

PAIRWISE ELIMINATION WITH INSTANCE-DEPENDENT GUARANTEES FOR BANDITS WITH COST SUBSIDY

Anonymous authors

Paper under double-blind review

ABSTRACT

Multi-armed bandits (MAB) are commonly used in sequential online decision-making when the reward of each decision is an unknown random variable. In practice, however, the typical goal of maximizing total reward may be less important than minimizing the total cost of the decisions taken, subject to a reward constraint. For example, we may seek to make decisions that have at least the reward of a reference “default” decision. This problem was recently introduced in the Multi-Armed Bandits with Cost Subsidy (MAB-CS) framework. MAB-CS is broadly applicable to problem domains where a primary metric (cost) is constrained by a secondary metric (reward), and there is an inability to explicitly determine the trade-off between these metrics. In our work, we first introduce the Pairwise-Elimination algorithm for a simplified variant of the cost subsidy problem with a known reference arm. We then generalize PE to PE-CS to solve the MAB-CS problem in the setting where the reference arm is the un-identified optimal arm. Next, we analyze the performance of both PE and PE-CS on the dual metrics of Cost and Quality Regret. Our instance-dependent analysis of PE and PE-CS reveals that both algorithms have an order-wise logarithmic upper bound on Cost and Quality Regret, making our policy the first with such a guarantee. Finally, experiments are conducted using the MovieLens 25M dataset for both PE and PE-CS and using a synthetic toy experiment for PE-CS revealing that our method invariably outperforms the ETC-CS baseline from the literature.

1 INTRODUCTION

Online sequential decision-making problems capture many applications where decisions must be made without knowing their outcomes in advance. After each decision, the resulting outcome or reward is observed, and an internal model is updated to improve future decisions. In clinical trials for example, the goal is to compare the therapeutic value of various drugs against an ailment. The decisions in this case represent administering a certain drug and the rewards are the apriori unknown efficacies of the candidate drugs. Communication networks are another example. Here decisions must be made about the communication channel to be employed. In this scenario the reward represents the success or lack thereof of communicating over a chosen channel. Multi-Armed Bandits (MABs) (Lattimore & Szepesvári, 2020) are a framework for *stateless* sequential decision making where the available decisions are abstracted as arms of a MAB problem instance. The stateless assumption implies that the distribution of rewards associated with an arm is not affected by the choices of past arms. The setting within MABs we work with is that of *stationary stochastic bandits* where the distribution of arm rewards does not evolve with time.

The generality of the assumptions imposed by the stationary stochastic bandits setting provides a wide net to capture a range of problem domains. However, in real-world applications there are often several competing objectives that go beyond the limited goal of maximizing reward. For instance consider the problem faced by a marketing agency where there are several communication modalities available to communicate the agency’s advertising message. Blindly maximizing the overall success rate (reward) in this case would be naive. Since such an approach would ignore the drastically different costs of using these modalities. Our example reveals that costs being associated with the sampling of any particular arm is a structure that appears quite naturally in applications.

In the marketing agency problem we know that the various communication modalities shall have unknown success rates for any brand new ad campaign. However, the cost of employing any modality will typically be known. These known arm sampling costs might manifest themselves in the form of a prescribed cost budget (Badanidiyuru et al., 2018), or as a metric whose cumulative value is to be minimized (Sinha et al., 2021), to specify two among many possible cost structures. We work with the latter among these settings. In particular our paper works with variants of the MAB with Cost-Subsidy (MAB-CS) framework introduced recently in Sinha et al. (2021).

1.1 MULTI-ARMED BANDITS WITH COST SUBSIDY (MAB-CS) SETTING

What makes the MAB-CS setting so interesting is that it requires the bandit policy minimize cumulative costs while obtaining cumulative reward that is *feasible* and not necessarily maximal. The problem therefore involves dual objectives: minimizing costs while ensuring that the reward meets or exceeds the feasibility threshold, denoted by μ_{CS} . For an MAB-CS problem instance with K arms we use c_i to denote the known cost associated with sampling any arm $i \in [K]$, and μ_i to denote the expectation of the reward received from sampling arm i . The *optimal arm* in our setting is then the least cost feasible arm. Mathematically if $S = \{i \in [K] : \mu_i \geq \mu_{CS}\}$ denotes the set of feasible arms, then the optimal arm $a^* = \arg \min_{i \in S} c_i$.

To honor the dual objectives of minimizing cost while maintaining feasible reward, cumulative cost and quality regret over a time horizon T , for instance ν , with policy π are defined as follows,

$$\begin{aligned} \text{Cost_Reg}(T, \nu, \pi) &= \sum_{t=1}^T \mathbb{E}_\pi \left[(c_{k_t} - c_{a^*})^+ \right] = \sum_{i=1}^K \Delta_{C,i}^+ \mathbb{E} [n_i(T)] \\ \text{Quality_Reg}(T, \nu, \pi) &= \sum_{t=1}^T \mathbb{E}_\pi \left[(\mu_{CS} - \mu_{k_t})^+ \right] = \sum_{i=1}^K \Delta_{Q,i}^+ \mathbb{E} [n_i(T)]. \end{aligned}$$

Where k_t is the arm sampled by bandit policy π in time slot t , and $n_i(T)$ is a random variable denoting the number times arm i was sampled over T time steps. The initial definitions of regret accumulated over horizon T in the center have been re-written in terms of the expected number of samples of each of the K arms at the right. Moreover we have used the definitions of cost gap $\Delta_{C,i} = (c_i - c_{a^*})$ and quality gap $\Delta_{Q,i} = (\mu_{CS} - \mu_i)$ in the decomposition. We highlight that as is standard for stationary stochastic bandits, the operator \mathbb{E}_π represents the expectation over the choices of arm k_t made by policy π during time slot t . In the remainder of this section, we discuss three variants of MAB-CS, each of which have a different specification for μ_{CS} . We end by motivating the importance of $\Delta_{C,i}^+, \Delta_{Q,i}^+$, where x^+ denotes $\max\{0, x\}$, in the regret definitions.

Take as an example the problem faced by a marketing agency. Consider that there are three modalities available for the agency to deliver their message, which are: (1) very expensive personalized door-step solicitation, (2) moderately expensive automated phone call, and (3) inexpensive email. Given these modalities, the agency's goal may be to achieve a prescribed sales rate with the minimum possible cost. Or, the goal may be that sales be at least a prescribed fraction of the sales of a certain communication modality. Finally, we may not have a reference mode in mind, and we may just desire a conversation rate that is (say) 80% as much as the highest unknown sales rate.

The first setting is captured by our novel contribution of the *fixed threshold* setting. In the fixed threshold setting we specify $\mu_{CS} = \mu_0$ for a known $\mu_0 \in \mathbb{R}$. We call the second setting the *known reference arm* setting which specifies $\mu_{CS} = (1 - \alpha)\mu_\ell$. Here ℓ is the index of the reference arm, and $\alpha \in [0, 1]$ is a known subsidy factor. The third setting was introduced in prior work and specifies $\mu_{CS} = (1 - \alpha)\mu^*$. Where μ^* is the largest among expected rewards. In particular, μ^* is the expected reward from sampling arm $i^* = \arg \max_{i \in [K]} \mu_i$. Similar to the *known reference arm* setting α is the subsidy factor. We refer to this third setting as the *subsidized best reward* setting.

With the MAB-CS framework, we target applications that are agnostic to the level of quality as long as the quality exceeds threshold μ_{CS} . This structure necessitates the zero-clipped operation inside the cost and quality regret definitions. Consider for example a problem where it is known that customers need to be provided a certain (possibly unknown) service quality level for them to continue their subscription. Any quality on top of the feasibility level μ_{CS} would not improve the performance of our solution. For the marketing communication example from earlier, the quality

108 threshold represents a sales conversation rate beyond which the profitability of the campaign is
 109 ensured. In this case too we would like decisions that are agnostic to sales success over μ_{CS} .

110
 111 In the absence of the zero-clipped structure, cost and quality regret that are sub-linear in horizon
 112 T could be achieved for our examples in an unintended fashion. An algorithm that samples sub-
 113 optimal arms $i \neq a^*$ a linear fraction of times but balances positive regret from infeasible decisions,
 114 by negative regret from stellar decisions would satisfy the un-clipped quality constraint but would
 115 be unsuitable for our example applications.

116 1.2 KEY CONTRIBUTIONS

117
 118 Our first contribution is to extend the MAB-CS framework to include two new settings. (1) The *fixed*
 119 *threshold* setting, and (2) The *known reference arm* setting. Second, we present instance dependent
 120 lower bounds on the expected number of pulls of any sub-optimal arm $i \neq a^*$ for our two novel
 121 settings as well as the third *subsidized best reward* setting from the literature.

122
 123 Next, we present an original order optimal algorithm Pairwise Elimination (PE) for regret minimiza-
 124 tion in the known reference arm setting. Under PE non-reference arms are pit against the reference
 125 arm in the ascending order of their costs. PE uses a principled elimination based regret minimiza-
 126 tion algorithm called Improved-UCB (Auer & Ortner, 2010) to determine whether an arm provides
 127 feasible rewards. We show that our PE has an instance dependent upper bound on both expected
 128 cost and quality regret that is $\mathcal{O}(\log T)$, and that PE only samples arms more expensive than the
 129 optimal action a^* at most a constant number of times under expectation.

130
 131 Next, we develop a generalization of PE for the subsidized best reward setting called PE-CS. We
 132 show that PE-CS too admits an $\mathcal{O}(\log T)$ instance dependent upper bound on both cost and quality
 133 regret that involves both notions of conventional sub-optimality gaps and quality gaps. Not only
 134 is PE-CS the first algorithm for the subsidized best reward setting with instance dependent upper
 135 bounds on regret, but also PE-CS offers a substantial improvement in performance over the only
 136 other algorithm from the literature for the subsidized best reward setting which admits a guarantee,
 137 namely ETC-CS. Although we find that PE-CS is not order optimal for all instances, our contribution
 138 includes characterizing the class of instances for which PE-CS is order optimal.

139 2 RELATED WORK

140
 141 **Structured Bandits:** There have been numerous works that impose additional structure onto the
 142 stationary stochastic bandits problem with the goal of better addressing specific application domains.
 143 This structure can sometimes come in the form of relationships imposed on the rewards of arms.
 144 These reward-relationships may be known (Kleinberg et al., 2008) or unknown (Gupta et al., 2021).
 145 Adding constraints that depend on the risks associated with sampling the rewards of an arm as in
 146 Wu et al. (2016) or Chen et al. (2022) is another form of the structured problem.

147
 148 **Bandits with Costs:** We contextualize the core contributions of this paper by comparing and con-
 149 trasting our setting and methods with related ones from the Literature. We build on the MAB-CS
 150 setting introduced in Sinha et al. (2021). A core component of the MAB-CS setting is that there
 151 is a known cost associated with sampling any arm that is specified as part of the problem instance.
 152 There have been numerous works within the MAB literature that include the notion that a price has
 153 to be paid for sampling an arm. Notably the Bandits with Knapsacks (Badanidiyuru et al., 2018)
 154 line of work also considers a setting with known costs. However, in Badanidiyuru et al. (2018) there
 155 is a limited cost-budget and reward must be optimized while satisfying strict budget constraints. In
 156 MAB-CS and its variants on the other hand, the goal is to minimize cumulative costs without there
 157 being any explicit constraints on cost. In MAB-CS, the constraints are in fact on reward, and are
 158 referred to as quality constraints.

159
 160 **Bandit with Constraints:** The quality constraints in our work closely resemble the constraints on
 161 expected rewards that are imposed in the work on Conservative Bandits (Wu et al., 2016). In both
 our work and in Wu et al. (2016), there is a constraint that requires the accumulated reward to
 exceed a $(1 - \alpha)$ discounted version of the reward of a reference arm. The conservative bandits
 setting only considers the cases where either the return of the reference arm is a known constant μ_0 ,
 or the case where the reference arm is known, but its return is unknown. The primary difference in

our work is that in addition to satisfying a quality constraint, in the MAB-CS setting, we must work to minimize the cumulative cost. The notion of costs are completely absent in Wu et al. (2016), moreover, in addition to the cases with a known reward threshold μ_0 , and a known reference arm with an unknown return, which we consider as novel extensions to the MAB-CS framework, we also address the problem of the original MAB-CS framework where the reference arm is the unknown optimal arm. Another key difference is that Wu et al. (2016) imposes the reward constraints in a cumulative anytime manner, whereas we impose it at every time-step.

BAI and Improved UCB In our paper, we work with the notions of Cost Regret and Quality Regret which are identical to the ones introduced by Sinha et al. (2021), however unlike the setting in Sinha et al. (2021) which only considers the case where the reference reward comes from the so-far unidentified optimal arm, we consider the additional cases (1) Where there is a fixed known threshold to be exceeded (which we call the known threshold setting), and (2) When there is a known reference arm ℓ whose reward μ_ℓ has to be exceeded however μ_ℓ itself is unknown. In addition to our novel PE-CS algorithm for the setting from Sinha et al. (2021) we present novel algorithms for our new settings (1) and (2) as well. In Sinha et al. (2021) the authors present three novel algorithms for the MAB-CS setting, the former two among which construct a set of empirically feasible arms by interleaving exploration and exploitation. We build up our approach to optimizing for the regret objectives by first solving the known reference arm ℓ with unknown reward μ_ℓ setting using a successive elimination style algorithm that compares candidate arms one at a time against the reference arm to see if they are feasible. We call this approach Pairwise Elimination (PE), and we adapt the elimination based regret minimization algorithm Improved-UCB (Auer & Ortner, 2010) to develop it. Then we generalize PE to the case where the reference arm is the unknown optimal arm by prepending PE with a Best-Arm-Identification (BAI) stage (also based on Auer & Ortner (2010)). We call this latter algorithm PE-CS.

3 ALGORITHMS AND ANALYSIS

As discussed in Section 1, we introduce novel settings called: (1) Fixed threshold setting with $\mu_{CS} = \mu_0$, and (2) known reference arm setting with $\mu_{CS} = (1 - \alpha)\mu_\ell$. In interest of building up to our presentation on PE-CS we start with the known reference arm setting and relegate the discussion and analysis of the known threshold setting to Appendix H. For the known reference arm setting, we present the lower bound and our novel Pairwise-Elimination algorithm in Section 3.1.

Our PE algorithm builds upon Improved-UCB (Figure 1 in Auer & Ortner (2010)), a regret minimization algorithm for the stationary stochastic bandits setting. We choose to build on the method since its successive elimination approach to regret minimization readily incorporates our insight that arms be evaluated in the order of their costs. One of the core features of Improved-UCB is that the cadence of sampling and elimination is governed by *rounds*. Each round is a period where every active arm is sampled to the extent prescribed by a function of the *round number*. A round concludes when all active arms have been sampled sufficiently, and at the end of a round, unsatisfactory arms are eliminated. We inherit the use of these rounds and associated formulas from Improved-UCB.

Finally, our third setting has the reward threshold $\mu_{CS} = (1 - \alpha)\mu^*$. This subsidized best reward setting is strictly more challenging than the known reference arm ℓ setting since the optimal arm itself is unknown. To solve it, we extend the PE algorithm by pre-pending it with a Best-Arm-Identification (BAI) stage to develop the PE-CS algorithm. The lower bound for the setting and the PE-CS algorithm are presented in Section 3.2.

3.1 KNOWN REFERENCE ARM SETTING

The known reference arm setting specifies $\mu_{CS} = (1 - \alpha)\mu_\ell$ as the subsidized (unknown) return of a known reference arm. Without loss of generality, we assume that all MAB-CS setting bandit instances have arms indexed in the ascending order of their costs. In the known reference arm setting any arm with cost higher than c_ℓ is necessarily sub-optimal and can be pruned out. Moreover we can think of arm ℓ as being an arm with index $K + 1$. First we provide an instance dependent lower bound on the expected number of samples of sub-optimal arms for a class of consistent policies. The definition of consistent policies is available in Appendix E.

Theorem 3.1 (Lower bound for known reference arm setting). *Under any consistent policy π the expected number of samples of a low cost arm and of the reference arm ℓ are lower bounded as,*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_i(T)]}{\log T} \geq \frac{2}{\Delta_{Q,i}^2}, \quad \text{for arms } i < a^*, \quad \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_\ell(T)]}{\log T} \geq \max_{i \leq a^*} \frac{2(1-\alpha)^2}{\Delta_{Q,i}^2}.$$

Where T denotes the problem horizon and the rewards of all K arms are Gaussian distributed with variance $\sigma^2 = 1$. Low cost is a term used relative to the cost of optimal arm a^* .

The proof of Theorem 3.1 (available in Appendix E) uses analytic techniques and arguments similar to those used to establish lower bounds for the classical multi-armed bandit setting (Garivier et al., 2019). We now discuss our PE (Algorithm 1) for the known reference arm setting. Under this setting, quality regret is calibrated against the expected reward of the reference arm ℓ which we denote μ_ℓ . For jointly optimizing cost and quality regret, we take an approach where the feasibility of arms is evaluated in the order of the costs of the arms: cheapest first. This insight motivates a pairwise comparison between the reference arm ℓ and the non-reference candidate arms.

Function 1: Pairwise Elimination Function PE()

Function $PE(\mathbf{n}: \text{Sample Vector}, \hat{\boldsymbol{\mu}}: \text{Empirical Means}, \boldsymbol{\omega}: \text{Round Numbers}, T: \text{Horizon}, i: \text{Episode}, \ell: \text{Reference Arm}, \alpha: \text{Subsidy Factor})$:

```

1   $\tilde{\Delta} \leftarrow 2^{-\omega_i}$ 
2   $\tau \leftarrow \left\lceil \frac{2 \log(T \tilde{\Delta}^2)}{\tilde{\Delta}^2} \right\rceil$ 
3  for  $k \in \{i, \ell\}$  do
4    if  $n_k < \tau$  then
5      return  $k, \boldsymbol{\omega}, i$  //  $k_t = k$ , round numbers  $\boldsymbol{\omega}$  and ep.  $i$  unchanged
6   $\beta \leftarrow \sqrt{\frac{\log(T \tilde{\Delta}^2)}{2\tau}}$ 
7  if  $(1 - \alpha)(\hat{\mu}_\ell + \beta) < \hat{\mu}_i - \beta$  then
8    return  $i, \boldsymbol{\omega}, \text{None}$  // Declare  $i$  as winner, set episode to None
9  else if  $\hat{\mu}_i + \beta < (1 - \alpha)(\hat{\mu}_\ell - \beta)$  then
10   return  $i+1, \boldsymbol{\omega}, i+1$  // Sample next candidate arm, Rounds  $\boldsymbol{\omega}$ 
11   unchanged, Update episode to that of next candidate arm
12 else
13    $\omega_i \leftarrow \omega_i + 1$  // Increment round only for arm  $i$  being evaluated
14   return  $i, \boldsymbol{\omega}, i$  // Move to next round within same episode

```

Algorithm 1: Pairwise Elimination (PE) for a known reference arm ℓ

Inputs: Bandit Instance ν , Horizon T , Reference Arm ℓ .

Initialize: Samples $n_k = 0$, Empirical Means $\hat{\mu}_k = 0$, Current Rounds $\omega_k = 0$,
 $\forall k \in [K] \cup \{\ell\}$, PE Episode $i = 1$.

```

1 while  $t \leq T$  do
2   if  $i \notin \{\text{None}, \ell\}$  then
3      $k_t, \boldsymbol{\omega}, i \leftarrow PE(\mathbf{n}, \hat{\boldsymbol{\mu}}, \boldsymbol{\omega}, T, i, \ell, \alpha)$  // receive arm to be sampled,
4     updated round numbers, and updated episode number
5   else
6      $k_t \leftarrow k_{t-1}$  // sample winning arm for remaining budget
7    $\hat{\boldsymbol{\mu}}(t+1), \mathbf{n}(t+1), t \leftarrow \text{sample\_and\_update}(k_t, \hat{\boldsymbol{\mu}}(t), \mathbf{n}(t), t)$  // in Appendix B

```

PE tracks and updates three bookkeeping variables. The number of times each arm has been sampled \mathbf{n} , the empirical mean reward sampled from each arm $\hat{\boldsymbol{\mu}}$, and the round ongoing¹ by each arm $\boldsymbol{\omega}$.

¹Although the round number ω_k achieved by an arm k can be uniquely determined from its number of samples n_k , we track them separately for a more lucid presentation.

For PE we assess candidate arms in the ascending order of their cost by assigning an *episode* to each candidate arm. This is represented by the main loop in Algorithm 1 in which i is used to refer to both the index of the arm being evaluated for feasibility and to the episode number associated with that arm. $\text{PE}(\cdot)$ keeps getting invoked until all K candidate arms are evaluated and we reach episode ℓ , or until when a lower cost feasible arm has been identified.

Inside the $\text{PE}(\cdot)$ subroutine the sample prescription τ is computed using the current round number ω_i associated with candidate arm i (line 2). Once the sample prescription τ is met for both arm i and arm ℓ , we check for elimination. Since no samples are ever discarded, episodes that are further downstream will re-use samples of arm ℓ from prior episodes when $n_\ell > \tau$. If candidate arm i is able to eliminate arm ℓ then the optimal arm has been identified. Else if arm ℓ is able to eliminate arm i , then the episode is incremented and we proceed to evaluating the next cheapest arm. If no elimination occurs, then we simply increment the round number ω_i . For the PE algorithm, we show the upper bound on cumulative cost and quality regret stated in Theorem 3.2.

Theorem 3.2 (Instance dependent upper bound on cumulative cost and quality regret for PE). *For bandit instance ν , over horizon T , the expected cumulative cost regret $\mathbb{E}[\text{Cost_Reg}(T, \nu)]$ and quality regret $\mathbb{E}[\text{Quality_Reg}(T, \nu)]$ of the PE algorithm are upper bounded respectively as,*

$$\underbrace{\left(1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+}_{\text{Contribution from arm } \ell \text{ under nominal termination in PE episode } a^*} + \underbrace{\left(\sum_{i=1}^{a^*} \frac{43}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+}_{\text{Contribution from arm } \ell \text{ under mis-termination in PE episode } \leq a^*} + \underbrace{\frac{43}{\Delta_{Q,a^*}^2} \max_{a^* < i \leq \ell} \Delta_{C,i}^+}_{\text{Contribution from episodes } > a^* \text{ in case of mis-termination during ep } a^*},$$

$$\underbrace{\sum_{i=1}^{a^*-1} \left(\Delta_{Q,i} + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}}\right)}_{\text{Contribution from arms } i < a^* \text{ under nominal termination in PE episode } a^*} + \underbrace{\sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}}}_{\text{Contribution from arms } i < a^* \text{ under mis-termination in PE episode } \leq a^*} + \underbrace{\frac{43}{\Delta_{Q,a^*}^2} \max_{a^* < i < \ell} \Delta_{Q,i}^+}_{\text{Contribution from episodes } > a^* \text{ in case of mis-termination during ep } a^*}.$$

There are two dimensions to the execution of PE being nominal. First within an episode the right arm should be eliminated within a reasonable number of samples. The second is the execution across episodes should terminate in episode a^* (with arm ℓ eliminated). To prove Theorem 3.2 we separately bound the expected number of samples of arms with cost lower and higher than the optimal action a^* , and the reference arm ℓ . To bound these samples we condition on the desirable and anticipated events at both the intra and inter episode levels and show that the probability of these desirable events not occurring is small. The details are available in Appendix F.

Overall, we find that not only does PE achieve cost and quality regret that are $\mathcal{O}(\log T)$, but also PE matches up to constant factors the lower bound on cost and quality regret implied by Theorem 3.1. To see this clearly, we breakdown Theorem 3.2 by contribution. For cost regret, the $\mathcal{O}(\log T)$ dependence on arm ℓ is contained in the first term, and the dependence on higher cost arms is a constant captured by the third. For quality regret the first term captures the $\mathcal{O}(\log T)$ dependence on low cost arms, and the entire contribution of higher cost arms is captured by a constant.

Practical Extensions of PE. We highlight that although PE makes comparisons between the candidate arms and the reference arm in a pairwise manner, samples of the reference arm are re-used across episodes. Since during most episodes the samples accrued shall be limited to ones of the candidate arm undergoing evaluation, this sample reuse is a key feature of the PE algorithm and endows it with good sample efficiency. In PE as presented in Algorithm 1, during an arbitrary episode evaluating the candidacy of arm i , the reference arm ℓ shall only ever have to be sampled if the samples of arm i start to exceed the samples of reference arm ℓ that were already available. In practice, we can implement another version of PE called *asymmetric-PE*. Asymmetric-PE allows for a mismatch between the number of samples of arms i and ℓ that go into computing the exploration bonus terms β . The details of the variant and an example comparing it to vanilla PE are available in Appendix C.

3.2 SUBSIDIZED BEST REWARD SETTING

Theorem 3.3 (Lower bound for subsidized best reward setting). *Under any consistent policy π , the expected number of samples of low cost arms, high cost arms, and the best reward arm i^* are lower*

324 bounded respectively as,

$$\begin{aligned}
325 & \\
326 & \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_i(T)]}{\log T} \geq \frac{2}{\Delta_{Q,i}^2}, \quad i < a^*. \quad \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_i(T)]}{\log T} \geq \frac{2}{(\frac{\mu_{a^*}}{1-\alpha} - \mu_i)^2}, \quad i > a^*, i \neq i^*. \\
327 & \\
328 & \\
329 & \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_{i^*}(T)]}{\log T} \geq 2(1-\alpha)^2 \max \left\{ \frac{1}{\Delta_{Q,a^*}^2}, \max_{i < a^*} \frac{1}{\Delta_{Q,i}^2} \mathbf{1}[\min_{i < a^*} \Delta_{Q,i} \leq (1-\alpha)\Delta_{\min}] \right\}. \\
330 & \\
331 &
\end{aligned}$$

332 Where $\Delta_{\min} = \mu^* - \max_{i \neq i^*} \mu_i$ is the smallest conventional gap, T denotes the horizon and the
333 rewards of all K arms are Gaussian distributed with $\sigma^2 = 1$.

334 The proof of Theorem 3.3 closely parallels the proof of Theorem 3.1 barring two salient differences.
335 First, in Theorem 3.3 there is a non-trivial lower bound on the number of samples of high-cost arms.
336 This is because unlike the known reference arm setting the arm i^* and its reward μ^* are unidentified.
337 Second, the lower bound on samples of i^* is not only a function of quality gaps $\Delta_{Q,i}, i \leq a^*$, but
338 also depends on the relationship between $\min_{i \leq a^*} \Delta_{Q,i}$ and the smallest conventional gap Δ_{\min} .
339 Next we discuss the extension of PE to the subsidized best reward setting to develop PE-CS.

341 **Algorithm 2:** Pairwise Elimination for Cost Subsidy Problem (PE-CS).

342 **Inputs:** Bandit Instance ν , Horizon T , Subsidy Factor α .

343 **Initialize:** Samples $n_k = 0$, Empirical Means $\hat{\mu}_k = 0$, Current Rounds $\omega_k = 0, \forall k \in [K]$,
344 BAI Candidates (active arms) $\mathcal{A} = [K]$, Arm $\ell = \text{None}$, PE Episode $i = 1$.

```

345 1 while  $t \leq T$  do
346 2   if  $\text{len}(\mathcal{A}) > 1$  then
347 3      $k_t, \omega, \mathcal{A} \leftarrow \text{BAI}(\mathbf{n}, \hat{\boldsymbol{\mu}}, \omega, T, \mathcal{A})$  // receive arm to be sampled,
348     updated round numbers, and updated active arms
349 4     if  $\text{len}(\mathcal{A}) = 1$  then
350 5        $\ell \leftarrow \mathcal{A}[0]$  // set  $\ell$  to be identified best arm
351 6       continue // ignore sample recommendation  $k_t$ 
352 7     else if  $i \notin \{\text{None}, \ell\}$  then
353 8        $k_t, \omega, i \leftarrow \text{PE}(\mathbf{n}, \hat{\boldsymbol{\mu}}, \omega, T, i, \ell, \alpha)$  // receive arm to be sampled,
354       updated round numbers, and updated episode number
355 9     else
356 10       $k_t \leftarrow k_{t-1}$  // sample winning arm for remaining budget
357 11       $\hat{\boldsymbol{\mu}}(t+1), \mathbf{n}(t+1), t \leftarrow \text{sample\_and\_update}(k_t, \hat{\boldsymbol{\mu}}(t), \mathbf{n}(t), t)$  // in Appendix B

```

361 When solving the subsidized best reward setting we face the additional challenge that the arm i^* is
362 unidentified. To solve this problem, PE-CS (Algorithm 2) comprises two stages, the *BAI stage* and
363 the *PE stage*. The $\text{BAI}()$ subroutine uses the same *round number* based cadence of sampling and
364 elimination leveraged in $\text{PE}()$ and is specified in Appendix B, Function 3. Under $\text{BAI}()$ once the set
365 of active arms collapses to a single arm, sampling decisions are passed over to the PE stage.

366 The PE stage in the PE-CS algorithm works identically to the PE algorithm described in Algorithm 1.
367 In fact, due to the modular and phased nature of PE-CS we use precisely the same subroutine Func-
368 tion 1 for $\text{PE}()$. For the PE-CS algorithm, there is all the more reason to make best use of the
369 accumulated samples of the reference arm since not only can n_ℓ accumulate during episodes of the
370 PE stage, but also can accrue in the BAI stage where the empirical best reward arm ℓ is the last sur-
371 viving and therefore most sampled arm. In passing control from the BAI stage to the PE stage, a key
372 role is played by the ongoing rounds ω . Used in both $\text{BAI}()$ and $\text{PE}()$, ω is responsible for tracking
373 the round up to which any arm has been sampled at any point in time. Once the time comes for
374 passing over control from $\text{BAI}()$ to $\text{PE}()$, the PE stage is able to pick up sampling and elimination
375 checks in any of its pairwise comparison episodes right where the BAI stage left-off. For PE-CS we
376 show upper bounds on expected cumulative cost and quality regret in Theorem 3.4.

377 **Theorem 3.4** (Instance dependent upper bound on cumulative cost and quality regret for PE-CS).
For bandit instance ν , over horizon T , the expected cumulative cost regret $\mathbb{E}[\text{Cost_Reg}(T, \nu)]$ and

quality regret $\mathbb{E} [\text{Quality_Reg}(T, \nu)]$ of the PE-CS algorithm are upper bounded respectively as,

$$\begin{aligned} & \Delta_{C,i^*}^+ \left(1 + \max_{\Delta \in \mathbf{\Delta}} \left\{ \frac{32 \log(T\Delta^2)}{\Delta^2} \right\} \right) + \Delta_{C,i^*}^+ \left(\sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) \right) \\ & + \max_{i>a^*, i \in [K]} \Delta_{C,i}^+ \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) + \sum_{i>a^*, i \in [K] \setminus \{i^*\}} \Delta_{C,i}^+ \left(1 + \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \right) \\ & + \max_{i \in [K]} \Delta_{C,i}^+ \left(\frac{11}{\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{\Delta_j^2} \right). \end{aligned}$$

$$\begin{aligned} & \sum_{i=1}^{a^*-1} \left(\Delta_{Q,i} + \frac{32 \log(T\Delta_{Q,i}^2)}{\Delta_{Q,i}} \right) + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}} + \max_{i>a^*, i \in [K]} \Delta_{Q,i}^+ \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) \\ & + \sum_{i>a^*, i \in [K] \setminus \{i^*\}} \Delta_{Q,i}^+ \left(1 + \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \right) + \max_{i \in [K]} \Delta_{Q,i}^+ \left(\frac{11}{\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{\Delta_j^2} \right). \end{aligned}$$

Where $\Delta_i = \mu^* - \mu_i$ are conventional gaps, $\Delta_{\min} = \mu^* - \max_{i \neq i^*} \mu_i$ is the smallest conventional gap, and $\mathbf{\Delta} = \{\Delta_{\min}\} \cup \{\Delta_{Q,j}\}_{j \leq a^*}$ is defined especially to bound of arm i^* under expectation.

We bound the expected number of samples of any arm under PE-CS by first conditioning on the outcome of the BAI-stage being either proper ($\ell = i^*$) or improper ($\ell \neq i^*$). The phased nature of PE-CS admits a modular analysis where conditioned on the outcome of BAI being proper, the bounds on samples shown in Theorem 3.2 for PE hold with slight modifications. In particular the three leading terms in the PE-CS bounds correspond directly (in order) to the three terms in the PE bounds for cost and quality regret. The fourth term in both the PE-CS cost and quality regret bounds corresponds to samples of high cost arms accrued during the BAI stage and the fifth term is the constant contribution to regret from an improper BAI outcome. Proof is available in Appendix G.

Finally, as with Theorem 3.2, we see that both the expected cost and quality regret are bounded by a quantity that is $\mathcal{O}(\log T)$. On comparing the upper bounds in Theorem 3.4 to the regret lower bounds implied by Theorem 3.3 we find a gap between the two. This gap arises from the class of bandit instances where precisely identifying arm i^* using BAI detracts from the actual goal of sampling cheap feasible arms. A more detailed discussion is available in Appendix E.

4 EXPERIMENTS

While in Section 3 we presented the PE and PE-CS algorithms to be the ideal choice for the MAB-CS setting with known and unknown reference arms respectively, the validation for our approach so far has been based exclusively on theoretical upper bounds on cumulative expected cost and quality regret. In this section, we complement our theoretical analysis with a study of the empirical performance of our methods. In particular, we compare PE and PE-CS with baselines from Sinha et al. (2021) on a problem instance derived from a real-world dataset. The real-world dataset we use is the MovieLens 25M dataset (Harper & Konstan, 2015). Next we describe how we make use of this dataset for our experiments.

The MovieLens 25M dataset consists of 25 million ratings for 62,000 movies rated by 162,000 users. It is a popular dataset for studying the performance of recommendation systems. The movie ratings in the dataset are from numerous users and on a 5 point scale. Moreover, every movie for which ratings are available in the dataset are tagged with one or more genres. Our MovieLens experiment simulates a scenario where a movie streaming website such as Netflix attempts to recommend movies to its users. The available movie selection is organized into genres and the movie streaming website deploys a recommendation system to decide which genre of movie should be served to its user. Whenever a genre is chosen by the recommendation system, a random movie tagged with that genre is drawn with replacement. Under these conditions, bandit arms become a natural abstraction for movie genres. The cost associated with pulling the arm corresponding to a certain genre is simply the average of the royalties that must be paid to the movie producer for every new streaming of a

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

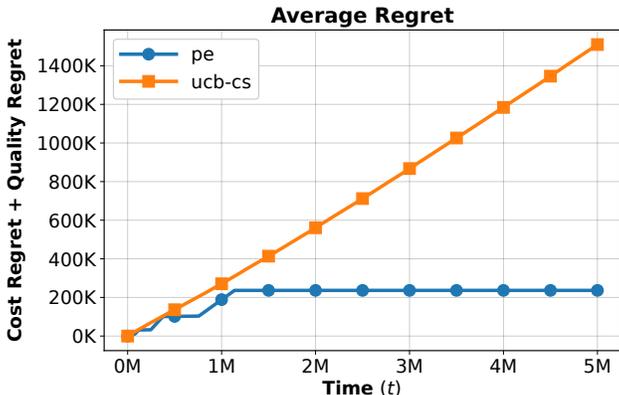


Figure 1: Fixed ℓ MovieLens Experiment. In the experiment: $\ell = 11$, $\mu_\ell = 0.703$, $\mu_{CS} = 0.668$.

movie. Since royalty data is unavailable as part of MovieLens 25M, we sample the costs associated with sampling any bandit arm (genre) to lie uniformly at random between 0 and 1.

For every genre we first obtain the mean 5-point scale rating of all movies tagged with that genre and then divide this rating by 5 so that it lies between 0 and 1. We then treat this fractional rating as the expected reward return from that genre. Through this process we end up with a bandit instance consisting of 20 arms corresponding to the 20 distinct genres the details of which are available in Appendix A.

Table 1. In all the experiments discussed in Section 4 we plot the summed together values of the cost regret and the quality regret. Looking at the summed regret allows us to view the performance of our method as a whole. The only way for an algorithm to perform well on summed regret is if it locks on satisfactorily to the best action a^* .

Number of Arms	Reward Spread	Cost Spread	Subsidy Factor α
20	0.659 - 0.785	0.02 - 0.964	0.05

Table 1: Information about the MovieLens Bandit Instance.

4.1 EVALUATING OUR PAIRWISE ELIMINATION (PE) ALGORITHM

To understand the effectiveness of PE empirically, we compare PE to a natural variant of the UCB-CS algorithm from Sinha et al. (2021). In the specification of UCB-CS (Algorithm 5 defined in Appendix B), the target reference arm (whose reward determines μ_{CS}) is the optimal arm i^* . UCB-CS estimates the index of arm i^* as the arm with the largest UCB-index in any time-slot. To develop a comparison with PE for the known reference arm ℓ setting, we simply replace this estimate with the true known fixed index of the reference arm while keeping the rest of the Algorithm the same. We call this variant of UCB-CS as UCB-CS Known ℓ and compare it to PE on the MovieLens Bandit Instance described earlier.

In Figure 1 we plot the performance of our novel PE algorithm and compare it against UCB-CS Known ℓ . Since there is no notion of a reference arm inherent to the MovieLens dataset, we arrange the 20 arms in the MovieLens bandit instance in ascending cost order and assign the 11th arm in the sequence to be the reference arm. To achieve sub-linear summed regret, an algorithm must primarily sample the best action and sample the other actions at most a sub-linear in horizon T number of times. Through the results of Figure 1 we find that while PE, through its cost-ordered arm-elimination scheme achieves this goal, UCB-CS known- ℓ is unsuccessful at it.

From the results in Section 4.1 we see that an interleaved exploration and exploitation approach such as UCB-CS known ℓ does not work for the known ℓ setting. Next in Section 4.2 we go into the experiments for PE-CS in the full cost-subsidy setting.

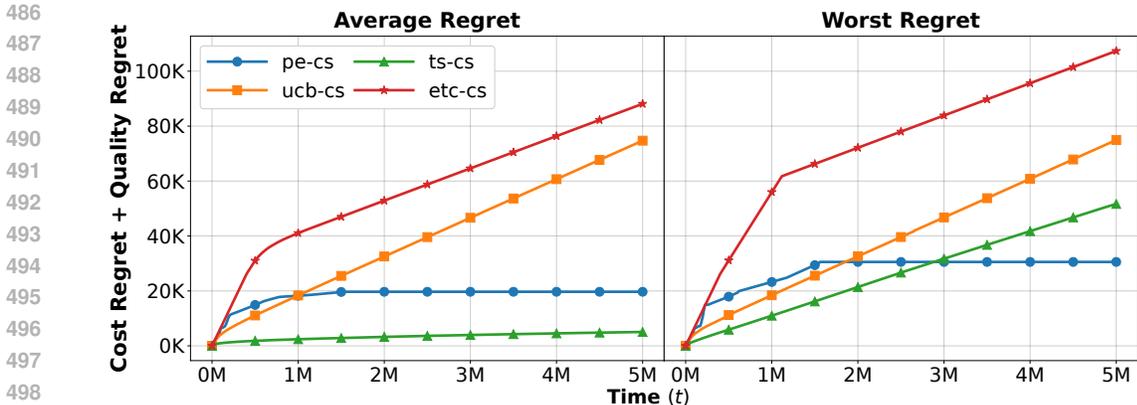


Figure 2: Unknown i^* MovieLens Experiment. In the experiment: $\mu^* = 0.785$, $\mu_{CS} = 0.746$.

4.2 EVALUATING OUR PE-CS (PAIRWISE ELIMINATION COST SUBSIDY) ALGORITHM

In Section 3, we saw that PE-CS admitted logarithmic instance dependent guarantees on expected cumulative cost and quality regret. The only algorithm for the full cost-subsidy problem from the literature that has an upper bound guarantee on expected cumulative regret is the ETC-CS algorithm from Sinha et al. (2021). Moreover their work also prescribes the UCB-CS and TS-CS algorithms which are approaches to solving the subsidized best reward problem that interleave exploration and exploitation that lack any performance guarantees. The three algorithms ETC-CS, UCB-CS, and TS-CS comprise all the algorithms from the literature and are specified in Appendix B. We compare PE-CS against all three of these approaches on the MovieLens bandit instance of Table 1.

In Figure 2 we have simulated PE-CS, UCB-CS, TS-CS, and ETC-CS on the MovieLens bandit instance for a horizon of 5 Million samples. We plot the summed cost and quality regret for an average over 100 independent runs on the left, and the worst performing case among the 100 runs for each algorithm on the right. The specifications of the baselines are available in Appendix B. Here the notion of worst is in terms of the summed terminal regret. ETC-CS consists of a pure exploration phase where each arm is sampled $\mathcal{O}(T^{2/3})$ times followed by an exploitation phase. As the only other algorithm with a regret guarantee ETC-CS is our primary competitor. We find that not only does the fixed budget τ end up sampling infeasible and high cost arms a lot more than PE-CS, but also it fails to exclude cheap infeasible arms from its empirically feasible set as evidenced by its linear regret trend in Figure 2.

Moreover PE-CS also outperforms the UCB-CS baseline and the latter algorithm has a linear regret trend arising from a persistent mis-identification of the best action. Although we find that the average of the regret over the 100 independent runs for TS-CS is lower than PE-CS, a closer examination of the regret trend reveals the problem with the performance of TS-CS. While initially it takes PE-CS more exploration to lock onto the best action, it does so in a consistent and reliable way and once it does, there is no further incremental regret. This is seen from both the average and worst-case regret trends in Figure 2 and from our upper bounds in Theorem 3.4. Whereas for TS-CS, while interleaving exploration and exploitation leads to lower regret at the outset, there is a distinct slow-but steady upward trend in regret observed for the method. The worst case summed regret trend reveals that a consistent failure to identify and exploit the best-action occurs for the worst performing case of TS-CS exhibiting its unreliability. The worst and other similar traces of TS-CS contributing linear regret disproportionately contribute to the upward regret trend for TS-CS.

A close examination of the Algorithm descriptions for UCB-CS and TS-CS reveals the source of their unreliability that we see in Figure 2. Both these algorithms work by constructing a set of empirically feasible arms and choose to sample the cheapest arm in this empirical set. This approach is vulnerable to a feasible arm being consistently excluded as a result of an initially poor performance. In our PE-CS on the other hand there is a systematic comparison between the arms where they are evaluated in the order of their costs and eliminated only when they have been sampled to be excluded with sufficient confidence.

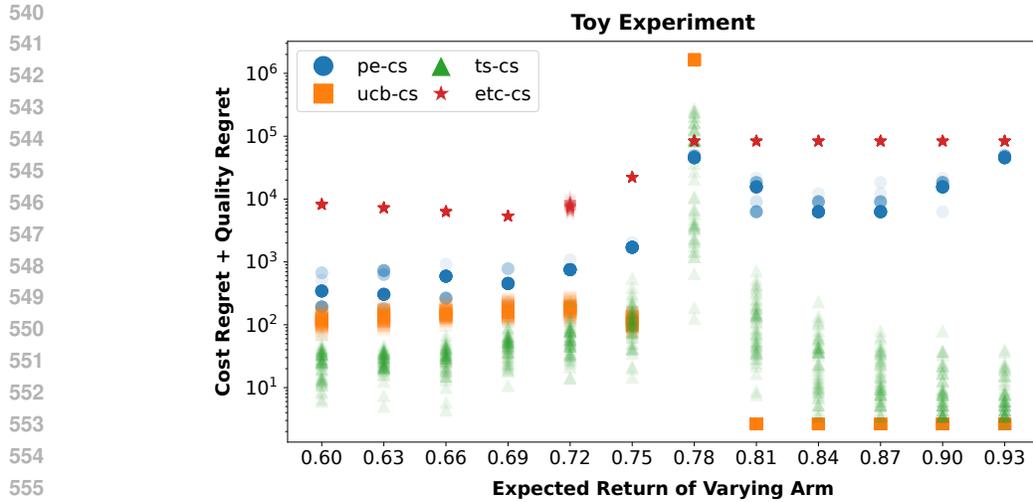


Figure 3: Toy Experiment performed using a family of four armed bandit instances. Expected rewards: $\mu_1 = 0.6, \dots, 0.93, \mu_2 = 0.81, \mu_3 = 0.95, \mu_4 = 0.8$, and costs: $c_1 = 0.05, c_2 = 0.9, c_3 = 0.9, c_4 = 1.0$, subsidy factor: $\alpha = 0.2$.

To take a closer look at the sensitivity of PE-CS and all three baselines, we perform an additional synthetic experiment in the known reference arm setting which we call the toy experiment. For the toy experiment, we create a family of 12 subsidized best reward problem instances each with four arms. Then we run all four algorithms over 50 independent runs of each instance. The expected reward of the first arm in the instance varies uniformly in the range 0.6 through 0.93 over the 12 instances in the family. Since the optimal return in all instances is $\mu^* = 0.95$ and the subsidy factor is $\alpha = 0.2$, the reward threshold μ_{CS} for all the instances is $0.8 \times 0.95 = 0.76$.

In Figure 3 we plot the results from the toy experiment on a scatter plot. On the y-axis is the summed terminal cost and quality regret (in log scale) and on the x-axis is the value of the varying expected return of the first arm of the instance family. Firstly, we find that on almost all instances, PE-CS performs either similar to or better than our primary comparator ETC-CS. Among the 12 instances tested here, the ones where the return of an arm is close to $\mu_{CS} = 0.76$ are the most challenging. We find that most runs of UCB-CS and several runs of TS-CS are unsuccessful at satisfactorily solving the $\mu_1 = 0.78$ case. Moreover, from Figures 2 and 3 we conclude that among the four algorithms tested here, PE-CS arguably offers the best balance between performance and reliability.

REFERENCES

- Peter Auer and Ronald Ortner. Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47:235–256, 2002.
- Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. Bandits with knapsacks. *J. ACM*, 65(3), mar 2018. ISSN 0004-5411. doi: 10.1145/3164539. URL <https://doi.org/10.1145/3164539>.
- Tianrui Chen, Aditya Gangrade, and Venkatesh Saligrama. Strategies for safe multi-armed bandits with logarithmic regret and risk. In *International Conference on Machine Learning*, pp. 3123–3148. PMLR, 2022.
- Aurélien Garivier, Pierre Ménard, and Gilles Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 44(2):377–399, 2019.
- Samarth Gupta, Shreyas Chaudhari, Gauri Joshi, and Osman Yağan. Multi-armed bandits with correlated arms. *IEEE Transactions on Information Theory*, 67(10):6711–6732, 2021.

594 F. Maxwell Harper and Joseph A. Konstan. The movielens datasets: History and context. *ACM*
595 *Trans. Interact. Intell. Syst.*, 5(4), dec 2015. ISSN 2160-6455. doi: 10.1145/2827872. URL
596 <https://doi.org/10.1145/2827872>.
597

598 Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-armed bandits in metric spaces. In
599 *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pp. 681–690, 2008.
600

601 Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
602

603 Deeksha Sinha, Karthik Abinav Sankararaman, Abbas Kazerouni, and Vashist Avadhanula. Multi-
604 armed bandits with cost subsidy. In Arindam Banerjee and Kenji Fukumizu (eds.), *Proceed-*
605 *ings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130
606 of *Proceedings of Machine Learning Research*, pp. 3016–3024. PMLR, 13–15 Apr 2021. URL
<https://proceedings.mlr.press/v130/sinha21a.html>.

607 Yifan Wu, Roshan Shariff, Tor Lattimore, and Csaba Szepesvari. Conservative bandits. In Maria Flo-
608 rina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference*
609 *on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1254–1262,
610 New York, New York, USA, 20–22 Jun 2016. PMLR. URL [https://proceedings.mlr.](https://proceedings.mlr.press/v48/wul16.html)
611 [press/v48/wul16.html](https://proceedings.mlr.press/v48/wul16.html).
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

A SUPPLEMENT TO EXPERIMENTS

B APPENDIX ALGORITHMS

In this Section we collect all the Algorithm blocks relevant to our work that couldn't find space in the main paper. First we present the `sample_and_update()` subroutine used in Algorithm 1 and Algorithm 2 and also used in the baselines in the current Appendix section.

Function 2: Sample and Update Function to update variables in-place and return them

Function `sample_and_update(k_t : Chosen Arm, $\hat{\mu}$: Empirical Means, \mathbf{n} : Sample Vector, t :`

`Time-Step):`

```

1   $r_t \leftarrow \text{sample}(k_t)$  // sample reward for arm  $k_t$ 
2   $\hat{\mu}_{k_t}(t+1) \leftarrow (\hat{\mu}_{k_t}(t) \times n_{k_t}(t) + r_t) / (n_{k_t}(t) + 1)$  // update mean
3   $n_{k_t}(t+1) \leftarrow n_{k_t}(t) + 1$  // update number of samples
4  return  $\hat{\mu}(t+1), \mathbf{n}(t+1), t+1$  // return updated means, samples, and
   time-step

```

Next we have the specification for the sub-routine `BAI()` used in PE-CS Algorithm 2.

Function 3: Best Arm Identification `BAI()`

Function `BAI(\mathbf{n} : Sample Vector, $\hat{\mu}$: Empirical Means, ω : Round Numbers, T : Horizon, \mathcal{A} :`

`Active Arms):`

```

1   $\hat{\Delta} \leftarrow 2^{-\max_{i \in [K]} \omega_i}$ 
2   $\tau \leftarrow \left\lceil \frac{2 \log(T \hat{\Delta}^2)}{\hat{\Delta}^2} \right\rceil$ 
3  for  $k \in \mathcal{A}$  do
4  |   if  $n_k < \tau$  then
5  |   |   return  $k, \omega, \mathcal{A}$  // next arm to be sampled, unchanged  $\omega$  and  $\mathcal{A}$ 
6   $\beta \leftarrow \sqrt{\frac{\log(T \hat{\Delta}^2)}{2\tau}}$ 
7  for  $i \in \mathcal{A}$  do
8  |    $\mu_i^{\text{UCB}}, \mu_i^{\text{LCB}} \leftarrow \hat{\mu}_i + \beta, \hat{\mu}_i - \beta$ 
9   $\mathcal{A}^+ \leftarrow \{i \in \mathcal{A} : \mu_i^{\text{UCB}} \geq \max_{j \in \mathcal{A}} \mu_j^{\text{LCB}}\}$  // update set of active arms
10  $k_t \leftarrow \text{Uniform}(\mathcal{A}^+)$  // tentatively, the next arm to be sampled
11 for  $i \in \mathcal{A}^+$  do
12 |    $\omega_i \leftarrow \omega_i + 1$  // increment round number for still active arms
13 return  $k_t, \omega, \mathcal{A}^+$ 

```

Next, we provide a precise specification of algorithms from prior work that we compare our methods against. In particular these are the ETC-CS, TS-CS, and UCB-CS Algorithms introduced by Sinha et al. (2021) and specified in Algorithms 3, 4, and 5 respectively. For the upper bound on regret provided in Sinha et al. (2021) to hold, we require the exploration budget τ to satisfy $\tau = c \left(\frac{T}{K}\right)^{\frac{2}{3}}$ where c is some unspecified proportionality constant. Based on a few trial runs and examples we pick $c = 5$ since in the average case, the above seemed to give the best performance overall for the ETC-CS approach.

Algorithm 3: Cost-Subsidized Explore-Then-Commit (ETC-CS)**Inputs:** Bandit instance ν , Cost Vector \mathbf{c} , Horizon T , Exploration Budget τ .**Initialize:** Empirical means $\hat{\mu}_k = 0$, Number of Samples $n_k = 0$, $\forall k \in [K]$.

```

702 1 while  $t \leq K\tau$  do
703   2    $k_t \leftarrow t \bmod K$ 
704   3    $\hat{\boldsymbol{\mu}}(t+1), \mathbf{n}(t+1), t \leftarrow \text{sample\_and\_update}(k_t, \hat{\boldsymbol{\mu}}(t), \mathbf{n}(t), t)$ 
705
706 4 while  $K\tau < t \leq T$  do
707   5   for  $i \in [K]$  do
708     6      $\beta_i(t) \leftarrow \sqrt{\frac{2 \log T}{n_i(t)}}$ 
709     7      $\mu_i^{\text{UCB}} \leftarrow \min\{\hat{\mu}_i(t) + \beta_i(t), 1\}$ 
710     8      $\mu_i^{\text{LCB}} \leftarrow \max\{\hat{\mu}_i(t) - \beta_i(t), 0\}$ 
711   9    $\ell_t \leftarrow \arg \max_{i \in [K]} \mu_i^{\text{LCB}}(t)$ 
712   10   $\text{Feas}(t) \leftarrow \{i : \mu_i^{\text{UCB}}(t) \geq (1 - \alpha)\mu_{\ell_t}^{\text{LCB}}(t)\}$ 
713   11   $k_t \leftarrow \arg \min_{i \in \text{Feas}(t)} c_i$ 
714   12   $\hat{\boldsymbol{\mu}}(t+1), \mathbf{n}(t+1), t \leftarrow \text{sample\_and\_update}(k_t, \hat{\boldsymbol{\mu}}(t), \mathbf{n}(t), t)$ 

```

Algorithm 4: Cost-Subsidized Thompson Sampling with Beta Priors (TS-CS)**Inputs:** Bandit Instance ν , Cost Vector \mathbf{c} , Subsidy Factor α , Beta Priors and Binomial Likelihood (Bernoulli Rewards).**Initialize:** Successes $S_k = 0$, Failures $F_k = 0 \forall k \in [K]$.

```

722 1 while  $t \leq K$  do
723   2    $k_t \leftarrow t$ 
724   3    $r_t \leftarrow \text{sample}(k_t)$ 
725   4    $S_{k_t}(t+1) \leftarrow S_{k_t}(t) + r_t$ 
726   5    $F_{k_t}(t+1) \leftarrow F_{k_t}(t) + 1 - r_t$ 
727   6    $t \leftarrow t + 1$ 
728
729 7 while  $K < t \leq T$  do
730   8   for  $i \in [K]$  do
731     9      $\theta_i(t) \sim \text{Beta}(S_i(t) + 1, F_i(t) + 1)$ 
732   10   $\ell_t \leftarrow \arg \max_{i \in [K]} \theta_i(t)$ 
733   11   $\text{Feas}(t) \leftarrow \{i : \theta_i(t) \geq (1 - \alpha)\theta_{\ell_t}(t)\}$ 
734   12   $k_t \leftarrow \arg \min_{i \in \text{Feas}(t)} c_i$ 
735   13   $r_t \leftarrow \text{sample}(k_t)$ 
736   14   $S_{k_t}(t+1) \leftarrow S_{k_t}(t) + r_t$ 
737   15   $F_{k_t}(t+1) \leftarrow F_{k_t}(t) + (1 - r_t)$ 

```

C ASYMMETRIC PAIRWISE ELIMINATION

Algorithms 1 and 2 as described in Section 3 use Function 1 PE() as a sub-routine. While in PE() the round number ω_i is used to determine the stipulated number of samples for both the candidate arm i and the reference arm ℓ , this does not have to be the case. By the time we commence episode i to evaluate the candidacy of arm i , we would have already accrued numerous samples of arm ℓ . In particular, we denote the number of samples of arm ℓ as $n_\ell(t) = \tau_{\omega_\ell}$, where the vector $\boldsymbol{\omega}$, first introduced in Section 3, is a vector recording highest round up to which the samples of each arm have evolved. Consequently, ω_ℓ is the highest round number reached for the samples of reference arm ℓ .

Based on this observation about the greater progressed round number ω_ℓ we create a variant of PE called Asymmetric-PE in Function 4 that replaces Function 1 in Algorithm 1.

Asymmetric-PE described in Function 4 has all the same inputs that conventional PE did. In addition, it has an input κ called the Maximum round deviation. While we have no bound on how much

Algorithm 5: Cost-Subsidized UCB (UCB-CS)**Inputs:** Bandit Instance ν , Cost Vector \mathbf{c} , Horizon T , Subsidy factor α .**Initialize:** Empirical means $\hat{\mu}_k = 0$, Number of Samples $n_k = 0 \forall k \in [K]$.

```

1 while  $t \leq K$  do
2    $k_t \leftarrow t$ 
3    $\hat{\boldsymbol{\mu}}(t+1), \mathbf{n}(t+1), t \leftarrow \text{sample\_and\_update}(k_t, \hat{\boldsymbol{\mu}}(t), \mathbf{n}(t), t)$ 
4 while  $K < t \leq T$  do
5   for  $i \in [K]$  do
6      $\beta_i(t) \leftarrow \sqrt{\frac{2 \log T}{n_i(t)}}$ 
7      $\mu_i^{\text{UCB}} \leftarrow \min \{ \hat{\mu}_i(t) + \beta_i(t), 1 \}$ 
8    $\ell_t \leftarrow \arg \max_{i \in [K]} \mu_i^{\text{UCB}}(t)$ 
9    $\text{Feas}(t) \leftarrow \{ i \in [K] : \mu_i^{\text{UCB}} \geq (1 - \alpha) \times \mu_{\ell_t}^{\text{UCB}} \}$ 
10   $k_t \leftarrow \arg \min_{i \in \text{Feas}(t)} c_i$ 
11   $\hat{\boldsymbol{\mu}}(t+1), \mathbf{n}(t+1), t \leftarrow \text{sample\_and\_update}(k_t, \hat{\boldsymbol{\mu}}(t), \mathbf{n}(t), t)$ 

```

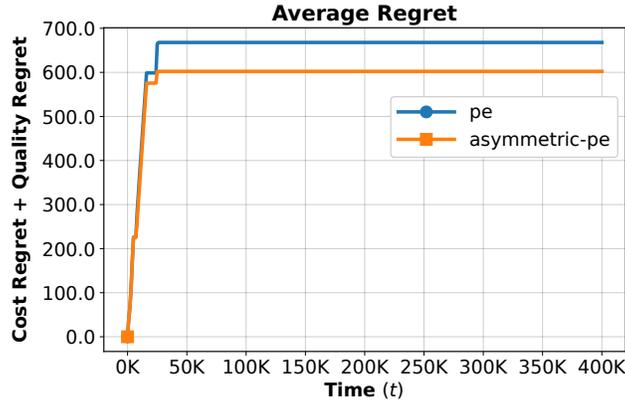


Figure 4: An experiment that illustrates potential savings in regret from the asymmetric-pe optimization. Bandit instance with Expected rewards: $\mu_1 = 0.74, \mu_2 = 0.5, \mu_3 = 0.8, \mu_4 = 0.75$, and costs: $c_1 = 0.15, c_2 = 0.2, c_3 = 0.21, c_4 = 0.25$, subsidy factor: $\alpha = 0.0, \ell = 3$.

larger ω_ℓ is compared to ω_i , when it comes to inferring the gap $\tilde{\Delta}_\ell$ corresponding to arm ℓ in Line 2 of Function 4, we restrict the round number we use to be at most κ larger than the round ω_i .

Effectively, the use of the further advanced round number ω_ℓ trades performance for tightness of the upper bound (as we shall see in the analysis of Section F). We get an improvement in performance since $\beta_\ell \leq \beta_i$ potentially leading to a resolution in lines 7 or 9 of Function 4 with a smaller value of β_i thereby requiring fewer samples of candidate arm i .

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

Function 4: Asymmetric Pairwise Elimination Function `Asymmetric_PE()`

Function `Asymmetric_PE`($\hat{\mu}$: Empirical Means, ℓ : Reference Arm, \mathbf{n} : Sample Vector, T :

Horizon, ω : Round Numbers, i : Episode, α : Subsidy Factor, κ : Max Round Deviation):

```

1  for  $k \in \{i, \ell\}$  do
2  |    $\tilde{\Delta}_k \leftarrow 2^{-\max\{\omega_i + \kappa, \omega_k\}}$ 
3  |    $\tau_k \leftarrow \left\lceil \frac{2 \log(T \tilde{\Delta}_k^2)}{\tilde{\Delta}_k^2} \right\rceil$ 
4  |   if  $n_k < \tau_k$  then
5  |   |   return  $k, \omega, i$            //  $k_t = k$ , round numbers  $\omega$  and episode  $i$ 
5  |   |   unchanged
6  |   |    $\beta_k \leftarrow \sqrt{\frac{\log(T \tilde{\Delta}_k^2)}{2\tau_k}}$ 
7  |   if  $(1 - \alpha)(\hat{\mu}_\ell + \beta_\ell) < \hat{\mu}_i - \beta_i$  then
8  |   |   return  $i, \omega, \text{None}$      // Declare  $i$  as winner, further episodes are
8  |   |   None
9  |   else if  $\hat{\mu}_i + \beta_i < (1 - \alpha)(\hat{\mu}_\ell - \beta_\ell)$  then
10 |   |   return  $i + 1, \omega, i + 1$  // Sample next candidate arm, Rounds  $\omega$ 
10 |   |   unchanged, Update episode to that of next candidate arm
11 |   else
12 |   |    $\omega_i \leftarrow \omega_i + 1$  // Increment round only for arm  $i$  being evaluated
13 |   |   return  $i, \omega, i$        // Move to next round within same episode

```

D PRELIMINARIES

In this section of the appendix, we collate the preliminary results required for the analysis of our cost-subsidy framework algorithms that follow in the forthcoming sections. The results stated without proof are standard results from the MAB literature.

Definition D.1 (Subgaussian Random Variable). *We say that X is σ -subgaussian if for any $\epsilon \geq 0$,*

$$\Pr(X - \mathbb{E}[X] \geq \epsilon) \leq \exp\left(\frac{-\epsilon^2}{2\sigma^2}\right). \quad (1)$$

Lemma D.1 (Bounded random variables are Subgaussian, Example 5.6(c) in Lattimore & Szepesvári (2020)). *If Random Variable $X \in [a, b]$ almost surely, then X is $\frac{b-a}{2}$ subgaussian.*

Lemma D.2 (Hoeffding Bound, Section 5.4 in Lattimore & Szepesvári (2020)). *Let X_1, X_2, \dots, X_n be n independent random variables, each bounded within the interval $[a, b]$: $a \leq X_i \leq b$. The empirical mean of these variables is given by,*

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i. \quad (2)$$

Then Hoeffding's inequality states that,

$$\Pr(\bar{X} - \mathbb{E}[X] \geq t) \leq \exp\left(-\frac{2nt^2}{(b-a)^2}\right). \quad (3)$$

Lemma D.3 (Iterated expectation over mutually exclusive and exhaustive events). *Let X be any integrable random variable over probability space $(\Omega, \mathcal{F}, \Pr)$, and let $\{E_i\}_{i=1}^n$ be a collection of mutually exclusive and exhaustive measurable events. That is $\bigcup_{i=1}^n E_i = \Omega$ and $E_i \cap E_j = \phi, \forall i, j \in [n], i \neq j$. Then the following identity holds,*

$$\mathbb{E}[X] = \sum_{i=1}^n \mathbb{E}[X | E_i] \Pr(E_i). \quad (4)$$

As a special case if the events are just some E and its complement E^c , then,

$$\mathbb{E}[X] = \mathbb{E}[X | E] \Pr(E) + \mathbb{E}[X | E^c] \Pr(E^c). \quad (5)$$

Proof. Define a sub σ -algebra of \mathcal{F} , $\mathcal{G} = \{\phi, E_1, E_2, \dots, E_n, \Omega\}$. Then,

$$\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X | \mathcal{G}]] \quad (\text{Because } \mathcal{G} \subset \mathcal{F}) \quad (6)$$

$$= \sum_{i=1}^n \mathbb{E}[X | E_i] \Pr(E_i). \quad (7)$$

□

Lemma D.4 (Expectation is at most equal to larger of the conditioned expectations). *Let X be any integrable random variable over probability space $(\Omega, \mathcal{F}, \Pr)$, and let $\{E_i\}_{i=1}^n$ be a collection of mutually exclusive and exhaustive measurable events. That is $\bigcup_{i=1}^n E_i = \Omega$ and $E_i \cap E_j = \phi, \forall i, j \in [n], i \neq j$. Then,*

$$\mathbb{E}[X] \leq \max_{i \in [n]} \{\mathbb{E}[X | E_i]\}. \quad (8)$$

Proof. Lemma D.4 can be considered a Corollary to Lemma D.3 as is illustrated by the following proof,

$$\mathbb{E}[X] = \sum_{i=1}^n \mathbb{E}[X | E_i] \Pr(E_i) \quad (\text{From the proof of Lemma D.3}). \quad (9)$$

$$\leq \left(\sum_{i=1}^n \Pr(E_i) \right) \cdot \left(\max_{i \in [n]} \{\mathbb{E}[X | E_i]\} \right) \quad (10)$$

$$= \max_{i \in [n]} \{\mathbb{E}[X | E_i]\}. \quad (11)$$

923 □

924 **Lemma D.5** (Regret Decomposition Lemma, Lemma 4.5 in Lattimore & Szepesvári (2020)). *For*
 925 *any policy π and stochastic bandit environment ν with K arms, for horizon T , the Expected Cumu-*
 926 *lative Regret $\text{Reg}_\pi(T, \nu)$ of policy π in ν satisfies,*
 927

$$\mathbb{E}[\text{Reg}_\pi(T, \nu)] = \sum_{i \in [K]} \Delta_i \mathbb{E}[n_i(T)]. \quad (12)$$

930 *This result may be trivially generalized to other notions of regret where the gap determining the*
 931 *incremental regret due to arm i is some arbitrary $\Delta_{X,i}$. In this case, the regret decomposition shall*
 932 *be,*
 933

$$\mathbb{E}[\text{Reg}_\pi^X(T, \nu)] = \sum_{i \in [K]} \Delta_{X,i} \mathbb{E}[n_i(T)]. \quad (13)$$

936 *In particular in our problem we have Cost and Quality regret which are,*
 937

$$\mathbb{E}[\text{Cost_Reg}(T, \nu)] = \sum_{i \in [K]} \Delta_{C,i} \mathbb{E}[n_i(T)] \quad (14)$$

$$\mathbb{E}[\text{Quality_Reg}(T, \nu)] = \sum_{i \in [K]} \Delta_{Q,i}^+ \mathbb{E}[n_i(T)]. \quad (15)$$

943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

E INSTANCE DEPENDENT LOWER BOUNDS FOR MAB-CS

LOWER BOUNDS FOR THE KNOWN REFERENCE ARM SETTING

We consider the case where the reference arm, denoted by arm ℓ , is known. Let $\nu = \{\nu_i\}_{i=1,\dots,K,\ell}$ denote the vector of reward distributions associated with all arms $i = 1, \dots, K, \ell$, where the index of the reference arm can be thought of as $\ell = K + 1$. Let \mathcal{M} denote the *model* under consideration that is a collection of all possible instances ν . Recall that without loss of generality we assume the arm indices are cost ordered, i.e., $c_1 < c_2 < \dots < c_K < c_\ell$, and that a^* denotes the optimal action defined as the cheapest arm whose mean reward meets the *feasibility* threshold μ_{CS} . In the known reference arm setting, the feasibility threshold is defined as $\mu_{CS} = (1 - \alpha)\mu_\ell$ for some $\alpha \in [0, 1]$ where μ_ℓ is the mean reward of the reference arm ℓ , which yields $a^* = \arg \min_{\{i: \mu_i \geq (1-\alpha)\mu_\ell\}} c_i$.

Definition E.1 (Consistent Policies). *A policy π is consistent if for all bandit instances ν and for all arms $i \neq a^*$, $\mathbb{E}[n_i(T)] = o(T^\gamma)$ for all $0 < \gamma \leq 1$.*

Definition E.2 (The key quantity D_{inf}). *Let KL denote the Kullback-Leibler divergence between two probability distributions. Given a distribution $\nu_i \in \mathcal{M}$ and a real number x , we define*

$$D_{\text{inf}}(\nu_i, x) = \inf \{KL(\nu_i, \nu'_i) : \nu'_i \in \mathcal{M} \text{ and } E[\nu'_i] > x\},$$

where $E[\nu'_i]$ denotes the mean of the distribution ν'_i . We also need a slightly modified version of this key quantity defined as follows:

$$\tilde{D}_{\text{inf}}(\nu_i, x) = \inf \{KL(\nu_i, \nu'_i) : \nu'_i \in \mathcal{M} \text{ and } E[\nu'_i] \leq x\}.$$

Although not always explicitly stated, we will be using $D_{\text{inf}}(\nu_i, x)$ for distributions ν_i with $\mu_i < x$ and $\tilde{D}_{\text{inf}}(\nu_i, x)$ for distributions ν_i with $\mu_i > x$.

Proofs of both Theorem E.2 and Theorem E.1 follow arguments similar to those used in establishing the well-known lower bounds for the classical multi-armed bandit setting. Here, we adopt the approach presented in Garivier et al. (2019) that is based on the following fundamental inequality (Garivier et al. (2019)[F]): For any ν, ν' ,

$$\sum_{i \in \{1, \dots, K, \ell\}} \mathbb{E}_\nu[n_i(T)] KL(\nu_i, \nu'_i) \geq \text{kl}(\mathbb{E}_\nu[Z], \mathbb{E}_{\nu'}[Z]) \quad (16)$$

where kl denotes the Kullback-Leibler divergence for Bernoulli distributions, i.e.,

$$\forall p, q \in [0, 1]^2, \quad \text{kl}(p, q) = p \log \frac{p}{q} + (1 - p) \log \frac{1 - p}{1 - q},$$

and where Z is any random variable with values in $[0, 1]$ (that is measurable with respect to the probability space that all reward random variables are defined on).

We will use (16) with $Z = \frac{n_i(T)}{T}$ for some arm i , and will also rely on the bound

$$\text{kl}(p, q) = p \log \frac{p}{q} + (1 - p) \log \frac{1 - p}{1 - q} \geq (1 - p) \log \frac{1}{1 - q} - \log 2 \quad (17)$$

Theorem E.1 (Lower bounds for the number of pulls of low-cost arms). *For any bandit instance ν over horizon T , for all consistent strategies, we have*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu[n_i(T)]}{\log T} \geq \frac{1}{D_{\text{inf}}(\nu_i, (1 - \alpha)\mu_\ell)}, \quad \text{for arms } i = 1, \dots, a^* - 1.$$

Proof. Given any bandit instance ν and any arm $i = 1, \dots, a^* - 1$, we consider a modified instance ν' where i is the unique best action: $\nu'_j = \nu_j$ for all $j \in \{1, \dots, K, \ell\} - \{i\}$ and ν'_i is any distribution in \mathcal{M} such that its expectation μ'_i satisfies $\mu'_i \geq (1 - \alpha)\mu_\ell$. The key idea here is that any consistent policy should have $\mathbb{E}_\nu[n_i(T)] = o(T^\gamma)$ while $\mathbb{E}_{\nu'}[n_i(T)] = T - o(T^\gamma)$. We now apply the fundamental inequality (16) with $Z = \frac{n_i(T)}{T}$. Noting that all KL terms except that for arm i is zero on the left hand side and using (17), we get

$$\mathbb{E}_\nu[n_i(T)] KL(\nu_i, \nu'_i) \geq \text{kl} \left(\frac{\mathbb{E}_\nu[n_i(T)]}{T}, \frac{\mathbb{E}_{\nu'}[n_i(T)]}{T} \right)$$

$$\geq \left(1 - \frac{\mathbb{E}_\nu[n_i(T)]}{T}\right) \log \left(\frac{T}{T - \mathbb{E}_{\nu'}[n_i(T)]}\right) - \log 2. \quad (18)$$

Given that arm $i \neq a^*$ in instance ν and $i = a^*$ in ν' , any consistent policy must have $\mathbb{E}_\nu[n_i(T)] = o(T^\gamma)$ and $\mathbb{E}_{\nu'}[n_i(T)] = T - o(T^\gamma)$ which leads to,

$$\liminf_{T \rightarrow \infty} \frac{1}{\log T} \left(1 - \frac{\mathbb{E}_\nu[n_i(T)]}{T}\right) \log \left(\frac{T}{T - \mathbb{E}_{\nu'}[n_i(T)]}\right) \geq 1 - \gamma,$$

for all $0 < \gamma \leq 1$. Thus, we get

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu[n_i(T)]}{\log T} \geq \frac{1}{\text{KL}(\nu_i, \nu'_i)} \quad (19)$$

and taking the supremum in the right-hand side over all distributions $\mu'_i \geq (1 - \alpha)\mu_\ell$, we get the bound of Theorem E.1.

Remark E.1. *The line of arguments used in the proof given above would not yield a similar lower bound for high-cost arms, i.e., arms $i = a^* + 1, \dots, K$. This is because even if the reward distribution of an arm $i \in \{a^* + 1, \dots, K\}$ is modified to make its mean exceed the feasibility threshold $(1 - \alpha)\mu_\ell$, arm a^* would remain to be the optimal action and thus we would not be able to state that $\mathbb{E}_{\nu'}[n_i(T)] = T - o(T^\gamma)$. This is consistent with the upper bounds established in Lemma F.7 for our PE algorithm where the number of samples for high-cost arms was seen to be upper bounded by a $O(1)$ term. In fact, the upper bounds for the PE algorithm will be seen to match within a constant factor the lower bounds for all sub-optimal arms; see Remark E.2 below.*

Theorem E.2 (Lower bounds for the number of pulls of the reference arm ℓ). *For any bandit instance ν over horizon T , for all consistent strategies, we have*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_\ell(T)]}{\log T} \geq \max \left\{ \frac{1}{D_{\text{inf}}\left(\nu_\ell, \frac{\mu_{a^*}}{1-\alpha}\right)}, \max_{i=1, \dots, a^*-1} \frac{1}{\tilde{D}_{\text{inf}}\left(\nu_\ell, \frac{\mu_i}{1-\alpha}\right)} \right\}. \quad (20)$$

Proof. Proof of Theorem E.2 follows similar arguments to that of Theorem E.1 but requires a more subtle application of the fundamental inequality (16). We will first establish the second part of (20) (associated with modifications that make an arm $i < a^*$ optimal) followed by the first part (that arises from modifications that make an arm $i > a^*$ optimal).

Consider any bandit instance ν with arm a^* being the optimal action. If $a^* = \ell$, then any consistent policy must choose it at least $T - o(T^\gamma)$ times and the bound given in the Theorem follows immediately. Next, we assume $a^* \neq \ell$. Consider any arm $i = 1, \dots, a^* - 1$ and note that we must have $\mu_i < \mu_{a^*}$ for all such arms. Now, consider a modified instance ν' , where $\nu'_j = \nu_j$ for all $j = 1, \dots, K$ and ν'_ℓ is any distribution in \mathcal{M} such that its expectation μ'_ℓ satisfies $(1 - \alpha)\mu'_\ell \leq \mu_i$. This means that arm i is feasible in the modified instance ν' , and given that it has lower cost than a^* , a^* is no longer the optimal action. In fact, in ν' , the optimal action is guaranteed to be among the arms $\{1, \dots, i\}$. Let $j \in \{1, \dots, i\}$ be the new optimal action. We will apply the fundamental inequality (16) with $Z = \frac{n_j(T)}{T}$ and note that any consistent policy should have $\mathbb{E}_\nu[n_j(T)] = o(T^\gamma)$ while $\mathbb{E}_{\nu'}[n_j(T)] = T - o(T^\gamma)$. Proceeding similarly with the proof of Theorem E.1, we get

$$\begin{aligned} \mathbb{E}_\nu[n_\ell(T)] \text{KL}(\nu_\ell, \nu'_\ell) &\geq \text{kl} \left(\frac{\mathbb{E}_\nu[n_j(T)]}{T}, \frac{\mathbb{E}_{\nu'}[n_j(T)]}{T} \right) \\ &\geq \left(1 - \frac{\mathbb{E}_\nu[n_j(T)]}{T}\right) \log \left(\frac{T}{T - \mathbb{E}_{\nu'}[n_j(T)]}\right) - \log 2. \end{aligned} \quad (21)$$

Noting that $\mathbb{E}_\nu[n_j(T)] = o(T^\gamma)$ and $\mathbb{E}_{\nu'}[n_j(T)] = T - o(T^\gamma)$ and using similar arguments as in the proof of Theorem E.1 we get

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu[n_\ell(T)]}{\log T} \geq \frac{1}{\text{KL}(\nu_\ell, \nu'_\ell)} \quad (22)$$

Taking the supremum in the right-hand side over all distributions $\nu'_\ell \in \mathcal{M}$ with $\mu'_\ell \leq \frac{\mu_i}{1-\alpha}$, we get

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_\ell(T)]}{\log T} \geq \frac{1}{\tilde{D}_{\text{inf}}\left(\nu_\ell, \frac{\mu_i}{1-\alpha}\right)} \quad (23)$$

Finally, since these arguments can be applied for any arm $i = 1, \dots, a^* - 1$, we get

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_\ell(T)]}{\log T} \geq \max_{i=1, \dots, a^*-1} \frac{1}{\tilde{D}_{\text{inf}}\left(\nu_\ell, \frac{\mu_i}{1-\alpha}\right)}. \quad (24)$$

To complete the proof of (20), we now consider a modified distribution that would make a^* sub-optimal. Consider any bandit instance ν with arm $a^* \neq \ell$ being the optimal action. Consider a modified instance ν' , where $\nu'_j = \nu_j$ for all $j = 1, \dots, K$ and ν'_ℓ is any distribution in \mathcal{M} such that its expectation μ'_ℓ satisfies $(1-\alpha)\mu'_\ell > \mu_{a^*}$. With this modification, the arm a^* is no longer feasible and thus not optimal. The optimal action should then be in the set $\{a^* + 1, \dots, K, \ell\}$. Let $j \in \{a^* + 1, \dots, K, \ell\}$ be the new optimal action. We will apply the fundamental inequality (16) with $Z = \frac{n_j(T)}{T}$ and note that any consistent policy should have $\mathbb{E}_\nu[n_j(T)] = o(T^\gamma)$ (since in the original instance $a^* \neq j$ was optimal) while $\mathbb{E}_{\nu'}[n_j(T)] = T - o(T^\gamma)$. Using the same ideas as before, we then get

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_\ell(T)]}{\log T} \geq \frac{1}{\text{KL}(\nu_\ell, \nu'_\ell)} \quad (25)$$

and taking the supremum in the right-hand side over all distributions $\nu'_\ell \in \mathcal{M}$ with $\mu'_\ell > \frac{\mu_{a^*}}{1-\alpha}$, we get

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_\ell(T)]}{\log T} \geq \frac{1}{D_{\text{inf}}\left(\nu_\ell, \frac{\mu_{a^*}}{1-\alpha}\right)} \quad (26)$$

Finally, we combine 24 and 26 to complete the proof of 20.

Remark E.2. When the reward distributions are Gaussian, i.e., the family \mathcal{M} of distributions contain all Gaussian distributions with a common variance σ^2 , the key terms $D_{\text{inf}}(\nu_i, x)$ and $\tilde{D}_{\text{inf}}(\nu_i, x)$ take the form of

$$D_{\text{inf}}(\nu_i, x) = \tilde{D}_{\text{inf}}(\nu_i, x) = \frac{(\mu_i - x)^2}{2\sigma^2}. \quad (27)$$

More generally, it was shown in Lattimore & Szepesvári (2020) that $D_{\text{inf}}(\nu_i, x) = O((\mu_i - x)^2)$ for most distribution families. Using this with $\sigma^2 = 1$, we obtain from Theorem E.1 and Theorem E.2 that

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_i(T)]}{\log T} \geq \frac{2}{\Delta_{Q,i}^2}, \quad \text{for arms } i = 1, \dots, a^* - 1.$$

and

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_\ell(T)]}{\log T} \geq \max_{i \leq a^*} \frac{2(1-\alpha)^2}{\Delta_{Q,i}^2}.$$

where $\Delta_{Q,i}$ is defined as before, i.e., $\Delta_{Q,i} = (1-\alpha)\mu_\ell - \mu_i$. These results match within a constant factor the upper bounds of our PE algorithm established in Lemmas F.2 and F.3 showing that the PE algorithm is order-wise optimal.

LOWER BOUNDS FOR THE SUBSIDIZED BEST REWARD SETTING

We now present lower bounds for the subsidized best reward setting. The core ideas in establishing these results are the same with those used in the *known reference arm* case. Namely, to establish a lower bound on the number of pulls needed from a given arm, we are seeking a modification in that arm's reward distribution that leads to a change in the optimal arm, i.e., arm a^* being no

longer optimal. Among the modified distributions that change the optimal arm, the one that has the minimum KL-divergence with the original distribution leads to the tightest lower bound on that arm's number of pulls. We will repeatedly use this argument in the following discussion but the repetitive parts of the proofs are omitted for brevity.

Let \mathcal{M} denote the *model* under consideration that contains all bandit instances ν , where ν specifies the reward distributions ν_i of all arms $i = 1, \dots, K$. Recall that the arm indices are cost ordered, i.e., $c_1 < \dots < c_{a^*} < \dots < c_{i^*} < \dots < c_K$, where a^* is the optimal action and i^* is the arm with the largest mean reward; namely $i^* = \arg \max_{i \in \{1, \dots, K\}} \mu_i$ and $a^* = \arg \min_{\{i: \mu_i \geq (1-\alpha)\mu^*\}} c_i$ where $\mu^* = \max_{i \in \{1, \dots, K\}} \mu_i = \mu_{i^*}$. The definition of a consistent policy and key quantities D_{inf} , \tilde{D}_{inf} remain as in the previous section.

LOWER BOUND ON THE NUMBER OF SAMPLES FOR ARMS $i < a^*$

Theorem E.3 (Lower bounds for the number of pulls of *low-cost* arms). *For any bandit instance ν over horizon T , for all consistent strategies, we have*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_i(T)]}{\log T} \geq \frac{1}{D_{\text{inf}}(\nu_i, (1-\alpha)\mu^*)}, \quad \text{for arms } i = 1, \dots, a^* - 1.$$

Proof. For arms cheaper than the optimal action a^* , i.e., arms $i = 1, \dots, a^* - 1$, our approach is to modify their reward distribution so that their mean reward is now larger than the feasibility threshold $(1-\alpha)\mu^*$. Since these arms have a lower cost than a^* , a^* will no longer be optimal in the modified instance, paving the way to establishing a lower bound for the number of arms $i = 1, \dots, a^* - 1$. Formally, given any bandit instance ν and any arm $i = 1, \dots, a^* - 1$, we consider a modified instance ν' where i is the unique best action: $\nu'_j = \nu_j$ for all $j \in \{1, \dots, K\} - \{i\}$ and ν'_i is any distribution in \mathcal{M} such that its expectation μ'_i satisfies $\mu'_i > (1-\alpha)\mu_{i^*}$. In the modified instance ν' , the feasibility threshold is either $(1-\alpha)\mu_{i^*}$ as before, or it changes to $(1-\alpha)\mu'_i$ by virtue of arm i being the arm with the largest mean reward. In either case, μ'_i exceeds this threshold and since $c_i < c_{a^*}$, arm a^* becomes a sub-optimal action. This proves that

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_i(T)]}{\log T} \geq \frac{1}{\text{KL}(\nu_i, \nu'_i)} \quad (28)$$

and taking the supremum in the right-hand side over all distributions $\nu'_i \in \mathcal{M}$ with $\mu'_i > (1-\alpha)\mu^*$, we get the bound of Theorem E.3.

Remark E.3. *For arms $i = 1, \dots, a^* - 1$, recall that the quality gaps are defined as $\Delta_{Q,i} = (1-\alpha)\mu^* - \mu_i$. Recalling Remark E.2, we see from Theorem E.3 that when the reward distributions are Gaussian with $\sigma^2 = 1$*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_i(T)]}{\log T} \geq \frac{2}{\Delta_{Q,i}^2}, \quad \text{for arms } i = 1, \dots, a^* - 1.$$

matching the upper bound of our PE-CS algorithm given in Lemma G.3.

LOWER BOUND ON THE NUMBER OF SAMPLES FOR ARMS $i > a^*, i \neq i^*$

Theorem E.4 (Lower bounds for the number of pulls of *high-cost* arms). *For any bandit instance ν over horizon T , for all consistent strategies, we have*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_i(T)]}{\log T} \geq \frac{1}{D_{\text{inf}}\left(\nu_i, \frac{\mu_{a^*}}{(1-\alpha)}\right)}, \quad \text{for arms } i \in \{a^* + 1, \dots, K\} - \{i^*\}.$$

Proof. For arms more expensive than a^* but with mean reward less than μ^* , the only modification we can make in their reward distribution to render a^* sub-optimal would be to increase their mean reward to a level where they will be the highest-mean arm and $(1-\alpha)$ times their new mean exceeds μ_{a^*} (so that a^* is no longer feasible and thus can not be optimal). Formally, given any bandit instance

ν and any arm $i \in \{a^* + 1, \dots, K\} - \{i^*\}$, we consider a modified instance ν' such that $\nu'_j = \nu_j$ for all $j \in \{1, \dots, K\} - \{i\}$ and ν'_i is any distribution in \mathcal{M} such that its expectation μ'_i satisfies $(1 - \alpha)\mu'_i > \mu_{a^*}$. In the modified instance ν' , we have $\mu'_i > \mu_{a^*}/(1 - \alpha) \geq \mu_{i^*}$ with the second inequality following from $\mu_{a^*} \geq (1 - \alpha)\mu_{i^*}$ since a^* was feasible in the original instance. Thus, in the modified instance arm i has the highest mean and the feasibility threshold becomes $(1 - \alpha)\mu'_i$ which exceeds the mean reward μ_{a^*} of arm a^* by construction. Thus, in the modified instance ν' , arm a^* is no longer the optimal action leading to

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_i(T)]}{\log T} \geq \frac{1}{\text{KL}(\nu_i, \nu'_i)} \quad (29)$$

and taking the supremum in the right-hand side over all distributions $\nu'_i \in \mathcal{M}$ with $\mu'_i > \frac{\mu_{a^*}}{1 - \alpha}$, we get the bound of Theorem E.3.

Remark E.4. As before, when reward distributions are Gaussian with $\sigma^2 = 1$, Theorem E.4 gives

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E} [n_i(T)]}{\log T} \geq \frac{2}{\left(\frac{\mu_{a^*}}{1 - \alpha} - \mu_i\right)^2}, \quad \text{for arms } i \in \{a^* + 1, \dots, K\} - \{i^*\}.$$

The corresponding upper bound for the number of samples from the high-cost arms for our PE-CS algorithm, given in Lemma G.5, scales as $\frac{1}{\Delta_i^2}$. Since arm a^* is feasible, we have

$$\frac{\mu_{a^*}}{1 - \alpha} - \mu_i > \mu^* - \mu_i = \Delta_i.$$

This suggests that there is potential for improvement in the PE-CS algorithm regarding the number of samples drawn from high-cost arms. For instance, if the two arms with the highest mean rewards have a small gap, it may not be critical to determine which one is superior (making Δ_i less relevant), provided that there exists a cheaper arm whose mean reward significantly exceeds the feasibility threshold corresponding to either of the two arms in contention to be the one with the largest mean reward.

LOWER BOUND ON THE NUMBER OF SAMPLES FOR ARM i^*

Theorem E.5 (Lower bound on the expected number of samples of the best arm). *For any bandit instance ν over horizon T , for all consistent strategies, we have*

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_{i^*}(T)]}{\log T} \geq \frac{1}{D_{\text{inf}}\left(\nu_{i^*}, \frac{\mu_{a^*}}{1 - \alpha}\right)}. \quad (30)$$

Also, with $\mu_j = \max_{i \in \{1, \dots, a^* - 1\}} \mu_i$ and $\mu_2 = \max_{i \in \{1, \dots, K\} - i^*} \mu_i$, we have

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_{i^*}(T)]}{\log T} \geq \frac{1}{\tilde{D}_{\text{inf}}\left(\nu_{i^*}, \frac{\mu_j}{1 - \alpha}\right)} \quad \text{if} \quad \mu_j \geq (1 - \alpha)\mu_2. \quad (31)$$

Proof. For the arm i^* that has the largest mean reward, establishing a lower bound is more subtle since we need to consider modifications in two different manners. First, we shall consider modifications to its reward distribution to make its mean larger such that arm a^* is no longer feasible and thus not optimal. Formally, given any bandit instance ν , we consider a modified instance ν' such that $\nu'_i = \nu_i$ for all $i \in \{1, \dots, K\} - \{i^*\}$ and ν'_{i^*} is any distribution in \mathcal{M} such that its expectation μ'_{i^*} satisfies $(1 - \alpha)\mu'_{i^*} > \mu_{a^*}$. In the modified instance ν' , arm i^* is still the one with the largest mean reward and the feasibility threshold has changed to $(1 - \alpha)\mu'_{i^*}$. Since $\mu_{a^*} < (1 - \alpha)\mu'_{i^*}$ by construction, arm a^* is no longer feasible and thus can not be optimal. This leads to

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_{i^*}(T)]}{\log T} \geq \frac{1}{\text{KL}(\nu_{i^*}, \nu'_{i^*})} \quad (32)$$

and taking the supremum in the right-hand side over all distributions $\nu'_{i^*} \in \mathcal{M}$ with $\mu'_{i^*} > \frac{\mu_{a^*}}{1 - \alpha}$, we get the first bound 30 of Theorem E.5.

Next, we consider ways of reducing the mean reward of arm i^* that would render a^* sub-optimal by making one of the cheaper arms $i = 1, \dots, a^* - 1$ feasible. Here, we shall notice that by reducing the mean reward of arm i^* , we can reduce the feasibility threshold to no lower than $(1 - \alpha)\mu_2$ where μ_2 is the second largest mean reward among all arms, i.e., $\mu_2 = \max_{i \in \{1, \dots, K\} - \{i^*\}} \mu_i$. Even then, this modification would make a^* sub-optimal *only if* one of the arms $i = 1, \dots, a^* - 1$ (that all have smaller cost than a^*) has a mean reward that would make it feasible at this threshold. Namely, with $\mu_j = \max_{i \in \{1, \dots, a^* - 1\}} \mu_i$ (i.e., μ_j is the largest mean reward among arms $1, \dots, a^* - 1$, and second largest among arms $1, \dots, a^*$), this would be possible only if $\mu_j \geq (1 - \alpha)\mu_2$. If $\mu_j < (1 - \alpha)\mu_2$, then no matter how much the mean reward of arm i^* is reduced, the arms $i = 1, \dots, a^* - 1$ would not be feasible (with feasibility threshold being at least $(1 - \alpha)\mu_2$), and arm a^* would remain optimal. In such cases, the only lower bound we can establish for the number of samples of arm i^* is that given by (30).

Now, consider any bandit instance ν for which $\mu_j \geq (1 - \alpha)\mu_2$. Then, consider a modified instance ν' such that $\nu'_i = \nu_i$ for all $i \in \{1, \dots, K\} - \{i^*\}$ and ν'_{i^*} is any distribution in \mathcal{M} such that its expectation μ'_{i^*} satisfies $(1 - \alpha)\mu'_{i^*} \leq \mu_j$. Since arm j was not feasible in the original instance ν , we must have $(1 - \alpha)\mu_{i^*} > \mu_j$, so $\mu'_{i^*} < \mu_{i^*}$. In the modified instance ν' , the arm with the largest mean reward is either still i^* or it becomes the arm with mean μ_2 . Thus, the feasibility threshold changes to $(1 - \alpha) \max\{\mu'_{i^*}, \mu_2\}$. By construction $\mu_j \geq (1 - \alpha)\mu'_{i^*}$ and due to our initial assumption on the instance ν we have $\mu_j \geq (1 - \alpha)\mu_2$ (which holds in the instance ν' since the distribution of all arms but i^* remain the same). Thus, we get $\mu_j \geq (1 - \alpha) \max\{\mu'_{i^*}, \mu_2\}$, meaning that the arm with mean reward μ_j is feasible in instance ν' . Since this arm is in $\{1, \dots, a^* - 1\}$, its cost is less than the cost of arm a^* meaning that arm a^* is not optimal in ν' ; instead the arm with mean reward μ_j is optimal. This leads to

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}_\nu [n_{i^*}(T)]}{\log T} \geq \frac{1}{\text{KL}(\nu_{i^*}, \nu'_{i^*})} \quad (33)$$

and taking the supremum in the right-hand side over all distributions $\nu'_{i^*} \in \mathcal{M}$ with $\mu'_{i^*} \leq \frac{\mu_j}{1 - \alpha}$, we get the second bound 31 of Theorem E.5.

Remark E.5. As before, when reward distributions are Gaussian, Theorem E.5 gives us

$$\begin{aligned} \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[n_{i^*}(T)]}{\log T} &\geq \max \left\{ \frac{2(1 - \alpha)^2}{((1 - \alpha)\mu_{i^*} - \mu_{a^*})^2}, \mathbf{1}[\mu_j \geq (1 - \alpha)\mu_2] \frac{2(1 - \alpha)^2}{((1 - \alpha)\mu_{i^*} - \mu_j)^2} \right\} \\ &= 2(1 - \alpha)^2 \max \left\{ \frac{1}{\Delta_{Q, a^*}^2}, \max_{i < a^*} \frac{1}{\Delta_{Q, i}^2} \mathbf{1}[\mu_j \geq (1 - \alpha)\mu_2] \right\} \\ &= 2(1 - \alpha)^2 \max \left\{ \frac{1}{\Delta_{Q, a^*}^2}, \max_{i < a^*} \frac{1}{\Delta_{Q, i}^2} \mathbf{1}[\min_{i < a^*} \Delta_{Q, i} \leq (1 - \alpha)\Delta_{\min}] \right\} \end{aligned}$$

where $\Delta_{Q, i} = (1 - \alpha)\mu^* - \mu_i$ is the quality gap of arm i , $\Delta_{\min} = \mu^* - \mu_2$ is the smallest of the regular gaps, and the last equality in the indicator function follows from the fact that

$$\min_{i < a^*} \Delta_{Q, i} = (1 - \alpha)\mu^* - \mu_j = (1 - \alpha)\Delta_{\min} + (1 - \alpha)\mu_2 - \mu_j$$

Remark E.6. The lower bound established in Theorem E.5, particularly the need to consider the case $\mu_j \geq (1 - \alpha)\mu_2$ separately can be explained as follows. For a policy to determine that a^* is the optimal action, it needs to be able to decide that arms $1, \dots, a^* - 1$ are not feasible but arm a^* is. For the latter, the gap between μ_{a^*} and the feasibility threshold $(1 - \alpha)\mu_{i^*}$ is important and affects the number of samples needed from the arm i^* , and this is seen in the bound (30). For the former, i.e., to decide that arms $1, \dots, a^* - 1$ are not feasible, a policy may not need any samples from arm i^* if it can determine their non-feasibility through another arm. For example, if the mean reward of arms $1, \dots, a^* - 1$ are all less than $(1 - \alpha)\mu_2$, then they can be determined to be not feasible based on the samples from the arm with mean reward μ_2 (i.e., the arm with the second largest mean reward among all). This is what gives rise to (30) being the only lower bound on the samples from arm i^* when $\mu_j < (1 - \alpha)\mu_2$. On the other hand, when $\mu_j \geq (1 - \alpha)\mu_2$, then at least one arm in $1, \dots, a^* - 1$ can not be deemed non-feasible by comparing it with the arm with mean reward μ_2 and samples from i^* would be necessary. The number of samples needed from i^* would be tied to the gaps between the mean reward of arms $1, \dots, a^* - 1$ and the feasibility threshold $(1 - \alpha)\mu^*$ and the smallest of these gaps is given by $(1 - \alpha)\mu_{i^*} - \mu_j$ which is consistent with the bound (31).

1296 **Remark E.7.** We now compare the lower bound established in Theorem E.5 with the upper bound
1297 for our PE-CS algorithm given in Lemma G.4. We can see that if $\mu_j \geq (1 - \alpha)\mu_2$, the dominant
1298 gap term in the established lower bound will be given by $\min_{i \leq a^*} |\Delta_{Q,i}|$, which in this case is also
1299 smaller than $\Delta_{min} = \mu^* - \mu_2$. Thus, the upper bound of the PE-CS algorithm given in Lemma
1300 G.4 matches the lower bound in this case within a constant factor. When $\mu_j < (1 - \alpha)\mu_2$, as
1301 also explained in the previous remark, the lower bound on the number of samples from i^* is only
1302 governed by Δ_{Q,a^*} , while the PE-CS upper bound depends additionally on $\min_{i < a^*} \Delta_{Q,i}$ and Δ_{min} .
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

F ANALYSIS FOR PE IN THE KNOWN REFERENCE ARM SETTING

In this Section we build up to the proof of Theorem 3.2 by upper bounding the expected number of samples of low-cost infeasible arms, the reference arm ℓ , and of high-cost arms under the Asymmetric-PE setting with maximum round-deviation κ . We then particularize these results to the $\kappa = 0$ case corresponding to conventional PE to obtain an upper bound on the expected regret for Algorithm 1 PE.

F.1 DEFINITIONS AND SETUP REQUIRED FOR ANALYSIS

As discussed in the main paper, the PE algorithm is inspired by the Improved-UCB successive elimination approach where sampling of arms occurs in un-interrupted batches called rounds. In Improved-UCB, a set of active arms is maintained and at the end of every round, arms in the active set are re-tested for their candidacy using an elimination criteria. Since in Pairwise-Elimination, we inherit the un-interrupted round based sampling scheme and elimination-criteria first used in Improved-UCB, to prove Theorem 3.2 we build on the analysis from Auer & Ortner (2010).

For $i \leq a^*$, define round number ρ_i to be,

$$\rho_i = \min \left\{ \omega_i \mid \tilde{\Delta}_{\omega_i} < \frac{|\Delta_{Q,i}|}{2} \right\}, \quad (34)$$

Intuitively, ρ_i is the PE round number during episode i by which we expect to either identify low-cost infeasible arm $i < a^*$ as being unsuitable or identify the best action a^* as being suitable. Moreover from Function 4, we know that for any arm the required number of samples to be drawn by round ω_i is given by,

$$\tau_{\omega_i} = \left\lceil \frac{2 \log \left(T \tilde{\Delta}_{\omega_i}^2 \right)}{\tilde{\Delta}_{\omega_i}^2} \right\rceil. \quad (35)$$

Table 2: Probabilistic Events Descriptions and Symbols for PE Analysis

Symbol	Event Description
$G_{1,i}, \forall i \leq a^*$	Episode i is executed to evaluate arm i
$G_{2,i}, \forall i < a^*$	Arm i is eliminated by arm ℓ by when round $\omega_i = \rho_i$, during episode i
G_{2,a^*}	Arm ℓ is eliminated by arm a^* by when $\omega_{a^*} = \rho_{a^*}$, during episode a^*
$G_{3,i}, \forall i < a^*$	Arm ℓ is not eliminated by arm i by when round $\omega_i = \rho_i - 1$, during episode i
G_{3,a^*}	Arm a^* is not eliminated by arm ℓ by when round $\omega_{a^*} = \rho_{a^*} - 1$, during episode a^*
$E_i, \forall i \leq a^*$	Available samples ran out during episode i before the sampling for round ρ_i could conclude and before an arm could be eliminated

To lay the groundwork for the forthcoming analysis we introduce notation $n_i(t_1, t_2)$ for the random variable denoting the number of samples of arm i accrued between time steps t_1 and t_2 both inclusive. Further, we use t_i to denote the final time-step in episode i . Thereby, the variable $n_\ell(1, t_i)$ denotes the number of samples of reference arm ℓ accrued by the end of episode i .

At the outset of our analysis, we define a large collection of probabilistic events needed for developing Theorem 3.2’s intermediate results in Table 2. Each event in Table 2 is a subset of the sample space Ω associated with a run of PE. In the descriptions of events in Table 2, when we say that an elimination event occurs *by a round*, we are including the round being mentioned. For example: “Arm i is eliminated by round ρ_i ” means that the arm i was eliminated in round ω_i such that $\omega_i \leq \rho_i$.

Remark F.1. In the event descriptions, whenever we refer to the reference arm ℓ available to the algorithm, we are really referring to the action whose expected return is $(1 - \alpha)\mu_\ell$, where α is the subsidy factor.

Arm ℓ not being eliminated by arm i during round ρ_i is covered by $G_{2,i}$, hence the rounds range up to $\rho_i - 1$ in the definition of $G_{3,i}$. Unlike the episodes $i < a^*$ where the nominal outcome is for the reference arm to win over the candidate arm, for episode a^* , nominally the arm a^* shall be the winner and therefore the events G_{2,a^*}, G_{3,a^*} are defined separately to account for this reality.

Lastly, we define compound event G_i using events in Table 2 as,

$$G_i = G_{1,i} \cap ((G_{2,i} \cap G_{3,i}) \cup E_i). \quad (36)$$

In English the event $G_i, i \leq a^*$ is the event that episode i occurs (event $G_{1,i}$) and that either an accurate and timely elimination of one arm by another is made (event $G_{2,i} \cap G_{3,i}$), or we run out of samples before a decision could be made and prior to the conclusion of round ρ_i (event E_i). Conditioning the samples $n_i(T)$ for PE on event G_i gives us the following key result,

$$\Pr(n_i(T) \leq \tau_{\rho_i} \mid G_i) = 1. \quad (37)$$

We leverage Equation 37 in conjunction with Lemma D.3 to prove Theorem 3.2. To use this procedure, we require Lemma F.1 that partitions the probability space Ω into mutually exclusive and exhaustive events including G_i .

Lemma F.1 (Partition of Ω with G_i). *The events $G_i, B_{1,i} = G_{1,i}^c$, and $B_i = G_{1,i} \cap (G_{2,i}^c \cup G_{3,i}^c) \cap E_i^c$ are exhaustive and pairwise exclusive $\forall i \leq a^*$.*

Proof. We can prove that the events stated in Lemma F.1 are mutually exclusive and exhaustive by showing that $G_i^c = B_{1,i} \cup B_i$. Since G_i and G_i^c are exhaustive showing so will show that the three events are exhaustive. Moreover, since G_i and G_i^c are mutually exclusive, and since $B_{1,i}$ and B_i are mutually exclusive by construction, we would also have all the events being pairwise mutually exclusive in addition to being exhaustive.

$$B_{1,i} \cup B_i = G_{1,i}^c \cup (G_{1,i} \cap (G_{2,i}^c \cup G_{3,i}^c) \cap E_i^c) \quad (38)$$

$$= (G_{1,i}^c \cup G_{1,i}) \cap (G_{1,i} \cap ((G_{2,i}^c \cup G_{3,i}^c) \cap E_i^c)) \quad (\cup \text{ distributes over } \cap) \quad (39)$$

$$= G_{1,i}^c \cup ((G_{2,i}^c \cup G_{3,i}^c) \cap E_i^c) \quad (40)$$

$$= G_i^c. \quad (41)$$

Where Equation 41 follows from the expression obtained using De Morgan's laws for G_i^c using the definition of G_i in Equation 36. \square

BOUND SAMPLES FOR THE CASE $i < a^*$

Lemma F.2 (Bound on the number of samples of a low-cost infeasible arm under Pairwise-Elimination). *When the maximum round deviation $\kappa = 0$, the expected number of samples of a low-cost arm with index $i < a^*$ over horizon T is upper bounded by,*

$$\mathbb{E}[n_i(T)] < 1 + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \frac{43}{\Delta_{Q,i}^2}.$$

Proof. The expected number of pulls $\mathbb{E}[n_i(T)]$ are bound by using the Iterated Expectation Lemma D.3 and conditioning on the event collection $G_i, B_{1,i}, B_i$ which are mutually exclusive and exhaustive per Lemma F.1.

$$\mathbb{E}[n_i(T)] = \mathbb{E}[n_i(T) \mid G_i] \cdot \Pr(G_i) + \mathbb{E}[n_i(T) \mid B_{1,i}] \cdot \Pr(B_{1,i}) + \mathbb{E}[n_i(T) \mid B_i] \cdot \Pr(B_i) \quad (42)$$

$$\leq \mathbb{E}[n_i(T) \mid G_i] + T \cdot \Pr(B_i) \quad (43)$$

$$= \mathbb{E}[n_i(T) \mid G_i] + T \cdot \Pr(G_{1,i} \cap (G_{2,i}^c \cup G_{3,i}^c) \cap E_i^c) \quad (B_i \text{ from Lemma F.1}) \quad (44)$$

$$\leq \mathbb{E}[n_i(T) \mid G_i] + T \cdot \Pr(G_{1,i} \cap (G_{2,i}^c \cup G_{3,i}^c)) \quad (45)$$

$$= \mathbb{E}[n_i(T) \mid G_i] + T \cdot \Pr((G_{1,i} \cap G_{2,i}^c) \cup (G_{1,i} \cap G_{3,i}^c)) \quad (\text{distributivity of } \cap) \quad (46)$$

$$\leq \mathbb{E}[n_i(T) \mid G_i] + T \cdot \Pr(B_{2,i} \cup B_{3,i}) \quad (\text{simplifying notation}) \quad (47)$$

$$= \mathbb{E}[n_i(T) \mid G_i] + T \cdot (\Pr(B_{2,i}) + \Pr(B_{3,i})) \quad (\text{using the union bound}). \quad (48)$$

Where Equation 43 follows from the fact that $n_i(T) \mid B_{1,i} = 0$ since there can be no samples of arm i if episode i never occurs, and we define $G_{1,i} \cap G_{2,i}^c$ and $G_{1,i} \cap G_{3,i}^c$ as $B_{2,i}$ and $B_{3,i}$ respectively for notational simplicity.

First we bound the number of samples of arm i conditioned on the good event G_i . Since during episode $i < a^*$, we either make the correct decision of eliminating arm i by episode $\omega_i = \rho_i$ as captured by $G_{2,i} \cap G_{3,i}$. Alternatively, under G_i we run out of samples as captured by E_i . In either case we will not have more than τ_{ρ_i} samples of arm i , where τ_{ρ_i} is given by,

$$\tau_{\rho_i} = \left\lceil \frac{2 \log(T \tilde{\Delta}_{\rho_i}^2)}{\tilde{\Delta}_{\rho_i}^2} \right\rceil. \quad (49)$$

By construction of the round ρ_i , for all $i \leq a^*$, we have,

$$\frac{|\Delta_{Q,i}|}{4} \leq \tilde{\Delta}_{\rho_i} < \frac{|\Delta_{Q,i}|}{2}. \quad (50)$$

Plugging in the bounds in Equation 50,

$$\tau_{\rho_i} \leq \left\lceil \frac{32 \log\left(\frac{T \Delta_{Q,i}^2}{4}\right)}{\Delta_{Q,i}^2} \right\rceil \quad (51)$$

$$< 1 + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2}. \quad (52)$$

Therefore, we have,

$$\mathbb{E}[n_i(T) \mid G_i] \leq \tau_{\rho_i} < 1 + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2}. \quad (53)$$

Next we bound $\Pr(B_{2,i})$ and $\Pr(B_{3,i})$ in order. Since $B_{2,i} = G_{1,i} \cap G_{2,i}^c$, from the specification of $G_{2,i}^c$ and the fact that intersecting with $G_{1,i}$ puts us in the sub-space of Ω where episode i occurs, $B_{2,i} \forall i < a^*$ is the event: “Arm i is **not** eliminated by arm ℓ by when round $\omega_i = \rho_i$, during episode i ”. Similarly $B_{3,i} \forall i < a^*$ is the event that “Arm ℓ is **eliminated** by arm i by when round $\omega_i = \rho_i - 1$, during episode i ”

Along the lines of the proof composition in Auer & Ortner (2010) for the Improved UCB algorithm, we construct three clauses on the empirical returns of arms i and ℓ in Equations 54, 55, and 56.

$$\hat{\mu}_i \leq \mu_i + \sqrt{\frac{\log(T \tilde{\Delta}_{\omega_i}^2)}{2\tau_{\omega_i}}} \quad (54)$$

$$\hat{\mu}_\ell \geq \mu_\ell - \sqrt{\frac{\log(T \tilde{\Delta}_{\omega_i}^2)}{2\tau_{\omega_i}}} \quad (55)$$

$$\hat{\mu}_\ell \geq \mu_\ell - \sqrt{\frac{\log(T \tilde{\Delta}_{\omega_\ell}^2)}{2\tau_{\omega_\ell}}}. \quad (56)$$

Clauses 54 and 55 holding when $\omega_i = \rho_i$ lead to the elimination of arm i by arm ℓ as is shown in the work that follows,

$$\sqrt{\log(T \tilde{\Delta}_{\rho_i}^2)} / 2\tau_{\rho_i} \leq \tilde{\Delta}_{\rho_i} / 2 < \Delta_{Q,i} / 4. \quad (57)$$

Therefore,

$$\hat{\mu}_i + \sqrt{\frac{\log(T \tilde{\Delta}_{\rho_i}^2)}{2\tau_{\rho_i}}} \leq \mu_i + 2\sqrt{\frac{\log(T \tilde{\Delta}_{\rho_i}^2)}{2\tau_{\rho_i}}} \quad (\text{From clause 54, and } \omega_i = \rho_i) \quad (58)$$

1566 clause itself we know that we must have $\rho < \rho_i$.

$$1567 \Pr(B_{3,i}) \leq \sum_{\rho=0}^{\rho_i-1} \frac{4^\kappa + 1}{T \tilde{\Delta}_\rho^2} = \sum_{\rho=0}^{\rho_i-1} \frac{(4^\kappa + 1) \cdot 4^\rho}{T} \text{ (Using } \tilde{\Delta}_m = 2^{-m} \text{)} \quad (69)$$

$$1570 < \frac{4^\kappa + 1}{3T} \cdot 4^{\rho_i} \text{ (Using the formula for the sum of a Geometric Series)} \quad (70)$$

$$1571 = \frac{4^\kappa + 1}{3T \tilde{\Delta}_{\rho_i}^2} \text{ (Since } \tilde{\Delta}_m = 2^{-m} \text{)} \quad (71)$$

$$1572 = \frac{16(4^\kappa + 1)}{3T \Delta_{Q,i}^2} \text{ (Because } \tilde{\Delta}_{\rho_i} \geq \frac{\Delta_{Q,i}}{4} \text{)} \quad (72)$$

$$1573 < \frac{11}{T \Delta_{Q,i}^2} \text{ (Since we impose in Lemma F.2).} \quad (73)$$

1580 Plugging in the bounds in Expressions 53, 68 and 73 into Equation 120 we get the overall bound on
1581 the number of samples stated in Lemma F.2. \square

1582 BOUNDING SAMPLES OF REFERENCE ARM ℓ

1583 Let the random variable Z denote the final episode in the run of PE. Then $\{Z = z\}$ constitutes
1584 a measurable event in the sample space Ω indicating that the terminal PE episode was z . This
1585 construction lets us establish F.3.

1586 **Lemma F.3** (Bound on Samples of Reference arm ℓ under Pairwise-Elimination). *Samples of ref-*
1587 *erence arm ℓ emanate from the episodes of candidate arms being compared to arm ℓ . When $\kappa = 0$,*
1588 *we show the bound,*

$$1589 \mathbb{E}[n_\ell(T)] < 1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \sum_{i=1}^{a^*} \frac{43}{\Delta_{Q,i}^2} + \mathbb{E}[n_\ell(T) \mid \{Z > a^*\}] \Pr(Z > a^*).$$

1590 Where Z is the terminal episode of PE.

1591 *Proof.* Under Pairwise-Elimination the nominal outcome is for episodes $i = 1, \dots, a^* - 1$ to result
1592 in the candidate arm i being eliminated by arm ℓ , followed arm a^* eliminating arm ℓ during episode
1593 a^* . To prove Lemma F.3 we condition on this nominal sequence and upper bound the probability of
1594 the outcome deviating from this sequence by a factor proportional to $\frac{1}{T}$. Throughout the episodes
1595 $i = 1, \dots, a^*$ the number of samples $n_\ell(T)$ are equal to the number of the most sampled
1596 candidate arm $i \leq a^*$. This is because in PE we re-use samples of arm ℓ across episodes and only
1597 further sample ℓ to keep up with the samples of a candidate arm. Motivated by this fact about n_ℓ we
1598 begin by defining a compound high-probability good event G .

1599 **Definition F.1** (Event G). *The compound good event G is the event that for each episode $i \leq a^*$*
1600 *that was executed, the episode satisfied the episode-wise good event G_i . Mathematically we define,*

$$1601 G = \bigcup_{z=1}^{a^*} \left(\{Z = z\} \cap \bigcap_{i=1}^z G_i \right). \quad (74)$$

1602 The complement of G , namely G^c can be written out using the definition of G and De Morgan's
1603 laws as,

$$1604 G^c = \bigcap_{z=1}^{a^*} \left(\{Z = z\}^c \cup \bigcup_{i=1}^z G_i^c \right) \quad (75)$$

$$1605 \subseteq \bigcup_{i=1}^{a^*} G_i^c \cup \{Z \neq a^*\} \quad \text{(picking } z = a^* \text{ from the iterated intersection)} \quad (76)$$

$$1606 = \bigcup_{i=1}^{a^*} G_i^c \cup \{Z < a^*\} \cup \{Z > a^*\} \quad (77)$$

$$= \bigcup_{i=1}^{a^*} (B_{1,i} \cup B_i) \cup \bigcup_{i=1}^{a^*} B_{1,i} \cup \{Z > a^*\} \quad (78)$$

$$= \bigcup_{i=1}^{a^*} (B_{1,i} \cup B_i) \cup \{Z > a^*\} \quad (79)$$

$$= \bigcup_{i=1}^{a^*} G_i^c \cup \{Z > a^*\}. \quad (80)$$

Where the first term in Equation 78 follows from the equivalence between G_i^c and $B_{1,i} \cup B_i$ shown in the proof of Lemma F.1. The second term in Equation 78 is based on the equivalence between the event $\{Z < a^*\}$, meaning that the final episode precedes a^* , and the event $\bigcup_{i=1}^{a^*} B_{1,i}$ which means that some episode $i = 1, \dots, a^*$ was not executed.

We leverage event G to bound the expected number of samples of arm ℓ . In particular we consider three events $G, G^c \cap \{Z \leq a^*\}$, and $G^c \cap \{Z > a^*\}$ that are mutually exclusive and exhaustive by construction. We decompose $\mathbb{E}[n_\ell(T)]$ by conditioning on this collection.

$$\begin{aligned} \mathbb{E}[n_\ell(T)] &= \mathbb{E}[n_\ell(T) | G] \Pr(G) + \mathbb{E}[n_\ell(T) | G^c \cap \{Z \leq a^*\}] \Pr(G^c \cap \{Z \leq a^*\}) \\ &\quad + \mathbb{E}[n_\ell(T) | G^c \cap \{Z > a^*\}] \Pr(G^c \cap \{Z > a^*\}) \end{aligned} \quad (81)$$

$$\leq \mathbb{E}[n_\ell(T) | G] + T \cdot \Pr(G^c \cap \{Z \leq a^*\}) + \mathbb{E}[n_\ell(T) | \{Z > a^*\}] \Pr(Z > a^*). \quad (82)$$

Where bound 82 follows from the simplified form of G^c developed in Equation 80.

Since samples of arm ℓ are reused between episodes with further sampling of arm ℓ only occurring to match the demand from a higher round number, we have,

$$\mathbb{E}[n_\ell(T) | G] = \mathbb{E}\left[\left(\max_{i \leq Z} n_i(1, t_Z)\right) | G\right] \quad (83)$$

$$\leq \mathbb{E}\left[\max_{i \leq a^*} n_i(1, t_{a^*}) | G\right] \quad (\because Z | G \leq a^*) \quad (84)$$

$$\leq \max_{i \leq a^*} \tau_{\rho_i} \quad (\text{using Equation 37, and by construction of } G \text{ as } \cap G_i) \quad (85)$$

$$= \max_{i \leq a^*} \left\lceil \frac{2 \log(T \tilde{\Delta}_{\rho_i}^2)}{\tilde{\Delta}_{\rho_i}^2} \right\rceil \quad (86)$$

$$< 1 + \max_{i \leq a^*} \frac{2 \log(T \tilde{\Delta}_{\rho_i}^2)}{\tilde{\Delta}_{\rho_i}^2} \quad (87)$$

$$\leq 1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2/4)}{\Delta_{Q,i}^2} \quad (\text{Using the ordering in Relation 50}) \quad (88)$$

$$< 1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2}. \quad (89)$$

To complete the bound $\mathbb{E}[n_\ell(T)]$, we develop a bound on the term $\Pr(G^c \cap \{Z \leq a^*\})$ in Lemma F.6. However, to prove Lemma F.6 we first require two intermediate results in the form of Lemmas F.4 and F.5.

Lemma F.4. *In all the outcomes contained in G^c ending in some episode $Z = i \leq a^*$, some $B_j, j \leq i$ must have held.*

$$G^c \cap \{Z = i\} \subseteq \bigcup_{j=1}^i B_j \quad \forall i \leq a^*. \quad (90)$$

1674

Proof.

1675

$$1676 \quad G^c \cap \{Z = i\} = ((G_i^c \cup G_i) \cap \{Z = i\}) \cap G^c \quad (91)$$

$$1677 \quad = ((G_i^c \cap \{Z = i\}) \cup (G_i \cap \{Z = i\})) \cap G^c \quad (\cap \text{ distributes over } \cup) \quad (92)$$

$$1678 \quad = (G_i^c \cap \{Z = i\}) \cup (G_i \cap \{Z = i\} \cap G^c) \quad (\because G_i^c \cap \{Z = i\} \subseteq G^c) \quad (93)$$

$$1679 \quad = ((B_{1,i} \cup B_i) \cap \{Z = i\}) \cup (G_i \cap \{Z = i\} \cap G^c) \quad (\because G_i^c = B_{1,i} \cup B_i) \quad (94)$$

$$1680 \quad \subseteq B_i \cup (G_i \cap \{Z = i\} \cap G^c) \quad (\because B_{1,i} \cap \{Z = i\} = \phi). \quad (95)$$

1681

1682 Now consider just the event $G_i \cap \{Z = i\} \cap G^c$ from Equation 95. We can find an event it is
1683 subsumed within in the following way,

1684

$$1685 \quad G_i \cap \{Z = i\} \cap G^c = G_i \cap \left(\bigcap_{j=1}^{i-1} G_j \cup \left(\bigcap_{j=1}^{i-1} G_j \right)^c \right) \cap \{Z = i\} \cap G^c \quad (96)$$

$$1688 \quad = \left(\bigcap_{j=1}^i G_j \cap \{Z = i\} \cap G^c \right) \cup \left(G_i \cap \bigcup_{j=1}^{i-1} G_j^c \cap \{Z = i\} \cap G^c \right) \quad (97)$$

$$1691 \quad = G_i \cap \bigcup_{j=1}^{i-1} (G_j^c \cap \{Z = i\}) \cap G^c \quad \left(\bigcap_{j=1}^i G_j \cap \{Z = i\} \subseteq G, \forall i \leq a^* \right) \quad (98)$$

$$1695 \quad = G_i \cap \bigcup_{j=1}^{i-1} ((B_{1,j} \cup B_j) \cap \{Z = i\}) \cap G^c \quad (\text{from Lemma F.1}) \quad (99)$$

$$1698 \quad = G_i \cap \bigcup_{j=1}^{i-1} (B_j \cap \{Z = i\}) \cap G^c \quad (\because B_{1,j} \cap \{Z = i\} = \phi, \forall j \leq i) \quad (100)$$

$$1701 \quad \subseteq \bigcup_{j=1}^{i-1} B_j. \quad (101)$$

1703

1704 Plugging in Equation 101 into Equation 95 gives us the result stated in Lemma F.4. \square

1705

1706 **Lemma F.5.** *Within the space of events that constitute G^c , if an episode $i \leq a^*$ does not occur, then*
1707 *there exists $j < i$ such that the event B_j occurred. Mathematically,*

$$1708 \quad B_{1,i} \cap G^c \subseteq \bigcup_{j=1}^{i-1} B_j, \forall i \leq a^*. \quad (102)$$

1711

1712

1713 *Proof.* We can prove the result by induction. First note that $B_{1,1} = \phi$ since Episode 1 always occurs.
1714 Let $i = 2$ represent the base case. Then,

$$1715 \quad B_{1,2} \cap G^c = \{Z = 1\} \cap G^c \quad (103)$$

$$1716 \quad \subseteq B_1 \quad (\text{Using Lemma F.4}). \quad (104)$$

1717

1718 Now say that the statement in Lemma F.5 holds true for some $i = k < a^*$. That is,

$$1719 \quad B_{1,k} \cap G^c \subseteq \bigcup_{j=1}^{k-1} B_j. \quad (105)$$

1721

1722 To prove Lemma F.5 we need to show this result for $i = k + 1$.

1723

$$1724 \quad B_{1,k+1} \cap G^c = (B_{1,k} \cup \{Z = k\}) \cap G^c \quad (106)$$

$$1725 \quad = (B_{1,k} \cap G^c) \cup (\{Z = k\} \cap G^c) \quad (\cap \text{ over } \cup) \quad (107)$$

1726

$$1727 \quad \subseteq \bigcup_{j=1}^{k-1} B_j \cup \bigcup_{j=1}^k B_j \quad (\text{using the induction hypothesis and Lemma F.4}). \quad (108)$$

$$= \bigcup_{j=1}^k B_j. \quad (109)$$

Where Equation 106 uses the identity $B_{1,k+1} = B_{1,k} \cup \{Z = k\}$ which breaks down the event of episode $k + 1$ not occurring into the event that episode k did not occur, or the event that episode k was the final episode Z . Equation 109 is the required result from the statement of Lemma F.5. \square

Lemma F.6 (Bound on $\Pr(G^c \cap \{Z \leq a^*\})$). *We can upper bound $\Pr(G^c \cap \{Z \leq a^*\})$ as,*

$$\Pr(G^c \cap \{Z \leq a^*\}) \leq \sum_{i=1}^{a^*} \Pr(B_i) \quad (110)$$

$$\leq \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + \Pr(B_{a^*}). \quad (111)$$

Where $B_i, \forall i \leq a^*$ is as defined in Lemma F.1, and using the bound on $\Pr(B_i)$ shown in the proof of Lemma F.2.

Proof. We begin by introducing the expanded expression for G^c developed in Equation 80.

$$\Pr(G^c \cap \{Z \leq a^*\}) = \Pr(G^c \cap \{Z \leq a^*\} \cap G^c) \quad (112)$$

$$\leq \Pr\left(\left(\bigcup_{i=1}^{a^*} G_i^c \cup \{Z > a^*\}\right) \cap \{Z \leq a^*\} \cap G^c\right) \quad (\text{from Equation 80}) \quad (113)$$

$$\leq \Pr\left(\left(\bigcup_{i=1}^{a^*} G_i^c \cap G^c\right)\right) \quad (\cap \text{ distributes over } \cup) \quad (114)$$

$$\leq \Pr\left(\bigcup_{i=1}^{a^*} (G_i^c \cap G^c)\right) \quad (\text{union bound and } \Pr(A, B) \leq \Pr(A)) \quad (115)$$

$$= \Pr\left(\bigcup_{i=1}^{a^*} ((B_{1,i} \cup B_i) \cap G^c)\right) \quad (\text{by definition of } G_i) \quad (116)$$

$$= \Pr\left(\bigcup_{i=1}^{a^*} ((B_{1,i} \cap G^c) \cup (B_i \cap G^c))\right) \quad (\cap \text{ distributes over } \cup) \quad (117)$$

$$\leq \Pr\left(\bigcup_{i=1}^{a^*} \left(B_i \cup \bigcup_{j=1}^{i-1} B_j\right)\right) \quad (\text{using Lemma F.5}) \quad (118)$$

$$= \Pr\left(\bigcup_{i=1}^{a^*} B_i\right) \quad (119)$$

$$\leq \sum_{i=1}^{a^*} \Pr(B_i) \quad (\text{Union bound, and } \{Z > a^*\} \implies B_{a^*}). \quad (120)$$

\square

To complete the upper bound on the expected number of samples $\mathbb{E}[n_\ell(T)]$ from Equation 82, we upper bound $\Pr(B_{a^*})$ along the same lines as $\Pr(B_i), i < a^*$ in the proof of Lemma F.2. The key difference being that the roles of candidate arm a^* and reference arm ℓ are reversed as compared to the earlier procedure since $\mu_{a^*} \geq \mu_\ell$. Moreover, just like the earlier proof we can define $B_{2,a^*} = G_{1,a^*} \cap G_{2,i}^c$ and $B_{3,a^*} = G_{1,a^*} \cap G_{3,i}^c$. In words B_{2,a^*} is the event that ‘‘Arm ℓ is **not** eliminated by arm a^* by when $\omega_{a^*} = \rho_{a^*}$, during episode a^* ’’. Similarly, B_{3,a^*} is the event ‘‘Arm a^* is **eliminated** by arm ℓ by when round $\omega_{a^*} = \rho_{a^*} - 1$, during episode a^* ’’. \square

We bound $\Pr(B_{2,a^*})$ and $\Pr(B_{3,a^*})$ separately by constructing clauses on $\hat{\mu}_\ell$, and $\hat{\mu}_{a^*}$ as before.

$$\hat{\mu}_\ell \leq \mu_\ell + \sqrt{\frac{\log(T\tilde{\Delta}_{\omega_{a^*}}^2)}{2\tau_{\omega_{a^*}}}} \quad (121)$$

$$\hat{\mu}_\ell \leq \mu_\ell + \sqrt{\frac{\log(T\tilde{\Delta}_{\omega_\ell}^2)}{2\tau_{\omega_\ell}}} \quad (122)$$

$$\hat{\mu}_{a^*} \geq \mu_{a^*} - \sqrt{\frac{\log(T\tilde{\Delta}_{\omega_{a^*}}^2)}{2\tau_{\omega_{a^*}}}}. \quad (123)$$

Clauses 121 and 123 holding when $\omega_{a^*} = \rho_{a^*}$ lead to the elimination of arm ℓ by arm a^* as is shown in the following steps,

$$\sqrt{\log(T\tilde{\Delta}_{\rho_{a^*}}^2)}/2\tau_{\rho_{a^*}} < \tilde{\Delta}_{\rho_{a^*}}/2 < |\Delta_{Q,a^*}|/4. \quad (124)$$

Therefore using $\omega_{a^*} = \rho_{a^*}$,

$$(1 - \alpha) \left(\hat{\mu}_\ell + \sqrt{\frac{\log(T\tilde{\Delta}_{\omega_\ell}^2)}{2\tau_{\omega_\ell}}} \right) \leq (1 - \alpha) \left(\hat{\mu}_\ell + \sqrt{\frac{\log(T\tilde{\Delta}_{\rho_{a^*}}^2)}{2\tau_{\rho_{a^*}}}} \right) \quad (125)$$

$$\leq (1 - \alpha) \left(\mu_\ell + 2\sqrt{\frac{\log(T\tilde{\Delta}_{\rho_{a^*}}^2)}{2\tau_{\rho_{a^*}}}} \right) \quad (\text{Equation 121}) \quad (126)$$

$$\leq (1 - \alpha)\mu_\ell + 2\sqrt{\frac{\log(T\tilde{\Delta}_{\rho_{a^*}}^2)}{2\tau_{\rho_{a^*}}}} \quad (127)$$

$$< (1 - \alpha)\mu_\ell + |\Delta_{Q,a^*}| - 2\sqrt{\frac{\log(T\tilde{\Delta}_{\rho_{a^*}}^2)}{2\tau_{\rho_{a^*}}}} \quad (\text{From 124}) \quad (128)$$

$$= \mu_{a^*} - 2\sqrt{\frac{\log(T\tilde{\Delta}_{\rho_{a^*}}^2)}{2\tau_{\rho_{a^*}}}} \quad (129)$$

$$\leq \hat{\mu}_{a^*} - \sqrt{\frac{\log(T\tilde{\Delta}_{\rho_{a^*}}^2)}{2\tau_{\rho_{a^*}}}} \quad (\text{From Equation 123}). \quad (130)$$

Here Equation 130 is the criteria for arm ℓ being eliminated by arm a^* in PE. Since Clause 121 and Clause 123 being true and applicable at round $\omega_i = a^*$ imply arm ℓ being eliminated, we can upper bound $\Pr(B_{2,a^*})$ as,

$$\Pr(B_{2,a^*}) \leq \Pr\left(\hat{\mu}_\ell > \mu_\ell + \sqrt{\frac{\log(T\tilde{\Delta}_{\omega_{a^*}}^2)}{2\tau_{\omega_{a^*}}}}\right) + \Pr\left(\hat{\mu}_{a^*} < \mu_{a^*} - \sqrt{\frac{\log(T\tilde{\Delta}_{\omega_{a^*}}^2)}{2\tau_{\omega_{a^*}}}}\right) \quad (131)$$

$$\leq \frac{2}{T\tilde{\Delta}_{\omega_{a^*}}^2} \quad (\text{Similar to the bounds in 63 and 64}). \quad (132)$$

Plugging in round number $\omega_{a^*} = \rho_{a^*}$ in Equation 132, and then plugging in the lower bound on $\tilde{\Delta}_{\rho_{a^*}}$ from Ordering 50, we have,

$$\Pr(B_{2,a^*}) \leq \frac{2}{T\tilde{\Delta}_{\rho_{a^*}}^2} \leq \frac{32}{T\Delta_{Q,a^*}^2}. \quad (133)$$

To complete the bound on $\Pr(B_{a^*})$ we must bound $\Pr(B_{3,a^*})$ which is the the probability of arm a^* being eliminated by arm ℓ by round $\omega_{a^*} = \rho_{a^*}$. Similar to the arguments in the Proof of Lemma F.2, the clauses in Equations 121 and 123 holding simultaneously preclude arm a^* from being removed by arm ℓ regardless of the round number, and we shall have,

$$\Pr(B_{3,a^*}) < \frac{16(4^\kappa + 1)}{3T\Delta_{Q,a^*}^2} \quad (134)$$

$$< \frac{11}{T\Delta_{Q,a^*}^2} \quad (\text{When we impose } \kappa = 0). \quad (135)$$

Using steps identical to those that lead up to Equation 48 we shall have,

$$\Pr(B_{a^*}) \leq \Pr(B_{2,a^*}) + \Pr(B_{3,a^*}) \quad (136)$$

$$< \frac{43}{T\Delta_{Q,a^*}^2} \quad (\text{from upper bounds in Equations 133 and 135}). \quad (137)$$

Combining the upper bound on the expected number of samples of arm ℓ in Equation 89, with the bound on $\Pr(G^c \cap \{Z \leq a^*\})$ in Lemma F.6, and the bound on $\Pr(B_{a^*})$ in Equation 137 all into 82 we reach the upper bound stated in Lemma F.3. \square

BOUNDING SAMPLES FOR HIGH COST ARMS $a^* < i < \ell$

Lemma F.7 (Bound on the expected number of samples of all high cost arms under Pairwise-Elimination). *When $\kappa = 0$ the expected number of samples of all the higher cost, non-reference arms, that is, arms with index in the range $a^* < i < \ell$ is upper bounded by a constant given by,*

$$\sum_{a^* < i < \ell} \mathbb{E}[n_i(T)] \leq \mathbb{E} \left[\sum_{a^* < i < \ell} n_i(T) \mid \{Z > a^*\} \right] \cdot \Pr(Z > a^*).$$

Where Z is the terminal PE episode.

Proof. We can upper bound the expected number of samples $\sum_{a^* < i < \ell} \mathbb{E}[n_i(T)]$ by conditioning on the event $\{Z \leq a^*\}$ and its complement.

$$\begin{aligned} \sum_{a^* < i < \ell} \mathbb{E}[n_i(T)] &= \mathbb{E} \left[\sum_{a^* < i < \ell} n_i(T) \right] \quad (\text{from the linearity of Expectation operator}) \quad (138) \\ &= \mathbb{E} \left[\sum_{a^* < i < \ell} n_i(T) \mid \{Z \leq a^*\} \right] \cdot \Pr(\{Z \leq a^*\}) \\ &\quad + \mathbb{E} \left[\sum_{a^* < i < \ell} n_i(T) \mid \{Z > a^*\} \right] \cdot \Pr(\{Z > a^*\}) \quad (\text{from Lemma D.3}). \end{aligned} \quad (139)$$

Where $\mathbb{E} \left[\sum_{i=a^*+1}^{\ell} n_i(T) \mid \{Z \leq a^*\} \right] = 0$ follows from the fact that there cannot be any samples of an arm $i > a^*$ in the case when the final episode Z is less than a^* . \square

Finally, we are in a position to Prove Theorem 3.2.

Proof of Theorem 3.2. First, we apply Regret decomposition in Lemma D.5 to Cost Regret,

$$\mathbb{E}[\text{Cost_Reg}(T, \nu)] \leq \sum_{i=a^*+1}^{\ell} \Delta_{C,i}^+ \mathbb{E}[n_i(T)] \quad (\text{because } \Delta_{C,i} \leq 0 \text{ for } i \leq a^*) \quad (140)$$

$$= \Delta_{C,\ell}^+ \mathbb{E}[n_\ell(T)] + \sum_{a^* < i < \ell} \Delta_{C,i}^+ \mathbb{E}[n_i(T)] \quad (141)$$

$$\leq \Delta_{C,\ell}^+ \mathbb{E}[n_\ell(T)] + \max_{a^* < i < \ell} \Delta_{C,i}^+ \sum_{a^* < i < \ell} \mathbb{E}[n_i(T)] \quad (142)$$

$$\begin{aligned}
&< \left(1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \sum_{i=1}^{a^*} \frac{43}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+ \\
&+ \Delta_{C,\ell}^+ \mathbb{E}[n_\ell(T) \mid \{Z > a^*\}] \cdot \Pr(Z > a^*) \\
&+ \max_{a^* < i < \ell} \Delta_{C,i}^+ \mathbb{E} \left[\sum_{a^* < i < \ell} n_i(T) \mid \{Z > a^*\} \right] \cdot \Pr(Z > a^*) \quad (143) \\
&\leq \left(1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \sum_{i=1}^{a^*} \frac{43}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+ \\
&+ \max_{a^* < i \leq \ell} \Delta_{C,i}^+ \cdot T \cdot \Pr(Z > a^*) \quad (144) \\
&\leq \left(1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \sum_{i=1}^{a^*} \frac{43}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+ \\
&+ \max_{a^* < i \leq \ell} \Delta_{C,i}^+ \cdot T \cdot \Pr(B_{a^*}) \quad (\because \{Z > a^*\} \implies B_{a^*}) \quad (145) \\
&\leq \left(1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \sum_{i=1}^{a^*} \frac{43}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+ + \frac{43}{\Delta_{Q,a^*}^2} \max_{a^* < i \leq \ell} \Delta_{C,i}^+. \quad (146)
\end{aligned}$$

Where the final inequality follows from the bound in 137. Similarly for Quality Regret we have,

$$\begin{aligned}
\mathbb{E}[\text{Quality_Reg}(T, \nu)] &= \sum_{i=1}^{\ell} \Delta_{Q,i}^+ \mathbb{E}[n_i(T)] \quad (147) \\
&= \sum_{i=1}^{a^*-1} \Delta_{Q,i} \mathbb{E}[n_i(T)] + \sum_{i=a^*+1}^{\ell-1} \Delta_{Q,i}^+ \mathbb{E}[n_i(T)] \quad (148) \\
&\leq \sum_{i=1}^{a^*-1} \Delta_{Q,i} \mathbb{E}[n_i(T)] + \max_{a^* < i < \ell} \Delta_{Q,i}^+ \sum_{i=a^*+1}^{\ell-1} \mathbb{E}[n_i(T)] \quad (149) \\
&< \sum_{i=1}^{a^*-1} \left(\Delta_{Q,i} + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}} + \frac{43}{\Delta_{Q,i}} \right) + \frac{43}{\Delta_{Q,a^*}^2} \max_{a^* < i < \ell} \Delta_{Q,i}^+. \quad (150)
\end{aligned}$$

Which are the bounds stated in Theorem 3.2. \square

For an improved understanding of these upper bounds, we provide a description of the terms.

First for Cost Regret,

$$\underbrace{\left(1 + \max_{i \leq a^*} \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+}_{\text{Contribution from } \ell \text{ under nominal termination in PE episode } a^*} + \underbrace{\left(\sum_{i=1}^{a^*} \frac{43}{\Delta_{Q,i}^2}\right) \Delta_{C,\ell}^+}_{\text{Contribution from } \ell \text{ under mis-termination in PE episode } \leq a^*} + \underbrace{\frac{43}{\Delta_{Q,a^*}^2} \max_{a^* < i \leq \ell} \Delta_{C,i}^+}_{\text{Contribution from episodes } > a^* \text{ in case of mis-termination during ep } a^*}.$$

Next for Quality Regret,

$$\underbrace{\sum_{i=1}^{a^*-1} \left(\Delta_{Q,i} + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}} \right)}_{\text{Contribution from } i < a^* \text{ under nominal termination in PE episode } a^*} + \underbrace{\sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}}}_{\text{Contribution from } i < a^* \text{ under mis-termination in PE episode } \leq a^*} + \underbrace{\frac{43}{\Delta_{Q,a^*}^2} \max_{a^* < i < \ell} \Delta_{Q,i}^+}_{\text{Contribution from episodes } > a^* \text{ in case of mis-termination during ep } a^*}.$$

G ANALYSIS FOR PE-CS IN THE FULL COST-SUBSIDY SETTING

We now turn towards proving Theorem 3.4 that establishes an upper bound on Expected Cumulative Cost and Quality regret for PE-CS. The PE-CS algorithm operates in the unknown reference arm setting, which we also refer to from time-to-time to be the Full Cost-Subsidy Setting Sinha et al. (2021). We have already shown upper bounds on Cost and Quality regret for Pairwise-Elimination (PE) for its operation in the known reference arm setting. The principle hurdle in generalizing the PE analysis to an analysis for PE-CS is working through the uncertainty associated with the identification of the best arm in the BAI stage of PE-CS. To perform the PE-CS analysis, not only do we need the definitions, notation, and setup from the analysis of PE, but also we require some additional constructs spelled out in the following Section.

G.1 PE-CS AND UNKNOWN REFERENCE ARM SETTING SPECIFIC DEFINITIONS

As described in Algorithm 2, PE-CS operates in two phases, a Best-Arm-Identification (BAI) phase and a Pairwise Elimination (PE) phase. As discussed in Section 3, the BAI phase of PE-CS is the Improved UCB algorithm Auer & Ortner (2010) terminated once there is a single active arm remaining. The single remaining arm is assigned to be the empirical reference arm denoted by ℓ . As we show in the subsequent work, by the manner in which PE-CS is setup, the event that the identified reference arm gets arm i^* in the line $\ell \leftarrow \mathcal{A}[0]$ (Line 6, Algorithm 2) is a high probability event. The core idea behind the analysis is to condition the expected number of samples on this desirable and likely outcome occurring during the BAI stage.

In addition to the notation defined in the main paper, we define more constructs that are specific to the analysis of PE-CS. For arm i , define round number σ_i within the Best-Arm-Identification (BAI) phase of PE-CS to be,

$$\sigma_i = \min \left\{ m \mid \tilde{\Delta}_m < \frac{\Delta_i}{2} \right\}. \quad (151)$$

Intuitively, round σ_i is the round number by which we expect arm i to be eliminated by arm i^* in the BAI stage of PE-CS.

We let the final round during which arm i was sampled during the BAI stage be denoted by the random variable Σ_i . To apportion the contributions of the BAI stage and the PE stage to the total number of samples $n_i(T)$, we introduce the variable t_{BAI} to denote the final time-step t in the BAI stage. As a consequence of these two definitions $n_i(1, t_{\text{BAI}}) \leq \tau_{\Sigma_i}$. The highest round number reached during the BAI stage overall for any and all active arms,

$$\Sigma_f = \max_{i \neq i^*} \Sigma_i. \quad (152)$$

And to denote the set of active arms at the last time-step of the BAI stage we use,

$$\mathcal{A}_f = \mathcal{A}(t_{\text{BAI}}). \quad (153)$$

Next, we define a collection of useful events in Table 3 for the analysis that follows. These events are contingent on outcomes occurring during the BAI stage alone.

Table 3: Probabilistic Events Descriptions and Symbols for PE-CS Analysis

Symbol	Event Description
$\Gamma_i, \forall i \neq i^*$	Arm i is eliminated by when round $m = \sigma_i$ during the BAI-stage
$\beta_{1,i}, \forall i \neq i^*$	Arm i is not eliminated by when round $m = \sigma_i$ and Arm $i^* \in \mathcal{A}_{\sigma_i}$ during the BAI-stage
$\beta_{2,i}, \forall i \neq i^*$	Arm i is not eliminated by when round $m = \sigma_i$ and Arm $i^* \notin \mathcal{A}_{\sigma_i}$ during the BAI-stage
$F_i, \forall i \neq i^*$	Final BAI round $\Sigma_f < \sigma_i$ and Arm $i \in \mathcal{A}_f$

Remark G.1 (The event F_i). *The event F_i in words is the event that samples run out before the validity of the event Γ_i could be checked and before the arm i could be eliminated. Equivalent to the description in Table 3, we can write $F_i = \{\Sigma_f < \sigma_i\} \cap \{i \in \mathcal{A}_f\}$. Using De Morgan’s rules, its complement is given by $F_i^c = \{\Sigma_f \geq \sigma_i\} \cup \{i \notin \mathcal{A}_f\}$. Consequently, $\Gamma_i \subseteq F_i^c$, and $\Gamma_i = \Gamma_i \cap F_i^c$. Since event F_i requires that $i \in \mathcal{A}_f$, the event $\{t_{\text{BAI}} = T\} \subset F_i \forall i \neq i^*$.*

Remark G.2 (Events are proper subsets of Ω). *Similar to the events defined in Table 2 for the analysis of PE, the events in Table 3 are proper subsets of the sample space Ω . From the setup inherited from Improved-UCB, Γ_i is a desirable and likely fate for Arm i during the BAI-stage.*

To analyze the outcome of the BAI stage of the PE-CS algorithm we intuit and validate a three way partition of the sample space Ω into the events Γ , β , and F . Γ is the event conditioning on which ensures that $\ell = i^*$ by requiring that the events Γ_i held for each arm $i \neq i^*$. Additionally we require that there are samples remaining at the end of the BAI stage by intersecting with $\{t_{\text{BAI}} < T\}$. This structure makes the downstream analysis of the PE-stage tractable.

$$\Gamma = \{t_{\text{BAI}} < T\} \cap \bigcap_{i \neq i^*} \Gamma_i. \quad (154)$$

Let A denote the set of sub-optimal arms, we know that $|A| = K - 1$. We use $\mathcal{P}(A)$ to denote the power set of A , that is the collection of all the possible sub-sets of A . Let $S \in \mathcal{P}(A)$ denote an arbitrary subset of A . We define an event $F(S) \subseteq \Omega$ parameterized by the set S as,

$$F(S) = \{t_{\text{BAI}} = T\} \cap \left(\bigcap_{i \in S} F_i \cap \bigcap_{j \in A \setminus S} (\Gamma_j \cap F_j^c) \right) \quad (155)$$

Intuitively, the set S consists of arms $i \in S$ for which F_i held thereby making Γ_i unverifiable. In contrast the arms j contained in $j \in A \setminus S$ are those for which Γ_j held. An inspection of the definition of $F(S)$ in Equation 155 reveals that $F(S_1) \cap F(S_2) = \phi \forall S_1, S_2 \in \mathcal{P}(A), S_1 \neq S_2$. Taking a union over all possible $F(S)$ we get the compound event F ,

$$F = \bigcup_{S \in \mathcal{P}(A)} F(S). \quad (156)$$

Finally, the event β is the event that Γ_i did not hold for some arm $i \neq i^*$ despite it being verifiable (F_i^c holding).

$$\beta = \bigcup_{i \neq i^*} (\Gamma_i^c \cap F_i^c). \quad (157)$$

Lemma G.1 (A partition of Ω using Γ). *The events Γ , F , and β as defined in Equations 154, 156, and 157 respectively, form a mutually exclusive and exhaustive partition of the sample space Ω .*

Proof. First we show that the sets are mutually exclusive by showing that their pairwise intersections namely $\Gamma \cap F$, $F \cap \beta$, and $\beta \cap \Gamma$ are all ϕ . Starting off, it is easy to see that $\Gamma \cap F = \phi$ since,

$$\Gamma \cap F \subseteq \{t_{\text{BAI}} < T\} \cap \{t_{\text{BAI}} = T\} = \phi \text{ (definitions from Equations 154 and 156)}. \quad (158)$$

Next, to show that $F \cap \beta = \phi$, it is sufficient to show that $F(S) \cap \beta = \phi$ for arbitrary $S \in \mathcal{P}(A)$.

$$F(S) \cap \beta = \{t_{\text{BAI}} = T\} \cap \left(\bigcap_{i \in S} F_i \cap \bigcap_{j \in A \setminus S} (\Gamma_j \cap F_j^c) \right) \cap \bigcup_{k \neq i^*} (\Gamma_k^c \cap F_k^c) \quad (159)$$

Since both $F_i \cap (\Gamma_i^c \cap F_i^c) = \phi$ and $(\Gamma_i \cap F_i^c) \cap (\Gamma_i^c \cap F_i^c) = \phi$, we shall have $F(S) \cap \beta = \phi$.

Lastly, for the pair $\beta \cap \Gamma$ we have,

$$\beta \cap \Gamma = \bigcup_{i \neq i^*} (\Gamma_i^c \cap F_i^c) \cap \bigcap_{j \neq i^*} \Gamma_j \quad \text{(since the indexing for } \beta \text{ and } \Gamma \text{ need not coincide)} \quad (160)$$

$$2052 \quad = \bigcup_{i \neq i^*} (\Gamma_i^c \cap F_i^c) \cap \bigcap_{j \neq i^*} (\Gamma_j \cap F_j^c) \quad (\text{since } \Gamma_j = F_j^c \cap \Gamma_j) \quad (161)$$

$$2053 \quad = \bigcup_{i \neq i^*} \bigcap_{j \neq i^*} ((\Gamma_i^c \cap F_i^c) \cap (\Gamma_j \cap F_j^c)) = \phi. \quad (162)$$

2057 Next we show that $\Gamma \cup F \cup \beta = \Omega$, that is, the event collection considered in Lemma G.1 is
2058 exhaustive.

$$2059 \quad \Gamma \cup F \supset \Gamma \cup F(\phi) = \bigcap_{i \neq i^*} (\Gamma_i \cap F_i^c) \quad (\text{plugging } F(S) \text{ when } S = \phi \text{ per Equation 155}) \quad (163)$$

$$2061 \quad \implies \Gamma \cup F \cup \beta \supset \bigcap_{i \neq i^*} (\Gamma_i \cap F_i^c) \cup \bigcup_{j \neq i^*} (\Gamma_j^c \cap F_j^c) \quad (164)$$

$$2063 \quad = \bigcup_{j \neq i^*} \bigcap_{i \neq i^*} ((\Gamma_j^c \cap F_j^c) \cup (\Gamma_i \cap F_i^c)) \quad (165)$$

$$2066 \quad \supseteq \bigcap_{i \neq i^*} ((\Gamma_i^c \cap F_i^c) \cup (\Gamma_i \cap F_i^c)) = \bigcap_{i \neq i^*} F_i^c. \quad (166)$$

2069 We shall now show that $F \cup \beta \supseteq \bigcup_{i \neq i^*} F_i$ which combined with Equation 166 completes the check
2070 on the exhaustive criteria. To do this we show that $F_i \subseteq F \cup \beta \forall i \neq i^*$.

2071

2072 PROOF THAT $F_i \subseteq F \cup \beta$:

2073 To prove this result, we start with the event $F(\{i\})$,

$$2074 \quad F(\{i\}) = \{t_{\text{BAI}} = T\} \cap \left(F_i \cap \bigcap_{j \in A, j \neq i} (\Gamma_j \cap F_j^c) \right) \quad (167)$$

$$2078 \quad = F_i \cap \bigcap_{j \in A, j \neq i} (\Gamma_j \cap F_j^c) \quad (\because F_i \subseteq \{t_{\text{BAI}} = T\}). \quad (168)$$

2081 Without loss of generality, let the set of remaining sub-optimal arms $A \setminus \{i\} = \{p, q, \dots\}$. The
2082 idea behind this proof is to identify sub-events $F(\cdot)$ such that iteratively taking their union with one
2083 another and with events lying in β reveals that $F_i \subseteq F \cup \beta$. Since $\beta = \bigcup_{k \neq i^*} (F_k^c \cap \Gamma_k^c)$ we have,

$$2084 \quad (F_p^c \cap \Gamma_p^c) \subseteq \beta \implies F_i \cap (\Gamma_p^c \cap F_p^c) \cap \bigcap_{j \in A \setminus \{i, p\}} (\Gamma_j \cap F_j^c) \subseteq \beta \quad (169)$$

2086 (intersecting with sets keeps us inside β)

$$2088 \quad \implies F(\{i\}) \cup F_i \cap (\Gamma_p^c \cap F_p^c) \cap \bigcap_{j \in A \setminus \{i, p\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta \quad (\because F(\{i\}) \subseteq F) \quad (170)$$

$$2091 \quad \implies F_p^c \cap F_i \cap \bigcap_{j \in A \setminus \{i, p\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta \quad (171)$$

2093 ($\because (\Gamma_p^c \cap F_p^c) \cup (\Gamma_p \cap F_p^c)$ from $\beta, F(\{i\})$)

$$2095 \quad \implies F(\{i, p\}) \cup F_p^c \cap F_i \cap \bigcap_{j \in A \setminus \{i, p\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta \quad (172)$$

$$2097 \quad \implies F_i \cap \bigcap_{j \in A \setminus \{i, p\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta. \quad (173)$$

2100 In reaching Equation 173, we have removed the dependence on Arm p for the event on the left. With
2101 the next series of equations, we further remove the dependence on Arm q .

$$2102 \quad F_i \cap \bigcap_{j \in A \setminus \{i, p\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta \quad (174)$$

$$2105 \quad \implies F_i \cap (\Gamma_q \cap F_q^c) \cap \bigcap_{j \in A \setminus \{i, p, q\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta. \quad (175)$$

2106 Similar to what we saw in the first iteration of this procedure, we have,
2107

$$2108 \quad F_i \cap (\Gamma_q^c \cap F_q^c) \cap \bigcap_{j \in A \setminus \{i,p,q\}} (\Gamma_j \cap F_j^c) \subseteq \beta \quad (176)$$

$$2109 \quad \implies F_i \cap F_q^c \cap \bigcap_{j \in A \setminus \{i,p,q\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta \quad (\text{combining with Equation 175}). \quad (177)$$

2110 Similar to how we reached the result in Equation 173 in starting from $F(\{i\})$, if instead we had
2111 started with the set $F(\{i,q\})$, and then eliminated the dependence on p , we would have shown,
2112

$$2113 \quad F_i \cap F_q \cap \bigcap_{j \in A \setminus \{i,p,q\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta. \quad (178)$$

2114 Combining Equations 177 and 178, we have,
2115

$$2116 \quad F_i \cap \bigcap_{j \in A \setminus \{i,p,q\}} (\Gamma_j \cap F_j^c) \subseteq F \cup \beta. \quad (179)$$

2117 Repeating this procedure iteratively, it is clear that the event on the left can be pruned down to simply
2118 F_i , and therefore,
2119

$$2120 \quad F_i \subseteq F \cup \beta. \quad (180)$$

2121 Since no assumptions were made on the choice of i , we have,
2122

$$2123 \quad \bigcup_{i \neq i^*} F_i \subseteq F \cup \beta. \quad (181)$$

2124 As mentioned earlier, we can combine $\Gamma \cup F \cup \beta \supseteq \bigcap_{i \neq i^*} F_i^c$ from Equation 166 and $F \cup \beta \supseteq$
2125 $\bigcup_{i \neq i^*} F_i$ from Equation 181 to obtain $\Gamma \cup F \cup \beta = \Omega$. \square
2126

2127 **Corollary G.1** (Corollary to Lemma G.1). *The events $\Gamma, \{F(S)\}_{S \in \mathcal{P}(A)}, \beta$ form a mutually exclu-*
2128 *sive and exhaustive partition over the sample space Ω . This result follows trivially from Lemma G.1*
2129 *and the fact that $S_1 \neq S_2 \implies F(S_1) \cap F(S_2) = \phi$.*
2130

2131 Next we prove an upper bound on the probability of the event β which we will need repeatedly in
2132 proving subsequent results.
2133

2134 **Lemma G.2** (Bound on $\Pr(\beta)$). *To bound the expected number of samples in all the cases pertinent*
2135 *to PE-CS we show the following bound on $\Pr(\beta)$,*
2136

$$2137 \quad \Pr(\beta) < \frac{11}{T \Delta_{\min}^2} + \sum_{i \neq i^*} \frac{32}{T \Delta_i^2}. \quad (182)$$

2138 *Proof.*
2139

$$2140 \quad \Pr(\beta) = \Pr\left(\bigcup_{i \neq i^*} (\Gamma_i^c \cap F_i^c)\right) \quad (183)$$

$$2141 \quad = \Pr\left(\bigcup_{i \neq i^*} (\beta_{1,i} \cup \beta_{2,i})\right) \quad (\text{since } \Gamma_i^c \cap F_i^c = \beta_{1,i} \cup \beta_{2,i}) \quad (184)$$

$$2142 \quad \leq \Pr\left(\bigcup_{i \neq i^*} \beta_{1,i}\right) + \Pr\left(\bigcup_{i \neq i^*} \beta_{2,i}\right) \quad (\text{Union Bound}) \quad (185)$$

$$2143 \quad \leq \sum_{i \neq i^*} \Pr(\beta_{1,i}) + \Pr\left(\bigcup_{i \neq i^*} \{\text{Arm } i^* \notin \mathcal{A}_{\sigma_i}\}\right) \quad (\text{Union Bound, Latter clause of } \beta_{2,i}) \quad (186)$$

$$= \sum_{i \neq i^*} \Pr(\beta_{1,i}) + \Pr(\{\text{Arm } i^* \notin \mathcal{A}_{\sigma_{\max}}\}). \quad (187)$$

Where $\sigma_{\max} = \max_{i \neq i^*} \sigma_i$, and Equation 187 follows from the fact that $i^* \notin \mathcal{A}_{\sigma_1} \implies i^* \notin \mathcal{A}_{\sigma_2}$ when $\sigma_2 > \sigma_1$.

The event $\beta_{1,i}$ is the event that Arm i is not eliminated by round σ_i while Arm i^* is active at the end of the sampling for round σ_i . The term $\Pr(\beta_{1,i})$ therefore can be bounded in a manner analogous to the way the probability of low-cost infeasible arm i not being eliminated by reference arm ℓ was bound in Equation 67. The difference being that the round number ρ_i is replaced by the round σ_i , or equivalently, the gap $\Delta_{Q,i}$ is replaced by the gap Δ_i . Therefore,

$$\Pr(\beta_{1,i}) \leq \frac{32}{T\Delta_i^2}. \quad (188)$$

The problem of analyzing $\Pr(\beta_{2,i})$ is analogous to the analysis of $\Pr(B_{3,i})$ in the proof of Lemma F.2 for the PE algorithm. By applying the Hoeffding bound (Lemma D.2) to the clauses in Equations 54 and 56 we were able to establish that the probability of the event $\{\text{Arm } \ell \text{ eliminated by infeasible arm } i \text{ after the sampling for an arbitrary round } \rho \text{ concludes}\}$ is upper bounded by $\frac{4^\kappa + 1}{T\Delta_\rho^2}$. Since for the BAI setting the samples of all the arms are always matched ($\kappa = 0$) here we shall have,

$$\Pr(\{\text{Arm } i^* \notin \mathcal{A}_{\sigma_{\max}}\}) \leq \sum_{\rho=0}^{m_{\sigma_{\max}}-1} \frac{2}{T\Delta_\rho^2} \leq \frac{2}{T} \sum_{\rho=0}^{m_{\sigma_{\max}}-1} 4^\rho \quad (189)$$

$$< \frac{2 \cdot 4^{\sigma_{\max}}}{3T} \quad (190)$$

$$\leq \frac{2}{3T\Delta_{\sigma_{\max}}^2} \quad (191)$$

$$\leq \frac{32}{3T\Delta_{\min}^2} \quad (192)$$

$$< \frac{11}{T\Delta_{\min}^2}. \quad (193)$$

Combining the bounds shown in Equations 188 and 193 we obtain the overall bound stated in Lemma G.2. \square

We now move on to analyzing the evolution of samples in the PE stage of PE-CS. The pieces needed from the analysis of the BAI stage are the partition over Ω from Lemma G.1, the bound on $\Pr(\beta)$ shown in Lemma G.2, and the Iterated Expectation Lemma D.3. The key difference between the analysis of PE and the PE-stage in PE-CS is the possibility that the round number Σ_i to which the samples of an arbitrary arm $i \neq i^*$ advance during the BAI-stage exceeds the round number ρ_i defined in Equation 34. Our modular proof technique sequesters both the pathological (event F) and the unlikely (event β) outcomes of the BAI stage away from the PE stage. In our approach the $\Sigma_i > \rho_i$ case in the analysis of the PE-stage of PE-CS only surfaces for episode a^* . Moreover, analyzing samples accrued during episode a^* is only called for when bounding the expected number of samples of the best arm i^* .

Remark G.3. *Similar to Remark F.1 we note here that Arm i^* during the PE stage of PE-CS really refers to a hypothetical Bandit Arm with expected return $(1 - \alpha)\mu_*$.*

Using all the definitions and constructs introduced in this section, we are now in a position to show an upper bound on the expected number of samples of low-cost arms in Lemma G.3, the best arm i^* in Lemma G.4, and high-cost arms in Lemma G.5.

2214 UPPER BOUND ON THE NUMBER OF SAMPLES FOR ARMS $i < a^*$

2215 **Lemma G.3** (Bound on the expected number of samples of a low-cost arm under PE-CS). *For any*
 2217 *low cost infeasible arm $i < a^*$, its expected number of samples accrued is upper bounded by,*

$$2218 \mathbb{E}[n_i(T)] < 1 + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \frac{43}{\Delta_{Q,i}^2} + \mathbb{E}[n_i(T) | \beta] \cdot \left(\frac{11}{T \Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T \Delta_j^2} \right).$$

2222 *Proof.* We begin the analysis by applying the Iterated Expectation Lemma D.3 to the partition de-
 2224 veloped in Lemma G.1.

$$2226 \mathbb{E}[n_i(T)] = \mathbb{E}[n_i(T) | \Gamma] \Pr(\Gamma) + \sum_{S \in \mathcal{P}(A)} \mathbb{E}[n_i(T) | F(S)] \Pr(F(S)) + \mathbb{E}[n_i(T) | \beta] \Pr(\beta)$$

$$2229 \leq \max \left\{ \mathbb{E}[n_i(T) | \Gamma], \{ \mathbb{E}[n_i(T) | F(S)] \}_{S \in \mathcal{P}(S)} \right\} + \mathbb{E}[n_i(T) | \beta] \Pr(\beta) \quad (195)$$

$$2231 (\because \Pr(\Gamma) + \sum \Pr(F(S)) < 1)$$

$$2233 \leq \max \{ \mathbb{E}[n_i(T) | \Gamma], \tau_{\sigma_i} \} + \mathbb{E}[n_i(T) | \beta] \Pr(\beta). \quad (196)$$

2235 Where Equation 196 follows from the fact that conditioned on any $F(S)$, the maximum round up
 2236 to which arm i can be sampled is σ_i both in the case when Γ_i holds ($i \notin S$), and in the case when
 2237 $i \in S$, and $F_i = \{\Sigma_f < \sigma_i\} \cap \{i \in \mathcal{A}_f\}$ holds instead. We now proceed by separately bounding the
 2238 $\mathbb{E}[n_i(T) | \Gamma]$ term by further conditioning on the cases where $\Sigma_i > \rho_i$ and $\Sigma_i \leq \rho_i$.

$$2239 \mathbb{E}[n_i(T) | \Gamma] = \mathbb{E}[n_i(T) | \Sigma_i \leq \rho_i, \Gamma] \Pr(\Sigma_i \leq \rho_i | \Gamma)$$

$$2240 + \mathbb{E}[n_i(T) | \Sigma_i > \rho_i, \Gamma] \Pr(\Sigma_i > \rho_i | \Gamma) \quad (\text{using Lemma D.3 on } \{\Sigma_i \leq \rho_i\})$$

$$2241 \quad (197)$$

$$2242 \leq \mathbb{E}[n_i(T) | \Sigma_i \leq \rho_i, \Gamma] = \mathbb{E}_1[n_i(T)] \quad (\text{introducing shorthand notation } \mathbb{E}_1).$$

$$2243 \quad (198)$$

2245 Where ρ_i is as defined in Equation 34. In writing Equation 198 we leverage the fact that for a low
 2246 cost arm with index $i < a^*$, $\Delta_{Q,i} = (1 - \alpha)\mu_* - \mu_i$ is necessarily a smaller gap than $\Delta_i = \mu_* - \mu_i$
 2247 since by construction, each of these low cost arms has a return $\mu_i < \mu_{CS} = (1 - \alpha)\mu_*$. It follows
 2248 that $\Pr(\Sigma_i > \rho_i, \Gamma) = \Pr(\Sigma_i > \rho_i | \Gamma) = 0$ because the largest value that the random variable Σ_i
 2249 can take under Γ is σ_i , and $\Delta_{Q,i} < \Delta_i \implies \rho_i > \sigma_i$.

2251 BOUND ON $\mathbb{E}_1[n_i(T)]$

2252 Since we enter the PE stage of the algorithm with a round number $\Sigma_i \leq \rho_i$, we use the same event
 2253 construction of the compound event $G_i \forall i < a^*$, defined and used in the Proof of Lemma F.2. There-
 2254 fore, just as before we work towards a bound by conditioning on the partition with G_i introduced
 2255 in Lemma F.1³. Let Λ_i be the random variable denoting the highest round number corresponding to
 2256 which sampling was performed for arm i during the run of PE-CS.

$$2258 \mathbb{E}_1[n_i(T)] = \mathbb{E}_1[n_i(T) | G_i] \Pr(G_i) + \mathbb{E}_1[n_i(T) | B_{1,i}] \Pr(B_{1,i}) + \mathbb{E}_1[n_i(T) | B_i] \Pr(B_i)$$

$$2260 \leq \max \{ \mathbb{E}_1[n_i(T) | G_i], \mathbb{E}_1[n_i(T) | B_{1,i}] \} + T \cdot \Pr(B_i) \quad (200)$$

$$2262 (\text{Since } \Pr(G_i) + \Pr(B_{1,i}) < 1)$$

$$2264 = \max \{ \mathbb{E}_1[n_i(T) | G_i], \mathbb{E}_1[n_i(1, t_{\text{BAI}}) + n_i(t_{\text{BAI}} + 1, T) | B_{1,i}] \} + T \cdot \Pr(B_i) \quad (201)$$

$$2266 \leq \max \{ \mathbb{E}_1[\tau_{\Lambda_i} | G_i], \mathbb{E}_1[\tau_{\Sigma_i} + 0 | B_{1,i}] \} + T \cdot \Pr(B_i) \quad (202)$$

³The probability operator \Pr is also for the conditional distribution conditioned on $\{\Sigma_i \leq \rho_i, \Gamma\}$.

2268 $(\tau_m \text{ as defined in Equation 35})$

2269 $\leq \max \{ \tau_{\rho_i}, \tau_{\sigma_i} \} + T \cdot \Pr(B_i)$ (203)

2270 $= \tau_{\rho_i} + T \cdot \Pr(B_i)$ (since $(1 - \alpha) \mu_* - \mu_i = \Delta_{Q,i} < \Delta_i = \mu_* - \mu_i \forall i < a^*$) (204)

2271 $< 1 + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + T \cdot \Pr(B_i)$ (Using bound on τ_{ρ_i} in Equation 53). (205)

2272
2273
2274
2275 Where the treatment of the random variable $n_i(t_{\text{BAI}} + 1, T)$ is based on the expression for the addi-
2276 tional rounds for which arm i is sampled during the PE stage of PE-CS. In Equation 202 conditioned
2277 on G_i there may be more samples, however conditioned on $B_{1,i}$ there shall be no further samples
2278 since the episode corresponding to i is never initiated. Now to bound $\Pr(B_i \mid \Sigma_i \leq \rho_i, \Gamma)$ we use
2279 arguments similar to the ones in Proof of Lemma F.2.

2280 $\Pr(B_i \mid \Sigma_i \leq \rho_i, \Gamma) = \Pr(B_{1,i} \cup B_{2,i} \mid \Sigma_i \leq \rho_i, \Gamma).$ (206)

2281 Where $B_{1,i} = G_{1,i} \cap G_{2,i}^c$ and $B_{2,i} = G_{1,i} \cap G_{3,i}^c$ as in the proof of Lemma F.2. The Probability
2282 $\Pr(B_{1,i})$ in Lemma F.2 was bound by the probability of either Clause 54 or Clause 55 being violated
2283 during round ρ_i . Due to the parallel nature of the construction here, and the possibility of round ρ_i
2284 being conducted, we can upper bound $\Pr(B_{1,i} \mid \Sigma_i \leq \rho_i, \Gamma)$ identically as,

2285 $\Pr(B_{1,i} \mid \Sigma_i \leq \rho_i, \Gamma) \leq \frac{32}{T \Delta_{Q,i}^2}.$ (207)

2287
2288 Now we move on to bounding $\Pr(B_{2,i} \mid \Sigma_i \leq \rho_i, \Gamma)$ by bounding the probability of arm ℓ being
2289 eliminated in any round $\omega_i = \rho$ lying in the range $\Sigma_i \leq \rho \leq \rho_i$. Just like in the proof of Lemma
2290 F.2, even here Clauses 54 and 56 holding simultaneously preclude arm ℓ from being eliminated by
2291 arm i regardless of round ρ . Therefore using work identical to the one that goes into establishing
2292 Equations 72 and 73 we have,

2293 $\Pr(B_{2,i} \mid \Sigma_i \leq \rho_i, \Gamma) \leq \sum_{\rho=\Sigma_i}^{\rho_i-1} \frac{4^\kappa + 1}{T \bar{\Delta}_\rho^2}$ (208)

2294 $\leq \sum_{\rho=0}^{\rho_i-1} \frac{4^\kappa + 1}{T \bar{\Delta}_\rho^2}$ (Because $\Sigma_i \geq 0$) (209)

2295 $< \frac{11}{T \Delta_{Q,i}^2}$ (From Equation 73 in Lemma F.2, $\kappa = 0$). (210)

2300 Applying the Union bound to Equation 206, and plugging in the bounds in 207 and 210, we have,

2301 $\Pr(B_i \mid \Sigma_i \leq \rho_i, \Gamma) < \frac{43}{T \Delta_{Q,i}^2}.$ (211)

2302 Substituting the bounds in Equations 211, 205, and Lemma G.2 into Equation 196 we obtain,

2303 $\mathbb{E}[n_i(T)] \leq \max \{ \tau_{\rho_i} + T \cdot \Pr(B_i), \tau_{\sigma_i} \} + \mathbb{E}[n_i(T) \mid \beta] \Pr(\beta)$ (212)

2304 $< 1 + \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \frac{43}{\Delta_{Q,i}^2} + \mathbb{E}[n_i(T) \mid \beta] \Pr(\beta).$ (213)

2305 Which when combined with Lemma G.2 is the bound stated in Lemma G.3. \square

2311 BOUND ON THE NUMBER OF SAMPLES FOR ARM i^*

2312
2313 **Lemma G.4** (Bound on the expected number of samples of the best arm). *For the best arm $i^* =$*
2314 *$\arg \max_{i \in [K]} \mu_i$, the expected number of samples accrued is upper bounded as,*

2315 $\mathbb{E}[n_{i^*}(T)] < 1 + \max \left\{ \frac{32 \log(T \Delta_{\min}^2)}{\Delta_{\min}^2}, \left\{ \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} \right\}_{i \leq a^*} \right\}$

2316 $+ \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + 2 \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) + \mathbb{E}[n_{i^*}(T) \mid \beta] \cdot \left(\frac{11}{T \Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T \Delta_j^2} \right).$

2322 *Proof.* Just like in the proof of Lemma G.3 we sequester away outcomes of the BAI stage that make
 2323 analysis of the PE stage intractable.

$$2324 \mathbb{E}[n_{i^*}(T)] = \mathbb{E}[n_{i^*}(T) | \Gamma] \Pr(\Gamma) + \sum_{S \in \mathcal{P}(A)} \mathbb{E}[n_{i^*}(T) | F(S)] \Pr(F(S)) + \mathbb{E}[n_{i^*}(T) | \beta] \Pr(\beta)$$

$$2327 \leq \max \left\{ \mathbb{E}[n_{i^*}(T) | \Gamma], \left\{ \mathbb{E}[n_{i^*}(T) | F(S)] \right\}_{S \in \mathcal{P}(S)} \right\} + \mathbb{E}[n_{i^*}(T) | \beta] \Pr(\beta)$$

$$2330 (\because \Pr(\Gamma) + \Pr(F) < 1)$$

$$2332 \leq \max \{ \mathbb{E}[n_{i^*}(T) | \Gamma], \tau_{\sigma_{\max}} \} + \mathbb{E}[n_{i^*}(T) | \beta] \Pr(\beta).$$

2334 Equation 216 results from the observation that the maximum round number up till which any arm is
 2335 sampled during the BAI stage under $F(S) \forall S \in \mathcal{P}(A)$ is σ_{\max} . Conditioned on Γ , the reference
 2336 arm ℓ is identified correctly to be i^* . Consequently, the expected number of samples $\mathbb{E}[n_{i^*}(T) | \Gamma]$
 2337 can be analyzed in a manner that closely parallels the analysis of $\mathbb{E}[n_{\ell}(T)]$ in the Proof of Lemma
 2338 F.3.

2339 The key difference from Lemma F.3 is that we grapple with the case $\Sigma_{a^*} > \rho_{a^*}$. This possi-
 2340 bility arises because our bandit instance may have $\Delta_{Q,a^*} < \Delta_{a^*}$, which in turn implies that
 2341 $\Pr(\rho_{a^*} < \Sigma_{a^*} < \sigma_{a^*} | \Gamma) > 0$. The outcome $\Sigma_{a^*} > \rho_{a^*}$ skips the checks associated with the
 2342 events $G_{2,i}$ and $G_{3,i}$. Mathematically, this means that for the event G_{a^*} as defined in Equation
 2343 36 we shall have $\Pr(G_{1,a^*} \cap ((G_{2,a^*} \cap G_{3,a^*}) \cup E_{a^*}) | \Sigma_{a^*} > \rho_{a^*}, \Gamma) = 0$. This motivates us to
 2344 expand the scope of the good event G_{a^*} .

2345 First we define new event $G_{4,a^*} \subset \Omega$ and then we augment the definition of G_{a^*} using this new
 2346 event.

2347 $G_{4,a^*} : \{ \text{Arm } i^* \text{ is eliminated by arm } i \text{ in round } \Sigma_{a^*}, \text{ during episode } a^* \text{ of the PE stage of PE-CS} \}$

2348 **Remark G.4.** Since arm a^* enters the PE stage of PE-CS with samples corresponding to Σ_{a^*}
 2349 number of rounds already accrued, checking whether the event G_{4,a^*} holds does not involve any
 2350 further sampling of arms.

2353 The definition of G_{a^*} is now changed to the one in Equation 217 which supersedes the prior generic
 2354 G_{a^*} (Equation 36).

$$2355 G_{a^*} = \underbrace{G_{1,a^*}}_{\text{Ep. } a^* \text{ is executed}} \cap \left(\underbrace{((G_{2,a^*} \cap G_{3,a^*}) \cup E_{a^*})}_{\text{Prior } G_{a^*} \text{ Clause}} \cup \underbrace{(G_{4,a^*} \cap \{\Sigma_{a^*} > \rho_{a^*}\})}_{\text{New Clause}} \right).$$

2356 **Remark G.5** (Implicit event in Prior G_{a^*} clause). We remark here that the event $\{\Sigma_{a^*} \leq \rho_{a^*}\}$
 2360 is implicit in the event $G_{1,a^*} \cap ((G_{2,a^*} \cap G_{3,a^*}) \cup E_{a^*})$ i.e. $G_{1,a^*} \cap ((G_{2,a^*} \cap G_{3,a^*}) \cup E_{a^*}) \subseteq$
 2361 $\{\Sigma_{a^*} \leq \rho_{a^*}\}$. This is because the event is contingent on correct eliminations happening leading up
 2362 to and during ρ_{a^*} .

2364 While retaining the definition of G from Equation 74, and the definition of B_{1,a^*} from Lemma F.1
 2365 we redefine B_{a^*} so that the relation $G_{a^*}^c = B_{1,a^*} \cup B_{a^*}$ continues to hold.

$$2366 G_{a^*} = G_{1,a^*} \cap (((G_{2,a^*} \cap G_{3,a^*}) \cup E_{a^*}) \cup (G_{4,a^*} \cap \{\Sigma_{a^*} > \rho_{a^*}\}))$$

$$2367 = G_{1,a^*} \cap (((((G_{2,a^*} \cap G_{3,a^*}) \cup E_{a^*})) \cap \{\Sigma_{a^*} \leq \rho_{a^*}\}) \cup (G_{4,a^*} \cap \{\Sigma_{a^*} > \rho_{a^*}\}))$$

2370 (due to Remark G.5)

$$2371 \Rightarrow G_{a^*}^c = G_{1,a^*}^c \cup \left(\underbrace{(((G_{2,a^*}^c \cup G_{3,a^*}^c) \cap E_{a^*}^c))}_{B_{a^*}^{\text{original}}} \cup \{\Sigma_{a^*} > \rho_{a^*}\} \right) \cap (G_{4,a^*}^c \cup \{\Sigma_{a^*} \leq \rho_{a^*}\})$$

$$\begin{aligned}
&= B_{1,a^*} \cup \left(\left(\left(B_{a^*}^{\text{original}} \cap \{\Sigma_{a^*} \leq \rho_{a^*}\} \right) \cup \{\Sigma_{a^*} > \rho_{a^*}\} \right) \right. \\
&\quad \left. \cap \left(\left(G_{4,a^*}^c \cap \{\Sigma_{a^*} > \rho_{a^*}\} \right) \cup \{\Sigma_{a^*} \leq \rho_{a^*}\} \right) \right) \quad (221)
\end{aligned}$$

$$\begin{aligned}
&= B_{1,a^*} \cup \underbrace{\left(\left(B_{a^*}^{\text{original}} \cap \{\Sigma_{a^*} \leq \rho_{a^*}\} \right) \cup \left(G_{4,a^*}^c \cap \{\Sigma_{a^*} > \rho_{a^*}\} \right) \right)}_{B_{a^*}}. \quad (222)
\end{aligned}$$

Where Equation 221 follows from $A \cup B = (A \cap B^c) \cup B$ being applied to both the left and right clauses. Equation 222 gives us the updated definition of the event B_{a^*} . Armed with the updated event B_{a^*} , we can now develop a bound for $\Pr(G^c | \Gamma)$ using the result in Lemma F.6. By trivially generalizing the bound on $\Pr(G^c)$ developed in Lemma F.6 to the scenario for this Proof where we condition on Γ we shall have,

$$\Pr(G^c | \Gamma) \leq \sum_{i=1}^{a^*} \Pr(B_j | \Gamma) + \Pr(B_{a^*} | \Gamma) \quad (223)$$

$$\leq \sum_{i=1}^{a^*-1} \frac{43}{T\Delta_{Q,i}^2} + 2 \cdot \Pr(B_{a^*} | \Gamma), \quad (224)$$

\because Equation 211, and $\Pr(\Sigma_i > \rho_i | \Gamma) = 0 \forall i < a^*$. We need now only develop a bound for $\Pr(B_{a^*} | \Gamma)$ under its definition in Equation 222.

$$\Pr(B_{a^*} | \Gamma) = \Pr\left(\left(B_{a^*}^{\text{original}} \cap \{\Sigma_{a^*} \leq \rho_{a^*}\}\right) \cup \left(G_{4,a^*}^c \cap \{\Sigma_{a^*} > \rho_{a^*}\}\right) \mid \Gamma\right) \quad (225)$$

$$\leq \Pr\left(B_{a^*}^{\text{original}} \cap \{\Sigma_{a^*} \leq \rho_{a^*}\} \mid \Gamma\right) + \Pr\left(G_{4,a^*}^c \cap \{\Sigma_{a^*} > \rho_{a^*}\} \mid \Gamma\right) \quad (226)$$

(Union Bound)

$$\leq \frac{43}{T\Delta_{Q,a^*}^2} + \Pr\left(G_{4,a^*}^c \cap \{\Sigma_{a^*} > \rho_{a^*}\} \mid \Gamma\right) \quad (\text{From Equation 137}). \quad (227)$$

From Equation 132 in the Proof of Lemma F.2, we have the probability of arm ℓ not being eliminated by arm a^* during round ω_{a^*} of episode a^* as being upper bounded by $\frac{2}{T\Delta_{\omega_{a^*}}^2}$. Therefore since elimination under G_{4,a^*} happens during episode a^* in round Σ_{a^*} ,

$$\Pr\left(G_{4,a^*}^c \cap \{\Sigma_{a^*} > \rho_{a^*}\} \mid \Gamma\right) \leq \frac{2}{T\tilde{\Delta}_{\sigma_{a^*}}^2} \quad (228)$$

(since $\Pr(\Sigma_{a^*} > \sigma_{a^*} | \Gamma) = 0$, and $\tilde{\Delta}_m$ decreases with m).

$$\leq \frac{32}{T\Delta_{a^*}^2} \quad (\text{since by definition of } \sigma_i, \tilde{\Delta}_{\sigma_i} \geq \frac{\Delta_i}{4}). \quad (229)$$

Combining the bounds in Equations 227 and 229 with the Expression 224 we have,

$$\Pr(G^c | \Gamma) \leq \sum_{i=1}^{a^*-1} \frac{43}{T\Delta_{Q,i}^2} + 2 \left(\frac{43}{T\Delta_{Q,a^*}^2} + \frac{32}{T\Delta_{a^*}^2} \right). \quad (230)$$

Armed with the result in Equation 230, we are now in a position to bound $\mathbb{E}[n_{i^*}(T) | \Gamma]$.

$$\mathbb{E}[n_{i^*}(T) | \Gamma] = \mathbb{E}[n_{i^*}(T) | G, \Gamma] \Pr(G | \Gamma) + \mathbb{E}[n_{i^*}(T) | G^c, \Gamma] \Pr(G^c | \Gamma) \quad (231)$$

$$\leq \mathbb{E}[n_{i^*}(T) | G, \Gamma] + T \cdot \Pr(G^c | \Gamma) \quad (232)$$

$$\leq \mathbb{E}[n_{i^*}(T) | G, \Gamma] + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + 2 \left(\frac{43}{\Delta_{Q,a^*}^2} + \frac{32}{\Delta_{a^*}^2} \right). \quad (233)$$

As in the proof of Lemma F.3, let Z denote the last PE episode. To bound $\mathbb{E}[n_{i^*}(T) | G, \Gamma]$ we leverage the result $\Pr(n_{i^*}(T) > \max_{i \neq i^*} n_i(1, t_Z) | G, \Gamma) = 0$. Conditioned on G, Γ there can be no further sampling of the best arm i^* beyond time t_Z , since under Γ the best arm i^* is the reference arm. Consequently,

$$\mathbb{E}[n_{i^*}(T) | G, \Gamma] \leq \mathbb{E}\left[\max_{i \neq i^*} n_i(1, t_Z) \mid G, \Gamma\right] \quad (234)$$

$$\leq \mathbb{E} \left[\max_{i \neq i^*} n_i(1, t_{a^*}) \mid G, \Gamma \right] \quad (\text{because } \Pr(Z > a^* \mid G, \Gamma) = 0) \quad (235)$$

$$= \mathbb{E} \left[\max \left\{ \{n_i(1, t_{a^*})\}_{i < a^*}, n_{a^*}(1, t_{a^*}), \{n_i(1, t_{a^*})\}_{i > a^*} \right\} \mid G, \Gamma \right] \quad (236)$$

$$\leq \mathbb{E} \left[\max \left\{ \{\tau_{\Lambda_i}\}_{i < a^*}, \tau_{\Lambda_{a^*}}, \{\tau_{\Sigma_i}\}_{i > a^*} \right\} \mid G, \Gamma \right] \quad (237)$$

$$\quad (\text{using definitions of } \tau, \Lambda, \text{ and } \Sigma)$$

$$\leq \max \left\{ \{\tau_{\rho_i}\}_{i < a^*}, \max \{\tau_{\rho_{a^*}}, \tau_{\sigma_{a^*}}\}, \{\tau_{\sigma_i}\}_{i > a^*} \right\} \quad (238)$$

$$\quad (\text{because } \Pr(\Lambda_{a^*} > \max\{\rho_{a^*}, \sigma_{a^*}\}) = 0)$$

$$\leq \max \left\{ \{\tau_{\rho_i}\}_{i \leq a^*}, \tau_{\sigma_{\max}} \right\}. \quad (239)$$

Combining Equations 239 and 233 and using the result in Equation 216 we obtain the upper bound stated in Lemma G.4,

$$\mathbb{E}[n_{i^*}(T)] \leq \max \left\{ \{\tau_{\rho_i}\}_{i \leq a^*}, \tau_{\sigma_{\max}} \right\} + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + 2 \left(\frac{43}{\Delta_{Q,a^*}^2} + \frac{32}{\Delta_{a^*}^2} \right) + \mathbb{E}[n_{i^*}(T) \mid \beta] \Pr(\beta) \quad (240)$$

$$< 1 + \max \left\{ \frac{32 \log(T \Delta_{\min}^2)}{\Delta_{\min}^2}, \left\{ \frac{32 \log(T \Delta_{Q,i}^2)}{\Delta_{Q,i}^2} \right\}_{i \leq a^*} \right\} + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} \quad (241)$$

$$+ 2 \left(\frac{43}{\Delta_{Q,a^*}^2} + \frac{32}{\Delta_{a^*}^2} \right) + \mathbb{E}[n_{i^*}(T) \mid \beta] \left(\frac{11}{T \Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T \Delta_j^2} \right). \quad (242)$$

Where Equation 242 follows directly from from Equation 37 and Lemma G.2. \square

UPPER BOUND ON THE NUMBER OF SAMPLES FOR ARMS $i > a^*, i \neq i^*$

Lemma G.5 (Bound on the expected number of samples of high-cost arms). *For any high cost arm with $i > a^*, i \neq i^*$, its expected number of samples are upper bounded as,*

$$\mathbb{E}[n_i(T)] < 1 + \frac{32 \log(T \Delta_i^2)}{\Delta_i^2} + \mathbb{E}[n_i(T) \mid \{Z > a^*\}, \Gamma] \cdot \left(\frac{32}{T \Delta_{a^*}^2} + \frac{43}{T \Delta_{Q,a^*}^2} \right) + \mathbb{E}[n_i(T) \mid \beta] \cdot \left(\frac{11}{T \Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T \Delta_j^2} \right).$$

Proof. To prove the bound stated in Lemma G.5 we proceed with initial steps identical to the ones that go into proving Lemmas G.3 and G.4.

$$\mathbb{E}[n_i(T)] = \mathbb{E}[n_i(T) \mid \Gamma] \Pr(\Gamma) + \sum_{S \in \mathcal{P}(A)} \mathbb{E}[n_i(T) \mid F(S)] \Pr(F(S)) + \mathbb{E}[n_i(T) \mid \beta] \Pr(\beta) \quad (243)$$

$$\leq \max \left\{ \mathbb{E}[n_i(T) \mid \Gamma], \max_{S \in \mathcal{P}(A)} \mathbb{E}[n_i(T) \mid F(S)] \right\} + \mathbb{E}[n_i(T) \mid \beta] \Pr(\beta) \quad (244)$$

$$\quad (\because \Pr(\Gamma) + \Pr(F) < 1)$$

$$\leq \max \left\{ \mathbb{E}[n_i(T) \mid \Gamma], \tau_{\sigma_i} \right\} + \mathbb{E}[n_i(T) \mid \beta] \left(\frac{11}{T \Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T \Delta_j^2} \right). \quad (245)$$

Where the final bound is from Lemma G.2. Now we must bound the expectation term $\mathbb{E}[n_i(T) \mid \Gamma]$. For this we recognize that during the PE-stage, the critical event which determines the number of

2484 samples further accrued for a high-cost arm is whether the final PE episode $Z > a^*$ or not. In the
 2485 case when $Z \leq a^*$, there are no further samples of arm i accrued beyond the BAI-stage.
 2486

$$2487 \mathbb{E}[n_i(T) | \Gamma] = \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \Pr(Z > a^* | \Gamma) \\
 2488 + \mathbb{E}[n_i(T) | \{Z \leq a^*\}, \Gamma] \Pr(Z \leq a^* | \Gamma) \quad (246)$$

$$2489 \leq \tau_{\sigma_i} + \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \Pr(B_{a^*} | \Gamma) \quad (\text{since } \{Z > a^*\} \subseteq B_{a^*}) \quad (247)$$

$$2491 \leq \tau_{\sigma_i} + \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \cdot \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right). \quad (248)$$

2493 Where the final line follows by using the bound developed in Equation 227. Returning to bounding
 2494 $\mathbb{E}[n_i(T)]$ we obtain the bound stated in Lemma G.5,
 2495

$$2496 \mathbb{E}[n_i(T)] \leq \tau_{\sigma_i} + \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \cdot \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right) \\
 2497 + \mathbb{E}[n_i(T) | \beta] \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right) \quad (249)$$

$$2502 < 1 + \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} + \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \cdot \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right)$$

$$2505 + \mathbb{E}[n_i(T) | \beta] \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right). \quad (250)$$

□

2509
 2510 Finally we have all the pieces needed to prove Theorem 3.4. We combine the results obtained in
 2511 Lemmas G.3, G.4, and G.5 by adding together the contributions to regret of all three categories of
 2512 arms while collating like terms.
 2513

2514 *Proof of Theorem 3.4.* Using Equation 13 from the Regret decomposition Lemma D.5 we can ex-
 2515 press and bound the Expected Cumulative Cost Regret as,
 2516

$$2517 \mathbb{E}[\text{Cost_Reg}(T, \nu)] \quad (251)$$

$$2518 = \sum_{i=1}^K \Delta_{C,i}^+ \mathbb{E}[n_i(T)] \quad (252)$$

$$2521 = \sum_{i < a^*} \Delta_{C,i}^+ \mathbb{E}[n_i(T)] + \sum_{i > a^*, i \neq i^*} \Delta_{C,i}^+ \mathbb{E}[n_i(T)] + \Delta_{C,i^*}^+ \mathbb{E}[n_{i^*}(T)] \quad (253)$$

$$2524 = \sum_{i > a^*, i \neq i^*} \Delta_{C,i}^+ \mathbb{E}[n_i(T)] + \Delta_{C,i^*}^+ \mathbb{E}[n_{i^*}(T)] \quad (\text{because } \Delta_{C,i} = 0, \forall i \leq a^*) \quad (254)$$

$$2526 < \sum_{i > a^*, i \neq i^*} \Delta_{C,i}^+ \left(1 + \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} + \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \cdot \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right) \right)$$

$$2529 + \mathbb{E}[n_i(T) | \beta] \cdot \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right)$$

$$2533 + \Delta_{C,i^*}^+ \left(1 + \max \left\{ \frac{32 \log(T\Delta_{\min}^2)}{\Delta_{\min}^2}, \left\{ \frac{32 \log(T\Delta_{Q,i}^2)}{\Delta_{Q,i}^2} \right\}_{i \leq a^*} \right\} + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} \right)$$

$$2536 + 2 \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) + \mathbb{E}[n_{i^*}(T) | \beta] \cdot \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right) \quad (255)$$

$$\begin{aligned}
&= \sum_{i>a^*} \Delta_{C,i} \left(1 + \mathbb{E}[n_i(T) | \beta] \cdot \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right) \right) + \sum_{i>a^*, i \neq i^*} \Delta_{C,i} \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \\
&+ \Delta_{C,i^*}^+ \left(\max_{\Delta \in \{\Delta_{\min}\} \cup \{\Delta_{Q,j}\}_{j \leq a^*}} \left\{ \frac{32 \log(T\Delta^2)}{\Delta^2} \right\} + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + 2 \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) \right) \\
&+ \sum_{i>a^*, i \neq i^*} \Delta_{C,i} \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \cdot \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right) \tag{256}
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i>a^*} \Delta_{C,i} + \Delta_{C,\max} \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right) \sum_{i>a^*} \mathbb{E}[n_i(T) | \beta] + \sum_{i>a^*, i \neq i^*} \Delta_{C,i} \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \\
&+ \Delta_{C,i^*}^+ \left(\max_{\Delta \in \{\Delta_{\min}\} \cup \{\Delta_{Q,j}\}_{j \leq a^*}} \left\{ \frac{32 \log(T\Delta^2)}{\Delta^2} \right\} + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + 2 \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) \right) \\
&+ \max_{i>a^*, i \neq i^*} \Delta_{C,i} \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right) \sum_{i>a^*, i \neq i^*} \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \tag{257}
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i>a^*} \Delta_{C,i} + \Delta_{C,\max} \left(\frac{11}{\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{\Delta_j^2} \right) + \sum_{i>a^*, i \neq i^*} \Delta_{C,i} \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \\
&+ \Delta_{C,i^*}^+ \left(\max_{\Delta \in \{\Delta_{\min}\} \cup \{\Delta_{Q,j}\}_{j \leq a^*}} \left\{ \frac{32 \log(T\Delta^2)}{\Delta^2} \right\} + \sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + 2 \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) \right) \\
&+ \max_{i>a^*, i \neq i^*} \Delta_{C,i} \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) \text{ (linearity of expectation, total samples } T). \tag{258}
\end{aligned}$$

Where $\Delta_{C,\max} = \max_{i \in [K]} \Delta_{C,i}$ is the largest cost gap among all arms. Equation 258 is the upper bound on Expected Cumulative Cost Regret stated in Theorem 3.4.

Proceeding identically, for Quality Regret we shall have,

$$\mathbb{E}[\text{Quality_Reg}(T, \nu)] \tag{259}$$

$$= \sum_{i=1}^K \Delta_{Q,i}^+ \mathbb{E}[n_i(T)] \tag{260}$$

$$= \sum_{i < a^*} \Delta_{Q,i} \mathbb{E}[n_i(T)] + \sum_{i > a^*, i \neq i^*} \Delta_{Q,i}^+ \mathbb{E}[n_i(T)] \tag{261}$$

$$< \sum_{i < a^*} \Delta_{Q,i} \left(1 + \frac{32 \log(T\Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \frac{43}{\Delta_{Q,i}^2} + \mathbb{E}[n_i(T) | \beta] \cdot \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right) \right) \tag{262}$$

$$+ \sum_{i > a^*, i \neq i^*} \Delta_{Q,i}^+ \left(1 + \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} + \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] \cdot \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right) \right)$$

$$+ \mathbb{E}[n_i(T) | \beta] \cdot \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right)$$

$$\leq \sum_{i < a^*} \left(\frac{32 \log(T\Delta_{Q,i}^2)}{\Delta_{Q,i}^2} + \frac{43}{\Delta_{Q,i}^2} \right) + \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right) \Delta_{Q,\max}^+ \sum_{i=1}^K \mathbb{E}[n_i(T) | \beta]$$

$$+ \left(\frac{32}{T\Delta_{a^*}^2} + \frac{43}{T\Delta_{Q,a^*}^2} \right) \max_{i > a^*} \Delta_{Q,i}^+ \sum_{i > a^*, i \neq i^*} \mathbb{E}[n_i(T) | \{Z > a^*\}, \Gamma] + \sum_{i=1}^K \Delta_{Q,i}^+ \tag{263}$$

$$+ \sum_{i>a^*, i \neq i^*} \Delta_{Q,i}^+ \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \quad (264)$$

$$\begin{aligned} &\leq \sum_{i<a^*} \left(\frac{32 \log(T\Delta_{Q,i}^2)}{\Delta_{Q,i}} + \frac{43}{\Delta_{Q,i}} \right) + \sum_{i>a^*, i \neq i^*} \Delta_{Q,i}^+ \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \\ &+ \max_{i>a^*} \Delta_{Q,i}^+ \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) + \Delta_{Q,\max}^+ \left(\frac{11}{T\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{T\Delta_j^2} \right) + \sum_{i=1}^K \Delta_{Q,i}^+. \end{aligned} \quad (265)$$

Where Equation 265 is the upper bound on Expected Cumulative Quality Regret states in Theorem 3.4 and follows from the Linearity of the Expectation operator and the total sample budget being T . \square

Similar to the description that followed the proof of Theorem 3.2, we have for cost regret,

$$\begin{aligned} &\underbrace{\Delta_{C,i^*}^+ \left(1 + \max_{\Delta \in \Delta} \left\{ \frac{32 \log(T\Delta^2)}{\Delta^2} \right\} \right)}_{\text{Contribution from } i^* \text{ under nominal termination in PE-stage episode } a^*} + \underbrace{\Delta_{C,i^*}^+ \left(\sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}^2} + \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right) \right)}_{\text{Contribution from } i^* \text{ under mis-termination in PE-stage episode } \leq a^*} \\ &+ \underbrace{\sum_{i>a^*, i \in [K] \setminus \{i^*\}} \Delta_{C,i}^+ \left(1 + \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \right)}_{\text{Contribution from high-cost arms with a proper end to the BAI-stage}} + \underbrace{\max_{i>a^*, i \in [K]} \Delta_{C,i}^+ \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right)}_{\text{Contribution from PE-stage episodes } > a^* \text{ in case of mis-termination during ep } a^*} \\ &+ \underbrace{\max_{i \in [K]} \Delta_{C,i}^+ \left(\frac{11}{\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{\Delta_j^2} \right)}_{\text{Contribution from improper end to BAI stage}}. \end{aligned}$$

And for quality regret,

$$\begin{aligned} &\underbrace{\sum_{i=1}^{a^*-1} \left(\Delta_{Q,i} + \frac{32 \log(T\Delta_{Q,i}^2)}{\Delta_{Q,i}} \right)}_{\text{Contribution from } i < a^* \text{ under nominal termination in PE-stage episode } a^*} + \underbrace{\sum_{i=1}^{a^*-1} \frac{43}{\Delta_{Q,i}}}_{\text{Contribution from } i < a^* \text{ under mis-termination in PE-stage episode } \leq a^*} + \underbrace{\max_{i>a^*, i \in [K]} \Delta_{Q,i}^+ \left(\frac{32}{\Delta_{a^*}^2} + \frac{43}{\Delta_{Q,a^*}^2} \right)}_{\text{Contribution from PE-stage episodes } > a^* \text{ in case of mis-termination during ep } a^*} \\ &+ \underbrace{\sum_{i>a^*, i \in [K] \setminus \{i^*\}} \Delta_{Q,i}^+ \left(1 + \frac{32 \log(T\Delta_i^2)}{\Delta_i^2} \right)}_{\text{Contribution from PE-stage episodes } > a^* \text{ in case of mis-termination during ep } a^*} + \underbrace{\max_{i \in [K]} \Delta_{Q,i}^+ \left(\frac{11}{\Delta_{\min}^2} + \sum_{j \neq i^*} \frac{32}{\Delta_j^2} \right)}_{\text{Contribution from improper end to BAI stage}}. \end{aligned}$$

Where $\Delta_i = \mu^* - \mu_i$ are conventional gaps, $\Delta_{\min} = \mu^* - \max_{i \neq i^*} \mu_i$ is the smallest conventional gap, and $\Delta = \{\Delta_{\min}\} \cup \{\Delta_{Q,j}\}_{j \leq a^*}$ is defined especially to bound of arm i^* under expectation.

H ANALYSIS FOR MINIMUM TOLERATED REWARD SETTING

Algorithm 1: MINIMUM TOLERATED REWARD (MTR) UCB

```

2650 1 Input: Number of arms  $K$ , Known costs  $c_k$  for each arm  $k \in [K]$ , Minimum Tolerated Reward
2651    $\mu_0$ .
2652 2 Initialize: Empirical Means  $\hat{\mu}_k = 0$ , Pulls of arm  $k : n_k = 0, \forall k \in [K]$ , Time step  $t = 0$ .
2653 3 for Each Round  $t$  do
2654 4   if  $t \leq K$  then
2655 5      $k_t \leftarrow t$ 
2656 6   else
2657 7      $C_t \leftarrow \left\{ k : \hat{\mu}_k(t) + \sqrt{\frac{2 \log t}{n_k(t)}} \geq \mu_0 \right\}$ 
2658 8     if  $C_t \neq \phi$  then
2659 9        $k_t \leftarrow \arg \min_{i \in C_t} c_i$ 
2660 10    else
2661 11       $k_t \leftarrow \text{Uniform}(K)$ 
2662 12     $X_{k_t}(t) \leftarrow \text{Ber}(\mu_{k_t})$ 
2663 13     $\hat{\mu}_{k_t}(t) \leftarrow \frac{\hat{\mu}_{k_t}(t-1)n_{k_t}(t-1) + X_{k_t}(t)}{n_{k_t}(t-1) + 1}$ 
2664 14     $n_{k_t}(t) \leftarrow n_{k_t}(t-1) + 1$ 
2665 15     $t \leftarrow t + 1$ 

```

H.1 GUARANTEES FOR MTR-UCB

Combining the results from the previous two subsections and the formula for Quality and Cost regrets the upper bound for Algorithm 1, is stated in Theorem H.1.

Theorem H.1 (Overall Regret Upper Bound for MTR-UCB).

$$\text{Quality_Reg}(n, \phi, \mu_0) = \sum_{i \in C^-} \mathbb{E}[T_i(n)] \Delta_i^\mu + \sum_{j \in C^+} \mathbb{E}[T_j(n)] \Delta_j^\mu \quad (266)$$

$$\leq \sum_{i \in C^-} \frac{8 \log n}{\Delta_i^\mu} + \left(1 + \frac{\pi^2}{3}\right) \sum_{j \in [K], j \neq i^*} \Delta_j^\mu \quad (267)$$

$$\text{Cost_Reg}(n, \phi, \mu_0) = \sum_{j \in C^+} \mathbb{E}[T_j(n)] \Delta_j^c \quad (268)$$

$$\leq \left(1 + \frac{\pi^2}{3}\right) \sum_{j \in C^+} \Delta_j^c \quad (269)$$

H.1.1 ANALYSIS FOR MTR-UCB

Under this setting, the quality gaps are defined as the 0-clipped version of the gap between the known minimum tolerated reward $\mu_{\text{CS}} = \mu_0$ and the expected return of an arm $j : \Delta_j^q = \max\{0, \mu_{\text{CS}} - \mu_j\}$.

For a policy π , we have, Cumulative Regret,

$$\text{Quality_Reg}(T, \phi, \mu_0) = \mathbb{E}_\pi \left[\sum_{t=1}^T \max(\mu_0 - \mu_{\pi_t}, 0) \right] \quad (270)$$

$$\text{Cost_Reg}(T, \phi, \mu_0) = \mathbb{E}_\pi \left[\sum_{t=1}^T \max(c_{\pi_t} - c_{i^*}, 0) \right] \quad (271)$$

$$(272)$$

2700 For a policy π , we have, Incremental Regret,

$$2701 \text{Quality_Reg}_\pi^{\text{incr}}(t, \phi, \mu_0) = \mathbb{E}_\pi [\max \{\mu_0 - \mu_{k_t}, 0\}] \quad (273)$$

$$2702 \text{Cost_Reg}_\pi^{\text{incr}}(t, \phi, \mu_0) = \mathbb{E}_\pi [\max \{c_{k_t} - c_{i^*}, 0\}] \quad (274)$$

2703 Also define gaps,

$$2704 \Delta_i^\mu = \max \{\mu_0 - \mu_i, 0\} \quad (275)$$

$$2705 \Delta_i^c = \max \{c_i - c_{i^*}, 0\} \quad (276)$$

2706 Note from 10/25: A key takeaway from the advising meeting today was that the uncertainty of
2707 the decision making of a policy can be in the construction of the filtered set of arms or in the final
2708 decision taken about arm k_t (either or, or both).

2709 Moreover, under the sorted costs structure, for a Bandit Instance $\phi(\boldsymbol{\mu}, \mathbf{c})$, we have,

$$2710 \text{Quality_Reg}(n, \phi, \mu_0) = \sum_{i=1}^{i^*-1} \mathbb{E}[T_i(n)] \Delta_i^\mu + \sum_{j=i^*+1}^K \mathbb{E}[T_j(n)] \Delta_j^\mu \quad (277)$$

$$2711 \text{Cost_Reg}(n, \phi, \mu_0) = \sum_{j=i^*+1}^K \mathbb{E}[T_j(n)] \Delta_j^c \quad (278)$$

2712 To bound quality regret,

- 2713 1. $n_i(T)$ and $T_i(n)$ are used interchangeably to represent the number of pulls of arm i up to a
2714 horizon
- 2715 2. Bound $\Pr(k \in [i^* - 1] \text{ entering } C_t)$. This is enough since the contribution from the
2716 beyond i^* arms will be a constant term as we shall see soon.
- 2717 3. $\Pr\left(\hat{\mu}_k(t) + \sqrt{\frac{2 \log t}{n_k(t)}} \geq \mu_0\right)$
- 2718 4. Where, $\mu_0 > \mu_i$, $\mu_0 - \mu_i = \Delta_i^\mu$

$$2719 \Pr\left(\left\{\hat{\mu}_k(t) + \sqrt{\frac{2 \log t}{n_k(t)}} \geq \mu_0\right\}\right) \quad (279)$$

$$2720 \Pr\left(\left\{\hat{\mu}_k(t) + \sqrt{\frac{2 \log t}{n_k(t)}} \geq \mu_i + \Delta_i^\mu\right\}\right) \quad (280)$$

$$2721 \Pr\left(\left\{\hat{\mu}_k(t) - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{n_k(t)}}\right\}\right) \quad (281)$$

2722 To deal with the R.V. $\sqrt{\frac{2 \log t}{n_k(t)}}(n_k(t))$ in denominator, we do the sum/max/min trick in the original
2723 UCB proof in the Auer et al. paper.

2724 Goal: Bound the probability of the event,

$$2725 \hat{\mu}_k(t) + \sqrt{\frac{2 \log t}{n_k(t)}} \geq \mu_0, \quad (282)$$

2726 using the Auer et al. paper tricks.

Let $i \in [i^* - 1]$ be any arm with $\Delta_i^\mu > 0$, then,

$$T_i(n) = 1 + \sum_{t=K+1}^n \mathbb{I}\{k_t = i\} \quad (283)$$

$$\leq \ell + \sum_{t=K+1}^n \mathbb{I}\{k_t = i, n_i(t-1) \geq \ell\} \quad (284)$$

$$\leq \ell + \sum_{t=K+1}^n \mathbb{I}\left\{\hat{\mu}_i(t-1) + \sqrt{\frac{2 \log(t-1)}{n_k(t-1)}} \geq \mu_0, n_i(t-1) \geq \ell\right\} \quad (285)$$

At this point we do not have precise knowledge of the number of pulls $n_i(t-1)$, so we just sum over all the possibilities and modify the indexing $t' = t-1$. Moreover, we update the notation $\hat{\mu}_i(s)$ to represent the mean return for arm i over s samples.

$$T_i(n) \leq \ell + \sum_{t=1}^{\infty} \sum_{s=\ell}^{t-1} \mathbb{I}\left\{\hat{\mu}_i(s) + \sqrt{\frac{2 \log t}{s}} \geq \mu_0\right\} \quad (286)$$

$$\hat{\mu}_i + \sqrt{\frac{2 \log t}{s}} \geq \mu_0 \iff \hat{\mu}_i(t) - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{s}}.$$

Applying Lemma D.2 to equation 286 we get,

$$\Pr\left(\hat{\mu}_i(s) - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{s}}\right) \leq \exp\left(-2s \left(\Delta_i^\mu - \sqrt{\frac{2 \log t}{s}}\right)^2\right) \quad (287)$$

$$\leq \exp\left(-2s \left((\Delta_i^\mu)^2 + \frac{2 \log t}{s} - 2\Delta_i^\mu \sqrt{\frac{2 \log t}{s}}\right)\right) \quad (288)$$

$$\leq \exp\left(-2s \left(-4 \log t + 4\Delta_i^\mu \sqrt{2s \log t}\right)\right) \quad (289)$$

$$= \frac{\exp(4\Delta_i^\mu \sqrt{2s \log t})}{t^4} \quad (290)$$

- We want to make the thing inside $\exp(-2s(\cdot))$ smaller for upper bound.

- Make $\sqrt{\frac{2 \log t}{s}}$ term largest \implies plug in $s = \left\lceil \frac{8 \log n}{(\Delta_i^\mu)^2} \right\rceil$

Refine the steps from Friday 10/27:

- Bound on $\mathbb{E}[T_i(n)]$ for $i < i^*$
- Bound on $\mathbb{E}[T_j(n)]$ for $j > i^*$

H.1.2 BOUND ON (A)

For $i < i^*$, we have $\Delta_i^\mu > 0$.

$$T_i(n) = 1 + \sum_{K+1}^n \mathbb{I}\{k_t = i\} \quad (291)$$

$$\leq \ell + \sum_{K+1}^n \mathbb{I}\{k_t = i, T_i(t-1) \geq \ell\} \quad (292)$$

$$\leq \ell + \sum_{K+1}^n \mathbb{I} \left\{ \hat{\mu}_i(t-1) + \sqrt{\frac{2 \log(t-1)}{T_i(t-1)}} \geq \mu_0, T_i(t-1) \geq \ell \right\}. \quad (293)$$

Here the number of pulls s at time t is implicit in the notation. In the next step when we sum over it, we make it explicit.

Since there is no precise knowledge of the number of pulls $T_i(n)$, we resort to summing over all possibilities and for convenience, we change the indexing of time to be $t' = t - 1$, and take the sum to ∞ .

$$\leq \sum_{t=1}^{\infty} \sum_{s=\ell}^{t-1} \mathbb{I} \left\{ \hat{\mu}_i(s) + \sqrt{\frac{2 \log t}{s}} \geq \mu_0 \right\}. \quad (294)$$

From previous notes, eventually, on applying Lemma D.2 we get,

$$\Pr \left(\hat{\mu}_i(s) - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{s}} \right) \quad (295)$$

$$\leq \exp \left(-2s \left((\Delta_i^\mu)^2 + \frac{2 \log t}{s} - 2\Delta_i^\mu \sqrt{\frac{2 \log t}{s}} \right) \right) \quad (296)$$

$$\quad (297)$$

Plug in $s = \frac{8 \log n}{(\Delta_i^\mu)^2}$ for the red s , and leave the remaining s untouched to work towards the bound.

Key steps from the analysis from the morning session are recapped for continuity,

$$T_i(n) = 1 + \sum_{K+1}^n \mathbb{I} \{k_t = i\} \quad (298)$$

$$\leq \ell + \sum_{t=K+1}^n \mathbb{I} \left\{ \hat{\mu}_i(t-1) + \sqrt{\frac{2 \log(t-1)}{T_i(t-1)}} \geq \mu_0, T_i(t-1) \geq \ell \right\} \quad (299)$$

$$\leq \ell + \sum_{t=1}^{n-1} \sum_{s_i=\ell}^{t-1} \mathbb{I} \left\{ \hat{\mu}_i(t) + \sqrt{\frac{2 \log t}{s_i}} \geq \mu_0 \right\} \quad (300)$$

Take expectations on both sides,

$$\mathbb{E}[T_i(n)] \leq \ell + \sum_{t=1}^{n-1} \sum_{s_i=\ell}^{t-1} \Pr \left(\hat{\mu}_i(t) - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{s_i}} \right) \quad (301)$$

Applying Hoeffding Inequality, we have,

$$\Pr \left(\hat{\mu}_i - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{s_i}} \right) \leq \exp \left(-2s_i \left(\Delta_i^\mu - \sqrt{\frac{2 \log t}{s_i}} \right)^2 \right) \quad (302)$$

$$\exp \left(-2s_i \left((\Delta_i^\mu)^2 + \frac{2 \log t}{s_i} - 2\Delta_i^\mu \sqrt{\frac{2 \log t}{s_i}} \right) \right) \quad (303)$$

Plugging into Expected number of pulls expression,

$$\mathbb{E}[T_i(n)] \leq \ell + \sum_{t=1}^{n-1} \sum_{s_i=\ell}^{t-1} \exp \left(-2s_i \left(\frac{2 \log t}{s_i} + (\Delta_i^\mu)^2 - 2\Delta_i^\mu \sqrt{\frac{2 \log t}{s_i}} \right) \right) \quad (304)$$

We pick $\ell = \left\lceil \frac{8 \log n}{(\Delta_i^\mu)^2} \right\rceil$ as usual. (305)

$$\leq \ell + \sum_{t=1}^{n-1} \sum_{s_i=\ell}^{t-1} \exp \left(-4 \log t + 2s_i (\Delta_i^\mu)^2 \left(\sqrt{\frac{\log t}{\log n}} - 1 \right) \right) \quad (306)$$

since $t < n$, we have, (307)

$$\mathbb{E} [T_i(n)] \leq \frac{8 \log n}{(\Delta_i^\mu)^2} + 1 + \sum_{t=1}^{n-1} \sum_{s_i=\ell}^{t-1} \exp(-4 \log t) \quad (308)$$

$$\leq \frac{8 \log n}{(\Delta_i^\mu)^2} + 1 + \sum_{t=1}^{\infty} \frac{1}{t^3} \quad (309)$$

$$\leq \frac{8 \log n}{(\Delta_i^\mu)^2} + 1 + \frac{\pi^2}{6} \quad (310)$$

H.1.3 BOUND FOR TERM (B)

Initial Attempt on 10/30/23 corrections added on 11/3/23.

Further notational corrections added on 11/15/23.

Bound $\mathbb{E} [T_j(n)]$ for arms $j > i^*$, define, $\delta_0 = \mu_{i^*} - \mu_0$, $\delta_0 \geq 0$, and note that $\Delta_{i^*}^\mu = 0$

$$T_j(n) \leq 1 + \sum_{t=K+1}^n \mathbb{I} \{k_t = j\} \quad (311)$$

Just as a note and an aside, these arms j can satisfy, $\mu_j - \mu_0 \geq 0$ or $\mu_j - \mu_0 < 0$. (312)

The former incurs no quality regret but the latter does. (313)

$$\leq 1 + \sum_{t=K+1}^n \mathbb{I} \{i^* \notin C_t\} \quad \text{otherwise arm } i^* \text{ would have been pulled and not arm } j \quad (314)$$

$$\leq 1 + \sum_{t=K+1}^n \mathbb{I} \left\{ \hat{\mu}_{i^*}(t-1) + \sqrt{\frac{2 \log(t-1)}{n_{i^*}(t-1)}} < \mu_0 \right\} \quad (315)$$

$$= 1 + \sum_{t=K+1}^n \mathbb{I} \left\{ \hat{\mu}_{i^*}(t-1) - \mu_{i^*} < -\delta_0 - \sqrt{\frac{2 \log(t-1)}{n_{i^*}(t-1)}} \right\} \quad (316)$$

Since no direct handle on s is available, we sum over all its possible values (317)

Also NB: We update the notation $\hat{\mu}_i(s)$ to represent (318)

the mean return for arm i over s samples. (319)

$$T_j(n) \leq 1 + \sum_{t=K+1}^n \sum_{s=1}^{t-1} \mathbb{I} \left\{ \hat{\mu}_{i^*}(s) - \mu_{i^*} < -\delta_0 - \sqrt{\frac{2 \log(t-1)}{s}} \right\} \quad (320)$$

Taking expectations, and perform a shift of 1 in time-indexing, we have, (321)

$$\mathbb{E} [T_j(n)] \leq 1 + \sum_{t=1}^{n-1} \sum_{s=1}^{t-1} \Pr \left(\hat{\mu}_{i^*}(s) - \mu_{i^*} < -\delta_0 - \sqrt{\frac{2 \log(t-1)}{s}} \right) \quad (322)$$

$$\leq 1 + \sum_{t=1}^{\infty} \sum_{s=1}^{t-1} \exp \left(-2s \frac{2 \log t}{s} \right) \quad (323)$$

$$= 1 + \sum_{t=1}^{\infty} \sum_{s=1}^{t-1} \frac{1}{t^4} \quad (324)$$

$$\leq 1 + \frac{\pi^2}{6} \quad (325)$$

H.1.4 FINAL UPPER BOUND ANALYSIS

MAB instance: $\phi(\boldsymbol{\mu}, \mathbf{c})$

For this instance, without loss of generality, we assume,

$$\mathbf{c} = [c_1, c_2, \dots, c_K] \quad (326)$$

$$\boldsymbol{\mu} = [\mu_1, \mu_2, \dots, \mu_K], \text{ where,} \quad (327)$$

$$c_1 \leq c_2 \leq c_3 \leq \dots \leq c_{i^*-1} \leq c_{i^*} \leq c_{i^*+1} \leq \dots \leq c_K \quad (328)$$

In the MAB-CS with minimum tolerated reward setting, the optimal arm i^* is the least cost arm that has expected reward more than μ_0 . We also define C^* to be set of feasible arms per the following,

$$C^* = \{k : \mu_k \geq \mu_0, k \in [K]\} \quad (329)$$

$$i^* = \arg \min_{i \in C^*} c_i. \quad (330)$$

We assume C^* is non-empty, i.e. at least one arm has expected return higher than the minimum-tolerated-reward μ_0 .

Further, for the analysis, define,

$$C^- = \{k : c_k < c_{i^*}, k \in [K]\} \quad (331)$$

$$C^+ = \{k : c_k \geq c_{i^*}, k \in [K]\} \quad (332)$$

$$\text{Quality_Reg}(n, \phi, \mu_0) = \sum_{i \in C^-} \mathbb{E}[T_i(n)] \Delta_i^\mu + \sum_{j \in C^+} \mathbb{E}[T_j(n)] \Delta_j^\mu \quad (333)$$

$$\text{Cost_Reg}(n, \phi, \mu_0) = \sum_{j \in C^+} \mathbb{E}[T_j(n)] \Delta_j^c \quad (334)$$

H.1.5 BOUND ON THE EXPECTED NUMBER OF PULLS OF HIGH COST ARMS

- An arm $j \in C^+$ contributes to Quality Regret if $\mu_j < \mu_0$, otherwise $\Delta_j^\mu = 0$ and there is no contribution
- An arm $j \in C^+$ always contributes to Cost Regret

$$T_j(n) \leq 1 + \sum_{t=K+1}^n \mathbb{I}\{k_t = j\} \quad (335)$$

$$\leq 1 + \sum_{t=K+1}^n \mathbb{I}\{k_t = j, C_t \neq \phi\} + \sum_{t=K+1}^n \mathbb{I}\{k_t = j, C_t = \phi\} \quad (336)$$

$$\leq 1 + \sum_{t=K+1}^n \mathbb{I}\{i^* \notin C_t\} + \sum_{t=K+1}^n \mathbb{I}\{i^* \in C_t\} \quad (337)$$

$$\leq 1 + 2 \sum_{t=K+1}^n \mathbb{I}\left\{ \hat{\mu}_{i^*}(T_{i^*}(t-1)) + \sqrt{\frac{2 \log(t-1)}{T_{i^*}(t-1)}} < \mu_0 \right\} \quad (338)$$

$$= 1 + 2 \sum_{t=K+1}^n \mathbb{I}\left\{ \hat{\mu}_{i^*}(T_{i^*}(t-1)) - \mu_{i^*} < \mu_0 - \mu_{i^*} - \sqrt{\frac{2 \log(t-1)}{T_{i^*}(t-1)}} \right\} \quad (339)$$

$$\leq 1 + 2 \sum_{t=K+1}^n \sum_{s=1}^{t-1} \mathbb{I} \left\{ \hat{\mu}_{i^*}(s) - \mu_{i^*} < \mu_0 - \mu_{i^*} - \sqrt{\frac{2 \log(t-1)}{s}} \right\} \quad (340)$$

Taking expectations, and perform a shift of 1 in time-indexing, we have,

$$\mathbb{E}[T_j(n)] \leq 1 + 2 \sum_{t=1}^{\infty} \sum_{s=1}^t \Pr \left(\hat{\mu}_{i^*}(s) - \mu_{i^*} < \mu_0 - \mu_{i^*} - \sqrt{\frac{2 \log t}{s}} \right) \quad (341)$$

$$\leq 1 + 2 \sum_{t=1}^{\infty} \sum_{s=1}^t \exp \left(-2s \frac{2 \log t}{s} \right) \quad (342)$$

$$= 1 + 2 \sum_{t=1}^{\infty} \sum_{s=1}^t \frac{1}{t^4} \quad (343)$$

$$\leq 1 + \frac{\pi^2}{3} \quad (344)$$

Therefore, the expected number of pulls of a high-cost arm under MTR-UCB is given by,

$$\mathbb{E}[T_j(n)] \leq 1 + \frac{\pi^2}{3} \quad (345)$$

H.1.6 BOUND ON THE PULLS OF LOW-COST INFEASIBLE ARMS

- An arm $i \in C^-$ always contributes to Quality Regret and does not contribute to cost regret

$$T_i(n) = 1 + \sum_{t=K+1}^n \mathbb{I} \{k_t = i, C_t \neq \phi\} + \sum_{t=K+1}^n \mathbb{I} \{k_t = i, C_t = \phi\} \quad (346)$$

$$= 1 + \sum_{t=K+1}^n \mathbb{I} \{k_t = i, C_t \neq \phi\} + \sum_{t=K+1}^n \mathbb{I} \{i^* \notin C_t\} \quad (347)$$

$$= \ell + \sum_{t=K+1}^n \mathbb{I} \{k_t = i, C_t \neq \phi, T_i(t-1) \geq \ell\} + \sum_{t=K+1}^n \mathbb{I} \{i^* \notin C_t\} \quad (348)$$

$$= \ell + \sum_{t=K+1}^n \mathbb{I} \left\{ \hat{\mu}_i(T_i(t-1)) + \sqrt{\frac{2 \log(t-1)}{T_i(t-1)}} \geq \mu_0, T_i(t-1) \geq \ell \right\} + \sum_{t=K+1}^n \mathbb{I} \{i^* \notin C_t\} \quad (349)$$

$$\leq \ell + \sum_{t=1}^{\infty} \sum_{s=\ell}^t \mathbb{I} \left\{ \hat{\mu}_i(s) + \sqrt{\frac{2 \log t}{s}} \geq \mu_0 \right\} + \sum_{t=1}^{\infty} \mathbb{I} \{i^* \notin C_t\} \quad (350)$$

Taking expectations on both sides, we have,

$$\mathbb{E}[T_i(n)] \leq \ell + \sum_{t=1}^{\infty} \sum_{s_i=\ell}^t \Pr \left(\hat{\mu}_i(s_i) - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{s_i}} \right) + \sum_{t=1}^{\infty} \Pr(i^* \notin C_t) \quad (351)$$

Applying Hoeffding Inequality, we have,

$$\Pr \left(\hat{\mu}_i - \mu_i \geq \Delta_i^\mu - \sqrt{\frac{2 \log t}{s_i}} \right) \leq \exp \left(-2s_i \left(\Delta_i^\mu - \sqrt{\frac{2 \log t}{s_i}} \right)^2 \right) \quad (352)$$

$$\leq \exp \left(-2s_i \left((\Delta_i^\mu)^2 + \frac{2 \log t}{s_i} - 2\Delta_i^\mu \sqrt{\frac{2 \log t}{s_i}} \right) \right) \quad (353)$$

Plugging in the bound in 353 and recognizing that we have already bounded the term $\sum_{t=1}^{\infty} \Pr(i^* \notin C_t)$ in the previous subsection when we bounded the number of pulls of a high-cost arm, we have,

$$\mathbb{E}[T_i(n)] \leq \ell + \sum_{t=1}^{\infty} \sum_{s_i=\ell}^t \exp\left(-2s_i \left(\frac{2 \log t}{s_i} + (\Delta_i^\mu)^2 - 2\Delta_i^\mu \sqrt{\frac{2 \log t}{s_i}}\right)\right) + \frac{\pi^2}{6} \quad (354)$$

We pick $\ell = \left\lceil \frac{8 \log n}{(\Delta_i^\mu)^2} \right\rceil$ as in the proof technique in Auer et al. (2002)

$$\mathbb{E}[T_i(n)] \leq \ell + \sum_{t=1}^{\infty} \sum_{s_i=\ell}^t \exp\left(-4 \log t + 2s_i(\Delta_i^\mu)^2 \left(\sqrt{\frac{\log t}{\log n}} - 1\right)\right) + \frac{\pi^2}{6} \quad (355)$$

since $t < n$, we have,

$$\mathbb{E}[T_i(n)] \leq \frac{8 \log n}{(\Delta_i^\mu)^2} + 1 + \sum_{t=1}^{n-1} \sum_{s_i=\ell}^{t-1} \exp(-4 \log t) + \frac{\pi^2}{6} \quad (357)$$

$$\leq \frac{8 \log n}{(\Delta_i^\mu)^2} + 1 + \sum_{t=1}^{\infty} \frac{1}{t^3} + \frac{\pi^2}{6} \quad (358)$$

$$\leq \frac{8 \log n}{(\Delta_i^\mu)^2} + 1 + \frac{\pi^2}{3} \quad (359)$$

Therefore, the expected number of pulls of a low-cost (necessarily infeasible) arm under MTR-UCB is given by,

$$\mathbb{E}[T_i(n)] \leq \frac{8 \log n}{(\Delta_i^\mu)^2} + 1 + \frac{\pi^2}{3} \quad (360)$$