaro/cpi_AISTATS2025.

PROOFS FOR LARGE BUDGETS

C PROOFS FOR LARGE BUDGET UPPER BOUND

C.1 Proof for Theorem 1

Proof. Note that, since Algorithm 2 is an extension of Cheshire et al. (2021), proofs for some technical lemmas (e.g. Lemmas 1, 3) are similar to existing work, up to the inclusion of an infinite action space and different elimination/backtracking probabilities. We write the main steps here and prove lemmas below.

Firstly, for all j in $\{1, ..., J\}$, let G_j be the "good event"

$$G_{j} = \left\{ E_{P,j} \cap \left\{ x^{*} \in (0,1) \setminus [a_{1}^{j}, a_{3}^{j}] \right\} \right\} \cup \left\{ E_{P,j}^{C} \cap E_{R,j} \cap \left\{ x^{*} \in [a_{2}^{j}, a_{3}^{j}] \right\} \right\}$$

$$\cup \left\{ E_{P,j}^{C} \cap E_{L,j} \cap \left\{ x^{*} \in [a_{1}^{j}, a_{2}^{j}] \right\} \right\}$$

$$(6)$$

in which we make the correct decision in eliminating or backtracking in phase j. Namely, when the change point is not within the remaining action space in phase j, $x^* \in (0,1) \setminus [a_1^j, a_3^j)$, we correctly backtrack. When $x^* \in [a_2^j, a_3^j)$ we correctly zoom into the right and eliminate the left half of region. And when $x^* \in [a_1^j, a_2^j)$ we correctly zoom into the left and eliminate the right half of the remaining action space, see Figure 2 again for illustration of the regions.

The first thing to note is that we need not make the correct decision in every phase. In particular, we can incorrectly eliminate a region when it contains the change point or incorrectly backtrack when the remaining action space contains the change point a limited number of times. As long as we fail in less than 1/4 of the phases, we will still achieve our objective $|\hat{x}_T - x^*| < \eta$. This is shown in Lemma 1.

Lemma 1. Let $\eta < 1/4$. Under Algorithm 2, if we have $\sum_{j=1}^{J} \mathbb{1}\{G_j^C\} < J/4$, then our final estimate will be within η of the true change point, $|\hat{x}_T - x^*| < \eta$.

We then calculate a lower bound on the probability of making the correct decision in round j, given the actions and rewards from all previous phases 1, ..., j - 1.

Lemma 2. Let \mathcal{F}_{j-1} be the sigma algebra generated by all actions and rewards in the first j-1 phases. Under Algorithm 2 we have

$$\mathbb{P}(G_j|\mathcal{F}_{j-1}) \ge 1 - 8 \exp\left(-\frac{t_j \Delta^2}{30\sigma^2}\right).$$

Using Lemma 2, we upper bound the probability that we make the incorrect decision in more than 1/4 of the phases.

Lemma 3. Let $\eta < 1/4$ and set $C_1 = \frac{1}{600}$, $C_2 = 34$, then the following inequality holds under Algorithm 2.

$$\mathbb{P}\left(\sum_{j=1}^{J} \mathbb{1}\{G_j^C\} \ge J/4\right) \le \exp\left(-C_1 \frac{\Delta^2}{\sigma^2} T + C_2 \log(\frac{1}{2\eta})\right).$$

Putting Lemma 1 and Lemma 3 together we get the required result. We also assume $T > 60 \log(1/2\eta)$ so that we can even run the algorithm (playing each action in each phase at least once).

C.2 Proof of Lemma 1

Proof. Except for the fact that we are working in a continuous action space and not a finite one, the structure of our backtracking policy is similar to that of Cheshire et al. (2021). So the proof of this lemma is almost identical to what is seen in Cheshire et al. (2021).

In our case, the objective that we would like to satisfy is a sufficiently accurate estimate of the change point $|\hat{x}_T - x^*| < \eta$ with high probability. As long as we make the correct decision in $\log_2(1/2\eta)$ more phases than

we make incorrect decisions, then we will have made enough correct decisions to cancel out our bad ones and correctly zoomed in enough to achieve $|\hat{x}_T - x^*| < \eta$. The quantity $\log_2(1/2\eta)$ is important as we need to correctly zoom in enough times such that the final region in the action space has sufficiently small width, i.e., $a_3^{J+1} - a_1^{J+1} < 2\eta$ and contains the true change point $x^* \in [a_1^{J+1}, a_3^{J+1})$. This ensures that estimating the midpoint of this region guarantees $|\hat{x}_T - x^*| < \eta$.

Writing this explicitly, it is sufficient to satisfy equation (7) to achieve $|\hat{x}_T - x^*| < \eta$.

$$\sum_{j=1}^{J} \mathbb{1}\{G_j\} - \sum_{j=1}^{J} \mathbb{1}\{G_j^C\} > \log_2(1/2\eta)$$
(7)

Hence to prove this lemma, we just need to demonstrate that $\sum_{j=1}^{J} \mathbb{1}\{G_j^C\} < J/4$ implies equation (7) and therefore $|\hat{x}_T - x^*| < \eta$.

We first note that if we have $\sum_{j=1}^{J} \mathbb{1}\{G_j^C\} < J/4$, then we also have that $\sum_{j=1}^{J} \mathbb{1}\{G_j\} \ge 3J/4$. We can then plug this into the following equation.

$$\sum_{j=1}^{J} \mathbb{1}\{G_{j}^{C}\} - \sum_{j=1}^{J} \mathbb{1}\{G_{j}^{C}\} > \frac{3J}{4} - \frac{J}{4} = \frac{J}{2}$$

$$= \frac{1}{2} \left\lceil 6 \log\left(\frac{1}{2\eta}\right) \right\rceil$$

$$= \frac{1}{2} \left\lceil 6 \frac{\log_{2}(\frac{1}{2\eta})}{\log_{2}(e)} \right\rceil$$

$$\geq \frac{1}{2} \left(6 \frac{\log_{2}(\frac{1}{2\eta})}{\log_{2}(e)} - 1 \right)$$

$$> \log_{2}(1/2\eta)$$
(8)

Where equation (8) comes from the definition of $J = \left\lceil 6 \log \left(\frac{1}{2\eta}\right) \right\rceil$ and the final inequality comes from the assumption that $\eta < 1/4$.

C.3 Proof of Lemma 2

Proof. For this proof, we will additionally denote $\mathbb{P}_j(\cdot) = \mathbb{P}(\cdot|\mathcal{F}_{j-1})$. From the definition of G_j in equation (6), we have that

$$\mathbb{P}_{j}(G_{j}) \geq \min\left\{\mathbb{P}_{j}\left(E_{P,j} \cap x^{*} \in (0,1) \setminus [a_{1}^{j}, a_{3}^{j})\right), \\ \mathbb{P}_{j}\left(E_{P,j}^{C} \cap E_{L,j}^{C} \cap x^{*} \in [a_{2}^{j}, a_{3}^{j})\right), \\ \mathbb{P}_{j}\left(E_{P,j}^{C} \cap E_{R,j}^{C} \cap x^{*} \in [a_{1}^{j}, a_{2}^{j})\right)\right\}$$

Now, since the events $\{x^* \in (0,1) \setminus [a_1^j, a_3^j)\}$, $\{x^* \in [a_2^j, a_3^j)\}$ and $\{x^* \in [a_1^j, a_2^j)\}$ are determined by the actions and rewards from the first j-1 phases in running the policy we have

$$\mathbb{P}_{j}(G_{j}) \geq \min\{\mathbb{P}_{j}\left(E_{P,j} \mid x^{*} \in (0,1) \setminus [a_{1}^{j}, a_{3}^{j})\right) \mathbb{1}\{x^{*} \in (0,1) \setminus [a_{1}^{j}, a_{3}^{j})\}, \\ \mathbb{P}_{j}\left(E_{P,j}^{C} \cap E_{L,j}^{C} \mid x^{*} \in [a_{2}^{j}, a_{3}^{j})\right) \mathbb{1}\{x^{*} \in [a_{2}^{j}, a_{3}^{j})\}, \\ \mathbb{P}_{j}\left(E_{P,j}^{C} \cap E_{R,j}^{C} \mid x^{*} \in [a_{1}^{j}, a_{2}^{j})\right) \mathbb{1}\{x^{*} \in [a_{1}^{j}, a_{2}^{j})\}\} \\ \geq \min\{\mathbb{P}_{j}\left(E_{P,j} \mid x^{*} \in (0,1) \setminus [a_{1}^{j}, a_{3}^{j})\right), \\ \mathbb{P}_{j}\left(E_{P,j}^{C} \cap E_{L,j}^{C} \mid x^{*} \in [a_{2}^{j}, a_{3}^{j})\right), \\ \mathbb{P}_{j}\left(E_{P,j}^{C} \cap E_{R,j}^{C} \mid x^{*} \in [a_{1}^{j}, a_{2}^{j})\right)\}$$

We can now use the following three lemmas to complete our proof.

Lemma 4. If $x^* \in (0,1) \setminus [a_1^j, a_3^j)$, then

$$\mathbb{P}_{j}\left(E_{P,j}\right) \geq 1 - 3\exp\left(-\frac{t_{j}\Delta^{2}}{30\sigma^{2}}\right)$$

Proof. Let's first assume that $x^* \in [a_3^j, 1]$ and $\mu_1 < \mu_2$. Now, if we have events H, H_1, H_2, H_3 such that $H_1 \cap H_2 \cap H_3 \subset H$, then $\mathbb{P}(H^C) \le \mathbb{P}(\{H_1 \cap H_2 \cap H_3\}^C) = \mathbb{P}(H_1^C \cup H_2^C \cup H_3^C) \le \mathbb{P}(H_1^C) + \mathbb{P}(H_2^C) + \mathbb{P}(H_3^C)$. We

can then apply this as follows for equation (10). But firstly, by definition of $E_{P,j}$ we can write the following.

$$\begin{aligned} \mathbb{P}_{j}(E_{P,j}^{C}) = \mathbb{P}_{j}\left(\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} - \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2}\right| & (9) \\ &> \frac{3}{4}\max\left(\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} - \hat{\mu}_{1,t_{j}}\right|, \left|\hat{\mu}_{0,t_{j}} - \frac{\hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{3}\right|\right)\right) \\ \leq \mathbb{P}_{j}\left(\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} - \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2}\right| > \frac{3}{4}\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} - \hat{\mu}_{1,t_{j}}\right|\right) \\ \leq \mathbb{P}_{j}\left(\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} > \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2}\right) \\ &+ \mathbb{P}_{j}\left(\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} + \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2}\right) \leq \frac{3}{4}\left(-\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} + \hat{\mu}_{1,t_{j}}}\right) \\ \leq \exp\left(-\frac{t_{j}\Delta^{2}}{8\sigma^{2}}\right) + \exp\left(-\frac{t_{j}\Delta^{2}}{24\sigma^{2}}\right) + \exp\left(-\frac{t_{j}\Delta^{2}}{30\sigma^{2}}\right) \end{aligned}$$

$$(10)$$

As required. Where the penultimate inequality (11) comes from upper bounding each of the three probabilities in equation (10). To show these bounds hold, we will bound the first probability by rewriting it as shown below.

$$\mathbb{P}_{j}\left(\frac{\hat{\mu}_{0,t_{j}}+\hat{\mu}_{a_{1}^{j},t_{j}}}{2} > \frac{\hat{\mu}_{a_{3}^{j},t_{j}}+\hat{\mu}_{1,t_{j}}}{2}\right) = \mathbb{P}_{j}\left(\hat{\mu}_{0,t_{j}}+\hat{\mu}_{a_{1}^{j},t_{j}}-\hat{\mu}_{a_{3}^{j},t_{j}}+\hat{\mu}_{1,t_{j}}<0\right) \\
= \mathbb{P}_{j}\left(\hat{\mu}_{0,t_{j}}+\hat{\mu}_{a_{1}^{j},t_{j}}-\hat{\mu}_{a_{3}^{j},t_{j}}+\hat{\mu}_{1,t_{j}}-\Delta<-\Delta\right) \\
\leq \exp\left(-\frac{t_{j}\Delta^{2}}{8\sigma^{2}}\right)$$
(13)

Where the final inequality comes from noticing that, conditioning on the action rewards from the first j-1 phases \mathcal{F}_{j-1} , the empirical means $\hat{\mu}_{0,t_j}, \hat{\mu}_{a_1^j,t_j}, \hat{\mu}_{a_3^j,t_j}, \hat{\mu}_{1,t_j}$ are independent, each σ^2/t_j -sub-Gaussian, and with respective means $\mu_1, \mu_1, \mu_1, \mu_2$. Hence, the sum $\hat{\mu}_{0,t_j} + \hat{\mu}_{a_1^j,t_j} - \hat{\mu}_{a_3^j,t_j} + \hat{\mu}_{1,t_j} - \Delta$ is mean-zero, $4\sigma^2/t_j$ -sub-Gaussian. Plugging this into the Hoeffding inequality from Proposition 2.5 in Wainwright (2019), we attain the final inequality above. For the final two probabilities in equation (10) we can similarly calculate the upper bounds as following to attain equation (11).

$$\mathbb{P}_{j}\left(\frac{\hat{\mu}_{0,t_{j}}+\hat{\mu}_{a_{1}^{j},t_{j}}+\hat{\mu}_{a_{3}^{j},t_{j}}}{3}>\hat{\mu}_{1,t_{j}}\right) \leq \exp\left(-\frac{t_{j}\Delta^{2}}{24\sigma^{2}}\right)$$
$$\mathbb{P}_{j}\left(-\frac{\hat{\mu}_{0,t_{j}}+\hat{\mu}_{a_{1}^{j},t_{j}}}{2}+\frac{\hat{\mu}_{a_{3}^{j},t_{j}}+\hat{\mu}_{1,t_{j}}}{2}>\frac{3}{4}\left(-\frac{\hat{\mu}_{0,t_{j}}+\hat{\mu}_{a_{1}^{j},t_{j}}+\hat{\mu}_{a_{3}^{j},t_{j}}}{3}+\hat{\mu}_{1,t_{j}}\right)\right) \leq \exp\left(-\frac{t_{j}\Delta^{2}}{30\sigma^{2}}\right)$$

Hence, returning to equation (12) we can achieve the following final bound as required for the lemma.

$$\implies \mathbb{P}_j(E_{P,j}) > 1 - 3\exp\left(-\frac{t_j\Delta^2}{30\sigma^2}\right) \tag{14}$$

If we instead assume that $x^* \in [0, a_1^j)$ then almost identical arguments would lead us to the same final equation (14). Furthermore if we were to assume that $\mu_1 > \mu_2$ then we could again use almost identical arguments in both cases to attain (14).

Lemma 5. If $x^* \in [a_2^j, a_3^j)$, then

$$\mathbb{P}_j\left(E_{P,j}^C \cap E_{L,j}^C\right) \ge 1 - 8\exp\left(-\frac{t_j\Delta^2}{30\sigma^2}\right)$$

Proof. Note that by a union bound we have

$$\mathbb{P}_j(\{E_{P,j}^C \cap E_{L,j}^C\}^C) = \mathbb{P}_j(E_{P,j} \cup E_{L,j}) \le \mathbb{P}_j(E_{P,j}) + \mathbb{P}_j(E_{L,j})$$
(15)

So we will focus on bounding the two probabilities on the right of the above equation. Again, we will first assume that $\mu_1 < \mu_2$.

Starting with $\mathbb{P}_{j}(E_{P,j})$, we can use a union bound to get the following.

$$\mathbb{P}_{j}(E_{P,j}) = \mathbb{P}_{j}\left(\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} - \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2}\right|$$

$$< \frac{3}{4}\max\left(\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} - \hat{\mu}_{1,t_{j}}\right|, \left|\hat{\mu}_{0,t_{j}} - \frac{\hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{3}\right|\right)\right)$$

$$\leq \mathbb{P}_{j}\left(\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} - \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2}\right| < \frac{3}{4}\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} - \hat{\mu}_{1,t_{j}}}\right|\right)$$

$$+ \mathbb{P}_{j}\left(\left|\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} - \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2}\right| < \frac{3}{4}\left|\hat{\mu}_{0,t_{j}} - \frac{\hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{3}\right|\right)$$

$$(17)$$

We can then denote the final two probabilities in equation (17) as (A') and (B') respectively. We can firstly bound (A') using a similar idea to equation (10). Namely, if we have events H, H_1, H_2, H_3 such that $H_1 \cap H_2 \cap H_3 \subset H$, then $\mathbb{P}(H^C) \leq \mathbb{P}(\{H_1 \cap H_2 \cap H_3\}^C) = \mathbb{P}(H_1^C \cup H_2^C \cup H_3^C) \leq \mathbb{P}(H_1^C) + \mathbb{P}(H_2^C) + \mathbb{P}(H_3^C)$. We use this to get the below equation.

$$\begin{aligned} (A') \leq \mathbb{P}_{j} \left(\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} > \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2} \right) \\ &+ \mathbb{P}_{j} \left(\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} > \hat{\mu}_{1,t_{j}} \right) \\ &+ \mathbb{P}_{j} \left(-\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{1}^{j},t_{j}}}{2} + \frac{\hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{1,t_{j}}}{2} < \frac{3}{4} \left(-\frac{\hat{\mu}_{0,t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}} + \hat{\mu}_{a_{3}^{j},t_{j}}}{3} + \hat{\mu}_{1,t_{j}} \right) \right) \\ \leq 3 \exp \left(-\frac{t_{j} \Delta^{2}}{24\sigma^{2}} \right) \end{aligned}$$
(18)

Where the final line again comes from using the Hoeffding inequality for each of the probabilities in (18), similar to the previous lemma for equation (13).

Furthermore, we can similarly show that $(B') \leq 3 \exp\left(-\frac{t_j \Delta^2}{24\sigma^2}\right)$. Hence, plugging our bounds for (A') and (B') into equation (17) we have the following.

$$\implies \mathbb{P}_j(E_{P,j}) \le 6 \exp\left(-\frac{t_j \Delta^2}{24\sigma^2}\right) \tag{19}$$

Now let's consider $\mathbb{P}_{j}(E_{L,j})$. By noting that $\{\hat{\mu}_{a_{1}^{j},t_{j}} < \hat{\mu}_{a_{3}^{j},t_{j}}\} \cap \{\hat{\mu}_{a_{2}^{j},t_{j}} - \hat{\mu}_{a_{1}^{j},t_{j}} < \hat{\mu}_{a_{3}^{j},t_{j}} - \hat{\mu}_{a_{2}^{j},t_{j}}\} \implies \{|\hat{\mu}_{a_{2}^{j},t_{j}} - \hat{\mu}_{a_{3}^{j},t_{j}} - \hat{\mu}_{a_{3}^{j},t_{j}}|\}, \text{ we can bound the following probability.}$

$$\mathbb{P}_{j}(E_{L,j}) = \mathbb{P}(\{|\hat{\mu}_{a_{1}^{j},t_{j}} - \hat{\mu}_{a_{2}^{j},t_{j}}| < |\hat{\mu}_{a_{2}^{j},t_{j}} - \hat{\mu}_{a_{3}^{j},t_{j}}|\}^{C}) \\
\leq \mathbb{P}\left(\left\{\{\hat{\mu}_{a_{1}^{j},t_{j}} < \hat{\mu}_{a_{3}^{j},t_{j}}\} \cap \{\hat{\mu}_{a_{2}^{j},t_{j}} - \hat{\mu}_{a_{1}^{j},t_{j}} < \hat{\mu}_{a_{3}^{j},t_{j}} - \hat{\mu}_{a_{2}^{j},t_{j}}\}\right\}^{C}\right) \\
\leq \mathbb{P}(\{\hat{\mu}_{a_{1}^{j},t_{j}} < \hat{\mu}_{a_{3}^{j},t_{j}}\}^{C}) + \mathbb{P}(\{\hat{\mu}_{a_{2}^{j},t_{j}} - \hat{\mu}_{a_{1}^{j},t_{j}} < \hat{\mu}_{a_{3}^{j},t_{j}} - \hat{\mu}_{a_{2}^{j},t_{j}}\}^{C}) \\
\leq \exp\left(-\frac{t_{j}\Delta^{2}}{4\sigma^{2}}\right) + \exp\left(-\frac{t_{j}\Delta^{2}}{12\sigma^{2}}\right) \\
\leq 2\exp\left(-\frac{t_{j}\Delta^{2}}{12\sigma^{2}}\right) \tag{20}$$

Where the penultimate inequality again holds using Hoeffding inequality as in equation (13). Hence, plugging this and equation (19) into (15) give us the following.

$$\mathbb{P}_{j}(\{E_{P,j}^{C} \cap E_{L,j}^{C}\}^{C}) \leq \mathbb{P}_{j}(E_{P,j}) + \mathbb{P}_{j}(E_{L,j})$$
$$\leq 8 \exp\left(-\frac{t_{j}\Delta^{2}}{30\sigma^{2}}\right)$$

As required. We note that one could use an almost identical method to attain the same bound if $\mu_1 > \mu_2$.

Lemma 6. If $x^* \in [a_1^j, a_2^j)$, then

$$\mathbb{P}_j\left(E_{P,j}^C \cap E_{R,j}^C\right) \ge 1 - 8\exp\left(-\frac{t_j\Delta^2}{30\sigma^2}\right)$$

Proof. Similar to Lemma 5

Hence, regardless of the position of the change point, we get the required bound for the probability of zooming in/out towards the correct direction, given the actions and rewards in previous phases, concluding our proof of Lemma 2.

C.4 Proof of Lemma 3

Proof. Again, the structure of this proof will be very similar to the proofs seen in Cheshire et al. (2021) since the structure of the policies is very similar. Of course, the probability of a good event in each round will be different as well as constants throughout.

$$p_j := \mathbb{P}_j(G_j^C)$$
$$p_0 := 8 \exp\left(-\frac{t_j \Delta^2}{30\sigma^2}\right)$$

We can furthermore assume that

$$\Delta \ge 88\sqrt{\frac{\sigma^2 \log(1/2\eta)}{T}} \tag{21}$$

Otherwise, if we were to assume the contrary, then plugging this into the exponential term in Theorem 1, the result from Theorem 1 would trivially be true. Now, using Lemma 2 and substituting this in our definition of p_0

we have

$$p_j \le p_0 = 8 \exp\left(-\frac{t_j \Delta^2}{30\sigma^2}\right)$$
$$\le 8 \exp\left(-88^2 \log(1/2\eta) \frac{t_j}{30T}\right)$$
(22)

$$\leq 8 \exp\left(-88^2 \log(1/2\eta) \frac{T}{10J} \frac{1}{30T}\right) \tag{23}$$

$$= 8 \exp\left(-88^{2} \log(1/2\eta) \frac{1}{300J}\right)$$

$$\leq 8 \exp\left(-88^{2} \log(1/2\eta) - \frac{1}{1}\right)$$
(24)

$$\leq 8 \exp\left(-88^2 \log(1/2\eta) \frac{1}{300(6 \log(1/2\eta) + 1)}\right)$$
(24)

$$\leq \frac{1}{4} \tag{25}$$

Where equation (22) come from substituting assumption on Δ in (21). Equation (23) comes from using the definition of $t_j = \lfloor \frac{T}{5J} \rfloor \geq \frac{T}{10J}$ (since we assume that $\frac{T}{5J} \geq 1$ for the policy to have enough samples to run). Equation (24) comes from the definition of $J = \lceil 6 \log(1/2\eta) \rceil \leq 6 \log(1/2\eta) + 1$. Then, the final inequality holds true whenever $\eta < 1/4$.

Now, the quantity of interest in this lemma is the probability that we fail in more than 1/4 of the phases. We can bound this probability from above using Markov's inequality and any $\lambda \ge 0$.

$$\mathbb{P}\left(\sum_{j=1}^{J} \mathbb{1}\{G_j^C\} \ge J/4\right) \le \frac{\mathbb{E}\left[\exp\left(\lambda \sum_{j=1}^{J} \mathbb{1}\{G_j^C\}\right)\right]}{\exp(\lambda \cdot \frac{J}{4})}$$
(26)

Now, in order to bound the expectation on the right hand side of the above equation as in Cheshire et al. (2021), we will introduce the following function $\phi_p(\lambda) = \log(1 - p + pe^{\lambda})$ which is non-decreasing in p. Hence, since $p_j \leq p_0$ from Lemma 2, we have $\phi_{p_t}(\lambda) \leq \phi_{p_0}(\lambda)$. Hence we have the following, starting by using the tower rule.

$$\mathbb{E}\left[\exp\left(\lambda\sum_{j=1}^{J}\mathbb{1}\left\{G_{j}^{C}\right\}\right)\right] = \mathbb{E}\left[\mathbb{E}_{J}\left[\exp\left(\lambda\mathbb{1}\left\{G_{J}^{C}\right\}\right)\right]\exp\left(\lambda\sum_{j=1}^{J-1}\mathbb{1}\left\{G_{j}^{C}\right\}\right)\right]$$
$$= \mathbb{E}\left[\exp\left(\phi_{p_{J}}(\lambda)\right)\exp\left(\lambda\sum_{j=1}^{J-1}\mathbb{1}\left\{G_{j}^{C}\right\}\right)\right]$$
$$(27)$$

$$\leq \mathbb{E} \left[\exp\left(\phi_{p_0}(\lambda)\right) \exp\left(\lambda \sum_{j=1}^{J-1} \mathbb{1}\left\{G_j^C\right\}\right) \right]$$

$$\leq \mathbb{E} \left[\exp\left(J \cdot \phi_{p_0}(\lambda)\right) \right]$$
(28)

Where equation (27) comes from the definition $p_J = \mathbb{P}_J(G_J^C)$ and hence that $\mathbb{E}_J\left[\exp\left(\lambda \mathbb{1}\{G_J^C\}\right)\right] = \exp(\lambda)p_J + 1 - p_J = \exp(\phi_{p_J}(\lambda))$. Equation (28) comes from the monotonicity of ϕ and we can get the final inequality by simply repeating these steps.

We can then substitute this bound in expectation back into the inequality from (26), remembering that we can choose any $\lambda \geq 0$

$$\mathbb{P}\left(\sum_{j=1}^{J} \mathbb{1}\{G_{j}^{C}\}\right) \leq \exp\left(-J \cdot \sup_{\lambda \geq 0} \left(\lambda/4 - \phi_{p_{0}}(\lambda)\right)\right) \\
= \exp\left(-J \cdot kl(1/4, p_{0})\right) \tag{29}$$

$$\leq \exp\left(-J\frac{1}{4}\log(\frac{1}{p_0}) + J\log(2)\right)$$

$$\leq \exp\left(-\frac{Jt_j}{120}\frac{\Delta^2}{\sigma^2} + J\left(\log(2) + \frac{\log(8)}{4}\right)\right)$$

$$\leq \exp\left(-\frac{T\Delta^2}{600\sigma^2} + 13\log(1/2\eta)\right)$$
(30)

As required. Where equation (29) comes from noting that $\sup_{\lambda \ge 0} \lambda q - \phi_p(\lambda) = kl(q, p)$ when $q \ge p$, where we denote kl(p,q) as the divergence between two Bernoulli distributions of parameters p and q. We can then use this and the fact that $p_0 \le 1/4$ from (25) to attain equation (29). Then we can attain equation (30) from the trick $kl(a,b) \ge a \log(1/b) - \log(2)$.

	_	_	
- 1			

D PROOFS FOR LARGE BUDGET LOWER BOUND

D.1 Overview

Before going into technical details, we provide a brief informal overview of the techniques we use to prove the lower bound in Theorem 2. In particular, we also discuss how η appears through our analysis, unlike previous works (see Section 4.2 for discussion). We first simplify the problem and assume that the two means (μ_1, μ_2) are known and that the change point is at one of $K = \lfloor 1/2\eta \rfloor$ different positions, $\{x_j^*\}_j$ indexed from smallest to largest, each at least 2η away from each other. Thus our objective is to identify which of these K prospective environments $V' = \{v_j\}_j$ we are in. The proof then uses the following three ingredients.

- (i) Conditions on the policy: The piecewise constant structure makes it difficult to isolate the effect of the expected number of plays in each region of \mathcal{A} with respect to the failure probability (whereas this is possible for each arm in Best Arm Identification, see Carpentier and Locatelli (2016) equation (7), and in unstructured Thresholding Bandits, see Locatelli et al. (2016) appendix A.1). Indeed since the two means are known, in order to distinguish v_i from all other environments in V', it is sufficient to sample exclusively on either side adjacent to the change point of v_i , x_i^* . However, this is infeasible in practice as the learner will not a priori know where the changepoint x_i^* occurs. To enforce some exploration across the action space, we begin by making a reasonable assumption that the failure probability of our policy is at most C' = 1/2 regardless of the environment (See Theorem 6 statement). Note in Appendix D.3 we show that we can omit this condition when the budget is sufficiently large.
- (ii) Flat environment v_0 : We consider a flat environment v_0 with no change in mean. Here there will always be a region in which π estimates the change point with probability less than 1/(K-1) (Lemma 7). We then choose a challenging reference environment $v_i \in V'$ to have a change point within this low probability region.
- (iii) Change of measure and transportation lemma: Using a change of measure from v_K to v_i , we can relate the number of samples played to the right of the change point in environment v_i , denoted $T_{(x_i^*,1]}$, with the failure probability (see equations (38) - (42)). For a policy to be able to identify a change in mean with high probability, we require a sufficient number of samples to the right and left of the change (missing this is why some existing policies can suffer when the change is near the boundary, see Section 7). Therefore, using (i), we can upper bound $T_{(x_i^*,1]}$. This is firstly formalised with a change of measure from v_i to v_0 . Then we use a transportation lemma (see Lemma 10 proof) combined with the 1/(K-1) probability in (ii) and C' from (i) (see Lemma 10 and equations (43)-(44)) to produce our lower bound. Note that the definition of K here involves η .

D.2 Lower Bound for Reasonable Policies

In line with what is outlined above, we begin by considering a lower bound on any policy which explores sufficiently. In Appendix D.3 we will show how this assumption can be removed to get our final lower bound.

Theorem 6. Let $\overline{V} \subset V(\Delta, \sigma)$ be the set of environments with change in mean Δ and Gaussian random noise with variance σ^2 . Let $C' \in (0, 1)$ and let $\overline{\Pi} := \{\pi \in \Pi : \forall v \in \overline{V}, \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| < \eta) \ge C'\} \subset \Pi$. Then, denoting x_v^* as the change point in environment v, we have

$$\inf_{\pi \in \bar{\Pi}} \sup_{v \in \bar{V}} \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| \ge \eta) \ge \frac{1}{8} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + C' \log\left(\left\lfloor \frac{1}{2\eta} \right\rfloor - 1\right)\right).$$

Proof. Fix some $\pi \in \Pi$. Let $\bar{V}' := \{v_1, ..., v_K\}$ with $K := \lfloor 1/2\eta \rfloor$ be a finite subset of \bar{V} in which all environments have mean rewards μ_1, μ_2 and noise variance σ^2 . Furthermore, let environment v_j have change point $x_j^* := 2\eta(j-1)$. Finally we will drop the subscript for π as we will fix this policy π for the remained or the proof, unless otherwise stated (i.e. denote $\mathbb{P}_{v_j} = \mathbb{P}_{v_j,\pi}$). Then we have

$$\sup_{v\in\bar{V}} \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| \ge \eta) \ge \sup_{v_j\in\bar{V}'} \mathbb{P}_{v_j}(|\hat{x}_T - x_j^*| \ge \eta)$$
(31)

Now, let v_0 be an environment in which there is no change in mean across the space. Namely the mean reward function has constant value μ_2 . We will now introduce several helpful lemmas before continuing with equation (31).

Lemma 7. There exists an $i \in \{1, ..., K-1\}$ such that

$$\mathbb{P}_{v_0}(|\hat{x}_T - x_i^*| < \eta) \le \frac{1}{K - 1}$$

Proof. Suppose, for contradiction, that $\forall i \in \{1, .., K-1\}$

$$\mathbb{P}_{v_0}(|\hat{x}_T - x_i^*| < \eta) > \frac{1}{K - 1}$$
(32)

$$\implies 1 \ge \sum_{i=1}^{K-1} \mathbb{P}_{v_0}(|\hat{x}_T - x_i^*| < \eta) > \sum_{i=1}^{K-1} \frac{1}{K-1} = 1$$
(33)

Which is a contradiction, hence the proof is complete. Note the final inequality on the left holds because the sets $\{(x_i^* - \eta, x_i^* + \eta)\}_{i=1}^{K-1}$ are disjoint and \mathbb{P}_{v_0} is a probability measure.

Hence, we set i to be the same i as in the above Lemma for the remainder of the proof.

Lemma 8. Let T_B be the number of times we play an action in the set $B \subseteq [0,1]$ over the whole budget T. Then, with i chosen to satisfy Lemma 7,

$$KL(\mathbb{P}_{v_i}, \mathbb{P}_{v_0}) \le \frac{\Delta^2}{2\sigma^2} \mathbb{E}_{v_i}\left[T_{[0, x_i^*]}\right].$$

Proof. Denote P_{i,A_t} and P'_{0,A_t} as the reward distributions of the action played in round t in both the environment v_i and v_0 , respectively. Then, from Ex 15.8 (Lattimore and Szepesvári, 2020), we have the following equation.

$$KL(\mathbb{P}_{v_{i}}, \mathbb{P}_{v_{0}}) = \mathbb{E}_{v_{i}} \left[\sum_{t=1}^{T} KL(P_{i,A_{t}}, P_{0,A_{t}}) \right]$$
$$= \mathbb{E}_{v_{i}} \left[\sum_{t=1}^{T} KL(P_{i,A_{t}}, P_{0,A_{t}}) \cdot (\mathbb{1}\{A_{t} \in [0, x_{i}^{*}]\} + \mathbb{1}\{A_{t} \in (x_{i}^{*}, 1]\}) \right]$$
$$= \mathbb{E}_{v_{i}} \left[\sum_{t=1}^{T} KL(P_{i,A_{t}}, P_{0,A_{t}}) \mathbb{1}\{A_{t} \in [0, x_{i}^{*}]\} \right]$$
$$= \mathbb{E}_{v_{i}} \left[\sum_{t=1}^{T} \frac{\Delta^{2}}{2\sigma^{2}} \mathbb{1}\{A_{t} \in [0, x_{i}^{*}]\} \right]$$
(34)

$$=rac{\Delta^2}{2\sigma^2}\mathbb{E}_{v_i}\left[T_{[0,x^*_i]}
ight]$$

As required. Where equation (34) comes from noting that for $A_t \in (x_i^*, 1]$, we have $P_{i,A_t} = P_{0,A_t}$, hence $KL(P_{i,A_t}, P_{0,A_t}) = 0$. Furthermore (35) comes from noting that for $A_t \in [0, x_i^*]$ we are comparing P_{i,A_t}, P_{0,A_t} which are two Gaussian distributions with difference in mean Δ and same variance σ^2 .

Lemma 9. Under the same setup as Lemma 8,

$$KL(\mathbb{P}_{v_i}, \mathbb{P}_{v_K}) \leq \mathbb{E}_{v_i}[T_{(x_i^*, x_K^*)}] \frac{\Delta^2}{2\sigma^2}.$$

Proof. Similar to Lemma 8, noting instead that the reward distributions in environments v_i, v_K only differ in the region in the action space $(x_i^*, x_K^*]$.

Lemma 10.

$$KL(\mathbb{P}_{v_i}, \mathbb{P}_{v_0}) \ge C' \log\left((K-1)\right) - \log(2)$$

Proof. Firstly, we have from Lemma 1 in Garivier et al. (2016), that the following holds for any measurable function Z which maps to [0, 1].

$$KL(\mathbb{P}_{v_i}, \mathbb{P}_{v_0}) \ge kl(\mathbb{E}_{v_i}(Z), \mathbb{E}_{v_0}(Z))$$

Now, if we choose $Z = \mathbb{1}\{E\}$ and let the event $E := \{|\hat{x}_T - x_i^*| < \eta\}$, then this becomes the below equation.

$$KL(\mathbb{P}_{v_i}, \mathbb{P}_{v_0}) \ge kl(\mathbb{P}_{v_i}(E), \mathbb{P}_{v_0}(E))$$
$$\ge \mathbb{P}_{v_i}(E) \log\left(\frac{1}{\mathbb{P}_{v_0}(E)}\right) - \log(2)$$
(36)

$$\geq C' \log\left(K - 1\right) - \log(2) \tag{37}$$

Where equation (36) comes from using a trick to bound $kl(a,b) \ge a \log(\frac{1}{b}) - \log(2)$. The the final inequality holds by Lemma 7 that $\mathbb{P}_{v_0}(E) \le \frac{1}{K-1}$ and using the assumption made for π in Theorem 6 that $\forall v \in \overline{V}$, $\mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| < \eta) \ge C'$.

Now, we put Lemmas 7, 8, 9, and 10 together and return to equation (31).

$$\sup_{v \in \bar{V}} \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| \ge \eta) \ge \sup_{v_j \in \bar{V}'} \mathbb{P}_{v_j}(|\hat{x}_T - x_j^*| \ge \eta)$$
$$\ge \frac{1}{2} \mathbb{P}_{v_i}(|\hat{x}_T - x_i^*| \ge \eta) + \frac{1}{2} \mathbb{P}_{v_K}(|\hat{x}_T - x_K^*| \ge \eta)$$
(38)

$$\geq \frac{1}{2} \mathbb{P}_{v_i}(|\hat{x}_T - x_i^*| \geq \eta) + \frac{1}{2} \mathbb{P}_{v_K}(|\hat{x}_T - x_i^*| < \eta)$$
(39)

$$\geq \frac{1}{4} \exp\left(-KL(\mathbb{P}_{v_i}, \mathbb{P}_{v_K})\right) \tag{40}$$

$$\geq \frac{1}{4} \exp\left(-\mathbb{E}_{v_i}[T_{(x_i^*, x_K^*)}]\frac{\Delta^2}{2\sigma^2}\right) \tag{41}$$

$$\geq \frac{1}{4} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + \frac{\Delta^2}{2\sigma^2}\left(\mathbb{E}_{v_i}[T_{[0,x_i^*]}] + \mathbb{E}_{v_i}[T_{(x_K^*,1]}]\right)\right)$$
(42)

$$\geq \frac{1}{4} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + \frac{\Delta^2}{2\sigma^2}\mathbb{E}_{v_i}[T_{[0,x_i^*]}]\right)$$

$$\tag{43}$$

$$\geq \frac{1}{4} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + KL(\mathbb{P}_{v_i}, \mathbb{P}_{v_0})\right)$$

$$\geq \frac{1}{4} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + C'\log\left(K-1\right) - \log(2)\right)$$
(44)

As required. Inequality (38) holds as we are simply taking the average of two point in the set over which the supremum is acting. Inequality (40) holds since $|x_K^* - x_i^*| \ge 2\eta$ and therefore $\{|\hat{x}_T - x_i^*| < \eta\} \subset \{|\hat{x}_T - x_K^*| \ge \eta\}$. Inequality (40) comes from the Bretagnolle-Huber inequality (Lattimore and Szepesvári, 2020) Theorem 14.2. Inequality (41) comes from Lemma 9. Equation (42) is true since $T = T_{[0,x_i^*]} + T_{(x_i^*,x_K^*]} + T_{(x_{K}^*,1]}$. Inequality (43) holds since $T_{(x_{K}^*,1]} \ge 0$. Then the final two inequalities hold from Lemmas 8 and 10, respectively.

D.3 Proof of Theorem 2

We can extend Theorem 6 to the set of any policies in Theorem 7 for a sufficiently large budget $T \geq \frac{\sigma^2}{\Delta^2} \log \left(\frac{1}{2}(\lfloor 1/2\eta \rfloor - 1)\right)$, which is stated below. Theorem 2 can be seen as a consequence of Theorem 7 below since we have that

$$\frac{\sigma^2}{\Delta^2} \log\left(\frac{1}{2}(\lfloor 1/2\eta \rfloor - 1)\right) \le \frac{\sigma^2}{\Delta^2} (1.59 \log(\lfloor \frac{1}{2\eta} \rfloor) - 2 \log(2)).$$

Therefore, since Theorem 7 holds for $T \geq \frac{\sigma^2}{\Delta^2} \log \left(\frac{1}{2} (\lfloor 1/2\eta \rfloor - 1) \right)$, it will also hold for $T \geq \frac{\sigma^2}{\Delta^2} (1.59 \log(\lfloor \frac{1}{2\eta} \rfloor) - 2 \log(2))$, as required.

Theorem 7. Let $\bar{V} \subset V(\Delta, \sigma)$ be the set of environments with change in mean Δ and Gaussian random noise with variance σ^2 . Then, for $T \geq \frac{\sigma^2}{\Delta^2} \log \left(\frac{1}{2} (\lfloor 1/2\eta \rfloor - 1) \right)$, we have

$$\inf_{\pi \in \Pi} \sup_{v \in \bar{V}} \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| \ge \eta) \ge \frac{1}{8} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + \frac{1}{2}\log\left(\frac{1}{2}\left(\left\lfloor\frac{1}{2\eta}\right\rfloor - 1\right)\right)\right).$$

Proof. Fix some $\pi \in \Pi$. Let's then consider two cases.

First Case: Suppose that

$$\forall v \in V, \quad \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| < \eta) \ge 1/2$$

Then, from Theorem 6, under this assumption, we have

$$\sup_{v\in\bar{V}} \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| \ge \eta) \ge \frac{1}{8} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + \frac{1}{2}\log\left(\frac{1}{2}(\lfloor 1/2\eta \rfloor - 1)\right)\right),$$

as required.

Second Case: Now, suppose instead that

$$\exists v \in V, \quad \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| < \eta) < 1/2.$$

We can equivalently write that

$$\exists v \in V, \quad \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| \ge \eta) \ge 1/2$$

And if this is the case, then we have the following

$$\sup_{v \in \bar{V}} \mathbb{P}_{v,\pi}(|\hat{x}_T - x_v^*| \ge \eta) \ge \frac{1}{2}$$
$$\ge \frac{1}{8} \exp\left(-\frac{\Delta^2}{2\sigma^2}T + \frac{1}{2}\log\left(\frac{1}{2}(\lfloor 1/2\eta \rfloor - 1)\right)\right)$$

As required. Where the final inequality holds since we are assuming that $T \ge \frac{\sigma^2}{\Delta^2} \log \left(\frac{1}{2}(\lfloor 1/2\eta \rfloor - 1)\right)$.

PROOFS FOR SMALL BUDGETS

E PROOFS FOR SMALL BUDGET LOWER BOUND

E.1 Proof of Theorem 4

Lemma 11. Let Π be the set of policies with fixed budget T. Let $\overline{V} \subset V(\Delta, \sigma)$ be the set of environments with change in mean Δ and Gaussian random noise with variance σ^2 . Denote x_v^* as the change point in environment v. Then it holds that,

$$\inf_{\pi \in \Pi} \sup_{v \in \bar{V}} \mathbb{P}_{\pi,v}(|\hat{x}_T - x_v^*| > \eta) \ge 1 - \frac{\Delta^2 T / 2\sigma^2 + \log(2)}{\log(\lfloor \frac{1}{2\eta} \rfloor)}.$$

Proof. This proof follows closely to that demonstrated in Wainwright (2019). Also, as it is oftentimes less obvious in this section, we will include all subscripts to denote the policy in use and environment in question. Additionally, we add superscript for the random variable of the final estimate for the change point by policy π in environment v as $\hat{x}_{T}^{\pi,v}$.

Let $\eta \leq 0.5$ and fix some arbitrary policy $\pi \in \Pi$. We choose M environments $v_j \in \overline{V}$ such that the 2η neighborhoods around the change point in each environment, $x_{v_j}^*$, are pairwise disjoint. We can choose the covering number, M, to be at least $M = \lfloor 1/2\eta \rfloor$. Let $J \sim \text{Uniform}\{1, ..., M\}$. We can then show the following inequality, since we are taking a supremum on the left, which is greater than an average across any subset.

$$\sup_{v \in V} \mathbb{P}_{\pi,v}(|\hat{x}_T^{\pi,v} - x_v^*| > \eta) \ge \frac{1}{M} \sum_{j=1}^M \mathbb{P}_{\pi,v_j}(|\hat{x}_T^{\pi,v_j} - x_{v_j}^*| > \eta)$$
$$= \mathbb{P}_{J,\pi,v_J}(|\hat{x}_T^{\pi,v_J} - x_{v_J}^*| > \eta)$$
(45)

Where we denote \mathbb{P}_{J,π,v_J} as the joint measure between $J \sim \text{Uniform}\{1, ..., M\}$ and \mathbb{P}_{π,v_J} . Now,given our choice of v_j 's and by defining the test $\psi(\hat{x}_T^{\pi,v_J}) = \operatorname{argmin}_{i \in \{1,...,M\}} |\hat{x}_T^{\pi,v_J} - x_{v_i}^*|$, we can then lower bound (45) with

$$\sup_{v \in V} \mathbb{P}_{\pi,v}(|\hat{x}_T^{\pi,v} - x_v^*| > \eta) \ge \mathbb{P}_{J,\pi,v_J}(\psi(\hat{x}_T^{\pi,v_J}) \neq J)$$
$$\ge 1 - \frac{\frac{1}{M^2} \sum_{j,k=1}^M KL(\mathbb{P}_{\pi,v_j}, \mathbb{P}_{\pi,v_k}) + \log(2)}{\log(M)}$$
(46)

$$\geq 1 - \frac{\left(\frac{\Delta^2 T}{2\sigma^2}\right) + \log(2)}{\log(M)} \tag{47}$$

Where equation (46) holds from using Fano's inequality and the loose bound for mutual information shown in equations 15.31 and 15.34 in Wainwright (2019) respectively. Then equation (47) holds from noting that, for any two environments in the sum in (47), the biggest difference in mean between the reward distributions anywhere in the action space is at most Δ . Hence, using a similar idea to Lemma 8 and 9, we can use the divergence decomposition from Lattimore and Szepesvári (2020) to show that for any pair $i \neq k$ we have $KL(\mathbb{P}_{\pi,v_j},\mathbb{P}_{\pi,v_k}) \leq \frac{\Delta^2}{2\sigma^2}\mathbb{E}_{\pi,v_j}[T] = \frac{\Delta^2}{2\sigma^2}T$ in our fixed-budget setting.

The proof is completed by noticing that for a fixed value of η , we can choose a covering of [0, 1] with covering number at least $M = \lfloor 1/2\eta \rfloor$.

We then simply extend Lemma 11 by using the fact that $e^{-x} \le 1 - x/2$ when we have that $0 \le x \le 1.59$. Doing so gives us the required bound for Theorem 4

F PROOFS FOR SMALL BUDGET UPPER BOUND

F.1 Proof for Theorem 3

We can reuse a lot of the calculations we performed when analysing the backtracking algorithm, since we have defined the left and right elimination criteria in the same way. The different this time, is that we have to bound the probability that a good even occurs in every phase and we eliminate the correct half of the space. We define the good event in this case, event G'_i , as

$$G'_{j} = \left\{ E^{C}_{L,j} \cap \left\{ x^{*} \in [a^{j}_{2}, a^{j}_{3}) \right\} \right\} \cup \left\{ E^{C}_{R,j} \cap \left\{ x^{*} \in [a^{j}_{1}, a^{j}_{2}) \right\} \right\}.$$

First, note that using the previous Appendix C, we can attain the following upper bound on related events.

Lemma 12. Given the rewards and actions from previous phases and under Algorithm 3; the probability that we fail to eliminate the correct half of the action space when it is actually on the right, and respectively left, can be upper bounded with

$$\mathbb{P}_{j}\left(E_{L,j}\left|x^{*}\in[a_{2}^{j},a_{3}^{j})\right\right) \leq 2\exp\left(-\frac{t_{j}\Delta^{2}}{12\sigma^{2}}\right),\\ \mathbb{P}_{j}\left(E_{R,j}\left|x^{*}\in[a_{2}^{j},a_{3}^{j})\right\right) \leq 2\exp\left(-\frac{t_{j}\Delta^{2}}{12\sigma^{2}}\right).$$

Proof. As mentioned before, for the first equation, we can directly use the calculation in equation (20) and use an almost identical method for the second inequality. \Box

Then, under Algorithm 3, we can upper bound the probability that event G'_j fails to occur, given the actions and rewards from previous phases.

Lemma 13. Under Algorithm 3

$$\mathbb{P}_j(G'_j) \ge 1 - 2\exp\left(-\frac{t_j\Delta^2}{12\sigma^2}\right)$$

Proof. We can lower bound $\mathbb{P}_j(G'_j)$ with

$$\begin{split} \mathbb{P}_{j}(G'_{j}) &= \mathbb{P}_{j}\left(\left\{E^{C}_{L,j} \cap \left\{x^{*} \in [a^{j}_{2}, a^{j}_{3})\right\}\right\} \cup \left\{E^{C}_{R,j} \cap \left\{x^{*} \in [a^{j}_{1}, a^{j}_{2})\right\}\right\}\right) \\ &\geq \min\left(\mathbb{P}_{j}\left(E^{C}_{L,j} \cap \left\{x^{*} \in [a^{j}_{2}, a^{j}_{3})\right\}\right), \mathbb{P}_{j}\left(E^{C}_{R,j} \cap \left\{x^{*} \in [a^{j}_{1}, a^{j}_{2})\right\}\right)\right) \\ &\geq \min\left(\mathbb{P}_{j}\left(E^{C}_{L,j} \left|x^{*} \in [a^{j}_{2}, a^{j}_{3})\right\right), \mathbb{P}_{j}\left(E^{C}_{R,j} \left|x^{*} \in [a^{j}_{1}, a^{j}_{2})\right\right)\right) \\ &\geq 1 - 2\exp\left(-\frac{t_{j}\Delta^{2}}{12\sigma^{2}}\right) \end{split}$$

As required. Where the final line comes from using Lemma 12.

We know that, under the event that G'_j occurs in every phase $j \in \{1, ..., J\}$, our objective $|\hat{x}_T - x^*| < \eta$ is attained. Hence, we can take a union bound over all the phases, combined with Lemma 13, we attain the required bound for Theorem 3.

G CHANGE POINTS AND INTUITION FOR ELIMINATION CRITERIA

G.1 Estimation in Offline Change Point Analysis

Suppose we have a sequence of sample and observation pairs $\{x'_i, y'_i\}_{i=1}^n$, ordered such that we have $x'_1 \leq x'_2 \leq \dots \leq x'_n$. We also denote $\bar{y}_{i:j}$ as the empirical mean of observations y'_i, \dots, y'_j ,

$$\bar{y}_{i:j} = \sum_{k=i}^{j} y'_k.$$

Then, suppose we want to fit a piecewise constant model with two means μ_1 and μ_2 which change at some index r in the sequence of observations $\{y'_i\}_{i=1}^n$, namely $f(x) = \mu_1 \mathbb{I}\{x \le x'_r\} + \mu_2 \mathbb{I}\{x > x'_r\}$. To do so, we can write the residual sum of squares from fitting the model as

$$RSS(r) = \sum_{i=1}^{r} (y'_i - \hat{y}_1)^2 + \sum_{i=r+1}^{n} (y'_i - \hat{y}_2)^2.$$

From this, it is quick to check that the least squares estimators for the two means \hat{y}_1 and \hat{y}_2 are just $\bar{y}_{1:r}$ and $\bar{y}_{r+1:n}$ and we can therefore write the least squares estimator for the true change point index, to minimise this RSS, as

$$\hat{r} = \operatorname{argmin}_{r \in \{1,..,n-1\}} \sum_{i=1}^{r} (y'_i - \bar{y}_{1:r})^2 + \sum_{i=r+1}^{n} (y'_i - \bar{y}_{r+1:n})^2$$
(48)

It additionally turns out that this is equal to the Maximum Likelihood Estimator for the index of the change point in a sequence when we assume the underlying distributions of the observations are Gaussian (Chen and Gupta, 2012). Therefore we can use equation 2.7 in Chen and Gupta (2012) to equivalently rewrite our least squares estimator from (48) as

$$\hat{r} = \operatorname{argmax}_{r \in \{1,..,n-1\}} \frac{r(n-r)}{n} (\bar{y}_{1:r} - \bar{y}_{r+1:n})^2$$
(49)

G.2 Justification of Elimination Criteria

In Algorithm 2, at the beginning of each phase we are given empirical means from repeatedly playing the five actions $0 \le a_1 < a_2 < a_3 \le 1$ (we drop the *j* subscript for the actions in this section).

The first problem is to determine if a piecewise constant function fits better in terms of residual sum of squares when the change is between a_1 and a_2 or between a_2 and a_3 . To do so, for simplicity, we choose to look only at the empirical means from playing actions a_1, a_2, a_3 (Note that we do not consider rewards from actions 0, 1 in this step only for simplicity. We could have instead included these actions on the boundary, but this would actually only affect the upper bound we attain in Theorem 1 by at most a constant factor.). Hence by using equation (49) and when fitting a piecewise constant model to our observed rewards, having the change point between a_2 and a_3 is a better fit that having the change point between a_1 and a_2 when we have

$$\frac{t_j(3t_j - t_j)}{3t_j} \left(\hat{\mu}_{a_1} - \frac{\hat{\mu}_{a_2} + \hat{\mu}_{a_3}}{2}\right)^2 < \frac{2t_j(3t_j - 2t_j)}{3t_j} \left(\frac{\hat{\mu}_{a_1} + \hat{\mu}_{a_2}}{2} - \hat{\mu}_{a_3}\right)^2 \tag{50}$$

It turns out that this is equivalent to the event $E_{R,j}$, namely

$$\iff |\hat{\mu}_{a_1} - \hat{\mu}_{a_2}| < |\hat{\mu}_{a_2} - \hat{\mu}_{a_3}|,$$

which is the criteria we used in Algorithm 2 to decide that the change point is in $x^* \in [a_2, a_3)$ and to then eliminate the region $[a_1, a_2)$. We can use a similar argument to justify the construction of the complementary event $E_{L,j}$.

Now, the second problem is to determine if having the change *outside* of the region $[a_1, a_3)$ would actually be a better fit. This time, again for simplicity, we consider only rewards from actions $0, a_1, a_3, 1$ (where omitting a_2)

helps make the backtracking condition simpler and the additional inclusion of a_2 would only improve the upper bound we attain in Theorem 1 by at most a constant factor). In this case, comparing the regions $[0, a_1)$, $[a_1, a_3)$, $[a_3, 1)$, we have that $[a_1, a_3)$ is *not* the best fitting region for the change point when

$$\left|\frac{\hat{\mu}_{0,t_j} + \hat{\mu}_{a_1^j,t_j}}{2} - \frac{\hat{\mu}_{a_3^j,t_j} + \hat{\mu}_{1,t_j}}{2}\right| < \sqrt{\frac{3}{4}} \max\left(\left|\frac{\hat{\mu}_{0,t_j} + \hat{\mu}_{a_1^j,t_j} + \hat{\mu}_{a_3^j,t_j}}{3} - \hat{\mu}_{1,t_j}\right|, \left|\hat{\mu}_{0,t_j} - \frac{\hat{\mu}_{a_1^j,t_j} + \hat{\mu}_{a_3^j,t_j} + \hat{\mu}_{1,t_j}}{3}\right|\right),$$

which again comes from equation (49). We then define this as our criteria for backtracking in the definition of event $E_{P,j}$ except we modify the constant $\sqrt{\frac{3}{4}}$ to be $\frac{3}{4}$. We make this modification for simplicity as well as the fact that it allows us to construct a very slightly tighter upper bound in Theorem 1. This constant determines how strict we are with the backtracking procedure and it would be interesting to study exactly how this constant affects our upper bound for the failure probability. Perhaps in doing so we could find a way to optimise the strictness of our backtracking rule to further improve our upper bound. However this is beyong the scope of our work for now.

H PROOFS FOR ADAPTIVE ALGORITHM

H.1 Proof of Theorem 5

Throughout this proof, we denote $C_1 = 1/600$, $C_2 = 13$.

Lets consider an environment v with change in mean reward at the change point Δ and sub-Gaussian constant σ^2 . We note that the event $\{T \ge \tau\}$, with

$$\tau = \gamma \frac{\sigma^2}{\hat{\Delta}_L^2} \log\left(\frac{1}{\eta}\right),$$

is equivalent to the event $\{\hat{\Delta}_L \geq \theta_{\Delta}\}$, where

$$\theta_{\Delta} = \sqrt{\gamma \frac{\sigma^2}{T} \log\left(\frac{1}{2\eta}\right)}.$$
(51)

Now, we can use the Law of Total Probability to rewrite the failure probability from SHA as follows.

$$\mathbb{P}_{SHA,v}(|\hat{x}_{T}^{SHA} - x^{*}| > \eta) = \mathbb{P}_{SHA,v}(|\hat{x}_{T}^{SHA} - x^{*}| > \eta | \hat{\Delta}_{L} \ge \theta_{\Delta}) \mathbb{P}_{SHA,v}(\hat{\Delta}_{L} \ge \theta_{\Delta}) + \mathbb{P}_{SHA,v}(|\hat{x}_{T}^{SHA} - x^{*}| > \eta | \hat{\Delta}_{L} < \theta_{\Delta}) \mathbb{P}_{SHA,v}(\hat{\Delta}_{L} < \theta_{\Delta}) = \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SHB} - x^{*}| > \eta | \hat{\Delta}_{L} \ge \theta_{\Delta}) \mathbb{P}_{SHA,v}(\hat{\Delta}_{L} \ge \theta_{\Delta})$$
(52)

$$+ \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SH} - x^*| > \eta | \hat{\Delta}_L < \theta_\Delta) \mathbb{P}_{SHA,v}(\hat{\Delta}_L < \theta_\Delta)$$

$$=\mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SHB} - x^*| > \eta)\mathbb{P}_{SHA,v}(\hat{\Delta}_L \ge \theta_{\Delta})$$

$$+\mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SH} - x^*| > \eta)\mathbb{P}_{AHS,v}(\hat{\Delta}_L < \theta_{\Delta})$$
(53)

Where equation (52) holds since if the event $\{\hat{\Delta}_L \geq \theta_{\Delta}\}$ holds, then the event $\{T \geq \tau\}$ holds. Hence the SHA algorithm will run SHB for the final T - L rounds and the estimate \hat{x}_{T-L}^{SHB} will be returned (See Algorithm 4 for explicit algorithm statement). Similarly, if event $\{\hat{\Delta}_L < \theta_{\Delta}\}$ holds, SHA will use SH for the final T - L rounds. Equation (53) comes from the fact that the random variable $\hat{\Delta}_L$ from the first L rounds is independent to the estimate from SHB using the final T - L rounds \hat{x}_{T-L}^{SHB} and so the condition can be removed. Similarly, $\hat{\Delta}_L$ is independent to the estimate from SH \hat{x}_{T-L}^{SH} .

Hence from equation (53), the failure probability for SHA decomposes into the failure probability when using SHB for T - L rounds multiplied by the probability that SHB is chosen, plus the failure probability from using SH multiplied by the probability that SH is chosen.

We now want to show that the failure probability for the SHA algorithm is near optimal for all budgets both small and large. We therefore split our analysis of the failure probability of SHA into three cases. **Case 1:** small budgets, **Case 2:** very large budgets, and **Case 3:** moderate budgets, which we define later. We demonstrate that for small budgets the probability that SHA chooses to run SH for the final T - L rounds is high enough that the failure probability is of near optimal form. For very large budgets we show that SHA chooses to run SHB for the final T - L rounds with high probability, and therefore our failure probability is near optimal. For moderate budgets we show that the failure probability from using SH or SHA would be similar and therefore our failure probability is near optimal regardless of our choice between the algorithms.

The result from Case 1 gives us the first upper bound of Theorem 5 for $T < T_1$ and combining the results from Case 2 and 3 gives us the second upper bound from Theorem 5 for $T \ge T_1$.

Case 1: Small budget setting: Suppose we are in the small budget setting where $T < T_1$, which can be equivalently written as

$$\Delta < \sqrt{\frac{\sigma^2}{T}} \left(1.59 \log\left(\left\lfloor \frac{1}{2\eta} \right\rfloor \right) - 2 \log(2) \right).$$
(54)

We firstly note that by the assumption that the reward distributions for all of our actions are σ^2 sub-Gaussian, we know that the empirical means $\hat{\mu}_{1,L}$ and $\hat{\mu}_{0,L}$ are each $2\sigma^2/L$ sub-Gaussian. Hence, $\hat{\mu}_{1,L} - \hat{\mu}_{0,L}$ is $4\sigma^2/L$ sub-Gaussian. We can then use Hoeffding's inequality (from Proposition 2.5 in Wainwright (2019)) to show that, under the SHA algorithm our estimate for the size of the change in mean $\hat{\Delta}_L$ is concentrated as follows.

$$\mathbb{P}_{SHA,v}(\hat{\Delta}_{L} \ge \theta_{\Delta}) = \mathbb{P}_{SHA}(|\hat{\mu}_{1,L} - \hat{\mu}_{0,L}| \ge \theta_{\Delta}) \\
= \mathbb{P}_{SHA}(|\hat{\mu}_{1,L} - \hat{\mu}_{0,L} - \Delta + \Delta| \ge \theta_{\Delta}) \\
\le \mathbb{P}_{SHA}(|\hat{\mu}_{1,L} - \hat{\mu}_{0,L} - \Delta| + \Delta \ge \theta_{\Delta}) \\
= \mathbb{P}_{SHA}(|\hat{\mu}_{1,L} - \hat{\mu}_{0,L} - \Delta| \ge \theta_{\Delta} - \Delta) \\
\le 2 \exp\left(-\frac{L(\theta_{\Delta} - \Delta)^{2}}{8\sigma^{2}}\right)$$
(55)

Note that by assumption $\gamma > 1.59$, therefore the $\theta_{\Delta} - \Delta$ term above is positive and this allows Hoeffding's inequality to hold for equation (55). Hence, returning to the equation (53) we have that the failure probability of SHA under small budgets is as follows.

$$\begin{split} \mathbb{P}_{SHA,v}(|\hat{x}_{T}^{SHA} - x^{*}| > \eta) &\leq \mathbb{P}_{SHA,v}(\hat{x}_{T-L}^{SHA} - x^{*}| > \eta)\mathbb{P}_{SHA,v}(\hat{\Delta}_{L} \geq \theta_{\Delta}) \\ &+ \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SH} - x^{*}| > \eta)\mathbb{P}_{AHS,v}(\hat{\Delta}_{L} \geq \theta_{\Delta}) \\ &+ \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SH} - x^{*}| > \eta) \\ &\leq 2 \exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T - L) + C_{2}\log(1/2\eta) - \frac{L(\theta_{\Delta} - \Delta)^{2}}{8\sigma^{2}}\right) \quad (56) \\ &+ 2\left\lceil \log_{2}\left(\frac{1}{2\eta}\right)\right\rceil \exp\left(\frac{-(T - L)\Delta^{2}}{36\sigma^{2}\log_{2}(1/2\eta)}\right) \\ &\leq 2 \exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T - L) + C_{2}\log(1/2\eta) - B(\sqrt{\gamma} - \sqrt{1.59})^{2}\frac{\log\left(\frac{1}{2\eta}\right)}{8}\right) \quad (57) \\ &+ 2\left\lceil \log_{2}\left(\frac{1}{2\eta}\right)\right\rceil \exp\left(\frac{-(T - L)\Delta^{2}}{36\sigma^{2}\log_{2}(1/2\eta)}\right) \\ &\leq 2 \exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T - L)\right) \\ &+ 2\left\lceil \log_{2}\left(\frac{1}{2\eta}\right)\right\rceil \exp\left(\frac{-(T - L)\Delta^{2}}{36\sigma^{2}\log_{2}(1/2\eta)}\right) \\ &\leq 4\left\lceil \log_{2}\left(\frac{1}{2\eta}\right)\right\rceil \exp\left(\frac{-(T - L)\Delta^{2}}{\sigma^{2}}\min\left\{\frac{1}{36\log_{2}(1/2\eta)}, C_{1}\right\}\right) \\ &\leq 4\left\lceil \log_{2}\left(\frac{1}{2\eta}\right)\right\rceil \exp\left(-C_{1}\frac{(T - L)\Delta^{2}}{\sigma^{2}\log_{2}(1/2\eta)}\min\left\{\frac{1}{36}, C_{1}\right\}\right) \\ &\leq 4\left\lceil \log_{2}\left(\frac{1}{2\eta}\right)\right\rceil \exp\left(-C_{1}\frac{(T - L)\Delta^{2}}{\sigma^{2}\log_{2}(1/2\eta)}\right) \end{aligned}$$

Note that equation (56) comes from substituting in our upper bounds from Theorem 1, Theorem 3 and the concentration inequality from equation (55). Also, in order to attain (57), we consider lower bounding the quantity $(\theta_{\Delta} - \Delta)^2$. In Case 1 we have $\theta_{\Delta} > \Delta$ and an upper bound (54) for Δ . Hence, to upper bound equation (56) we substitute in the largest value Δ can take in Case 1, namely

$$\Delta = \sqrt{\frac{\sigma^2}{T}} \left(1.59 \log\left(\left\lfloor \frac{1}{2\eta} \right\rfloor \right) - 2 \log(2) \right),$$

as well as our definition of θ_{Δ} from (51) and the fact that L = BT in order to attain equation (57). Then equation (58) only holds when $B(\sqrt{\gamma} - \sqrt{1.59})^2/8 > C_2$, which is assumed in the statement of Theorem 5. The upper bound in equation (59) is of the form of the small budget lower bound in Theorem 4, hence SHA is nearly minimax optimal up to constants and loglog terms under small budgets.

Case 2: Very Large Budgets: Now assume that the budget is sufficiently large such that

 $\Delta > 2\theta_{\Delta},$

where θ_{Δ} is defined in equation (51). Note that in this case, we can make a similar concentration argument to Case 1, using Hoeffding's inequality to show that $\hat{\Delta}_L < \theta_{\Delta}$ occurs with small probability. In particular,

$$\mathbb{P}_{SHA,v}(\Delta_L < \theta_{\Delta}) = \mathbb{P}_{SHA}(|\hat{\mu}_{1,L} - \hat{\mu}_{0,L}| < \theta_{\Delta}) \\
= \mathbb{P}_{SHA}(|\hat{\mu}_{1,L} - \hat{\mu}_{0,L} - \Delta + \Delta| < \theta_{\Delta}) \\
\leq \mathbb{P}_{SHA}(|\Delta| - |\hat{\mu}_{1,L} - \hat{\mu}_{0,L} - \Delta| < \theta_{\Delta}) \\
= \mathbb{P}_{SHA}(|\hat{\mu}_{1,L} - \hat{\mu}_{0,L} - \Delta| > \Delta - \theta_{\Delta}) \\
\leq 2 \exp\left(-\frac{L(\theta_{\Delta} - \Delta)^2}{8\sigma^2}\right).$$
(60)

Note that, since in Case 2 we assume that $\Delta > 2\theta_{\Delta}$, and therefore $\Delta > \theta_{\Delta}$. Hence the $\Delta - \theta_{\Delta}$ term above is positive and this allows Hoeffding's inequality to hold for equation (60). We can now rewrite the failure probability for the SHA algorithm in this case as

$$\begin{split} \mathbb{P}_{SHA,v}(|\hat{x}_{T}^{SHA} - x^{*}| > \eta) &\leq \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SHB} - x^{*}| > \eta)\mathbb{P}_{SHA,v}(\hat{\Delta}_{L} \geq \theta_{\Delta}) \\ &+ \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SHB} - x^{*}| > \eta)\mathbb{P}_{AHS,v}(\hat{\Delta}_{L} < \theta_{\Delta}) \\ &\leq \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SHB} - x^{*}| > \eta)\mathbb{P}_{AHS,v}(\hat{\Delta}_{L} < \theta_{\Delta}) \\ &\leq \exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T - L) + C_{2}\log(1/2\eta)\right) \qquad (61) \\ &+ 4\left[\log_{2}\left(\frac{1}{2\eta}\right)\right]\exp\left(\frac{-(T - L)\Delta^{2}}{36\sigma^{2}\log_{2}(1/2\eta)} - \frac{L(\theta_{\Delta} - \Delta)^{2}}{8\sigma^{2}}\right) \\ &\leq \exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T - L) + C_{2}\log(1/2\eta)\right) \\ &+ 4\left[\log_{2}\left(\frac{1}{2\eta}\right)\right]\exp\left(\frac{-(T - L)\Delta^{2}}{36\sigma^{2}\log_{2}(1/2\eta)} - \frac{L(\Delta/2)^{2}}{8\sigma^{2}}\right) \\ &= \exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T - L) + C_{2}\log(1/2\eta)\right) \\ &+ 4\left[\log_{2}\left(\frac{1}{2\eta}\right)\right]\exp\left(\frac{-(T - L)\Delta^{2}}{36\sigma^{2}\log_{2}(1/2\eta)} - \frac{L\Delta^{2}}{32\sigma^{2}}\right) \\ &\leq \exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T - L) + C_{2}\log(1/2\eta)\right) \\ &+ 4\left[\log_{2}\left(\frac{1}{2\eta}\right)\right]\exp\left(-\frac{L\Delta^{2}}{32\sigma^{2}}\right) \\ &\leq 5\exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}B(1 - B)T + C_{2}\log(1/2\eta)\right). \end{aligned}$$

Here equation (61) comes from substituting our upper bounds for SH in Theorem 3 and SHB in Theorem 1, as well as using our upper bound in equation (60). Furthermore, equation (62) comes from the fact that $\Delta > 2\theta_{\Delta}$. Finally, equation (63) comes from recalling that L = BT. Note that equation (63) matches our lower bound from Theorem 2 up to constants, hence SHA is near optimal up to constants for these very large budgets.

Case 3 Moderate Budgets: Let's suppose instead that the budget is only moderately large and is bounded by

$$\Delta \in \left[\sqrt{\frac{\sigma^2}{T} \left(1.59 \log\left(\left\lfloor \frac{1}{2\eta} \right\rfloor\right) - 2 \log(2)\right)}, 2\theta_{\Delta}\right].$$

We can equivalently write this condition as

$$\Delta \in \left[\sqrt{\frac{\sigma^2}{T}} \left(1.59 \log\left(\left\lfloor \frac{1}{2\eta} \right\rfloor\right) - 2 \log(2)\right), 2\sqrt{\gamma \frac{\sigma^2}{T}} \log\left(\frac{1}{2\eta}\right)\right].$$
(64)

In this case, as long as we have $\gamma < 1800/(1-B)$, then for all such Δ in the moderate budget setting (64), we have a better guarantee for SH in Theorem 3 than for SHB in Theorem 1 with budget input T - L. This is because, combining $\gamma < 1800/(1-B)$ and our condition (64) for Case 3, implies that $T - L < 7200 \frac{\sigma^2}{\Delta^2} \log \left(\frac{1}{\eta}\right)$. Plugging in these values for the budget T - L into Theorems 3 and 1 shows that our gaurantee for SH outperforms SHB in this moderate budget setting. We note that $7200 \frac{\sigma^2}{\Delta^2} \log \left(\frac{1}{\eta}\right)$ is different to the definition of T_1 (see Section 5). The constant 7200 comes from the comparison of our two upper bounds for SH/SHB (which are loose in constants), whereas in practice the budget at which SHB begins to outperform SH is significantly smaller than this (see Section B).

Hence, we can upper bound our failure probability in this region as

$$\mathbb{P}_{SHA,v}(|\hat{x}_{T}^{SHA} - x^{*}| > \eta) \leq \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SHB} - x^{*}| > \eta)\mathbb{P}_{SHA,v}(\hat{\Delta}_{L} \geq \theta_{\Delta}) \\
+ \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SH} - x^{*}| > \eta)\mathbb{P}_{AHS,v}(\hat{\Delta}_{L} < \theta_{\Delta}) \\
\leq \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SH} - x^{*}| > \eta) \cdot 1 \\
+ \mathbb{P}_{SHA,v}(|\hat{x}_{T-L}^{SH} - x^{*}| > \eta) \cdot 1 \\
\leq 2\exp\left(-C_{1}\frac{\Delta^{2}}{\sigma^{2}}(T-L) + C_{2}\log(1/2\eta)\right)$$
(65)

$$= 2 \exp\left(-C_1 \frac{\Delta^2}{\sigma^2} (1-B)T + C_2 \log(1/2\eta)\right).$$
 (66)

Here equation (65) comes from the above discussion that for moderate budgets in Case 3, our gaurantee for SH is better than that of SHB. Equation (66) comes from recalling that L = BT. Equation (66) is of the form of our lower bound in Theorem 2, therefore SHA algorithm is nearly optimal under moderate budgets as well.