
Fair Exploration via Axiomatic Bargaining

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Motivated by the consideration of fairly sharing the cost of exploration between
2 multiple groups in learning problems, we develop the Nash bargaining solution in
3 the context of multi-armed bandits. Specifically, the ‘grouped’ bandit associated
4 with any multi-armed bandit problem associates, with each time step, a single group
5 from some finite set of groups. The utility gained by a given group under some
6 learning policy is naturally viewed as the reduction in that group’s regret relative to
7 the regret that group would have incurred ‘on its own’. We derive policies that yield
8 the Nash bargaining solution relative to the set of incremental utilities possible
9 under any policy. We show that on the one hand, the ‘price of fairness’ under such
10 policies is limited, while on the other hand, regret optimal policies are arbitrarily
11 unfair under generic conditions. Our theoretical development is complemented by
12 a case study on contextual bandits for warfarin dosing where we are concerned
13 with the cost of exploration across multiple races and age groups.

14 1 Introduction

15 Exploration in learning problems has an implicit cost, inasmuch that exploring actions that are
16 eventually revealed to be sub-optimal incurs regret. We study how this cost of exploration is shared
17 in a system with multiple stakeholders. At the outset, we present two motivating examples.

18 **Personalized Medicine and Adaptive Trials:** Multi-stage, adaptive designs [1, 2, 3, 4], are widely
19 viewed as a frontier in clinical trials. More generally, the ability to collect detailed patient level
20 data, and real time monitoring (eg. glucose monitoring for diabetes [5, 6]) has raised the specter
21 of learning personalized treatments. Among other formulations, such problems may be viewed as
22 contextual bandits. For instance, for the problem of optimal warfarin dosing [7], the context at each
23 time step corresponds to a patient’s covariates, arms correspond to different dosages of warfarin, and
24 the reward is the observed efficacy of the assigned dose. In examining such a study in retrospect, it is
25 natural to measure the regret incurred by distinct groups of patients (eg. by race or age). What makes
26 a profile of regret across such groups fair or unfair?

27 **Revenue Management for Search Advertising:** Ad platforms enjoy a tremendous amount of
28 flexibility in the the choice of ads served against search queries. Specifically, this flexibility exists
29 both in selecting a slate of advertisers to compete for a specific search, and then in picking a winner
30 from this slate. Now a key goal for the platform is learning the affinity of any given ad for a given
31 search. In solving such a learning problem – for which many variants have been proposed [8, 9] – we
32 may again ask the question of who bears the cost of exploration, and whether the profile of such costs
33 across various groups of advertisers is fair.

34 1.1 Bandits, Groups and Axiomatic Bargaining

35 Delaying a formal development to later, any bandit problem has an associated ‘grouped’ variant.
36 Specifically, we are given a finite set of groups (eg. races or age groups in the warfarin example), and

37 each group is associated with an arrival probability and a distribution over action sets. At each time
38 step, a group and an action set is drawn from this distribution from which the learning algorithm must
39 pick an action. Heterogeneity in groups is thus driven by differences in their respective distributions
40 over feasible action sets. In addition to measuring overall regret, we also care about the regret incurred
41 by specific groups, which we can view as the cost of exploration borne by that group.

42 In reasoning about ‘fair’ regret profiles we turn to the theory of axiomatic bargaining. There, a central
43 decision maker is concerned with the incremental utility earned by each group from collaborating,
44 relative to the utility the group would earn on its own. Here this incremental utility is precisely the
45 reduction in regret for any given group relative to the optimal regret that group would have incurred
46 ‘on its own’. A *bargaining solution* maximizes some objective function over the set of achievable
47 incremental utilities. The *utilitarian solution*, for instance, maximizes the sum of incremental utilities
48 which would reduce here to the usual objective of minimizing total regret. The *Nash bargaining*
49 *solution* maximizes an alternative objective, the Nash Social Welfare (SW) function. This latter
50 solution is the unique solution to satisfy a set of axioms any ‘fair’ solution would reasonably satisfy.
51 *This paper develops the Nash bargaining solution to the (grouped) bandit problem.*

52 1.2 Contributions

53 In developing the Nash bargaining solution, we focus primarily on what is arguably the simplest
54 non-trivial grouped bandit setting. Specifically, we consider the ‘grouped’ K -armed bandit model,
55 wherein each group corresponds to a subset of the K arms. We make the following contributions
56 relative to this problem:

57 *Regret Optimal Policies are Unfair (Theorem 3.1):* We show that all regret optimal policies for the
58 grouped K -armed bandit share a structural property that make them ‘arbitrarily unfair’ – in the sense
59 that the Nash SW is $-\infty$ for these solutions – under a broad set of conditions on the problem instance.

60 *Achievable Fairness (Theorem 3.2):* We derive an instance-dependent upper bound on the Nash SW
61 for the grouped K -armed bandit. This can be viewed as a ‘fair’ analogue to a regret lower bound
62 (e.g. [10]) for the problem, since a lower bound on achievable regret (forgoing any fairness concerns)
63 would in effect correspond to an upper bound on the utilitarian SW for the problem.

64 *Nash Solution (Theorem 4.1):* We produce a policy that achieves the Nash solution. Specifically, the
65 Nash SW under this policy achieves the upper bound we derive on the Nash SW for all instances of
66 the grouped K -armed bandit.

67 *Price of Fairness for the Nash Solution (Theorem 4.2):* We show that the ‘price of fairness’ for
68 the Nash solution is small: the Nash solution achieves at least $O(1/\sqrt{G})$ of the reduction in regret
69 achieved under a regret optimal solution relative to the regret incurred when groups operate separately.

70 Taken together, these results establish a rigorous framework for the design of bandit algorithms that
71 yield fair outcomes across groups at a low cost to total regret. As a final contribution, we extend our
72 framework beyond the grouped K -armed bandit and undertake an empirical study:

73 *Linear Contextual Bandits and Warfarin Dosing:* We extend our framework to grouped linear
74 contextual bandits, yielding a candidate Nash solution there. Applied to a real-world dataset on
75 warfarin dosing using race and age groups, we show (a) a regret optimal solution that ignores groups
76 is dramatically unfair, and (b) the Nash solution balances out reductions in regret across groups at the
77 cost of a small increase in total regret.

78 1.3 Related Literature

79 Two pieces of prior work have a motivation similar to our own. [11] studies a setting with multiple
80 agents with a common bandit problem, where each agent can decide which action to take at each time.
81 They show that ‘free-riding’ is possible — an agent that can access information from other agents
82 can incur only $O(1)$ regret in several classes of problems. This is consistent with our motivation.
83 [12] studies a very similar grouped bandit model to ours, and provides a ‘counterexample’ in which a
84 group can have a negative externality on another group. This example is somewhat pathological and
85 stems from considering an instance-specific fixed time horizon; instead, if $T \rightarrow \infty$, all externalities
86 become non-negative (details in Appendix A.1). Our grouped bandit model is also similar to *sleeping*
87 *bandits* [13], in which the set of available arms is adversarially chosen in each round. The known,
88 fixed group structure in our model allows us to achieve tighter regret bounds than [13].

89 There have also been a handful of papers [14, 15, 16, 17] that study ‘fairness in bandits’ in a
 90 completely different context. These works enforce a fairness criterion between *arms*, which is
 91 relevant in settings where a ‘pull’ represents some resource that is allocated to that arm, and these
 92 pulls should be distributed between arms in a fair manner. In these models, the decision maker’s
 93 objective (maximize reward) is distinct from that of a group (obtain ‘pulls’), unlike our setting (and
 94 motivating examples) where the groups and decision maker are aligned in their eventual objective.

95 Our upper bound on Nash SW borrows classic techniques from the regret lower bound results of [10]
 96 and [18]. Our policy follows a similar pattern to recent work on regret-optimal, optimization-based
 97 policies for structured bandits [19, 20, 21, 22]. Unlike those policies, our policy has no forced
 98 exploration. Further the optimization problem defining the Nash solution can generically have
 99 multiple solutions whereas the aforementioned approaches would require this solution to be unique;
 100 our approach does not require a unique solution. Nonetheless, we believe that the framework in the
 101 aforementioned works can be fruitfully leveraged to construct Nash solutions for general grouped
 102 bandits, and we provide such a candidate solution as an extension.

103 Our fairness framework is inspired by the literature on fairness in welfare economics — see [23, 24].
 104 Specifically, we study fairness in exploration through the lens of the axiomatic bargaining framework,
 105 first studied by [25], who showed that enforcing four desirable axioms induces a unique fair solution.
 106 [26] is an excellent textbook reference for this topic.

107 2 The Axiomatic Bargaining Framework for Bandits

108 Let $\theta \in \Theta$ be an unknown parameter and let \mathcal{A} be the action set. For every arm $a \in \mathcal{A}$, $(Y_n(a))_{n \geq 1}$
 109 is an i.i.d. sequence of rewards drawn from a distribution $F(\theta, a)$ parameterized by θ and a . We let
 110 $\mu(a) = \mathbb{E}[Y_1(a)]$ be the expected reward of arm a . In defining a *grouped* bandit problem, we let \mathcal{G}
 111 be a set of G groups. Each group $g \in \mathcal{G}$ is associated with a probability distribution P^g over $2^{\mathcal{A}}$,
 112 and a probability of arrival p_g ; $\sum_g p_g = 1$. The identity of the group arriving at time t , g_t , is chosen
 113 independently according to this latter distribution; \mathcal{A}_t is then drawn according to P^{g_t} . An instance of
 114 the grouped bandit problem is specified by $\mathcal{I} = (\mathcal{A}, \mathcal{G}, p, P, F, \theta)$, where all quantities except for θ
 115 are known. At each time t , a central decision maker observes g_t and \mathcal{A}_t , chooses an arm $A_t \in \mathcal{A}_t$
 116 to pull and observes the reward $Y_{N_t(A_t)+1}(A_t)$, where $N_t(a)$ is the total number of times arm a was
 117 pulled up to but not including time t . Let $A_t^* \in \operatorname{argmax}_{a \in \mathcal{A}_t} \mu(a)$ be an optimal arm at time t . Given
 118 an instance \mathcal{I} and a policy π , the *total regret*, and the *group regret* for group $g \in \mathcal{G}$ are respectively

$$R_T(\pi, \mathcal{I}) = \mathbb{E} \left[\sum_{t=1}^T (\mu(A_t^*) - \mu(A_t)) \right] \text{ and } R_T^g(\pi, \mathcal{I}) = \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(g_t = g) (\mu(A_t^*) - \mu(A_t)) \right],$$

119 where the expectation is over randomness in arrivals (g_t, \mathcal{A}_t) , rewards $Y_n(a)$, and the policy π .
 120 Finally, so that the notion of an optimal policy for some class of instances, \mathcal{I} , is well defined, we
 121 restrict attention to *consistent* policies which yield sub-polynomial regret for any instance in that
 122 class: $\Psi = \{\pi : R_T(\pi, \mathcal{I}) = o(T^b) \forall \mathcal{I} \in \mathcal{I}, \forall b > 0\}$.

123 2.1 Background: Axiomatic Bargaining

124 The axiomatic bargaining problem is specified by the number of agents n , a set of feasible utility
 125 profiles $U \subseteq \mathbb{R}^n$, and a disagreement point $d \in \mathbb{R}^n$, that represents the utility profile when agents
 126 cannot come to an agreement. A solution $f(\cdot, \cdot)$ to the bargaining problem selects an agreement
 127 $u^* = f(U, d) \in U$, in which agent i receives utility u_i^* . It is assumed that there is at least one point
 128 $u \in U$ such that $u > d$, and we assume U is compact and convex.

129 The bargaining framework proposes a set of axioms a fair solution u^* should ideally satisfy:

- 130 (a) *Pareto Optimality*: There is no $u \in U$ with $u \geq u^*$, $u \neq u^*$.
- 131 (b) *Invariance to Affine Transformations*: If $U' = \{a^\top u + b : u \in U\}$ and $d' = a^\top d + b$, then
 132 $f(U', d')_i = a_i u_i^* + b_i$ for any $a \in \mathbb{R}_+^n$, $b \in \mathbb{R}^n$.
- 133 (c) *Independence of Irrelevant Alternatives*: If $V \subseteq U$ where $u^* \in V$, then $f(V, d) = u^*$.
- 134 (d) *Symmetry*: If U and d are symmetric, $u_i^* = u_j^* \forall i, j$.

135 Now (b) implies that $f(U, d) = f(\{u - d : u \in U\}, 0) + d$. It is therefore customary to normalize
 136 the origin to the disagreement point, i.e. assume $d = 0$, and implicitly that U has been appropriately
 137 translated. So translated, U is interpreted as a set of feasible utility *gains* relative to the disagreement

138 point. The seminal work of [25] showed that there is a unique bargaining solution that satisfies the
 139 above four axioms, and it is the outcome that maximizes the *Nash social welfare (SW) function* [27]:

$$W(u) = \begin{cases} \sum_{i=1}^n \log(u_i) & u_i > 0 \forall i \in [n] \\ -\infty & \text{otherwise.} \end{cases}$$

140 We will interchangeably refer to $u^* = \operatorname{argmax}_{u \in U} W(u)$ as the *Nash solution* or as *proportionally*
 141 *fair*. If $u \in U$ such that $W(u) = -\infty$, we say that u is *unfair*.

142 2.2 Fairness Framework for Grouped Bandits

143 We now consider the Nash bargaining solution in the context of the grouped bandit problem. To
 144 do so, we need to appropriately define the utility gain under any policy. We begin by formalizing
 145 the rewards to a single group under a policy where no information was shared across groups, which
 146 represents the disagreement point. Specifically, let \mathcal{I}_g be the ‘single-group’ bandit instance obtained
 147 by considering the instance \mathcal{I} restricted to arrivals of group g so that in any period t in which $g_t \neq g$,
 148 we receive no reward under any action. Let us denote by π_g^* an optimal policy for instances of type
 149 \mathcal{I}_g (i.e. π_g^* is optimal in the non-grouped bandit setting) so that for any instance of type \mathcal{I}_g , and any
 150 other consistent policy π'_g for instances of that type,

$$(1) \quad \limsup_{T \rightarrow \infty} \frac{R_T(\pi_g^*, \mathcal{I}_g)}{\log T} \leq \liminf_{T \rightarrow \infty} \frac{R_T(\pi'_g, \mathcal{I}_g)}{\log T}.$$

151 Now letting $\tilde{R}_T^g(\mathcal{I}) \triangleq R_T(\pi_g^*, \mathcal{I}_g)$, we define, with a slight abuse of notation, the T -period utility
 152 earned by group g under π_g^* , and any other consistent policy π for instances of type \mathcal{I} respectively, as:

$$\mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(g_t = g) \mu(A_t^*) \right] - \tilde{R}_T^g(\mathcal{I}) \triangleq u_T^g(\pi_g^*) \text{ and } \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(g_t = g) \mu(A_t^*) \right] - R_T^g(\pi, \mathcal{I}) \triangleq u_T^g(\pi).$$

153 The T -period utility gain under a policy π is then $u_T^g(\pi) - u_T^g(\pi_g^*) = \tilde{R}_T^g(\mathcal{I}) - R_T^g(\pi, \mathcal{I})$. Since our
 154 goal is to understand long-run system behavior, we define asymptotic utility gain for any group g :

$$\text{UtilGain}^g(\pi, \mathcal{I}) = \liminf_{T \rightarrow \infty} \frac{\tilde{R}_T^g(\mathcal{I}) - R_T^g(\pi, \mathcal{I})}{\log T}.$$

155 Equipped with this definition, we may now identify the set of incremental utilities for an instance \mathcal{I} ,
 156 as $U(\mathcal{I}) = \{(\text{UtilGain}^g(\pi, \mathcal{I}))_{g \in \mathcal{G}} : \pi \in \Psi\}$. We can readily show that the Nash solution remains
 157 the unique solution satisfying the fairness axioms presented in Section 2.1 relative to $U(\mathcal{I})$. We finish
 158 up by finally defining the Nash solution to the grouped bandit problem. Since we find it convenient to
 159 associate a SW function with a policy (as opposed to a vector of incremental utilities), the Nash SW
 160 function for grouped bandits is equivalently defined as:

$$(2) \quad W(\pi, \mathcal{I}) = \begin{cases} \sum_{g \in \mathcal{G}} \log(\text{UtilGain}^g(\pi, \mathcal{I})) & \text{UtilGain}^g(\pi, \mathcal{I}) > 0 \forall g \in \mathcal{G} \\ -\infty & \text{otherwise.} \end{cases}$$

161 So equipped, we finish by defining the Nash solution to the grouped bandit problem.

162 **Definition 2.1.** Suppose a policy π^* satisfies $W(\pi^*, \mathcal{I}) = \sup_{\pi \in \Psi} W(\pi, \mathcal{I})$ for every instance
 163 $\mathcal{I} \in \mathcal{I}$. Then, we say that π^* is the Nash solution for \mathcal{I} and that it is *proportionally fair*.

164 2.3 Grouped K -armed Bandit Model

165 The grouped K -armed bandit is arguably the simplest non-trivial class of grouped bandits. Let
 166 $\mathcal{A} = [K]$. Denote by $\mathcal{A}^g \subseteq \mathcal{A}$ a subset of arms corresponding to group g and by \mathcal{G}_a a subset of groups
 167 corresponding to arm a . For each g , P^g places unit mass on \mathcal{A}^g so that the set of arms available at
 168 time t is $\mathcal{A}_t = \mathcal{A}^{g_t}$. Assume $\theta \in (0, 1)^K$, and the single period reward $Y_1(a) \sim \text{Bernoulli}(\theta(a))$.
 169 We assume that $\theta(a) \neq \theta(a')$ for all $a \neq a'$. Since the set of arms available at each time step only
 170 depends on the arriving group, we denote by $\text{OPT}(g) = \max_{a \in \mathcal{A}^g} \theta(a)$ the optimal mean reward for
 171 group g . We take π_g^* to be the KL-UCB policy of [28] since KL-UCB is optimal (in the sense of (1))
 172 for vanilla K -armed bandits. We may write the T -period regret in this model as

$$(3) \quad R_T(\pi, \mathcal{I}) = \sum_{g \in \mathcal{G}} \sum_{a \in \mathcal{A}^g} \Delta^g(a) \mathbb{E}[N_T^g(a)],$$

173 where $N_T^g(a)$ is the number of times that group g has pulled arm a after T time steps, and
 174 $\Delta^g(a) = \text{OPT}(g) - \theta(a)$. Lastly, we state a condition guaranteeing $U(\mathcal{I})$ contains a point $u > 0$;
 175 Proposition G.1 in Appendix G proves the following assumption is necessary and sufficient:

176 **Assumption 2.2.** *Every group g has at least one suboptimal arm that is shared with another group.*
 177 *That is, for every g , $\exists a \in \mathcal{A}^g$ such that $\mu(a) < \text{OPT}(g)$ and $|\mathcal{G}_a| \geq 2$.*

178 3 Fairness-Regret Trade-off

179 In this section, we prove that a regret-optimal policy for a generic grouped K -armed bandit must
 180 necessarily be unfair. We then turn to deriving an upper bound on achievable Nash SW.

181 3.1 Unfairness of Regret Optimal Policies

182 We first state the main result, which states that regret optimal policies are arbitrarily unfair. In fact,
 183 we show that perversely the most ‘disadvantaged’ group (in a sense we make precise shortly) bears
 184 the brunt of exploration in that it sees no improvement in regret relative to if it were ‘on its own’.

185 **Theorem 3.1.** *Let π be a regret optimal policy. Let \mathcal{I} be an instance of the grouped K -armed bandit
 186 where $g_{\min} \triangleq \text{argmin}_{g \in G} \text{OPT}(g)$ is unique. Then, $W_{\mathcal{I}}(\pi) = -\infty$ and $\text{UtilGain}^{g_{\min}}(\pi, \mathcal{I}) = 0$.*

187 *Proof.* We define regret optimality by proving tight lower and upper bounds on regret, and these
 188 bounds imply necessary properties of all regret optimal policies that yield the desired result.

189 We first lower bound the total number of pulls, $\mathbb{E}[N_T(a)]$, of a suboptimal arm. Denote by $\mathcal{A}_{\text{sub}}^g =$
 190 $\{a \in \mathcal{A}^g : \theta(a) < \text{OPT}(g)\}$ the suboptimal actions for group g , and denote by $\mathcal{A}_{\text{sub}} = \{a \in \mathcal{A} :$
 191 $a \in \mathcal{A}_{\text{sub}}^g \forall g \in \mathcal{G}_a\}$ the set of arms that are not optimal for any group. Now since a consistent policy
 192 for the grouped K -armed bandit is automatically consistent for the vanilla K -armed bandit obtained
 193 by restricting to any of its component groups g , the standard lower bound of [10] implies that for any
 194 $a \in \mathcal{A}_{\text{sub}}^g$, $\liminf_{T \rightarrow \infty} \mathbb{E}[N_T(a)]/\log T(g) \geq J^g(a)$ where $J^g(a) \triangleq 1/\text{KL}(\theta(a), \text{OPT}(g))$ and
 195 $T(g)$ is the number of arrivals of group g up to and including time T . Since this must hold for any
 196 group, and since $\lim_T \log T / \log T(g) = 1$ a.s.,

$$(4) \quad \liminf_{T \rightarrow \infty} \frac{\mathbb{E}[N_T(a)]}{\log T} \geq J(a)$$

197 for all $a \in \mathcal{A}_{\text{sub}}$ where $J(a) = \max_{g \in \mathcal{G}_a} J^g(a)$. Now, denote by $\Gamma(a) = \text{argmin}_{g \in \mathcal{G}_a} \text{OPT}(g)$ the
 198 set of groups that have the smallest optimal reward out of all groups that have access to a . Then the
 199 smallest regret incurred in pulling arm a is simply $\Delta^g(a)$ for any $g \in \Gamma(a)$. With a slight abuse, we
 200 denote this quantity by $\Delta^{\Gamma(a)}(a)$. (4) immediately implies that for any consistent policy π ,

$$(5) \quad \liminf_{T \rightarrow \infty} \frac{R_T(\pi, \mathcal{I})}{\log T} \geq \sum_{a \in \mathcal{A}_{\text{sub}}} \Delta^{\Gamma(a)}(a) J(a).$$

201 In fact, we show that the KL-UCB policy [28] (surprisingly) achieves this lower bound; the proof of
 202 this claim is somewhat involved and can be found in Appendix C. Consequently, any regret optimal
 203 policy must achieve the limit infimum in (5). In turn, this implies that a policy $\pi \in \Psi$ is regret optimal
 204 if and only if, the number of pulls of arms $a \in \mathcal{A}_{\text{sub}}$ achieve the lower bound (4), i.e.

$$(6) \quad \lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_T(a)]}{\log T} = J(a) \quad \forall a \in \mathcal{A}_{\text{sub}}$$

205 and further that any pulls of arm a from a group $g \notin \Gamma(a)$ must be negligible, i.e.

$$(7) \quad \lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_T^g(a)]}{\log T} = 0 \quad \forall a \in \mathcal{A}, g \notin \Gamma(a).$$

206 Now, turning our attention to g_{\min} , we have by assumption that g_{\min} is the only group in
 207 $\Gamma(a)$ for all $a \in \mathcal{A}^{g_{\min}}$. Consequently, by (7), we must have that for any optimal policy,
 208 $\lim_{T \rightarrow \infty} \mathbb{E}[N_T^{g_{\min}}(a)]/\log T = \lim_{T \rightarrow \infty} \mathbb{E}[N_T(a)]/\log T$ for all $a \in \mathcal{A}^{g_{\min}}$. And since $J(a) =$
 209 $J^{g_{\min}}(a)$ for all $a \in \mathcal{A}^{g_{\min}} \cap \mathcal{A}_{\text{sub}}$, (6) then implies that the regret for group g_{\min} is precisely

$$\lim_{T \rightarrow \infty} \frac{R_T^{g_{\min}}(\pi, \mathcal{I})}{\log T} = \sum_{a \in \mathcal{A}_{\text{sub}}^{g_{\min}}} \Delta^{g_{\min}}(a) J^{g_{\min}}(a).$$

210 But this is precisely $\lim_T \tilde{R}_T^{g_{\min}}(\mathcal{I})/\log T$. Thus, $\text{UtilGain}^{g_{\min}}(\pi, \mathcal{I}) = 0$, and $W_{\mathcal{I}}(\pi) = -\infty$. \square

211 The proof also illustrates that if $g_{\max} \triangleq \text{argmax}_{g \in G} \text{OPT}(g)$ is unique, then g_{\max} incurs no regret
 212 from *any* shared arm in a regret optimal policy. If all suboptimal arms for g_{\max} are shared with another
 213 group, then g_{\max} incurs zero (log-scaled) regret in an optimal policy. In summary, regret optimal
 214 policies are unfair, and achieve perverse outcomes with the most disadvantaged groups gaining
 215 nothing and the most advantaged groups gaining the most from sharing the burden of exploration.

216 3.2 Upper Bound on Nash SW

217 The preceding question motivates asking what is in fact possible with respect to fair outcomes. To
 218 that end, we derive an instance-dependent upper bound on the Nash SW. We may view this as a ‘fair’
 219 analogue to instance-dependent lower bounds on regret.

220 Recall the definition of $W(\pi, \mathcal{I})$ in (2), and let $W^*(\mathcal{I}) = \sup_{\pi \in \Psi} W(\pi, \mathcal{I})$. Fix an instance \mathcal{I} with
 221 unknown parameter vector θ . We first upper bound $W(\pi, \mathcal{I})$. Recall that KL-UCB is the policy π_g^*
 222 used to define $\tilde{R}_T^g(\mathcal{I})$. The fact that KL-UCB is optimal in the vanilla K -armed bandit implies:

$$(8) \quad \lim_{T \rightarrow \infty} \frac{\tilde{R}_T^g(\mathcal{I})}{\log T} = \sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) J^g(a).$$

223 Next, we re-write $R_T^g(\pi, \mathcal{I})/\log T$. Given a policy π , for any action a and group g , let $q_T^g(a, \pi) \in$
 224 $[0, 1]$ be the *percentage* of times that group g pulls arm a , out of the total number of times arm a is
 225 pulled. That is, $\mathbb{E}[N_T^g(a)] = q_T^g(a, \pi) \mathbb{E}[N_T(a)]$, where $\sum_{g \in G} q_T^g(a, \pi) = 1$ for all a . Then,

$$(9) \quad \frac{R_T^g(\pi, \mathcal{I})}{\log T} = \sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) q_T^g(a, \pi) \frac{\mathbb{E}[N_T(a)]}{\log T} \geq \sum_{a \in \mathcal{A}_{\text{sub}}^g \cap \mathcal{A}_{\text{sub}}} \Delta^g(a) q_T^g(a, \pi) \frac{\mathbb{E}[N_T(a)]}{\log T}.$$

226 Recalling $\text{UtilGain}^g(\pi, \mathcal{I}) = \liminf_{T \rightarrow \infty} \frac{\tilde{R}_T^g(\mathcal{I}) - R_T^g(\pi, \mathcal{I})}{\log T}$, combining (8), (9), and (4) yields:

$$\text{UtilGain}^g(\pi, \mathcal{I}) \leq \liminf_{T \rightarrow \infty} \sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) (J^g(a) - q_T^g(a, \pi) J(a) \mathbf{1}\{a \in \mathcal{A}_{\text{sub}}\}).$$

227 Using the definition of $W(\pi, \mathcal{I})$ and taking the lim inf outside of the sum gives

$$W(\pi, \mathcal{I}) \leq \liminf_{T \rightarrow \infty} \sum_{g \in G} \log \left(\sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) (J^g(a) - q_T^g(a, \pi) J(a) \mathbf{1}\{a \in \mathcal{A}_{\text{sub}}\}) \right)^+.$$

228 But since $\sum_{g \in G} q_T^g(a, \pi) = 1$ for every T, a , it must be that the limit infimum above is achieved for
 229 some vector $(q^g(a))$ satisfying $\sum_{g \in G} q^g(a) = 1$ for all a . This immediately yields an upper bound
 230 on $W^*(\mathcal{I})$: Let $Y^*(\mathcal{I})$ be the optimal value to the program $P(\theta)$, and let q_* be an optimal solution.

$$(P(\theta)) \quad \begin{aligned} & \max_{q \geq 0} \sum_{g \in G} \log \left(\sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) (J^g(a) - q^g(a) J(a)) \right)^+ \\ & \text{s.t.} \quad \sum_{g \in G} q^g(a) = 1 \quad \forall a \in \mathcal{A}_{\text{sub}} \\ & \quad \quad q^g(a) = 0 \quad \forall g \in G, a \notin \mathcal{A}_{\text{sub}} \cap \mathcal{A}_g. \end{aligned}$$

231 Then, we have shown:

232 **Theorem 3.2.** *For every instance \mathcal{I} of the grouped K -armed bandit, $W^*(\mathcal{I}) \leq Y^*(\mathcal{I})$.*

233 4 Nash Solution for Grouped K -armed Bandits

234 We turn our attention in this section to constructive issues: we first develop an algorithm that achieves
 235 the Nash SW upper bound of Theorem 3.2 and thus establish that this is the Nash solution for the
 236 grouped K -armed bandit. In analogy to the unfairness of a regret optimal policy, it is then natural to
 237 ask whether the regret under this Nash solution is large relative to optimal regret; we show thankfully
 238 that this ‘price of fairness’ is relatively small.

239 **4.1 The Nash Solution: PF-UCB**

240 The algorithm we present here ‘Proportionally Fair’ UCB (or PF-UCB) works as follows: at each
 241 time step it computes the set of arms that optimize the (KL) UCB for some group. Then, when a
 242 group arrives, it asks whether any arm from this set has been ‘under-explored’ where the notion of
 243 under-exploration is measured relative to an estimated optimal solution to $P(\theta)$. Such an arm, if
 244 available, is pulled. Absent the availability of such an arm, a greedy selection is made.

245 Specifically, let $\hat{\theta}_t$ be the empirical mean estimate of θ at time t . $P(\hat{\theta}_t)$ is then our approximation to
 246 $P(\theta)$ at time t and we denote by \hat{q}_t the optimal solution to this program with smallest euclidean norm.
 247 Note that finding such a solution constitutes a tractable convex optimization problem. We define
 248 the standard KL-UCB for an arm, $\text{UCB}_t(a) = \max\{q : N_t(a)\text{KL}(\hat{\theta}_t(a), q) \leq \log t + 3 \log \log t\}$.
 249 Finally, we denote by $A_t^{\text{UCB}}(g) \in \arg\max_{a \in \mathcal{A}^g} \text{UCB}_t(a)$ the arm with the highest UCB for group g
 250 at time t , and by $\mathcal{A}_t^{\text{UCB}} = \{A_t^{\text{UCB}}(g) : g \in \mathcal{G}\}$ the set of arms that have the highest UCB for *some*
 251 group. PF-UCB then proceeds as follows. At time t :

- 252 1. If there is an available arm $a \in \mathcal{A}^{g_t} \cap \mathcal{A}_t^{\text{UCB}}$ such that $N_t^{g_t}(a) \leq \hat{q}_t^{g_t}(a)N_t(a)$, pull a . If
 253 there are multiple arms matching this criteria, pull one of them uniformly at random.
- 254 2. Otherwise, pull the greedy arm $A_t^{\text{greedy}}(g_t) \in \arg\max_{a \in \mathcal{A}^{g_t}} \hat{\theta}_t(a)$.

255 PF-UCB explores at time t by pulling an arm if it is the arm with the highest UCB for *some* group
 256 (not necessarily group g_t), *and* the current group g_t has not pulled it as many times as it should have
 257 according to the solution \hat{q}_t . PF-UCB constitutes a Nash solution for the grouped K -armed bandit.
 258 Specifically, we prove the following theorem in Appendix E:

259 **Theorem 4.1.** *For any instance \mathcal{I} of the grouped K -armed bandit, we have for all groups g ,*

$$\lim_{T \rightarrow \infty} \frac{R_T^g(\pi^{\text{PF-UCB}}, \mathcal{I})}{\log T} = \sum_{a \in \mathcal{A}^g} \Delta^g(a) q_*^g(a) J(a).$$

260 It is worth noting that relative to the existing optimization-based algorithms for structured bandits
 261 (e.g. [19, 20, 21, 22]), PF-UCB does no forced sampling. In addition, we make no requirement that
 262 the solution to the optimization problem $P(\theta)$ is unique as these existing policies require. In fact,
 263 optimal solutions to $P(\theta)$ are not unique, and the choice of a solution that has smallest euclidean
 264 norm is carefully shown to provide the necessary ‘stability’ while being computationally tractable.
 265 That said, the next section shows how we can fruitfully leverage an existing algorithm from [22] to
 266 construct a candidate Nash solution for a setting beyond the grouped K -armed bandit.

267 **4.2 Price of Fairness**

268 Whereas PF-UCB is proportionally fair, what price do we pay with respect to regret? To answer this
 269 question we compute in this section an upper bound on the ‘price of fairness’. Specifically, define

$$\text{SYSTEM}(\mathcal{I}) = \sum_{g \in \mathcal{G}} \text{UtilGain}^g(\pi^{\text{KL-UCB}}, \mathcal{I}) \text{ and } \text{FAIR}(\mathcal{I}) = \sum_{g \in \mathcal{G}} \text{UtilGain}^g(\pi^{\text{PF-UCB}}, \mathcal{I}).$$

270 $\text{UtilGain}^g(\pi^{\text{KL-UCB}}, \mathcal{I})$ is the reduction in group g ’s regret under a *regret optimal* policy in the
 271 grouped setting relative to the optimal regret it would have endured on its own; $\text{SYSTEM}(\mathcal{I})$
 272 aggregates this reduction in regret across all groups. Similarly, $\text{UtilGain}^g(\pi^{\text{PF-UCB}}, \mathcal{I})$ is the reduction
 273 in group g ’s regret under a *proportionally fair* policy, and $\text{FAIR}(\mathcal{I})$ aggregates this across groups.
 274 The price of fairness (PoF) asks what fraction of the optimal reduction in regret is lost to fairness:

$$\text{PoF}(\mathcal{I}) \triangleq \frac{\text{SYSTEM}(\mathcal{I}) - \text{FAIR}(\mathcal{I})}{\text{SYSTEM}(\mathcal{I})}.$$

275 Of course, $\text{PoF}(\mathcal{I})$ is a quantity between 0 and 1, where smaller values are preferable.

276 Now for an instance \mathcal{I} , let $s^g(\mathcal{I}) = \sup_{\pi \in \Psi^+(\mathcal{I})} \text{UtilGain}^g(\pi, \mathcal{I})$ be the maximum achievable utility
 277 gain (or equivalent, the largest reduction in regret possible) for group g , where $\Psi^+(\mathcal{I}) = \{\pi \in \Psi : \text{UtilGain}^g(\pi, \mathcal{I}) \geq 0 \forall g \in \mathcal{G}\}$. Then, $R(\mathcal{I}) = \min_{g \in \mathcal{G}} s^g(\mathcal{I}) / \max_{g \in \mathcal{G}} s^g(\mathcal{I})$ is a measure of the
 278 inherent asymmetry of the instance \mathcal{I} with respect to utility gain across groups. We show:
 279

280 **Theorem 4.2.** *For an instance \mathcal{I} of the grouped K -armed bandit, $\text{PoF}(\mathcal{I}) \leq 1 - R(\mathcal{I}) \frac{2\sqrt{G}-1}{G}$.*

281 The proof relies on an analysis of the price of fairness for general convex allocation problems in [29]
 282 and may be found in Appendix F. The key takeaway from this result is that, treating the inherent
 283 asymmetry $R(\mathcal{I})$ as a constant, the price of fairness grows *sub-linearly* in the number of groups G . It
 284 is unclear we can expect this with other fairness solution concepts: for instance, we would expect the
 285 price of fairness under a max-min solution to grown linearly with the number of groups [29]. Further,
 286 whereas the bound above depends on the topology of the instance only through $R(\mathcal{I})$, a topology
 287 specific analysis may well yield stronger results. For instance:

288 **Proposition 4.3.** *Let \mathcal{I} be an instance such that for every arm $a \in \mathcal{A}$, either $\mathcal{G}_a = \mathcal{G}$ or $|\mathcal{G}_a| = 1$.*
 289 *Then $\text{PoF}(\mathcal{I}) \leq \frac{1}{2}$.*

290 This result shows that for a specific class of topologies, the price of fairness is a constant independent
 291 of any parameters including the number of groups or the mean rewards. In Section 6 we study the
 292 price of fairness computationally in the context of random families of instances.

293 5 Extension to Grouped Contextual Linear Bandits

294 In this section, we introduce the grouped linear contextual bandit model and propose a candidate
 295 Nash solution by extending the regret optimal policy of [22] (without theory). We apply this model
 296 and the policies in Section 6 for an empirical case study.

297 **Grouped Linear Contextual Bandit Model:** Let $\theta \in \mathbb{R}^d$ and $\mathcal{A} \subseteq \mathbb{R}^d$. The reward for pulling arm
 298 a for the n 'th time is $Y_n(a) = \langle a, \theta \rangle + \varepsilon_{a,n}$, where $\varepsilon_{a,n}$ is distributed i.i.d. $N(0, 1)$. Let $\mathcal{M} \subseteq \mathbb{R}^d$
 299 be the set of contexts, where $|\mathcal{M}| = M < \infty$, and each $m \in \mathcal{M}$ is associated with an action set
 300 $\mathcal{A}(m) \subseteq \mathcal{A}$. Each group $g \in \mathcal{G}$ has a probability of arrival, p^g , and a distribution P^g over contexts
 301 $[M]$. At each time t , a group g_t is drawn independently from $(p^g)_g$, then a random context $m_t \sim P^{g_t}$
 302 is drawn. The action set at time t is $\mathcal{A}_t = \mathcal{A}(m_t)$. Let \mathcal{M}^g be the contexts in the support of P^g . Let
 303 $\text{OPT}(m) = \max_{a \in \mathcal{A}(m)} \langle a, \theta \rangle$ and $\Delta(m, a) = \text{OPT}(m) - \langle a, \theta \rangle$.

304 **Regret Optimal Policy:** [22] provides an instance-dependent lower bound for linear contextual
 305 bandits as the optimal value of the following optimization problem:

$$\begin{aligned} (L(\theta)) \quad Y(\mathcal{M}) &= \min_{Q \geq 0} \sum_{m \in \mathcal{M}} \sum_{a \in \mathcal{A}(m)} Q(m, a) \Delta(m, a) \\ &\text{s.t.} \quad Q(a) = \sum_{m: a \in \mathcal{A}(m)} Q(m, a) \quad \forall a \in \mathcal{A} \\ &\quad (Q(a))_{a \in \mathcal{A}} \in \mathcal{Q}, \end{aligned}$$

306 where \mathcal{Q} is the following polytope ensuring the consistency of the policy:

$$\mathcal{Q} = \left\{ (Q(a))_{a \in \mathcal{A}} : \|a\|_{H_Q^{-1}}^2 \leq \Delta(m, a)^2 / 2 \quad \forall m \in [M], a \in \mathcal{A}(m), H_Q = \sum_{a \in \mathcal{A}} Q(a) a a^\top \right\}.$$

307 The variable $Q(m, a)$ represents how often context m pulls arm a . [22] provides a policy (OAM)
 308 whose regret matches this lower bound. At a high level, like PF-UCB, OAM solves $L(\hat{\theta}_t)$ at each time
 309 step and ‘follows’ the solution; but it does not make use of a UCB and rather uses forced exploration.
 310 There are many details in the OAM policy and the full description can be found in Appendix A.2.

311 **Candidate Nash Solution:** We propose a policy which runs exactly OAM, except that the optimiza-
 312 tion problem solved at every time step is changed to the following:

$$\begin{aligned} (L^{\text{fair}}(\theta)) \quad \max_{Q \geq 0} &\quad \sum_{g \in \mathcal{G}} \log \left(Y(\mathcal{M}^g) - \sum_{m \in \mathcal{M}^g} \sum_{a \in \mathcal{A}(m)} Q^g(m, a) \Delta(m, a) \right)^+ \\ &\text{s.t.} \quad Q(a) = \sum_{g \in \mathcal{G}} \sum_{m \in \mathcal{M}^g: a \in \mathcal{A}(m)} Q^g(m, a) \quad \forall a \in \mathcal{A} \\ &\quad (Q(a))_{a \in \mathcal{A}} \in \mathcal{Q}. \end{aligned}$$

313 Compared to $(L(\theta))$, the objective is modified to maximize the Nash SW, and the new variable
 314 $Q^g(m, a)$ represents how often group g with context m should pull arm a .

315 We do not have a theoretical guarantee that this extension of OAM is indeed the Nash solution.
 316 This is not implied by [22] since there is an added group structure on the bandit model and OAM
 317 requires that the optimization problem has a unique solution, which $(L^{\text{fair}}(\theta))$ does not. Proving such
 318 a guarantee is a natural direction for future work.

319 **6 Experiments**

320 We consider two sets of experiments. The first seeks to understand the PoF in synthetic instances to
 321 shed further light on the impact of topology. The second is a real-world case study that returns to the
 322 Warfarin dosing example discussed in motivating the paper where we seek to understand unfairness
 323 under a regret optimal policy and the extent to which the Nash solution can mitigate this problem.

324 **Synthetic Grouped K -Armed Bandits:** We consider two generative models that differ in how
 325 the bipartite graph matching groups to available arms is generated. In ‘i.i.d.’, each edge appears
 326 independently with probability 0.5, and $K = 10$ is fixed. The mean reward of each arm is i.i.d.
 327 $U(0, 1)$. In ‘Skewed’, $K = G + 1$, and a group $g \in \{1, \dots, G - 1\}$ has access to arms $\{g, G\}$,
 328 while the last group $g = G$ has access to all arms. The rewards of arms $1, \dots, G - 1$ are equal, and
 329 $\mu(1) < \mu(G) < \mu(G + 1)$ are generated randomly by sorting three i.i.d. $U(0, 1)$ random variables.

330 Table 1 shows that the PoF is very small in the ‘i.i.d.’ setting, and contrary to Theorem 4.2 the PoF
 331 actually decreases as G gets large. This suggests an interesting conjecture for future research: the PoF
 332 may actually grow negligible in large random bandit instances. The ‘Skewed’ structure is motivated
 333 by our PoF analysis where we see that the PoF increase – albeit slowly – with G .

Table 1: The median and 95th percentile of the PoF for synthetic instances of the grouped K -armed bandit over 500 runs of each method.

G	i.i.d.				Skewed			
	3	5	10	50	3	5	10	50
Median	0.073	0.054	0.040	0.015	0.327	0.407	0.454	0.521
95th percentile	0.289	0.177	0.142	0.063	0.632	0.764	0.845	0.924

Table 2: Asymptotic disagreement point, regret, and utility gains for each group under the regret optimal and fair policies, where groups are either based on race or age. The numbers are derived from the optimal solution to $(L(\theta))$ and $(L^{\text{fair}}(\theta))$ for the regret optimal and fair policies respectively, for the grouped linear contextual bandit instance based on the warfarin dataset. As regret scales logarithmically as $T \rightarrow \infty$, these numbers represent the coefficient of $\log T$ term.

		Race				Age		
		A	B	C	Total	A	B	Total
Regret	Disagreement point	25.6	74.8	78.6	179.1	164.7	78.0	242.8
	Regret optimal	1.9	5.6	71.1	78.6	151.6	23.2	174.8
	Fair	0.0	25.4	54.0	79.4	149.3	29.3	178.7
Utility Gain	Regret optimal	23.7	69.2	7.6	100.4	13.1	54.9	68.0
	Fair	25.6	49.4	24.6	99.6	15.4	48.7	64.1

334 **Warfarin Dosing Case Study:** Warfarin is a common blood thinner whose optimal dose varies
 335 widely across patients. We use a publicly available dataset [30] to evaluate the effect of using a
 336 proportionally fair policy on learning the optimal personalized dose of warfarin. A detailed description
 337 of the experimental setup is deferred to Appendix A.3. The dataset contains covariates and the optimal
 338 dose of warfarin for 5700 patients. Both the age and race of patients are available and we use these
 339 to define groups. We use a linear contextual bandit setup with five features and an intercept; three
 340 actions (dose levels) are available to any arriving patient.

341 The results in Table 2 shows that for both groups based on race and age, the fair solution effectively
 342 ‘balances out’ the utility gains across groups with a small increase in regret. For race, we see that
 343 the disagreement point for groups B and C are very similar, but the regret optimal solution ends up
 344 benefitting B substantially more than C. The fair solution is able to ‘even out’ the utility gain between
 345 C to B for a small increase in regret. For age, the impact of fairness is smaller than with race which is
 346 potentially since there is less opportunity to learn across age groups than across race.

347 **References**

- 348 [1] Edward S Kim, Roy S Herbst, Ignacio I Wistuba, J Jack Lee, George R Blumenschein, Anne
349 Tsao, David J Stewart, Marshall E Hicks, Jeremy Erasmus, Sanjay Gupta, et al. The battle trial:
350 personalizing therapy for lung cancer. *Cancer discovery*, 1(1):44–53, 2011.
- 351 [2] Donald A Berry. Adaptive clinical trials in oncology. *Nature reviews Clinical oncology*,
352 9(4):199, 2012.
- 353 [3] Donald A Berry. The brave new world of clinical cancer research: adaptive biomarker-driven
354 trials integrating clinical practice with clinical research. *Molecular oncology*, 9(5):951–959,
355 2015.
- 356 [4] Hope S Rugo, Olufunmilayo I Olopade, Angela DeMichele, Christina Yau, Laura J van’t Veer,
357 Meredith B Buxton, Michael Hogarth, Nola M Hylton, Melissa Paoloni, Jane Perlmutter, et al.
358 Adaptive randomization of veliparib–carboplatin treatment in breast cancer. *New England
359 Journal of Medicine*, 375(1):23–34, 2016.
- 360 [5] Richard M Bergenstal, Mary Johnson, Rebecca Passi, Anuj Bhargava, Natalie Young, Davida F
361 Kruger, Eran Bashan, Stanley G Bisgaier, Deanna J Marriott Isaman, and Israel Hodish. Auto-
362 mated insulin dosing guidance to optimise insulin management in patients with type 2 diabetes:
363 a multicentre, randomised controlled trial. *The Lancet*, 393(10176):1138–1148, 2019.
- 364 [6] Revital Nimri, Tadej Battelino, Lori M Laffel, Robert H Slover, Desmond Schatz, Stuart A
365 Weinzimer, Klemen Dovc, Thomas Danne, and Moshe Phillip. Insulin dose optimization using
366 an automated artificial intelligence-based decision support system in youths with type 1 diabetes.
367 *Nature medicine*, 26(9):1380–1384, 2020.
- 368 [7] Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates.
369 *Operations Research*, 68(1):276–294, 2020.
- 370 [8] Thore Graepel, Joaquin Quinonero Candela, Thomas Borchert, and Ralf Herbrich. Web-scale
371 bayesian click-through rate prediction for sponsored search advertising in microsoft’s Bing
372 search engine. In *ICML*, 2010.
- 373 [9] Deepak Agarwal, Bo Long, Jonathan Traupman, Doris Xin, and Liang Zhang. Laser: A scalable
374 response prediction platform for online advertising. In *Proceedings of the 7th ACM international
375 conference on Web search and data mining*, pages 173–182, 2014.
- 376 [10] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Ad-
377 vances in applied mathematics*, 6(1):4–22, 1985.
- 378 [11] Christopher Jung, Sampath Kannan, and Neil Lutz. Quantifying the burden of exploration and
379 the unfairness of free riding. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium
380 on Discrete Algorithms*, pages 1892–1904. SIAM, 2020.
- 381 [12] Manish Raghavan, Aleksandrs Slivkins, Jennifer Vaughan Wortman, and Zhiwei Steven Wu.
382 The externalities of exploration and how data diversity helps exploitation. In *Conference on
383 Learning Theory*, pages 1724–1738. PMLR, 2018.
- 384 [13] Robert Kleinberg, Alexandru Niculescu-Mizil, and Yogeshwer Sharma. Regret bounds for
385 sleeping experts and bandits. *Machine learning*, 80(2):245–272, 2010.
- 386 [14] Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. Fairness in learning:
387 Classic and contextual bandits. *arXiv preprint arXiv:1605.07139*, 2016.
- 388 [15] Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalya Mandal, and David C Parkes.
389 Calibrated fairness in bandits. *arXiv preprint arXiv:1707.01875*, 2017.
- 390 [16] Stephen Gillen, Christopher Jung, Michael Kearns, and Aaron Roth. Online learning with an
391 unknown fairness metric. *arXiv preprint arXiv:1802.06936*, 2018.
- 392 [17] Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Y Narahari. Achieving fairness in the
393 stochastic multi-armed bandit problem. In *Proceedings of the AAAI Conference on Artificial
394 Intelligence*, volume 34, pages 5379–5386, 2020.

- 395 [18] Todd L Graves and Tze Leung Lai. Asymptotically efficient adaptive choice of control laws
396 incontrolled markov chains. *SIAM journal on control and optimization*, 35(3):715–743, 1997.
- 397 [19] Tor Lattimore and Csaba Szepesvari. The end of optimism? an asymptotic analysis of finite-
398 armed linear bandits. In *Artificial Intelligence and Statistics*, pages 728–737. PMLR, 2017.
- 399 [20] Richard Combes, Stefan Magureanu, and Alexandre Proutiere. Minimal exploration in structured
400 stochastic bandits. *arXiv preprint arXiv:1711.00400*, 2017.
- 401 [21] Bart Van Parys and Negin Golrezaei. Optimal learning for structured bandits. *Available at*
402 *SSRN 3651397*, 2020.
- 403 [22] Botao Hao, Tor Lattimore, and Csaba Szepesvari. Adaptive exploration in linear contextual
404 bandit. In *International Conference on Artificial Intelligence and Statistics*, pages 3536–3545.
405 PMLR, 2020.
- 406 [23] H Peyton Young. *Equity: in theory and practice*. Princeton University Press, 1995.
- 407 [24] Amartya Sen and James Eric Foster. *On Economic Inequality*. Oxford university press, 1997.
- 408 [25] John F Nash. The bargaining problem. *Econometrica: Journal of the econometric society*,
409 pages 155–162, 1950.
- 410 [26] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. *Microeconomic theory*,
411 volume 1. Oxford university press New York, 1995.
- 412 [27] Mamoru Kaneko and Kenjiro Nakamura. The nash social welfare function. *Econometrica:*
413 *Journal of the Econometric Society*, pages 423–435, 1979.
- 414 [28] Aurélien Garivier and Olivier Cappé. The kl-ucb algorithm for bounded stochastic bandits
415 and beyond. In *Proceedings of the 24th annual conference on learning theory*, pages 359–376,
416 2011.
- 417 [29] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. The price of fairness. *Operations*
418 *research*, 59(1):17–31, 2011.
- 419 [30] Michelle Whirl-Carrillo, Ellen M McDonagh, JM Hebert, Li Gong, K Sangkuhl, CF Thorn,
420 Russ B Altman, and Teri E Klein. Pharmacogenomics knowledge for personalized medicine.
421 *Clinical Pharmacology & Therapeutics*, 92(4):414–417, 2012.
- 422 [31] International Warfarin Pharmacogenetics Consortium. Estimation of the warfarin dose with
423 clinical and pharmacogenetic data. *New England Journal of Medicine*, 360(8):753–764, 2009.
- 424 [32] Fumihiko Takeuchi, Ralph McGinnis, Stephane Bourgeois, Chris Barnes, Niclas Eriksson,
425 Nicole Soranzo, Pamela Whittaker, Venkatesh Ranganath, Vasudev Kumanduri, William
426 McLaren, et al. A genome-wide association study confirms *vkorc1*, *cyp2c9*, and *cyp4f2*
427 as principal genetic determinants of warfarin dose. *PLoS Genet*, 5(3):e1000433, 2009.

428 Checklist

- 429 1. For all authors...
- 430 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
431 contributions and scope? [Yes]
- 432 (b) Did you describe the limitations of your work? [Yes]
- 433 (c) Did you discuss any potential negative societal impacts of your work? [Yes] Total
434 social welfare can decrease, and the extent of this is evaluated in Section 4.2 and is one
435 focus of the experiments in Section 6.
- 436 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
437 them? [Yes]
- 438 2. If you are including theoretical results...
- 439 (a) Did you state the full set of assumptions of all theoretical results? [Yes]

- 440 (b) Did you include complete proofs of all theoretical results? [Yes] All proofs are in
441 supplemental material.
- 442 3. If you ran experiments...
- 443 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
444 mental results (either in the supplemental material or as a URL)? [Yes] Supplemental
445 material.
- 446 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
447 were chosen)? [Yes]
- 448 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
449 ments multiple times)? [Yes] We report both median and 95 percentile.
- 450 (d) Did you include the total amount of compute and the type of resources used (e.g., type
451 of GPUs, internal cluster, or cloud provider)? [Yes]
- 452 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 453 (a) If your work uses existing assets, did you cite the creators? [Yes] [31]
- 454 (b) Did you mention the license of the assets? [Yes] In Appendix A.3.
- 455 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
456 Code for experiments in supplemental material.
- 457 (d) Did you discuss whether and how consent was obtained from people whose data you're
458 using/curating? [Yes] In Appendix A.3.
- 459 (e) Did you discuss whether the data you are using/curating contains personally identifiable
460 information or offensive content? [Yes] In Appendix A.3.
- 461 5. If you used crowdsourcing or conducted research with human subjects...
- 462 (a) Did you include the full text of instructions given to participants and screenshots, if
463 applicable? [N/A]
- 464 (b) Did you describe any potential participant risks, with links to Institutional Review
465 Board (IRB) approvals, if applicable? [N/A]
- 466 (c) Did you include the estimated hourly wage paid to participants and the total amount
467 spent on participant compensation? [N/A]

468 **A Deferred Descriptions**

469 **A.1 Negative Externality Example from [12]**

470 [12] provide an example of an instance where there exists a sub-population that is better off when
 471 UCB is run on that sub-population alone, compared to running UCB on the entire population. The
 472 example they provide depends on the total time horizon T . We claim that this does not occur when
 473 you fix an instance and consider asymptotic log-scaled regret, $\lim_{T \rightarrow \infty} \frac{R_T}{\log T}$.

474 Fix any time T_0 , and consider the two-armed instance according to $T = T_0$ from Definition 1 of [12].
 475 The population consists of three buckets that depend on their starting location: A, B, and C. The
 476 sub-population consisting of B and C is dubbed the “minority”, while A is the “majority”. Note that
 477 only B has access to both arms and hence it is the only bucket that can ever incur regret. Group B
 478 pulls the arm that has a higher UCB, defined as $\hat{\theta}_t(a) + \sqrt{\frac{\alpha \log T_0}{N_t(a)}}$ for some tuning parameter $\alpha > 0$.

479 We first summarize informally how the negative externality arises. Because arms 1 and 2 are so
 480 close together, even after $O(T_0)$ time steps, which arm has a higher UCB is not dominated by the
 481 difference between their empirical means, but it is dominated the second term of the UCB: $\sqrt{\frac{\alpha \log T_0}{N_t(a)}}$,
 482 which is just a function of the number of pulls $N_t(a)$. That is, group B essentially ends up pulling the
 483 arm that has fewer pulls. Therefore, when only the minority exists, since C only pulls arm 2, arm 1
 484 ends up having a higher UCB, and hence B ends up always pulling arm 1. However, if the majority
 485 group exists, arm 1 always has more pulls than arm 2 since there are more people from A than C.
 486 Then, B ends up essentially always pulling arm 2. If arm 2 is the arm that has a lower true reward
 487 than arm 1, then regret is higher when the majority group exists — therefore, the existence of the
 488 majority can have a “negative externality” on the minority.

489 However, if we fix this instance and let $T \rightarrow \infty$, then no matter which arms is better, from
 490 Theorem C.1, the total log-scaled regret is 0 from running KL-UCB. Moreover, when the majority
 491 does not exist, then the minority incurs non-zero log-scaled regret when $\theta_1 < \theta_2$. Therefore, the
 492 presence of the majority can only help the minority. Now, as explained in [12], it is true that the
 493 presence of the majority can negatively affect the minority in the early time steps (i.e. $t < T_0$). In
 494 the asymptotic regime, such a negative externality corresponds to adding $o(\log T)$ regret, which is
 495 deemed insignificant in our setting.

496 **A.2 Optimal Allocation Matching (OAM) Policy**

497 We describe the OAM algorithm from [22].

498 **Preliminaries:** Let $G_t = \sum_{s=1}^{t-1} A_s A_s^\top$ and let $\hat{\theta}_t = G_t^{-1} \sum_{s=1}^{t-1} A_s Y_s$ be the least squares estimate
 499 of θ at time t . Let $\hat{\Delta}_t^m(a) = \max_{a' \in \mathcal{A}(m)} \langle a' - a, \hat{\theta}_t \rangle$ be the corresponding estimate of $\Delta^m(a)$. Let
 500 $\hat{\Delta}_t^{\min} = \min_{m \in [M]} \min_{a \in \mathcal{A}(m), \hat{\Delta}_t(m, a) > 0} \hat{\Delta}_t(m, a)$ be the smallest nonzero instantaneous regret.
 501 Let

$$f_{T, \delta} = 2 \left(1 + \frac{1}{\log T} \right) \log \left(\frac{1}{\delta} \right) + cd \log(d \log T),$$

502 where c is an absolute constant. Let $f_T = f_{T, 1/T}$.

503 Define the following optimization problem that takes $\tilde{\Delta}(m, a)$ as input:

$$(K) \quad \begin{aligned} & \min \sum_{m \in \mathcal{M}} \sum_{a \in \mathcal{A}(m)} Q(m, a) \tilde{\Delta}(m, a) \\ & \text{s.t.} \quad \|a\|_{H_T^{-1}}^2 \leq \frac{\tilde{\Delta}(m, a)^2}{f_T} \quad \forall m \in \mathcal{M}, a \in \mathcal{A}(m) \\ & \quad \quad Q(m, a) \geq 0 \quad \forall m \in \mathcal{M}, a \in \mathcal{A}, \end{aligned}$$

504 where $H_T = \sum_{m \in \mathcal{M}} \sum_{a \in \mathcal{A}(m)} Q(m, a) a a^\top$ is invertible. Let $(\hat{Q}_t(m, a))_{m \in \mathcal{M}, a \in \mathcal{A}}$ be the solution
 505 to (K) using $\tilde{\Delta} = \hat{\Delta}_t$.

506 **Algorithm:** We are now ready to state the algorithm. At each time step t , observe context m_t and
 507 do the following. First, check whether

$$(10) \quad \|a\|_{G_t^{-1}}^2 \leq \frac{\hat{\Delta}_t(m_t, a)^2}{f_T} \quad \forall a \in \mathcal{A}(m_t).$$

508 If (10) is satisfied, we exploit; otherwise, we explore.

509 **Exploit:** Pull the greedy arm: $\operatorname{argmax}_{a \in \mathcal{A}(m_t)} \langle a, \hat{\theta}_t \rangle$.

510 **Explore:** Let $s(t)$ be the total number of exploration rounds so far. Solve the empirical optimization
 511 problem (K) to get solution $\hat{Q}_t(m_t, a)$.

- 512 1. Check whether $N_t^{m_t}(a) \geq \min(\hat{Q}_t(m_t, a), f_T/(\hat{\Delta}_t^{\min})^2)$ holds for all available arms $a \in$
 513 $\mathcal{A}(m_t)$. If so, pull the UCB arm $A_t = \operatorname{argmax}_{a \in \mathcal{A}(m_t)} \langle a, \hat{\theta}_t \rangle + \sqrt{f_T/1/s(t)^2} \|a\|_{G_t^{-1}}$.
- 514 2. Check whether there exists an available arm $a \in \mathcal{A}(m_t)$ such that $N_t(a) \leq \varepsilon_t s(t)$, where
 515 $\varepsilon_t = 1/\log \log t$. If there is, then pull $A_t = \operatorname{argmin}_{a \in \mathcal{A}(m_t)} N_t(a)$.
- 516 3. If the above two criteria are not true, then pull $A_t =$
 517 $\operatorname{argmin}_{a \in \mathcal{A}(m_t)} \frac{N_t(a)}{\min(\hat{Q}_t(m_t, a), f_T/(\hat{\Delta}_t^{\min})^2)}$.

518 A.3 Warfarin Experiment Details

519 We use a publicly available dataset for warfarin dosing that was collected by the Pharmacogenomics
 520 Knowledge Base (PharmGKB [30]), which is under a Creative Commons license¹. The dataset
 521 contains 5700 patients who were treated with warfarin from 21 research groups over 9 countries.
 522 Consent for all patients was obtained previously from each center, and no personally identifiable
 523 information was used. The dataset contains the optimal dose of warfarin for each patient, which
 524 was found by doctors through trial and error. It also includes many other covariates for each patient
 525 including demographics, clinical features, and genetic information.

526 **Groups:** We group the patients either by race or age. There were three distinct races in the dataset,
 527 which we label as A, B, and C. For age, we split the patients into two age groups, where the threshold
 528 age was 70.

529 **Contexts:** The OAM and PF-OAM policies assume a finite number of possible feature vectors, and
 530 the optimization problem ($L(\theta)$) scales with this number. Therefore, for tractability, we only use
 531 five features for the contexts of the patients, where we discretize each feature into two bins. We use
 532 the five features that are most correlated with the optimal warfarin dosage, and we use the empirical
 533 median of each feature to discretize them. The five features that we use are: age, weight, whether
 534 the patient was taking another drug (amiodarone), and two binary features capturing whether the
 535 patient has a particular genetic variant of genes Cyp2C9 and VKORC1, two genes that are known
 536 to affect warfarin dosage [32]. Out of $2^5 = 32$ different possible feature vectors, there were 21 that
 537 were present in the data.

538 **Rewards:** We bin the optimal dosage levels into three arms as was done in [7]: Low (under 3
 539 mg/day), Medium (3-7 mg/day), and High (over 7 mg/day). To ensure that the model is correctly
 540 specified, for each arm, we train a linear regression model using the entire dataset from the five
 541 contexts to the binary reward on whether the optimal dosage for that patient belongs in that bin. Let
 542 $\theta_a \in \mathbb{R}^6$ be the learned linear regression parameter for each arm ($d = 6$ to include the intercept).²
 543 To model this as grouped linear contextual bandits as described in Section 5, we let $d = 18$ and let
 544 $\theta = (\theta_1, \theta_2, \theta_3) \in \mathbb{R}^d$. When a patient with covariates $X \in \mathbb{R}^6$ arrives, the actions available are
 545 $\{(X, \mathbf{0}, \mathbf{0}), (\mathbf{0}, X, \mathbf{0}), (\mathbf{0}, \mathbf{0}, X)\}$, and their expected reward from arm a is $\langle X, \theta_a \rangle$ for $a \in \{1, 2, 3\}$.

546 **Algorithms:** We assume a patient is drawn i.i.d. from the dataset at each time step, and we compute
 547 the asymptotic group regret of the OAM policy ('Regret optimal') and the fair extension ('Fair') as
 548 described in Section 5:

¹<https://creativecommons.org/licenses/by-sa/4.0/>

²The linear regression step is done solely to remove model misspecification. The purpose of this study is not to show that the linear contextual bandit is a good fit for this dataset — this was already demonstrated in [7]. rather, the purpose is to demonstrate how incorporating fairness changes the outcome from a policy that does not take fairness into account on a bandit instance that approximates a real-world setting. rather, the purpose is to

- 549 • *Regret optimal:* Using the true values θ , we solve $(L(\theta))$ and obtain solution
550 $(Q(m, a))_{m \in [M], a \in \mathcal{A}}$. Then, the total (log-scaled) regret incurred by context m is
551 $\sum_{a \in \mathcal{A}} \Delta(m, a) Q(m, a)$. Since we assume the group arrivals are i.i.d., for each con-
552 text, we allocate the regret to groups in proportion to the group's frequency. That
553 is, for each m , let $(w^g(m))_{g \in \mathcal{G}}$, $\sum_{g \in \mathcal{G}} w^g(m) = 1$ be the empirical distribution of
554 groups among patients with context m . Then, the total regret assigned to group g is
555 $\sum_{m \in [M]} w^g(m) \sum_{a \in \mathcal{A}} \Delta(m, a) Q(m, a)$.
- 556 • *Fair:* Using the true values Δ , we solve $(L^{\text{fair}}(\theta))$ and obtain solu-
557 tion $(Q^g(m, a))_{g \in \mathcal{G}, m \in [M], a \in \mathcal{A}}$. The total regret assigned to group g is
558 $\sum_{m \in [M]} \sum_{a \in \mathcal{A}} \Delta(m, a) Q^g(m, a)$.

559 All experiments were run on a Macbook Pro with a 2.5 GHz Intel Core i7 processor.

560 B Proof Preliminaries

561 B.1 Notation

562 For all of the subsequent proofs, we assume that an instance \mathcal{I} is *fixed*. We often use big-O notation,
563 which is with respect to $T \rightarrow \infty$, unless otherwise specified. The big-O hides constants that may
564 depend on any other parameter other than T , including the instance \mathcal{I} . In general, when we introduce
565 a *constant*, it may depend on any other parameters other than T . We are usually not concerned
566 with the values of the constants as we are concerned with asymptotic results (though we do concern
567 ourselves with constants in front of the leading term, usually $\log T$). We sometimes re-use letters like
568 c for constants but they do not refer to the same value.

569 The UCB of an arm is defined as:

$$(11) \quad \text{UCB}_t(a) = \max\{q : N_t(a) \text{KL}(\hat{\theta}_t(a), q) \leq \log t + 3 \log \log t\}.$$

570 Let $\text{Pull}_t(a)$ be the indicator for arm a being pulled at time t , and let $\text{Pull}_t^g(a)$ be the indicator for
571 when arm a is pulled by group g . We define the class of log-consistent policies:

572 **Definition B.1.** A policy π for the grouped bandit problem is *log-consistent* for if for any instance
573 $(\theta, G, (p_g)_{g \in G}, (\mathcal{A}_g)_{g \in G})$, for any group g ,

$$(12) \quad \mathbb{E} \left[\sum_{a \in \mathcal{A}_{\text{sub}}(g)} N_T^g(a) \right] = O(\log T).$$

574 That is, the expected number of times that group g pulled a suboptimal arm by time t is logarithmic
575 in the number of arrivals of g .

576 B.2 Commonly Used Lemmas

577 We state a few lemmas that are used several times for both Theorem C.1 and Theorem 4.1. These
578 lemmas do not depend on the policy that is used. The first result shows that the number of times that
579 an arm's UCB is smaller than its true mean is small.

580 **Lemma B.2.** Let $\Lambda_t = \{\text{UCB}_t(a) \geq \theta(a) \forall a \in \mathcal{A}\}$ be the event that UCB for every arm is valid at
581 time t .

$$\sum_{t=1}^T \Pr(\bar{\Lambda}_t) = O(\log \log T).$$

582 *Proof.* For a fix arm a , $\sum_{t=1}^T \Pr(\text{UCB}_t(a) < \theta(a)) = O(\log \log T)$ follows from Theorem 10 of
583 [28], plugging in $\delta = \log t + 3 \log \log t$ as is done in the proof of Theorem 2 of [28]. The result
584 follows from a union bound over all actions $a \in \mathcal{A}$. \square

585 The second lemma states a relationship between the radius of the UCB of an arm and the number of
586 pulls of the arm.

587 **Lemma B.3.** Let $0 < \alpha < \beta < 1$. There exists a constant $c > 0$ such that if $\hat{\theta}_t(a) \leq \alpha$ and
588 $\text{UCB}_t(a) \geq \beta$, then $N_t(a) < c \log t$.

589 *Proof.* Suppose $\hat{\theta}_t(a) \leq \alpha$ and $\text{UCB}_t(a) \geq \beta$. Then, $\text{KL}(\hat{\theta}_t(a), \text{UCB}_t(a)) \geq \text{KL}(\alpha, \beta)$. Let
590 $c = \frac{4}{\text{KL}(\alpha, \beta)}$. By definition of the UCB (11), $N_t(a) \leq \frac{\log t + 3 \log \log t}{\text{KL}(\hat{\theta}_t(a), \text{UCB}_t(a))} \leq c \log t$. \square

591 This result essentially states that if the radius of the UCB of an arm is larger than a constant, then the
592 number of pulls of the arm is at most $O(\log t)$; this result follows simply from the definition of the
593 UCB (11). The next result states that if an arm a is pulled, then its empirical mean will be close to its
594 true mean.

595 **Lemma B.4.** For any group g and arm $a \in \mathcal{A}^g$, if $L < \theta(a) < U$,

$$\sum_{t=1}^T \Pr(\text{Pull}_t(a), \hat{\theta}_t(a) \notin [L, U]) = O(1).$$

596 where big- O hides constants that may depend on the instance and L, U .

597 *Proof.* Let $\hat{\theta}^n(a)$ be the empirical mean after n pulls of arm a . Let $E_{t,n}$ be the event that the number
598 of times arm 1 has been pulled before time t is exactly n .

$$\begin{aligned} & \sum_{t=1}^T \Pr(\text{Pull}_t(a), \hat{\theta}_t(a) \notin [L, U]) \\ &= \sum_{t=1}^T \sum_{n=1}^T \Pr(\text{Pull}_t(a), \hat{\theta}^n(a) \notin [L, U], E_{t,n}) \\ &= \sum_{n=1}^T \sum_{t=1}^T \Pr(\hat{\theta}^n(a) \notin [L, U] \mid \text{Pull}_t(a), E_{t,n}) \Pr(\text{Pull}_t(a), E_{t,n}) \end{aligned}$$

599 If $F_{t,n} = \{\text{Pull}_t(a), E_{t,n}\}$, then for any n , the events $F_{1,n}, \dots, F_{T,n}$ are disjoint. Then, by the law
600 of total probability, $\Pr(\hat{\theta}^n(a) \notin [L, U]) \geq \sum_{t=1}^T \Pr(\hat{\theta}^n(a) \notin [L, U] \mid F_{t,n}) \Pr(F_{t,n})$. Therefore,

$$\sum_{t=1}^T \Pr(\text{Pull}_t(a), \hat{\theta}_t(a) \notin [L, U]) \leq \sum_{n=1}^T \Pr(\hat{\theta}^n(a) \notin [L, U]) \leq \sum_{n=1}^T \exp(-\alpha n).$$

601 for some $\alpha > 0$ since the rewards of arm a are Bernoulli. Therefore, $\sum_{t=1}^T \Pr(\text{Pull}_t(a), \hat{\theta}_t(a) \notin$
602 $[L, U]) = O(1)$. \square

603 C Proof that KL-UCB is Regret Optimal

604 In this section, we prove that the KL-UCB policy is regret-optimal. At each time step, $\pi^{\text{KL-UCB}}$
605 chooses the arm with the highest UCB, defined as (11), out of all arms available.

606 **Theorem C.1.** For all instances \mathcal{I} of the grouped K -armed bandit,

$$(13) \quad \liminf_{T \rightarrow \infty} \frac{R_T(\pi^{\text{KL-UCB}}, \mathcal{I})}{\log T} \leq \sum_{a \in \mathcal{A}_{\text{sub}}} \Delta^{\Gamma(a)}(a) J(a).$$

607 The first step of the proof is to show that the number of pulls of a suboptimal arm is optimal:

608 **Proposition C.2.** Let $a \in \mathcal{A}_{\text{sub}}$ be a suboptimal arm. KL-UCB satisfies

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[N_T(a)]}{\log T} \leq J(a).$$

609 This result can be shown using the existing analysis of KL-UCB from [28]. The next step is to
610 analyze how these pulls are distributed across groups. In particular, we need to show that a group
611 never pulls a suboptimal arm a if $g \notin \Gamma(a)$. This is the result of the next theorem:

612 **Proposition C.3.** Let $a \in \mathcal{A}$. Let $g \in G_a$, $g \notin \Gamma(a)$ be a group that has access to the arm but is not
613 the group that has the smallest optimal out of G_a . Then, KL-UCB satisfies

$$\mathbb{E}[N_T^g(a)] = O(\log \log T),$$

614 where the big- O hides constants that depend on the instance.

615 This result implies that for any arm a , the regret incurred by group $g \notin \Gamma(a)$ pulling the arm
 616 is $o(\log T)$, and is equal to 0 when scaled by $\log T$. Theorem C.1 then follows from combining
 617 Proposition C.2 and Proposition C.3.

618 In this section, we prove Proposition C.3. Let $a \in \mathcal{A}$ and let $A \in \Gamma(a)$ be a group that has access to
 619 that arm with the smallest OPT. Let group $B \notin \Gamma(a)$ be another group that has access to arm a . Let
 620 θ^A, θ^B be the optimal arms for group A and B respectively. We use θ^A, θ^B to refer to both the arm
 621 and the arm means. Our goal is to show $\mathbb{E} [N_T^B(a)] = O(\log \log T)$.

$$\begin{aligned} \mathbb{E} [N_T^B(a)] &= \sum_{t=1}^T \Pr(\text{Pull}_t^B(a)) \\ &= \sum_{t=1}^T \Pr(\text{Pull}_t^B(a), \text{UCB}_t(\theta^B) \geq \theta^B) + \sum_{t=1}^T \Pr(\text{Pull}_t^B(a), \text{UCB}_t(\theta^B) < \theta^B). \end{aligned}$$

622 The second sum can be bounded by Lemma B.2, since $\sum_{t=1}^T \Pr(\text{Pull}_t^B(a), \text{UCB}_t(\theta^B) < \theta^B) \leq$
 623 $\sum_{t=1}^T \Pr(\bar{\Lambda}_t) = O(\log \log T)$. Therefore, our goal is to show

$$(14) \quad \sum_{t=1}^T \Pr(\text{Pull}_t^B(a), \text{UCB}_t(\theta^B) \geq \theta^B) = O(\log \log T).$$

624 We state a slightly more general result that implies (14).

625 **Lemma C.4.** *Suppose we run any log-consistent policy π . Let $r > 0$ be fixed. For any $a \in \mathcal{A}$,*

$$\sum_{t=1}^T \Pr(\text{Pull}_t(a), \text{UCB}_t(a) \geq \text{OPT}(\Gamma(a)) + r) = O(\log \log T),$$

626 *where the constant in the big-O may depend on the instance and r .*

627 The rest of this section proves Lemma C.4.

628 C.1 Probabilistic Lower Bound of $N_t(a)$ for Grouped Bandit

629 One of the main tools used in the proof of Lemma C.4 is a high probability lower bound on the
 630 number of pulls of a suboptimal arm. Let $W_t(g)$ be the number of arrivals of group g by time t .
 631 Let $R_t^g = \{W_t(g) \geq \frac{p_g t}{2}\}$ be the event that the number of arrivals of group g is at least half of the
 632 expected value. We condition on the event R_t^g to ensure that a group has arrived a sufficient number
 633 of times.

634 **Proposition C.5.** *Let g be a group, and let $a \in \mathcal{A}_{\text{sub}}^g$ be a suboptimal arm for group g . Fix $\varepsilon \in (0, 1)$.
 635 Suppose we run a log-consistent policy as defined in Definition B.1. Then,*

$$\Pr \left(N_t(a) \leq \frac{(1 - \varepsilon) \log t}{KL(\theta(a), \text{OPT}(g))} \mid R_t^g \right) = O \left(\frac{1}{\log t} \right),$$

636 *where the big-O notation is with respect to $t \rightarrow \infty$.*

637 The proof of this result can be found in Appendix D.3. For an arm $a \notin \mathcal{A}_{\text{sub}}$, we have the following
 638 stronger result:

639 **Proposition C.6.** *Let a be an arm that is optimal for some group g . Suppose we run a log-consistent
 640 policy. Then, for any $b > 0$,*

$$\Pr (N_t(a) \leq b \log t \mid R_t^g) = O \left(\frac{1}{\log t} \right),$$

641 *where the big-O notation is with respect to $t \rightarrow \infty$ and hide constants that depend on both b and the
 642 instance.*

643 **C.2 Proof of Lemma C.4**

644 **Outline:** Let $A \in \Gamma(a)$ be a group that has the smallest optimal out of all arms with access to a .
645 The main idea of this lemma is that group A does not “allow” the UCB of arm a to grow as large
646 as $\text{OPT}(A) + r$, as group A would pull arm a once the UCB is above $\text{OPT}(A)$. Proposition C.5
647 implies that $\text{UCB}_t(a)$ is not larger than $\text{OPT}(A)$ with high probability. If this occurs at time t , since
648 the radius of the UCB grows slowly (logarithmically), the earliest time that the UCB can grow to
649 $\text{OPT}(A) + r$ is t^γ , for some $\gamma > 1$. We divide the time steps into epochs, where if epoch k starts at
650 time s_k , it ends at s_k^γ . This exponential structure gives us $O(\log \log T)$ epochs in total, and we show
651 that the expected number of times that $\text{UCB}_t(a) > \text{OPT}(A) + r$ during one epoch is $O(1)$.

652 **Proof:** We denote by θ_a the true mean reward of arm a and by $\hat{\theta}_t$ the empirical mean reward of a at
653 the start of time t . Let $U = \text{OPT}(\Gamma(a)) + r$. Let $A \in \Gamma(a)$, and let $\theta^A = \text{OPT}(A)$. If $a \notin \mathcal{A}_{\text{sub}}$,
654 then let $\theta^A = \text{OPT}(A) + r/2$. Let $b > 0$ such that $\frac{\text{KL}(\theta_a, U)}{\text{KL}(\theta_a, \theta^A)} = 1 + b$. Define $\theta_u \in [\theta_a, \theta^A]$ such
655 that $\frac{\text{KL}(\theta_u, U)}{\text{KL}(\theta_u, \theta^A)} = 1 + \frac{b}{2}$. We have $\theta_a < \theta_u < \theta^A < U$. Define $\gamma \triangleq 1 + \frac{b}{4}$. Let $\varepsilon > 0$ such that
656 $\frac{1-\varepsilon}{1+\varepsilon} \cdot \frac{\text{KL}(\theta_u, U)}{\text{KL}(\theta_u, \theta^A)} = \gamma$.

657 By Lemma B.4, $\sum_{t=1}^T \Pr(\text{Pull}_t(a), \hat{\theta}_t(a) > \theta_u) = O(1)$. Therefore, we can assume $\hat{\theta}_t(a) \leq \theta_u$.
658 Denote the event of interest by $E_t = \{\text{Pull}_t(a), \text{UCB}_t(a) \geq \theta^A + r, \hat{\theta}_t(a) \leq \theta_u\}$. Our goal is to
659 show $\sum_{t=1}^T \Pr(E_t) = O(\log \log T)$.

660 Divide the time interval T into $K = O(\log \log T)$ epochs. Let epoch k start at $s_k \triangleq \lceil 2^{\gamma k} \rceil$ for $k \geq 0$.
661 Let $\mathcal{T}_k = \{s_k, s_k + 1, \dots, s_{k+1} - 1\}$ be the time steps in epoch k . This epoch structure satisfies the
662 following properties:

- 663 1. The total number of epochs is $O(\log \log T)$.
- 664 2. $\frac{\log s_{k+1}}{\log s_k} = \gamma$ for all $k \geq 0$.

665 We will treat each epoch separately. Fix an epoch k . Our goal is to bound $\mathbb{E}[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t)]$.
666 Lemma B.3 implies that there exists a constant $c > 0$ such that if E_t occurs, it must be that
667 $N_t(a) < c \log t$. Hence,

$$\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t) \leq c \log s_{k+1}.$$

668 Define the event $G_t = \left\{ N_t(a) \geq (1 - \varepsilon) \frac{\log t}{\text{KL}(\mu, \theta^A)} \right\}$. The following claim says that if G_{s_k} is true,
669 then E_t never happens during that epoch.

670 **Claim C.7.** *Suppose G_{s_k} is true. Let t_0 be such that if $t \geq t_0$, $\log \log t \leq \varepsilon \log t$. Then, if*
671 $s_k \geq t_0$, $\sum_{t=s_k}^{s_{k+1}} \mathbf{1}(E_t) = 0$.

672 This result follows from the fact that the event G_{s_k} implies that the radius of the UCB is “small” at
673 time s_k , and the epoch is defined so that the radius will not grow large enough that E_t can occur
674 during epoch k . Therefore, we have the following:

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t) \right] = \mathbb{E} \left[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t) \middle| \bar{G}_{s_k} \right] \Pr(\bar{G}_{s_k}) \leq c \log s_{k+1} \Pr(\bar{G}_{s_k}).$$

675 We can bound $\Pr(\bar{G}_{s_k})$ using the probabilistic lower bound of Proposition C.5.

676 **Claim C.8.** $\Pr(\bar{G}_{s_k}) \leq O\left(\frac{1}{\log s_k}\right)$.

677 Then, property 2 of the epoch structure implies $\mathbb{E}[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t)] = O(1)$. Since the number of
678 epochs is $O(\log \log T)$,

$$\mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(E_t) \right] \leq \sum_{k=1}^K \mathbb{E} \left[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t) \right] = O(\log \log T),$$

679 as desired.

680 **C.3 Proof of Claims**

681 *Proof of Claim C.7.* Let $t = s_k > t_0$ and let $t' \geq t$ such that $E_{t'}$ is true. By definition of KL-UCB,

$$N_{t'}(a) \leq \frac{\log t' + 3 \log t'}{\text{KL}(\hat{\theta}_{t'}, \text{UCB}_{t'}(\theta))}.$$

682 Since $E_{t'}$ implies $\text{UCB}_{t'}(a) > \theta^B$ and $\hat{\theta}_{t'} \leq \theta_u$, we have $N_{t'}(a) \leq \frac{\log t' + 3 \log t'}{\text{KL}(\theta_u, \theta^B)}$. Since G_{s_k} is true,
683 $N_{t'}(a) \geq (1 - \varepsilon) \frac{\log s_k}{\text{KL}(\theta_a, \theta^A)}$. Therefore, it must be that

$$\begin{aligned} (1 - \varepsilon) \frac{\log s_k}{\text{KL}(\theta_a, \theta^A)} &\leq \frac{\log t' + 3 \log \log t'}{\text{KL}(\theta_u, \theta^B)} \leq \frac{(1 + \varepsilon) \log t'}{\text{KL}(\theta_u, \theta^B)} \\ \Rightarrow \frac{1 - \varepsilon}{1 + \varepsilon} \cdot \frac{\text{KL}(\theta_u, \theta^B)}{\text{KL}(\theta_a, \theta^A)} \log s_k &\leq \log t' \\ \Rightarrow t' &\geq s_k^\gamma. \end{aligned}$$

684 This implies that t' is not in epoch k . □

685 *Proof of Claim C.8.* For group $g = A$, Proposition C.5 (or Proposition C.6 if $a \notin \mathcal{A}_{\text{sub}}$) states that

$$\Pr(\bar{G}_{s_k} \mid R_{s_k}^g) = O\left(\frac{1}{\log s_k}\right).$$

686 (We show in Appendix D.1 that KL-UCB is log-consistent.)

687 Now we need to bound $\Pr(\bar{R}_{s_k}^g) = \Pr(M_{s_k}(A) \leq \frac{p_A s_k}{2})$. Note that $M_s(A) = \sum_{t=1}^s Z_t^A$, where
688 $Z_t^A \stackrel{\text{iid}}{\sim} \text{Bern}(p_A)$. By Hoeffding's inequality,

$$\Pr\left(M_{s_k}(A) \leq \frac{p_A s_k}{2}\right) < \exp\left(-\frac{1}{2} p_A^2 s_k\right).$$

689 Combining, we have

$$\Pr(\bar{G}_k) \leq \Pr(\bar{R}_k) + \Pr(\bar{G}_k \mid R_k) \leq O\left(\frac{1}{\log s_k}\right).$$

690 □

691 **D Deferred Proofs for Theorem C.1**

692 For any $\varepsilon > 0$, let

$$K_\varepsilon^g(x) = \left\lceil \frac{1 + \varepsilon}{\text{KL}(\theta_a, \text{OPT}(g))} (\log x + 3 \log \log x) \right\rceil.$$

693 To show both Proposition C.2 and the fact that KL-UCB is log-consistent, we make use of the
694 following lemma.

695 **Lemma D.1.** Let $a \in \mathcal{A}$. Let $g \in G_a$ be a group in which a is suboptimal. For any $\varepsilon > 0$,

$$(15) \quad \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), N_t(a) \geq K_\varepsilon^g(T)) \right] = O(\log \log T).$$

696 *Proof.* Let $\varepsilon > 0$. Recall that A_g^* is the optimal arm for group g , and $\text{OPT}(g)$ is the mean reward of
697 A_g^* .

$$\begin{aligned} &\mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), N_t(a) \geq K_\varepsilon^g(T)) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), N_t(a) \geq K_\varepsilon^g(T), \text{UCB}_t(A_g^*) \geq \text{OPT}(g)) \right] + \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), \text{UCB}_t(A_g^*) < \text{OPT}(g)) \right] \end{aligned}$$

698 The second term is $O(\log \log T)$ from Lemma B.2. We will show that the first term is $O(1)$. Let
 699 $\hat{\theta}_s(a)$ be the empirical mean of a after s pulls. Consider the event $\{A_t = a, g_t = g, N_t(a) =$
 700 $s, \text{UCB}_t(A_g^*) \geq \text{OPT}(g)\}$, where $s \geq K_n$. Suppose this is true at time t . Then, it must be that
 701 $\text{UCB}_t(a) \geq \text{OPT}(g)$. For this to happen, by definition of KL-UCB, it must be that

$$(16) \quad s \text{KL}(\hat{\theta}_s(a), \text{OPT}(g)) \leq \log t + 3 \log \log t.$$

702 Since $s \geq K_\varepsilon^g(T)$ and $t \leq T$, we must have

$$(17) \quad \text{KL}(\hat{\theta}_s(a), \text{OPT}(g)) \leq \frac{\log T + 3 \log \log T}{K_\varepsilon^g(T)} = \frac{\text{KL}(\theta_a, \text{OPT}(g))}{1 + \varepsilon}.$$

703 Let $r > \theta_a$ such that $\text{KL}(r, \text{OPT}(g)) = \frac{\text{KL}(\theta_a, \text{OPT}(g))}{1 + \varepsilon}$. Then, for (17) to occur, it must be that
 704 $\hat{\theta}_s(a) \geq r$. Then, we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), N_t(a) \geq K_\varepsilon^g(n), \text{UCB}_t(A_g^*) \geq \text{OPT}(g)) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \sum_{s=K_n}^{\infty} \mathbf{1}(\text{Pull}_t^g(a), N_t(a) = s, \text{UCB}_t(A_g^*) \geq \text{OPT}(g)) \right] \\ &\leq \mathbb{E} \left[\sum_{t=1}^T \sum_{s=K_n}^{\infty} \mathbf{1}(\text{Pull}_t^g(a), N_t(a) = s, \hat{\theta}_s(a) \geq r) \right] \\ &= \mathbb{E} \left[\sum_{s=K_n}^{\infty} \mathbf{1}(\hat{\theta}_s(a) \geq r) \sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), N_t(a) = s) \right] \\ &\leq \sum_{s=K_n}^{\infty} \Pr(\hat{\theta}_s(a) \geq r). \end{aligned}$$

705 Since $r > \mu(a)$, there exists a constant $C_3 > 0$ that depends on ε and r such that $\Pr(\mu_s(a) \geq r) \leq$
 706 $\exp(-sC_3)$. Therefore, $\sum_{s=K_n}^{\infty} \Pr(\hat{\theta}_s(a) \geq r) = O(1)$ and we are done.

707 □

708 D.1 Proof that KL-UCB is log-consistent

709 This basically follows from Lemma D.1. Let $\varepsilon = 1/2$. Fix a group g , and let a be a suboptimal arm
 710 for g .

$$\begin{aligned} \mathbb{E}[N_T^g(a)] &= \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a)) \right] \\ &\leq K_\varepsilon^g(T) + \mathbb{E} \left[\sum_{t=1}^{t_g(n)} \mathbf{1}(\text{Pull}_t^g(a), N_t(a) \geq K_\varepsilon^g(T)) \right] \\ &= K_\varepsilon^g(T) + \log \log(T). \end{aligned}$$

711 We are done since $K_\varepsilon^g(T) = O(\log T)$.

712 D.2 Proof of Proposition C.2

713 Let $a \in \mathcal{A}_{\text{sub}}$ be a suboptimal arm. Let $\varepsilon > 0$. Let

$$K_T = \max_{g \in G_a} K_\varepsilon^g(T).$$

714 Clearly, the maximum is attained in the group g with the smallest $\text{OPT}(g)$, so.

$$K_T = \left\lceil \frac{1 + \varepsilon}{\text{KL}(\theta_a, \text{OPT}(\Gamma(a)))} (\log T + 3 \log \log T) \right\rceil.$$

$$\begin{aligned}
\mathbb{E}[N_T(a)] &= \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(A_t = a) \right] \\
&\leq K_T + \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(A_t = a, N_t(a) \geq K_T) \right] \\
&\leq K_T + \sum_{g \in G_a} \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), N_t(a) \geq K_T) \right] \\
&\leq K_T + \sum_{g \in G_a} O(\log \log T).
\end{aligned}$$

715 where the last inequality follows from Eq. (15) of Lemma D.1. Since this holds for any $\varepsilon > 0$, the
716 desired result holds.

717 D.3 Proof of Proposition C.5 and Proposition C.6

718 Let g be a group, and let j be a suboptimal arm for group g ; i.e. $\theta_j < \text{OPT}(g)$. Fix $\varepsilon > 0$. We
719 assume that the event $R_t^g = \{W_t(g) \geq \frac{p_g t}{2}\}$ holds. Fix $\delta > 0$ such that $\frac{1-\delta}{1+\delta} = 1 - \varepsilon$. Let $a = \delta/2$.
720 We construct another instance γ where arm j is replace with λ so that arm j is the optimal arm for g
721 in the same manner as the Lai-Robbins proof. Specifically, $\lambda > \theta_j$ such that

$$\text{KL}(\theta_j, \lambda) = (1 + \delta)\text{KL}(\theta_j, \text{OPT}(g)).$$

722 Our goal is to bound the probability of event $\left\{ N_t(j) \leq \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)} \right\}$, which we split into two events:

$$\begin{aligned}
C_t &= \left\{ N_t(j) \leq \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)}, L_{N_t(j)} \leq (1-a)\log t \right\}, \\
E_t &= \left\{ N_t(j) \leq \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)}, L_{N_t(j)} > (1-a)\log t \right\},
\end{aligned}$$

723 where $L_m = \sum_{i=1}^m \log \left(\frac{f(Y_i; \theta_j)}{f(Y_i; \lambda)} \right)$.

724 Assumption (12), there exists a constant c such that if t is large enough that $\Pr(R_t^g) \geq 1/2$,

$$\mathbb{E}_\gamma \left[\sum_{a \in \mathcal{A}_{\text{sub}}} N_t^g(a) \mid R_t^g \right] \leq c \log t.$$

725 Since j is the unique optimal arm under γ ,

$$\mathbb{E}_\gamma \left[W_t(g) - N_t^g(j) \mid R_t^g \right] \leq c \log t.$$

726 Using Markov's inequality and using the fact that $W_t(g) \geq \frac{p_g t}{2}$, we get

$$\begin{aligned}
\Pr_\gamma \left(N_t^g(j) \leq \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)} \mid R_t^g \right) &= \Pr_\gamma \left(W_t(g) - N_t^g(j) \geq W_t(g) - \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)} \mid R_t^g \right) \\
&\leq \Pr_\gamma \left(W_t(g) - N_t^g(j) \geq \frac{p_g t}{2} - \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)} \mid R_t^g \right) \\
&\leq \frac{\mathbb{E} [W_t(g) - N_t^g(j) \mid R_t^g]}{\frac{p_g t}{2} - \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)}} \\
&= O \left(\frac{\log t}{t} \right).
\end{aligned}$$

727 **Bounding** $\Pr(C_t \mid R_t^g)$: Following through with the same steps as the original proof, we can replace
 728 (2.7) with

$$\Pr_\theta(C_t \mid R_t^g) \leq t^{1-a} \Pr_\gamma(C_t \mid R_t^g) \leq t^{1-a} O\left(\frac{\log t}{t}\right) = O\left(\frac{\log t}{t^a}\right).$$

729 **Bounding** $\Pr(E_t \mid R_t^g)$: Next, we need to show a probabilistic result in lieu of (2.8) of [10]. Let
 730 $m = \frac{(1-\delta)\log t}{\text{KL}(\theta_j, \lambda)}$ and let $\alpha > 0$ such that $(1 + \alpha) = \frac{1-a}{1-\delta}$. We need to upper bound

$$\begin{aligned} \Pr_\theta\left(\max_{j \leq m} L_j > (1-a)\log t\right) &= \Pr_\theta\left(\max_{j \leq m} L_j > (1+\alpha)\text{KL}(\theta_j, \lambda)m\right) \\ &\leq \Pr_\theta\left(\max_{j \leq m}\{L_j - j\text{KL}(\theta_j, \lambda)\} > \alpha\text{KL}(\theta_j, \lambda)m\right). \end{aligned}$$

731 Let $Z_i = \log\left(\frac{f(Y_i; \theta_j)}{f(Y_i; \lambda)}\right) - \text{KL}(\theta_j, \lambda)$. We have $\mathbb{E}[Z_i] = 0$. Let $\text{Var}(Z_i) = \sigma^2$. Then, by Kol-
 732 mogorov's inequality, we have

$$\begin{aligned} \Pr_\theta\left(\max_{j \leq m} \sum_{i=1}^j Z_i > \alpha\text{KL}(\theta_j, \lambda)m\right) &\leq \frac{1}{\alpha^2\text{KL}(\theta_j, \lambda)^2 m^2} \text{Var}\left(\sum_{i=1}^m Z_i\right) \\ &= \frac{\sigma^2}{\alpha^2\text{KL}(\theta_j, \lambda)^2 m} \\ &= O\left(\frac{1}{\log t}\right), \end{aligned}$$

733 since $m = \Theta(\log t)$.

734 **Combine:** Combining, we have

$$\begin{aligned} \Pr_\theta\left(N_t(j) \leq \frac{(1-\delta)\log n}{\text{KL}(\theta_j, \lambda)} \mid R_t^g\right) &= \Pr_\theta(C_n \mid R_t^g) + \Pr_\theta(E_n \mid R_t^g) \\ &= O\left(\frac{\log t}{t^a}\right) + O\left(\frac{1}{\log t}\right). \end{aligned}$$

735 Since $\text{KL}(\theta_j, \lambda) \leq (1+\delta)\text{KL}(\theta_j, \text{OPT}(g))$ and $\frac{1-\delta}{1+\delta} = 1 - \varepsilon$, we have

$$\Pr_\theta\left(N_t(j) \leq \frac{(1-\varepsilon)\log t}{\text{KL}(\theta_j, \text{OPT}(g))} \mid R_t^g\right) \leq O\left(\frac{1}{\log t}\right)$$

736 as desired.

737 *Proof of Proposition C.6.* The proof of this result follows the same steps as Proposition C.5. Let
 738 $\varepsilon = 1/2$ and let $\theta^* > \theta_j$ so that $\frac{1-\varepsilon}{\text{KL}(\theta_j, \theta^*)} = b$. In the proof of Proposition C.5, replace $\text{OPT}(g)$ with
 739 θ^* . Then, the same proof goes through and we get $\Pr(N_t(j) \leq b \log n \mid R_t^g) = O\left(\frac{1}{\log t}\right)$. \square

740 E Proof of Theorem 4.1

741 To prove Theorem 4.1, our goal is to show that the total number of pulls of a suboptimal arm a is
 742 $J(a) \log T$, and those pulls are distributed amongst groups according to $q_*^g(a)$. The policy PF-UCB
 743 assigns arms in a way that the distribution of groups that have pulled arm a converges to $\hat{q}_t^g(a)$.
 744 Hence, our goal is to show that $\hat{q}_t^g(a)$ is usually “close” to $q_*^g(a)$.

745 Let $\delta_0 = \min_{a \neq a'} \frac{|\theta(a) - \theta(a')|}{4}$. For $\delta \in (0, \delta_0)$ let $H_t(\delta) = \{\hat{\theta}_t(a) \in [\theta(a) - \delta, \theta(a) + \delta] \forall a \in \mathcal{A}\}$
 746 be the event that all arms are within their “ δ -boundaries”. Since $\delta < \delta_0$, this implies that the ranking
 747 of the arms do not change if $H_t(\delta)$ is true (i.e. $\theta(a) < \theta(a') \Rightarrow \hat{\theta}_t(a) < \hat{\theta}_t(a')$). We first state a result
 748 pertaining to the program $(P(\theta))$, which states that if $H_t(\delta)$ is true, the approximate solution \hat{q}_t is
 749 also close to the true solution q_* .

750 **Proposition E.1.** For any $\varepsilon > 0$, there exists $\delta > 0$ such that if $H_t(\delta)$, then $\hat{q}_t^g(a) \in [q_*^g(a) -$
751 $\varepsilon, q_*^g(a) + \varepsilon]$ for all $a \in \mathcal{A}$ and $g \in \mathcal{G}$.

752 The proof of Proposition E.1 can be found in Appendix G.4. This result implies that when we have
753 good empirical estimates of θ (i.e. $H_t(\delta)$ is true), the policy of ‘following’ the solution $\hat{q}_t^g(a)$ will
754 give us the desired ‘split’ of pulls between groups. Therefore, our goal is to show that suboptimal
755 arms are pulled only when $H_t(\delta)$ is true.

756 For $a \in \mathcal{A}_{\text{sub}}^g$, there are two reasons why $\text{Pull}_t^g(a)$ would occur: (i) $a = A_t^{\text{UCB}}(g')$ for some group
757 g' , or (ii) $a = A_t^{\text{greedy}}(g)$. We show that the regret from (ii) is negligible:

758 **Proposition E.2.** Let g be a group, and let $a \in \mathcal{A}_{\text{sub}}^g$ be a suboptimal arm for g .

$$\sum_{t=1}^T \Pr(\text{Pull}_t^g(a), A_t^{\text{greedy}}(g) = a) = O(\log \log T).$$

759 Therefore, all of the regret stems from pulls of type (i), when an arm has the highest UCB. The next
760 result says that essentially all pulls occur when $H_t(\delta)$ is true:

761 **Proposition E.3.** Let $\delta > 0$. For any group g and action $a \in \mathcal{A}_{\text{sub}}^g$,

$$\sum_{t=1}^T \Pr(\text{Pull}_t^g(a), A_t^{\text{greedy}}(g) \neq a, \bar{H}_t(\delta)) = O(\log \log T).$$

762 Lastly, we show that the total number of times an arm $a \in \mathcal{A}_{\text{sub}}$ is pulled matches the lower bound:

763 **Proposition E.4.** Let $a \in \mathcal{A}_{\text{sub}}$.

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_T(a)]}{\log T} = J(a).$$

764 We now prove Theorem 4.1 using Propositions E.2-E.4.

765 *Proof of Theorem 4.1.* Fix a group g and an arm $a \in \mathcal{A}_{\text{sub}}^g$. Let $\varepsilon > 0$. Let $\delta \in (0, \delta_0)$ according to
766 Proposition E.1. Let $H_t = H_t(\delta)$.

$$\begin{aligned} \mathbb{E}[N_T^g(a)] &= \sum_{t=1}^T \Pr(\text{Pull}_t^g(a)) \\ &= \sum_{t=1}^T (\Pr(\text{Pull}_t^g(a), A_t^{\text{greedy}}(g) \neq a, H_t) \\ &\quad + \Pr(\text{Pull}_t^g(a), A_t^{\text{greedy}}(g) = a) + \Pr(\text{Pull}_t^g(a), A_t^{\text{greedy}}(g) \neq a, \bar{H}_t)) \\ (18) \quad &\leq \sum_{t=1}^T \Pr(\text{Pull}_t^g(a), a \in \mathcal{A}_t^{\text{UCB}}, H_t) + O(\log \log T). \end{aligned}$$

767 where the last step follows from Proposition E.3 and Proposition E.2.

768 First, assume that $a \notin \mathcal{A}_{\text{sub}}$. That is, there exists a group g' such that a is optimal for g' . We claim
769 that $\Pr(\text{Pull}_t^g(a) \mid a \in \mathcal{A}_t^{\text{UCB}}, H_t) = 0$. Notice that when H_t is true, a is not the greedy arm for g ,
770 and moreover, $a \notin \hat{\mathcal{A}}_{\text{sub}}$. Therefore, a is not involved in the optimization problem $(P(\theta))$, and a is
771 not the greedy arm for g , so g would not pull a when H_t is true. Therefore, $\text{Pull}_t^g(a) = 0$ when H_t is
772 true. This implies that if $a \notin \mathcal{A}_{\text{sub}}$,

$$(19) \quad \lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_T^g(a)]}{\log T} = 0.$$

773 Next, assume $a \in \mathcal{A}_{\text{sub}}$. By definition of the algorithm, if $\{\text{Pull}_t^g(a), a \in \mathcal{A}_t^{\text{UCB}}\}$ occurs, then
774 $N_t^g(a) \leq \hat{q}_t^g(a)N_t(a)$. If $H_t(\delta)$, then $\hat{q}_t^g(a) \leq q_t^g(a) + \varepsilon$. Therefore, $\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), a \in$

775 $\mathcal{A}_t^{\text{UCB}}, H_t(\delta)) \leq (q_t^g(a) + \varepsilon)N_T(a)$. Then, using (18), we can write

$$\begin{aligned} \limsup_{T \rightarrow \infty} \frac{\mathbb{E}[N_T^g(a)]}{\log T} &= \limsup_{T \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(\text{Pull}_t^g(a), a \in \mathcal{A}_t^{\text{UCB}}, H_t(\delta)) \right] + O(\log \log T)}{\log T} \\ &\leq \limsup_{T \rightarrow \infty} \frac{(q^g(a) + \varepsilon)\mathbb{E}[N_T(a)]}{\log T} \\ &\leq (q^g(a) + \varepsilon)J(a), \end{aligned}$$

776 where the last inequality follows from Proposition E.4. Since this holds for all $\varepsilon > 0$,

$$(20) \quad \limsup_{T \rightarrow \infty} \frac{\mathbb{E}[N_T^g(a)]}{\log T} \leq q^g(a)J(a).$$

777 Recall that Proposition E.4 states

$$(21) \quad \lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_T(a)]}{\log T} = J(a).$$

778 This implies that (20) must be an equality all g . If this weren't the case, then $\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[N_T(a)]}{\log T}$
779 would be strictly less than $J(a)$, which would be a contradiction.

780 Moreover, we claim that (20) and (21) implies $\lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_T^g(a)]}{\log T} = q^g(a)J(a)$ for all g . By contra-
781 diction, suppose there exists a $g' \in \mathcal{G}$ such that $\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[N_T^{g'}(a)]}{\log T} = q^{g'}(a)J(a) - \alpha$ for some
782 $\alpha > 0$. Then, (21) implies that $\limsup_{T \rightarrow \infty} \sum_{g \neq g'} \frac{\mathbb{E}[N_T^g(a)]}{\log T} \geq (1 - q^{g'}(a))J(a) + \alpha$, which is a
783 contradiction. Therefore, for every g ,

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}[N_T^g(a)]}{\log T} = q^g(a)J(a).$$

784 Combining with (19) yields the desired result:

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}[\text{Regret}_T^g(a)]}{\log T} = \lim_{T \rightarrow \infty} \frac{\sum_{a \in \mathcal{A}} \Delta^g(a) \mathbb{E}[N_T^g(a)]}{\log T} = \lim_{T \rightarrow \infty} \sum_{a \in \mathcal{A}_{\text{sub}}} \Delta^g(a) q^g(a) J(a).$$

785

□

786 E.1 Proof of Propositions E.2-E.4

787 *Proof of Proposition E.2.* Let $g \in \mathcal{G}$ and let $a \in \mathcal{A}_{\text{sub}}^g$. We bound $\sum_{t=1}^T \Pr(\text{Pull}_t^g(a), a =$
788 $A_t^{\text{greedy}}(g))$. We can assume that the events $\hat{\theta}_t(a) \in [\theta(a) - \delta, \theta(a) + \delta]$ and Λ_t occur using Lemma B.4,
789 and Lemma B.2 respectively. Since a is the greedy arm, it must be that $\hat{\theta}_t(a') \leq \theta(a) + \delta$ for all
790 $a' \in \mathcal{A}^g$.

791 Define the event

$$R_t = \{A_t^{\text{greedy}}(g) = a, \Lambda_t, \hat{\theta}_t(a) \leq \theta(a) + \delta, \hat{\theta}_t(a') \leq \theta(a) + \delta \forall a' \in \mathcal{A}^g\}.$$

792 Our goal is to bound $\sum_{t=1}^T \Pr(R_t)$.

793 For R_t to occur, $\hat{\theta}_t(a') \leq \theta(a) + \delta$ (since a is the greedy arm) and $\text{UCB}_t(a') \geq \text{OPT}(g)$ (since Λ_t)
794 for all $a' \in \mathcal{A}_{\text{opt}}^g$. By Lemma B.3 there exists a constant $c > 0$ such that if $N_t(a') > c \log t$ for some
795 $a' \in \mathcal{A}_{\text{opt}}^g$, R_t cannot happen. Moreover, for every $a' \in \mathcal{A}_{\text{opt}}^g$, $\Pr(N_t(a') < c \log t) < O\left(\frac{1}{\log t}\right)$
796 from Proposition C.6.

797 Divide the time period into epochs, where epoch k starts at time $s_k = 2^{2^k}$. Let \mathcal{T}_k be the time
798 steps in epoch k . Let $G_k = \{N_{s_k}(a) > 3c \log s_k \forall a \in \mathcal{A}_{\text{opt}}^g\}$ be the event that all optimal arms
799 were pulled at least $3c \log s_k$ times by the start of epoch k . If G_k occurs, since $s_k = \sqrt{s_{k+1}}$,

800 $N_{s_{k+1}}(a) > \frac{3}{2}r \log s_{k+1} > r \log s_{k+1}$, and hence R_t can never happen during epoch k . Moreover,
 801 $\Pr(\bar{G}_k) = O\left(\frac{1}{\log s_k}\right)$ for any k .

802 Suppose we are in a “bad epoch”, where G_k does not occur. We claim that R_t can’t occur more
 803 than $O(\log s_{k+1})$ times during epoch k . For R_t to occur, the arm j with the highest UCB satisfies
 804 $\text{UCB}_t(j) \geq \text{OPT}(g)$ and $\hat{\theta}_t(j) \leq \theta(a) + \delta$.

805 **Claim E.5.** For any action $j \in \mathcal{A}^g$, $\sum_{t=1}^s \Pr(A_t^{\text{UCB}}(g) = j, \text{UCB}_t(j) \geq \text{OPT}(g), \hat{\theta}_t(j) \leq \theta(a) +$
 806 $\delta \mid \bar{G}_k) = O(\log s)$.

807 Using Claim E.5 and taking a union bound over all actions j implies $\sum_{t \in \mathcal{T}_k} \Pr(R_t \mid \bar{G}_k) =$
 808 $\sum_{t \in \mathcal{T}_k} \sum_{j \in \mathcal{A}^g} \Pr(R_t, A_t^{\text{UCB}}(g) = j \mid \bar{G}_k) = O(\log s_{k+1})$. Since $\Pr(\bar{G}_k) = O\left(\frac{1}{\log s_k}\right)$,
 809 $\sum_{t \in \mathcal{T}_k} \Pr(R_t) = O(1)$. Since there are $O(\log \log T)$ epochs, $\sum_{t=1}^T \Pr(R_t) = O(\log \log T)$.

810 \square

811 *Proof of Proposition E.3.* Let $H_t = H_t(\delta)$. Fix a group g and an arm $a \in \mathcal{A}_{\text{sub}}^g$. For g to pull a
 812 when $A_t^{\text{greedy}}(g) \neq a$, it must be that $a \in \mathcal{A}_t^{\text{UCB}}$.

813 First, assume $a \notin \mathcal{A}_{\text{sub}}$. Then, there exist groups $G \subseteq \mathcal{G}$ in which a is optimal. If a is the greedy arm
 814 for some $g' \in G$, then $a \notin \hat{\mathcal{A}}_{\text{sub}}$, implying a is not considered in the optimization problem (\hat{P}_t) . In
 815 this case, group g would never pull arm a . Therefore, it must be that a is not the greedy arm for all
 816 groups in G . We show the following lemma, which proves the proposition for an arm $a \notin \mathcal{A}_{\text{sub}}$.

817 **Lemma E.6.** Let $a \notin \mathcal{A}_{\text{sub}}$, and let G be the set of groups in which a is optimal. Then,

$$\sum_{t=1}^T \Pr(\text{Pull}_t(a), A_t^{\text{greedy}}(g) \neq a \forall g \in G, a \in \mathcal{A}_t^{\text{UCB}}) = O(\log \log T).$$

818 Now assume $a \in \mathcal{A}_{\text{sub}}$. We assume that the events Λ_t and $\hat{\theta}_t(a) \in [\theta(a) - \delta, \theta(a) + \delta]$ hold using
 819 Lemma B.2 and Lemma B.4. Since $a \in \mathcal{A}_t^{\text{UCB}}$ and Λ_t , it must be that $\text{UCB}_t(a) \geq \text{OPT}(\Gamma(a))$. Let
 820 $E_t = \{\text{Pull}_t^g(a), \Lambda_t, \hat{\theta}_t(a) \in [\theta(a) - \delta, \theta(a) + \delta], \text{UCB}_t(a) \geq \text{OPT}(\Gamma(a))\}$ Our goal is to show

$$\mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(E_t, \bar{H}_t) \right] = O(\log \log T).$$

821 Divide the time interval into epochs, where epoch k starts at time $s_k = 2^{2^k}$. Let $K = O(\log \log T)$
 822 be the total number of epochs. Let \mathcal{T}_k be the time steps in epoch k .

823 Let $H_k = \cap_{t \in \mathcal{T}_k} H_t$. Clearly, if H_k is true, then by definition, $\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t, \bar{H}_t) = 0$. Therefore, we can
 824 write

$$\mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(E_t, \bar{H}_t) \right] = \sum_{k=1}^K \mathbb{E} \left[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t, \bar{H}_t) \right] = \sum_{k=1}^K \left(\mathbb{E} \left[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t, \bar{H}_t) \mid \bar{H}_k \right] \Pr(\bar{H}_k) \right)$$

825 We bound the expectation and the probability separately.

826 **1) Bounding** $\mathbb{E} \left[\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t, \bar{H}_t) \mid \bar{H}_k \right]$: If E_t occurs at some time step t , $\text{UCB}_t(a) \geq$
 827 $\text{OPT}(\Gamma(a))$ and $\hat{\theta}_t(a) \leq \theta(a) + \delta$. By Lemma B.3 it must be that $N_t(a) = O(\log t)$. Clearly,
 828 $N_s(a) \geq \sum_{t=1}^s \mathbf{1}(E_t)$, implying that $\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t) = O(\log s_{k+1})$. Therefore, $\sum_{t \in \mathcal{T}_k} \mathbf{1}(E_t, \bar{H}_t) \leq$
 829 $\sum_{t=1}^{s_{k+1}} \mathbf{1}(E_t) = O(\log s_{k+1})$

830 **2) Bounding** $\Pr(\bar{H}_k)$: For $a \in \mathcal{A}_{\text{sub}}$ let $c_a = \frac{0.9}{\text{KL}(\theta(a), \text{OPT}(\Gamma(a)))}$. For $a \notin \mathcal{A}_{\text{sub}}$, let $c_a = 1$. Let
 831 $F_k = \{\hat{\theta}_{s_k}(a) \in [\theta(a) - \delta/2, \theta(a) + \delta/2], N_{s_k}(a) \geq c_a \log s_k \forall a \in \mathcal{A}\}$ be the event that at time
 832 s_k , all arms a have been pulled $c_a \log s_k$ times and all arms are within an “inner” boundary (half as
 833 small as the boundary defined for H_t). We bound $\Pr(\bar{H}_k)$ by conditioning on the event F_k . Firstly,
 834 we bound $\Pr(\bar{F}_k)$ using the probabalistic lower bound of Proposition C.5-C.6:

835 **Lemma E.7.** For any k , $\Pr(\bar{F}_k) = O\left(\frac{1}{\log s_k}\right)$.

836 Next, we show that if F_k is true, then H_k occurs with probability at least $1 - O\left(\frac{1}{\log s_k}\right)$.

837 **Lemma E.8.** For any action a , $\Pr\left(\hat{\theta}_t(a) \notin [\theta(a) - \delta, \theta(a) + \delta] \text{ for some } t \in \mathcal{T}_k \mid F_k\right) \leq$
 838 $O\left(\frac{1}{\log s_k}\right)$.

839 Therefore,

$$\Pr(\bar{H}_k) \leq \Pr(\bar{F}_k) + \Pr(\bar{H}_k \mid F_k) = O\left(\frac{1}{\log s_k}\right).$$

840 **3) Combine:** Combining, we have

$$\begin{aligned} \mathbb{E}\left[\sum_{t=1}^T \mathbf{1}(E_t, \bar{H}_t)\right] &\leq \sum_{k=1}^K \left(O(\log s_{k+1}) O\left(\frac{1}{\log s_k}\right)\right) \\ &\leq \sum_{k=1}^K O(1) \\ &= O(\log \log T), \end{aligned}$$

841 where the last inequality follows due to the fact that $\frac{\log s_{k+1}}{\log s_k} = 2$ for any k . \square

842 *Proof of Proposition E.4.* Let $a \in \mathcal{A}_{\text{sub}}$. We need to show $\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[N_T(a)]}{\log T} \leq J(a)$, as the
 843 lower bound is implied by (4). By Proposition E.2, the number of times a is pulled when a is the
 844 greedy arm for some group g is $O(\log \log T)$. Therefore,

$$\mathbb{E}[N_T(a)] = \sum_{t=1}^T \Pr(\text{Pull}_t(a), a \in \mathcal{A}_t^{\text{UCB}}, H_t(\delta)) + O(\log \log T).$$

845 The rest of the proof relies on the same argument as Proposition C.2. The main idea is that after
 846 $J(a) \log T + o(\log T)$ pulls of a , the UCB of a will not be larger than $\text{OPT}(\Gamma(a))$, and therefore
 847 $a \notin \mathcal{A}_t^{\text{UCB}}$. \square

848 E.2 Deferred Proofs

849 *Proof of Claim E.5.* Recall that $G_k = \{N_{s_k}(a) > 3c \log s_k \ \forall a \in \mathcal{A}_{\text{opt}}^g\}$. We will show
 850 $\sum_{t=1}^T \Pr(A_t^{\text{UCB}} = j, \text{UCB}_t(j) \geq \text{OPT}(g), \hat{\theta}_t(j) \leq \theta(a) + \delta \mid \bar{G}_k) = O(\log \log T)$. From
 851 Lemma B.3, there exists a constant c' such that if $N_t(j) > c' \log T$ then, $\{\text{UCB}_t(j) \geq$
 852 $\text{OPT}(g), \hat{\theta}_t(j) \leq \theta(a) + \delta\}$ cannot occur.

$$\begin{aligned} &\sum_{t \in \mathcal{T}_k} \Pr(A_t^{\text{UCB}}(g) = j, \text{UCB}_t(j) \geq \text{OPT}(g), \hat{\theta}_t(j) \leq \theta(a) + \delta \mid \bar{G}_k) \\ &= \sum_{n=1}^{c' \log T} \sum_{t \in \mathcal{T}_k} \Pr(A_t^{\text{UCB}}(g) = j, \text{UCB}_t(j) \geq \text{OPT}(g), \hat{\theta}_t(j) \leq \theta(a) + \delta, N_t(a) = n \mid \bar{G}_k) \\ (22) \quad &\leq \sum_{n=1}^{c' \log T} \sum_{t \in \mathcal{T}_k} \Pr(A_t^{\text{UCB}}(g) = j, N_t(a) = n \mid \bar{G}_k). \end{aligned}$$

853 Our goal is to show that $\sum_{t \in \mathcal{T}_k} \Pr(A_t^{\text{UCB}}(g) = j, N_t(a) = n \mid \bar{G}_k) = O(1)$ for any n . Fix n , and
 854 write

$$\sum_{t \in \mathcal{T}_k} \Pr(A_t^{\text{UCB}}(g) = j, N_t(j) = n \mid \bar{G}_k) = \mathbb{E}\left[\sum_{t \in \mathcal{T}_k} \mathbf{1}(A_t^{\text{UCB}}(g) = j, N_t(j) = n) \mid \bar{G}_k\right]$$

855 Let $L_t = \mathbf{1}(A_t^{\text{UCB}}(g) = j, N_t(j) = n)$ be the indicator for the event of interest. Our goal is to count
856 the number of times L_t occurs. Let $Y_m = \{\exists t : \sum_{s=1}^t L_s = m\}$ be the event that L_s occurs at least
857 m times. Note that for Y_m to occur, it must be that Y_{m-1} occurred. Therefore, by expliciting writing
858 out the expectation, we have

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(A_t^{\text{UCB}}(g) = j, N_t(j) = n) \mid \bar{G}_k \right] &\leq \sum_{m \geq 1} m \Pr(Y_m \mid \bar{G}_k) \\ &= \sum_{m \geq 1} m \Pr(Y_m \mid Y_{m-1}, \bar{G}_k) \Pr(Y_{m-1} \mid \bar{G}_k). \end{aligned}$$

859 We claim that there exists a $\lambda \in (0, 1)$ such that $\Pr(Y_m \mid Y_{m-1}, \bar{G}_k) \leq \lambda$. Let τ be the time when
860 L_s occurred for the $m - 1$ 'th time, which exists since Y_{m-1} is true. For Y_m to occur, it must be that
861 arm j was not pulled at time τ , even though arm j is the UCB. Given that j is the UCB, there exists
862 a group g in which $N_\tau^g(a) \leq \hat{q}_\tau^g(a) N_\tau(a)$. If such a group arrives, it will pull j with probability at
863 least $\frac{1}{K}$. Therefore, at time τ , the probability that arm j will be pulled is at least $\min_{g \in G} \frac{p_g}{K}$. Then,
864 $\lambda = 1 - \min_{g \in G} \frac{p_g}{K}$ satisfies $\Pr(Y_m \mid Y_{m-1}, \bar{G}_k) \leq \lambda$.

865 Therefore,

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \mathbf{1}(A_t^{\text{UCB}} = j, N_t(j) = n) \mid \bar{G}_k \right] &= \sum_{m \geq 1} m \Pr(Y_m \mid Y_{m-1}, \bar{G}_k) \Pr(Y_{m-1} \mid \bar{G}_k) \\ &\leq \sum_{m \geq 1} m \lambda^m \\ &= O(1). \end{aligned}$$

866 Substituting back into (22) gives

$$\sum_{t=1}^T \Pr(A_t^{\text{UCB}} = j, \text{UCB}_t(j) \geq \text{OPT}(g), \hat{\theta}_t(j) \leq \theta(a) + \delta \mid \bar{G}_k) \leq \sum_{n=1}^{c' \log T} O(1) = O(\log T).$$

867 □

868 *Proof of Lemma E.6.* Let $a \notin \mathcal{A}_{\text{sub}}$, let G be the set of groups in which a is an optimal arm. We
869 condition on whether a is the UCB for some group in G .

870 First, suppose $a = A_t^{\text{UCB}}(g)$ for some group $g \in G$, implying $\theta(a) = \text{OPT}(g)$. We can assume
871 $\hat{\theta}_t(a) > \text{OPT}(g) - \delta$ from Lemma B.4. Then, if a is not the greedy arm for g , there exists a
872 suboptimal arm $j \in \mathcal{A}_{\text{sub}}^g$ with higher mean but lower UCB than a . This implies that the UCB radius
873 of j is smaller than the UCB radius of a , implying that j was pulled more times: $N_t(j) \geq N_t(a)$.
874 We show that this event cannot happen often. Let $E_t = \{\text{Pull}_t(a), A_t^{\text{greedy}}(g) \neq a, a \in \mathcal{A}_t^{\text{UCB}}, a =$
875 $A_t^{\text{UCB}}(g), \hat{\theta}_t(a) > \text{OPT}(g) - \delta\}$. For any $j \in \mathcal{A}_{\text{sub}}^g$,

$$\begin{aligned} &\sum_{t=1}^T \mathbf{1}(E_t, N_t(j) \geq N_t(a), \hat{\theta}_t(j) > \text{OPT}(g) - \delta) \\ &\leq \sum_{t=1}^T \sum_{n=1}^t \sum_{n_j=n}^t \mathbf{1}(E_t, \hat{\theta}_{n_j}(j) > \text{OPT}(g) - \delta, N_t(j) = n_j, N_t(a) = n) \\ &\leq \sum_{n_j=1}^T \mathbf{1}(\hat{\theta}_{n_j}(j) > \text{OPT}(g) - \delta) \sum_{n=1}^{n_j} \sum_{t=n}^T \mathbf{1}(E_t, N_t(a) = n) \\ &\leq \sum_{n_j=1}^T \mathbf{1}(\hat{\theta}_{n_j}(j) > \text{OPT}(g) - \delta) n_j, \end{aligned}$$

876 where the last inequality uses $\sum_{t=n}^T \mathbf{1}(E_t, N_t(a) = n) \leq 1$ (since pulling arm a increasing $N_t(a)$ by
877 1). Since $\Pr(\hat{\theta}_n(j) > \text{OPT}(g) - \delta) \leq \exp(-cn)$ for some constant $c > 0$, $\sum_{t=1}^T \Pr(E_t, N_t(j) \geq$
878 $N_t(a), \hat{\theta}_t(j) > \text{OPT}(g) - \delta) = O(1)$. Taking a union bound over actions $j \in \mathcal{A}_{\text{sub}}^g$ gives us the
879 desired result:

$$\sum_{t=1}^T \Pr(\text{Pull}_t(a), A_t^{\text{greedy}}(g) \neq a \forall g \in G, a \in A_t^{\text{UCB}}, \exists g \in G : a = A_t^{\text{UCB}}(g)) = O(\log \log T).$$

880 Now, suppose $a \notin A_t^{\text{UCB}}(g)$ for all $g \in G$. This means that there is another group h where
881 $a = A_t^{\text{UCB}}(h)$, but a is suboptimal for h . We assume Λ_t holds. Let a_h be an optimal arm for h . Since
882 $\Lambda_t, \text{UCB}_t(a_h) \geq \text{OPT}(h)$. Therefore, it must be that $\text{UCB}_t(a) \geq \text{OPT}(h)$. By Lemma C.4,

$$\sum_{t=1}^T \Pr(\text{Pull}_t(a), \text{UCB}_t(a) \geq \text{OPT}(h)) = O(\log \log T).$$

883 This finishes the proof. □

884 *Proof of Lemma E.7.* Fix $a \in \mathcal{A}$ and time t . We will show $\Pr(\hat{\theta}_{s_k}(a) \in [\theta(a) - \delta/2, \theta(a) +$
885 $\delta/2], N_{s_k}(a) \geq c_a \log s_k) \geq 1 - O\left(\frac{1}{\log t}\right)$. Then the result follows from taking a union bound over
886 actions. We first show that PF-UCB is log-consistent.

887 **Lemma E.9.** *PF-UCB is log-consistent.*

888 Let $g \in \Gamma(a)$. Since $\Pr(M_t(a) < \frac{p_a}{2}t) \leq \exp(-\frac{1}{2}p_g t)$, we can assume that there have been at least
889 $\frac{p_a}{2}t$ arrivals of g by time t . Then, using Proposition C.5 and Proposition C.6, we know that at time
890 t , $\Pr(N_t(a) < c_a \log t | M_t(a) \geq \frac{p_a}{2}t) \leq O\left(\frac{1}{\log t}\right)$. Next, we show that the probability of the event
891 $\hat{\theta}_t(a) \notin [\theta(a) - \delta/2, \theta(a) + \delta/2]$ given that we have more than $c_a \log t$ pulls of a is small.

$$\begin{aligned} & \Pr(\hat{\theta}_t(a) \notin [\theta(a) - \delta/2, \theta(a) + \delta/2] \mid N_t(a) \geq c_a \log t) \\ &= \sum_{n=c_a \log t}^t \Pr(\hat{\theta}_n(a) \notin [\theta(a) - \delta/2, \theta(a) + \delta/2] \mid N_t(a) = n) \Pr(N_t(a) = n) \\ &\leq \sum_{n=c_a \log t}^t \exp(-c_1 n) \Pr(N_t(a) = n) \\ &\leq c_3 \exp(-c_2 \log t) \\ &\leq \frac{c_3}{t^{c_2}}, \end{aligned}$$

892 for some constants $c_1, c_2, c_3 > 0$ that depends on the instance, a , and δ . Combining, we have that for
893 any action a , $\Pr(\hat{\theta}_{s_k}(a) \in [\theta(a) - \delta/2, \theta(a) + \delta/2], N_{s_k}(a) \geq c_a \log s_k) \geq 1 - O\left(\frac{1}{\log t}\right)$.

894 □

895 *Proof of Lemma E.8.* Let $U_a = \theta(a) + \delta$ and $U_a^I = \theta(a) + \delta/2$. Let $\eta = U_a - U_a^I$. Since F_k is
896 true, $N_{s_k}(a) \geq c_a \log s_k$. Let $n_1 = N_{s_k}(a)$. Let $\hat{\theta}^{n_1}(a)$ be the empirical average of arm a after n
897 pulls. We will bound

$$\Pr(\cup_{n_2=n_1+1}^{\infty} \{\hat{\theta}^{n_2}(a) \notin [L_a, U_a]\} \mid \hat{\theta}^{n_1}(a) \in [L_a^I, U_a^I]).$$

898 For any n_2 , $\hat{\theta}^{n_2}(a) > U_a$ implies $\hat{\theta}^{n_2}(a) > \hat{\theta}^{n_1}(a) + \eta$. Fix $n_2 > n_1$. Let $m = n_2 - n_1$.

$$\begin{aligned} \left\{ \hat{\theta}^{n_2}(a) > U_a \right\} &= \left\{ \sum_{i=1}^{n_2} X_i > n_2 U_a \right\} \\ &= \left\{ n_1 \hat{\theta}^{n_1}(a) + \sum_{i=n_1+1}^{n_2} X_i > n_2 U_a \right\} \\ &= \left\{ \sum_{j=1}^m X_{n_1+j} > n_1(U_a - \hat{\theta}^{n_1}(a)) + m U_a \right\} \\ &= \left\{ \sum_{j=1}^m (X_{n_1+j} - \mu) > n_1(U_a - \hat{\theta}^{n_1}(a)) + m(U_a - \mu) \right\} \end{aligned}$$

899 **Case $m \leq n_1$:** Since $U_a - \mu > 0$ and $U_a - \hat{\theta}^{n_1}(a) > \eta$ if F_k is true,

$$\begin{aligned} \Pr \left(\bigcup_{m=1}^{n_1} \{ \hat{\theta}^{n_1+m}(a) > U_a \} \mid F_k \right) &\leq \Pr \left(\bigcup_{m=1}^{n_1} \left\{ \sum_{j=1}^m (X_{n_1+j} - \mu) > n_1 \eta \right\} \mid F_k \right) \\ &\leq \Pr \left(\max_{m=1, \dots, n_1} S_m > n_1 \eta \mid F_k \right), \end{aligned}$$

900 where $S_m = \sum_{j=1}^m (X_{n_1+j} - \mu)$. Given that $X_{n_1+j} - \mu$ are zero mean independent random variables,
901 by Kolomogorov's inequality, we have

$$\begin{aligned} \Pr \left(\bigcup_{m=1}^{n_1} \{ \hat{\theta}^{n_1+m}(a) > U_a \} \mid F_k \right) &\leq \frac{1}{n_1^2 \eta^2} \text{Var}(S_{n_1}) \\ &= \frac{\sigma^2}{n_1 \eta^2} \\ &= \frac{\sigma^2}{\eta^2} \cdot \frac{1}{c_a \log s_k}, \end{aligned}$$

902 where $\sigma_2 = \text{Var}(X_1)$.

903 **Case $m > n_1$:**

$$\begin{aligned} \Pr \left(\bigcup_{m=n_1}^{\infty} \{ \hat{\theta}^{n_1+m}(a) > U_a \} \mid F_k \right) &\leq \Pr \left(\bigcup_{m=n_1}^{\infty} \left\{ \frac{\sum_{j=1}^m (X_{n_1+j} - \mu)}{m} > U_a - \mu \right\} \mid F_k \right) \\ &\leq \sum_{m=n_1}^{\infty} \Pr \left(\frac{\sum_{j=1}^m (X_{n_1+j} - \mu)}{m} > U_a - \mu \mid F_k \right) \\ &\leq \sum_{m=n_1}^{\infty} \exp(-mD) \\ &= \frac{\exp(-n_1 D)}{1 - \exp(-D)} \\ &= \frac{1}{s_k^{c_a D} (1 - \exp(-D))}, \end{aligned}$$

904 for a constant $D > 0$ that depends on $U_a - \mu$ and σ^2 .

905 Therefore,

$$\begin{aligned}
& \Pr \left(\bigcup_{m=1}^{\infty} \{ \hat{\theta}^{N_{s_k}(a)+m}(a) > U_a \} \mid F_k \right) \\
& \leq \Pr \left(\bigcup_{m=1}^{n_1} \{ \hat{\theta}^{N_{s_k}(a)+m}(a) > U_a \} \mid F_k \right) + \Pr \left(\bigcup_{m=n_1}^{\infty} \{ \hat{\theta}^{N_{s_k}(a)+m}(a) > U_a \} \mid F_k \right) \\
& \leq \frac{\sigma^2}{\eta^2} \cdot \frac{1}{c_a \log s_k} + \frac{1}{s_k^{c_a D} (1 - \exp(-D))} \\
& = O \left(\frac{1}{\log s_k} \right),
\end{aligned}$$

906 as desired. □

907 *Proof of Lemma E.9.* Fix a group g . At time t , if group g arrives, the PF-UCB pulls either the UCB
908 arm or the greedy arm. The original regret analysis of KL-UCB from [28] shows that

$$\sum_{t=1}^T \Pr(A_t \notin \mathcal{A}_{\text{opt}}^g, A_t = A_t^{\text{UCB}}, g_t = g) = O(\log T).$$

909 Proposition E.2 shows that the number of times the greedy arm is pulled and incurs regret is
910 $O(\log \log T)$. Combining, the total regret is $O(\log T)$. □

911 F Price of Fairness Proofs

912 F.1 Proof of Theorem 4.2

913 *Proof.* Consider the set of profiles $(s^g)_{g \in \mathcal{G}}$ that are in the feasible region of the polytope defined
914 by the constraints of $(P(\theta))$. Refer to this polytope as the “utility set”, in the language of [29]. This
915 utility set is compact and convex, and therefore we can apply Theorem 2 of [29], which gives us
916 the desired inequality. It is easy to see that the point in this utility set that maximizes total utility
917 corresponds to a regret-optimal policy, and the point in the utility set that maximizes proportional
918 fairness corresponds to PF-UCB (by definition, since PF-UCB maximizes proportional fairness within
919 this set). □

920 F.2 Proof of Proposition 4.3

921 *Proof.* In this proof, for convenience, we use subscripts instead of superscript to refer to groups g
922 since we do not need to refer to time steps.

923 Let $\{1, \dots, M\}$ be the set of shared arms, where $\theta_1 \leq \dots \leq \theta_M$. Let $\mathcal{G} = [G]$ be the set of
924 groups, where $\text{OPT}(1) \leq \dots \leq \text{OPT}(G)$. We assume that $\theta_M < \text{OPT}(1)$. (If there is a shared
925 arm whose reward is as large as $\text{OPT}(1)$, then neither policy will incur any regret from this arm,
926 and hence this arm is irrelevant.) In this case, all of the regret in the regret-optimal solution goes
927 to group 1, and the other groups incur no regret. Therefore, the total utility gain of the regret-
928 optimal solution is the sum of the regret at the disagreement point for groups 2 to G . Specifically,
929 $\lim_{T \rightarrow \infty} \text{SYSTEM}_T(\mathcal{I}) = \lim_{T \rightarrow \infty} \sum_{g=2}^G \frac{\tilde{R}_T^g(\pi^{\text{KL-UCB}})}{\log T}$.

930 We will show that for each group $g \geq 2$, the regret incurred from PF-UCB is less than half of the
931 regret at the disagreement point — i.e. $R_T^g(\pi^{\text{PF-UCB}}, \mathcal{I}) \leq \frac{1}{2} \tilde{R}_T^g(\mathcal{I})$. Then, the utility gain for the
932 group reduces by at most a half from the regret-optimal solution, which is our desired result.

933 Let $R_g = \lim_{T \rightarrow \infty} \frac{R_T^g(\pi^{\text{PF-UCB}}, \mathcal{I})}{\log T}$ and $\tilde{R}_g = \lim_{T \rightarrow \infty} \frac{\tilde{R}_T^g(\mathcal{I})}{\log T}$ for all $g \in \mathcal{G}$. Recall that the proportion-
 934 ally fair solution comes out of the optimal solution to the following optimization problem:

$$(P(\theta)) \quad \begin{aligned} & \max_{q \geq 0} \sum_{g \in \mathcal{G}} \log \left(\sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) (J^g(a) - q^g(a)J(a)) \right)^+ \\ & \text{s.t.} \quad \sum_{g \in \mathcal{G}} q^g(a) = 1 \quad \forall a \in \mathcal{A}_{\text{sub}} \\ & \quad \quad q^g(a) = 0 \quad \forall g \in G, a \notin \mathcal{A}_{\text{sub}} \cap \mathcal{A}_g. \end{aligned}$$

935 We first show a structural result of the optimal solution. Note that in terms of minimizing total regret,
 936 it is optimal for group 1 to pull all suboptimal arms. Therefore, if $q_g(a) > 0$ for some $g > 1$, we
 937 think of this as “transferring” pulls of arm a from group 1 to group g . This transfer increases the
 938 regret by a factor of $\frac{\Delta_g(a)}{\Delta_1(a)}$. We prove the following property that these transfers must satisfy:

939 **Claim F.1** (Structure of Optimal Solution). *For $g \in [M]$, let $b = \max\{a : q_g(a) > 0\}$. If $h < g$,
 940 then $q_h(a) = 0$ for all $a < b$.*

941 Writing out the KKT conditions of the optimization problem gives us the following result.

942 **Claim F.2** (KKT conditions). *Let $g, h \in \mathcal{G}$, $a \in \mathcal{A}$ such that $q_g(a) > 0$ and $h < g$. Then,
 943 $s_g \geq s_h \frac{\Delta_g(a)}{\Delta_h(a)}$. Moreover, if $q_1(a) > 0$, $s_g \leq \frac{\Delta_2(a)}{\Delta_1(a)} s_1$ for any $g > 1$.*

944 The next claim is immediate from Claim F.2.

945 **Claim F.3.** *If $h < g$ and there exists an arm a such that $q_g(a) > 0$, then $s_g \leq s_h$.*

946 Regret is minimized if $q_1(a) = 1$ for all a , in which case $s_1 = 0$. If $s_1 \neq 0$, then we think of this
 947 as pulls from group 1 that are re-allocated to other groups $g \neq 1$. This re-allocation increases total
 948 regret, since other groups incur more regret from pulling any arm compared to group 1.

949 Let $a_0 = \max\{a : q_g(a) \neq 1\}$. All pulls for any action $a > a_0$ come from group 1. We claim that
 950 $q_2(a_0) > 0$. Suppose not. Let $a' > 2$ such that $q_2(a_0) > 0$. Then, by Claim F.1, $q_2(a) = 0$ for
 951 all a . This implies that $s_2 = r_2 > r_{a'} \geq s_{a'}$, which contradicts Claim F.3. Then, by Claim F.2,
 952 $s_2 = s_1 \frac{\Delta_2(a_0)}{\Delta_1(a_0)}$.

953 Next, we claim that $s_2 \geq \frac{\tilde{R}_2}{2}$, which proves the desired result for $g = 2$. Note that s_1 represents the
 954 amount of regret that was “transferred” from group 1 to other groups, which increases the total regret.
 955 If all of this was transferred to group 2, the total regret from group 2 would be at most $s_1 \frac{\Delta_2(a_2)}{\Delta_1(a_2)} \leq s_2$.
 956 Therefore, $R_2 \leq s_2$. Since $R_2 + s_2 = \tilde{R}_2$, $s_2 \geq \frac{\tilde{R}_2}{2}$.

957 For $g > 2$, Claim F.2 shows $s_g \geq s_2$. Moreover, since $\text{OPT}(g) \geq \text{OPT}(2)$, $\tilde{R}_g \leq \tilde{R}_2$. Therefore,
 958 $s_g \geq s_2 \geq \frac{\tilde{R}_2}{2} \geq \frac{\tilde{R}_g}{2}$ as desired.

959 □

960 F.3 Proof of Claims

961 *Proof of Claim F.1.* Suppose not. Let $g \in \mathcal{G}$ and $b = \max\{a : q_g(a) > 0\}$. Let $a < b$ such that
 962 $q_h(a) > 0$. Then, since $\sum_{g'} q_{g'}(a) = 1$, $q_g(a) < 1$. By the ordering of arms and groups, we have

$$(23) \quad \frac{\Delta_h(a)}{\Delta_g(a)} > \frac{\Delta_h(b)}{\Delta_g(b)}.$$

963 We essentially show, using this inequality, that if we want to “transfer” pulls from group h to g , it
 964 is more efficient to do so using arm a rather than arm b , and hence it is a contradiction that $q_h(b)$ is
 965 positive.

966 We construct a “swap” that will strictly increase the objective function. Let $\varepsilon = \min\{q_h(a), q_g(b), 1 -$
 967 $q_g(a), 1 - q_h(b)\}$.

- 968 • Decrease $q_h(a)$ by ε , and increase $q_h(b)$ by $\frac{\Delta_h(a)J(a)}{\Delta_h(b)J(b)}\varepsilon \leq \varepsilon$, where the last inequality
969 follows from the convexity of $\text{KL}(\theta_b, \cdot)$. By construction, s_h does not change.
- 970 • Increase $q_g(a)$ by ε , and decrease $q_g(b)$ by $\frac{\Delta_h(a)J(a)}{\Delta_h(b)J(b)}\varepsilon$. The first operation decreases s_g
971 by $\Delta_g(a)J(a)\varepsilon$, while the second operation increases s_g by $\frac{\Delta_h(a)J(a)\Delta_g(b)}{\Delta_h(b)}\varepsilon$. By (23), this
972 strictly increases s_g overall.

973 This is a contradiction. □

974 *Proof of Claim F.2.* From the stationarity KKT condition, we have that

$$\begin{aligned} \frac{\Delta_g(a)J(a)}{s_g} + \lambda(a) - \mu_g(a) &= 0, \\ \frac{\Delta_h(a)J(a)}{s_h} + \lambda(a) - \mu_h(a) &= 0, \end{aligned}$$

975 for some $\lambda_a \in \mathbb{R}$ and $\mu_g(a), \mu_h(a) \geq 0$. From complementary slackness, $\mu_g(a)q_g(a) = 0$. Since
976 $q_g(a) > 0$, it must be that $\mu_g(a) = 0$. Since $\mu_h(a) \geq 0$, $\frac{\Delta_g(a)J(a)}{s_g} \leq \frac{\Delta_h(a)J(a)}{s_h}$. □

977 G Other Proofs

978 G.1 Proof that Nash Solution is Unique Under Grouped Bandit Model

979 The uniqueness of the Nash bargaining solution in the general bargaining problem requires that the set
980 U is convex. In the grouped bandit model, it is not clear that the set $U(\mathcal{I}) = \{(\text{UtilGain}^g(\pi, \mathcal{I}))_{g \in \mathcal{G}} : \pi \in \Psi\}$
981 is convex. In this section, we show that the uniqueness theorem still holds in the grouped
982 bandit setting.

983 Let G be the number of groups. Let $W(u) = \sum_{g \in \mathcal{G}} \log u_g$, and let $f(U) = \arg\max_{u \in U} W(u)$ for
984 $U \subseteq \mathbb{R}^G$. Fix a grouped bandit instance \mathcal{I} , and let $u^* = f(U(\mathcal{I}))$. We first show that u^* is unique (i.e.
985 $\arg\max_{u \in U(\mathcal{I})} W(u)$ is unique). Suppose there was another $u' \in U(\mathcal{I})$ with the same welfare. Then,
986 let $\bar{u} \in U(\mathcal{I})$ be the policy that runs u' with probability 50%, and u^* with probability 50%. Using
987 the fact that $\liminf_{T \rightarrow \infty} (a_T + b_T) \geq \liminf_{T \rightarrow \infty} a_T + \liminf_{T \rightarrow \infty} b_T$ implies that $\bar{u}_g \geq \frac{1}{2}(u_g^* + u'_g)$
988 for all g . Since \log is strictly concave, $\log \bar{u}_g > \frac{1}{2}(\log u_g^* + \log u'_g)$. This implies $W(\bar{u}) > W(u^*)$,
989 which is a contradiction.

990 Next, we show that f is the unique solution that satisfies the four axioms. Let $U = U(\mathcal{I})$. It is easy
991 to see that this solution satisfies the axioms. We need to show that no other solution satisfies them.
992 Suppose $g(\cdot)$ satisfies the axioms. We need to show $g(U) = f(U)$. Let $U' = \{(\alpha_g u_g)_{g \in \mathcal{G}} : u \in U; \alpha_g u_g^* = 1, \alpha_g > 0\}$. U' is the translated utility set so that u^* becomes the $\mathbf{1}$ vector. Then, the
993 optimal welfare is $W(\mathbf{1}) = 0$. We need to show $g(U') = \mathbf{1}$. We claim that there is no $v \in U'$ such
994 that $\sum_{g \in \mathcal{G}} v_g > G$. Assume that such a v exists. For $\lambda \in (0, 1)$, let t be the utilities from the policy
995 that runs the policy induced by v with probability λ , and the policy induced by $\mathbf{1}$ with probability
996 $1 - \lambda$. Then, by the same argument with \liminf to prove uniqueness, $t_g \geq \lambda v_g + (1 - \lambda)1$. If λ is
997 small enough, then $\sum_{g \in \mathcal{G}} \log t_g > 0$. This is a contradiction to $\mathbf{1}$ maximizing $W(\cdot)$.

999 Consider the symmetric set $U'' = \{u \in \mathbb{R}^G : u \geq 0, \sum_g u_g \leq G\}$. We have shown that $U' \subseteq U''$.
1000 By Pareto efficiency and symmetry, it must be that $g(U'') = \mathbf{1}$. By independence of irrelevant
1001 alternatives, $g(U') = \mathbf{1}$, and we are done.

1002 G.2 Proof that Assumption 2.2 is Sufficient

1003 **Proposition G.1.** *If an instance \mathcal{I} satisfies Assumption 2.2, then there exists a consistent policy π*
1004 *such that $f(\pi) > -\infty$. Otherwise, $f(\pi) = -\infty$ for all $\pi \in \Psi$.*

1005 *Proof.* First, suppose \mathcal{I} satisfies Assumption 2.2. We need to show that there exists a consistent
1006 policy such that $f(\pi) > -\infty$. We will construct a feasible solution to the optimization problem
1007 $(P(\theta))$ with a strictly positive objective value. This will imply that the objective value Y^* is strictly
1008 larger than 0, and hence the social welfare of PF-UCB is higher than $-\infty$.

1009 For each arm $a \in \mathcal{A}$, let $g(a) \in \Gamma(a)$. Start with $q^{g(a)}(a) = 1$ for all a and $q^g(a) = 0$ for $g \neq g(a)$.
 1010 We will modify these values for suboptimal arms \mathcal{A}_{sub} . For arm $a \in \mathcal{A}_{\text{sub}}$, let $g'(a) \neq g(a)$ be
 1011 another group with access to arm a . We will “split” the pulls of arm a between groups $g(a)$ and
 1012 $g'(a)$ in a way that both groups benefit from the disagreement point. Let $p(a) \in [0, 1]$ such that
 1013 $p(a)J(a) = J^{g'(a)}(a)$. Let $q^{g'(a)} = p(a)/2$ and $q^{g(a)} = 1 - p(a)/2$. Then, $J^g(a) - q^g(a)J(a) > 0$
 1014 for $g \in \{g(a), g'(a)\}$. This implies that $s^g > 0$ for all g , and therefore $Y^* > 0$. This proves the first
 1015 part of the proposition.

1016 For the second statement, suppose \mathcal{I} does not satisfy Assumption 2.2. Let g' be the group that
 1017 does not have a suboptimal arm that is shared with another group. First, suppose g' does not have
 1018 any suboptimal arms. Then, all arms available to group g' is optimal, so group g' will incur zero
 1019 regret regardless of the algorithm. Hence, the utility gain for group g' is exactly 0, and therefore
 1020 $W(\pi, \mathcal{I}) = -\infty$ for any π .

1021 Next, suppose g' does have a suboptimal arm but it is not shared. Let π be a consistent policy. Then
 1022 from the following upper bound on Nash SW from Section 3.2,

$$W(\pi, \mathcal{I}) \leq \liminf_{T \rightarrow \infty} \sum_{g \in \mathcal{G}} \log \left(\sum_{a \in \mathcal{A}^g} \Delta^g(a) (J^g(a) - q_T^g(a, \pi)J(a)) \right)^+.$$

1023 Since g' is the only group with access to arm a for every $a \in \mathcal{A}_{\text{sub}}^{g'}$, it must be that $q_T^{g'}(a, \pi) = 1$
 1024 for every $a \in \mathcal{A}_{\text{sub}}^{g'}$. Moreover, $J^{g'}(a) = J(a)$ for every $a \in \mathcal{A}_{\text{sub}}^{g'}$. This implies that the term
 1025 corresponding to g' in the sum equals $\log 0 = -\infty$. Therefore, $W(\pi, \mathcal{I}) = -\infty$ for any $\pi \in \Psi$. \square

1026 G.3 Omitted Details of Theorem 3.2

1027 We provide details on the two steps in Section 3.2 starting from (9). (4) implies that for every $\varepsilon > 0$,
 1028 there exists a T_ε such that if $T \geq T_\varepsilon$, then

$$\frac{\mathbb{E}[N_T(a)]}{\log T} \geq (1 - \varepsilon)J(a).$$

1029 Therefore, for large enough T , plugging into (9), we get

$$\frac{R_T^g(\pi, \mathcal{I})}{\log T} \geq \sum_{a \in \mathcal{A}_{\text{sub}}} \Delta^g(a) q_T^g(a, \pi) J(a) (1 - \varepsilon).$$

1030 This implies that

$$\limsup_{T \rightarrow \infty} \frac{R_T^g(\pi, \mathcal{I})}{\log T} \geq \limsup_{T \rightarrow \infty} (1 - \varepsilon) \sum_{a \in \mathcal{A}_{\text{sub}}} \Delta^g(a) q_T^g(a, \pi) J(a).$$

1031 Since this holds for every $\varepsilon > 0$ and the RHS is continuous in ε ,

$$(24) \quad \limsup_{T \rightarrow \infty} \frac{R_T^g(\pi, \mathcal{I})}{\log T} \geq \limsup_{T \rightarrow \infty} \sum_{a \in \mathcal{A}_{\text{sub}}} \Delta^g(a) q_T^g(a, \pi) J(a).$$

1032 Plugging in (24) into the definition of $\text{UtilGain}^g(\pi, \mathcal{I})$ gives

$$\text{UtilGain}^g(\pi, \mathcal{I}) \leq \liminf_{T \rightarrow \infty} \sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) (J^g(a) - q_T^g(a, \pi)J(a) \mathbf{1}\{a \in \mathcal{A}_{\text{sub}}\}).$$

1033 Using the definition of $W(\pi, \mathcal{I})$ and taking the \liminf outside of the sum gives

$$W(\pi, \mathcal{I}) \leq \liminf_{T \rightarrow \infty} \sum_{g \in \mathcal{G}} \log \left(\sum_{a \in \mathcal{A}_{\text{sub}}^g} \Delta^g(a) (J^g(a) - q_T^g(a, \pi)J(a) \mathbf{1}\{a \in \mathcal{A}_{\text{sub}}\}) \right)^+.$$

1034 G.4 Proof of Proposition E.1

1035 *Proof.* First, we prove the statement with respect to the variables $(s^g)_{g \in \mathcal{G}}$. Let $f_s(s) = \sum_{g \in \mathcal{G}} \log s^g$,
 1036 and let $s_*^g = \sum_{a \in \mathcal{A}^g} \Delta^g(a) (J^g(a) - q_*^g(a)J(a))$ and $\hat{s}_t^g = \sum_{a \in \mathcal{A}^g} \hat{\Delta}^g(a) (\hat{J}^g(a) - \hat{q}_t^g(a)\hat{J}(a))$.

1037 Since f_s is strictly concave with respect to s , s_*^g is unique. Define the event $H_t(\delta) = \{\hat{\theta}_t(a) \in$
 1038 $[\theta(a) - \delta, \theta(a) + \delta]\}$ for all $a \in \mathcal{A}$.

1039 **Lemma G.2.** For any $\varepsilon > 0$, there exists $\delta > 0$ such that if $H_t(\delta)$, then $\hat{s}_t^g \in [s_*^g - \varepsilon, s_*^g + \varepsilon]$ for all
1040 $g \in \mathcal{G}$.

1041 This shows that if $H_t(\delta)$, then the variables \hat{s}_t^g are close to s_*^g for all g . Next, we need to show that
1042 the corresponding q 's are also close. Let $\text{proj}(z, P)$ be the projection of point z onto a polytope P .

1043 Let $Q = \{q : \sum_{g \in G} q^g(a) = 1 \ \forall a \in \mathcal{A}_{\text{sub}}, q^g(a) = 0 \ \forall g \in G, a \notin \mathcal{A}_{\text{sub}}, q^g(a) \geq 0 \ \forall g \in$
1044 $G, a \in \mathcal{A}\}$ be the feasible space. Let $S^g(q, \tilde{\theta}) = \sum_{a \in \mathcal{A}^g} \tilde{\Delta}^g(a) (\tilde{J}^g(a) - q^g(a) \tilde{J}(a))$, where
1045 $\tilde{\Delta}^g(a)$, $\tilde{J}^g(a)$, and $\tilde{J}(a)$ are computed with $\tilde{\theta}$.

1046 Given $s = (s^g)_{g \in \mathcal{G}}$, let $Q(s, \tilde{\theta}) = \{q^g(a) \in Q : S^g(q, \tilde{\theta}) = s^g\}$ be the set of all feasible q 's that
1047 corresponds to the solution s under the parameters $\tilde{\theta}$. Note that $Q(s, \tilde{\theta})$ is a linear polytope, and we
1048 can write it as $Q(s, \tilde{\theta}) = \{q : A(\tilde{\theta})q = b(s), q \geq 0\}$ for a matrix $A(\tilde{\theta})$ and a vector $b(s)$. We are
1049 interested in the polytopes $Q(s, \theta)$ and $Q(\hat{s}_t, \hat{\theta}_t)$, which correspond the optimal solutions of $(P(\theta))$
1050 and (\hat{P}_t) respectively. The next two lemmas state that these polytopes are close together:

1051 **Lemma G.3.** Let $\varepsilon > 0$. There exists $\delta > 0$ such that if $H_t(\delta)$, for any $\hat{q} \in Q(\hat{s}_t, \hat{\theta}_t)$,
1052 $\|\text{proj}(\hat{q}, Q(s, \theta)) - \hat{q}\|_2 \leq \varepsilon$.

1053 **Lemma G.4.** Let $\varepsilon > 0$. There exists $\delta > 0$ such that if $H_t(\delta)$, for any $q \in Q(s, \theta)$,
1054 $\|\text{proj}(q, Q(\hat{s}_t, \hat{\theta}_t)) - q\|_2 \leq \varepsilon$.

1055 Let $q_* = \text{argmin}_{q \in Q(s, \theta)} \|q\|_2^2$, $\hat{q} = \text{argmin}_{q \in Q(\hat{s}_t, \hat{\theta}_t)} \|q\|_2^2$. Our goal is to show $\|q_* - \hat{q}\|_1 \leq \varepsilon$.
1056 Let $R(\eta) = \{q \in Q(s, \theta) : \|q\|_2 \leq \|q_*\|_2 + \eta\}$ for $\eta > 0$. Since the function $\|\cdot\|_2^2$ is strongly
1057 convex and q_* is minimizer, we have the following result:

1058 **Claim G.5.** For every $\varepsilon > 0$, there exists $\eta > 0$ such that if $q \in R(\eta)$, then $\|q - q_*\|_2 \leq \varepsilon$.

1059 First, assume $\|\hat{q}_t\|_2 \leq \|q_*\|_2$. Let $\eta > 0$ be from Claim G.5 using $\varepsilon = \frac{\varepsilon}{2}$. Let $\delta > 0$ be from
1060 Lemma G.3 using $\varepsilon = \min\{\frac{\varepsilon}{2}, \eta\}$. Let $q' = \text{proj}(\hat{q}, Q(s, \theta)) \in Q(s, \theta)$. From Lemma G.3,
1061 $\|\hat{q}_t - q'\|_2 \leq \eta$, implying $\|q'\|_2 \leq \|\hat{q}_t\|_2 + \eta \leq \|q_*\|_2 + \eta$. Therefore, $q' \in R(\eta)$. Claim G.5
1062 implies $\|q' - q_*\|_2 \leq \frac{\varepsilon}{2}$. Let $\delta > 0$ correspond to $\frac{\varepsilon}{2}$ from Lemma G.3, so that $\|\hat{q}_t - q'\|_2 \leq \frac{\varepsilon}{2}$. Then,

$$\|\hat{q}_t - q_*\|_2 \leq \|\hat{q}_t - q'\|_2 + \|q' - q_*\|_2 \leq \varepsilon.$$

1063 An analogous argument shows the same result in the case that $\|q_*\|_2 \leq \|\hat{q}_t\|_2$ using Lemma G.4.

1064 □

1065 G.4.1 Proof of Lemmas

1066 We first state an additional lemma:

1067 **Lemma G.6.** For any $\varepsilon > 0$ there exists a $\delta > 0$ such that if $H_t(\delta)$, then for any feasible solution q ,
1068 $|f(q) - \hat{f}(q)| < \varepsilon$.

1069 *Proof of Lemma G.6.* Let q be a feasible solution. Let $S^g(q, \tilde{\theta}) =$
1070 $\sum_{a \in \mathcal{A}^g} \tilde{\Delta}^g(a) (\tilde{J}^g(a) - q^g(a) \tilde{J}(a))$, where $\tilde{\Delta}^g(a)$, $\tilde{J}^g(a)$, and $\tilde{J}(a)$ are computed with
1071 $\tilde{\theta}$.

1072 For each g , let $\varepsilon_g > 0$ be such that if $|\tilde{s}^g - s_*^g| \leq \varepsilon_g$, then $|\log s_*^g - \log \tilde{s}^g| \leq \frac{\varepsilon}{G}$. $\Delta^g(a)$, $J^g(a)$, and
1073 $J(a)$ are all differentiable functions of θ with finite derivatives around θ_* . Then, it is possible to find
1074 $\delta_g > 0$ such that if $H_t(\delta_g)$, $|\hat{\Delta}^g(a) (\hat{J}^g(a) - q^g(a) \hat{J}(a)) - \Delta^g(a) (J^g(a) - q^g(a) J(a))| \leq \frac{\varepsilon_g}{|\mathcal{A}|}$.
1075 Summing over actions, $|S^g(q, \hat{\theta}_t) - S^g(q, \hat{\theta})| \leq \varepsilon_g$. Then, if $H_t(\delta_g)$, $|\log S^g(q, \hat{\theta}_t) - \log S^g(q, \hat{\theta})| \leq$
1076 $\frac{\varepsilon_g}{G}$. Take $\delta = \min_{g \in \mathcal{G}} \delta_g$. If $H_t(\delta)$ is true, $|f(q) - \hat{f}(q)| < \varepsilon$. □

1077 *Proof of Lemma G.2.* Let $\varepsilon > 0$. Let $S_\varepsilon = \{s : |s^g - s_*^g| \leq \varepsilon \ \forall g\}$ be the set around s_* of interest.
1078 Our goal is to show that $f_s(\hat{s}) \in S_\varepsilon$. Let $f_{\text{bd}} = \max\{f(s) : s \in \text{bd}(S_\varepsilon)\} < f^*$ be the largest f on
1079 the boundary of S_ε . Then, if $f_s(s) > f_{\text{bd}}$, it must be that $s \in S_\varepsilon$. (Since the entire line between
1080 s and s_* must have a value of f_s that is higher than $f_s(s)$ due to concavity, and it must cross the

1081 boundary.) Therefore, we need to show $f_s(\hat{s}_t) > f_{\text{bd}}$. Let \hat{q}_t be the corresponding solution to \hat{s}_t .
 1082 Then, $f_s(\hat{s}_t) = \hat{f}_t(\hat{q}_t)$. Let $\delta > 0$ as in Lemma G.6 with $\varepsilon = f^* - f_{\text{bd}}$. Then, if $H_t(\delta)$ is true,

$$f_s(\hat{s}_t) = \hat{f}_t(\hat{q}_t) \geq \hat{f}_t(q_*) \geq f(q_*) - (f^* - f_{\text{bd}}) = f_{\text{bd}},$$

1083 where the second inequality follows from Lemma G.6.

1084 □

1085 *Proof of Lemma G.3.* Let $\varepsilon > 0$. Let n be the dimension of q . We will make use of the following
 1086 closed form formula for the projection onto a linear subspace:

1087 **Fact G.7.** Let $P = \{x : Ax = b\}$. The orthogonal projection of z onto P is $\text{proj}(z, P) =$
 1088 $z - A^\top(AA^\top)^{-1}(Az - b)$.

1089 Let $Q = Q(s, \tilde{\theta})$, and let A, b be the corresponding parameters of the linear constraints; i.e. $Q = \{x :$
 1090 $Ax = b, x \geq 0\}$. Similarly, let $\hat{Q} = Q(\hat{s}_t, \hat{\theta}_t)$, and let \hat{A}, \hat{b} be defined similarly. Note that Fact G.7
 1091 only works with equality constraints.

1092 We define a distance between two linear polytopes. We use the notation $P(D, f) = \{x : Dx = f\}$.
 1093 Then, $Q = P(A, b)$, $\hat{Q} = P(\hat{A}, \hat{b})$.

1094 **Definition G.8.** For two polytopes $P(A, b)$ and $P(A', b')$, the distance is defined as
 1095 $d(P(A, b), P(A', b')) = \max\{\|A - A'\|_2, \|b - b'\|_2\}$.

1096 Note that for every $\alpha > 0$, there exists $\delta > 0$ such that $H_t(\delta)$ implies $d(Q, \hat{Q}) \leq \alpha$ using Lemma G.2.
 1097 For any $\mathcal{I} \in 2^{[n]}$, let $P_{\mathcal{I}} = P(A_{\mathcal{I}}, b_{\mathcal{I}}) = \{x : Ax = b, x_i = 0 \forall i \in \mathcal{I}\}$.

1098 **Claim G.9.** *There exists a constant $C \geq 1$ such that for any $\mathcal{I} \in 2^{[n]}$ and any \tilde{A}, \tilde{b} of same dimensions
 1099 as $A_{\mathcal{I}}, b_{\mathcal{I}}$, if $\tilde{q} \in P(\tilde{A}, \tilde{b})$ with $\tilde{q} \leq 1$ (for all elements), then $\|\tilde{q} - \text{proj}(\tilde{q}, P_{\mathcal{I}})\|_2 \leq Cd(P_{\mathcal{I}}, P(\tilde{A}, \tilde{b}))$.*

1100 *Proof of Claim G.9.* From Fact G.7, we have $\|\tilde{q} - \text{proj}(\tilde{q}, P_{\mathcal{I}})\|_2 = \|A_{\mathcal{I}}^\top(A_{\mathcal{I}}A_{\mathcal{I}}^\top)^{-1}(A_{\mathcal{I}}\tilde{q} - b_{\mathcal{I}})\|_2$.
 1101 Since $\tilde{q} \in P(\tilde{A}, \tilde{b})$, $\tilde{A}\tilde{q} = \tilde{b}$. Let $\lambda = \max_{\mathcal{I}} \|A_{\mathcal{I}}^\top(A_{\mathcal{I}}A_{\mathcal{I}}^\top)^{-1}\|_2$ and let $d = d(P_{\mathcal{I}}, P(\tilde{A}, \tilde{b}))$.
 1102 Therefore,

$$\begin{aligned} \|\tilde{q} - \text{proj}(\tilde{q}, P_{\mathcal{I}})\|_2 &\leq \lambda \|(A_{\mathcal{I}} - \tilde{A})\tilde{q} + (\tilde{b} - b_{\mathcal{I}})\|_2 \\ &\leq \lambda \left(\|A_{\mathcal{I}} - \tilde{A}\|_2 \|\tilde{q}\|_2 + \|\tilde{b} - b_{\mathcal{I}}\|_2 \right) \\ &\leq 2\lambda nd. \end{aligned}$$

1103 Therefore, $C = 2\lambda n$. □

1104 We now describe an iterative process to prove this result.

1105 Let $Q^0 = \{q : Aq = b\}$ (Q without the non-negativity constraint), and same with $\hat{Q}^0 = \{q : \hat{A}q = \hat{b}\}$.
 1106 Let $\alpha_0 = d(Q^0, \hat{Q}^0)$. Let $\tilde{q}^0 = \text{proj}(\hat{q}, Q^0)$. By Claim G.9, $\|\hat{q} - \tilde{q}^0\|_2 \leq C\alpha_0$. If $\tilde{q}^0 \geq 0$, then
 1107 STOP here.

1108 Otherwise, find an index i which violates the non-negativity constraint using the following method:

- 1109 • Let $q \in Q$ be an arbitrary feasible point ($q \geq 0$).
- 1110 • From the point \tilde{q}^0 , move along the direction towards q . Let p^0 be the first point on this line
 1111 where p^0 is non-negative.
- 1112 • Since Q is simply Q^0 with non-negativity constraints and both sets are convex, $p^0 \in Q$.
- 1113 • Let i be an index where $\tilde{q}_i^0 < 0$ and $p_i^0 = 0$ (the last index to become non-negative).

1114 Since $\hat{q} \geq 0$, it must be that $\hat{q}_i \leq C\alpha_0$ since $\|\hat{q}^0 - \hat{q}\| \leq C\alpha_0$.

1115 Let Q^1 be the same polytope as Q^0 , but with the additional constraint that $q_i = 0$ — call this constraint
 1116 C . Let A^1, b^1 be the corresponding equality constraints for Q^1 . Let \hat{Q}^1 be the same polytope as \hat{Q}^0 , but
 1117 with the additional equality constraint that $q_i = \hat{q}_i$ — call this constraint \hat{C} . Let \hat{A}^1, \hat{b}^1 be the equality
 1118 constraints for \hat{Q}^1 . Note that the only difference between constraints C and \hat{C} is the right hand side,

1119 which differ by at most $C\alpha_0$. Therefore, $d(Q^1, \hat{Q}^1) \leq d(Q^0, \hat{Q}^0) + C\alpha_0 \leq 2C\alpha_0$. Clearly, $\hat{q} \in \hat{Q}^1$.
 1120 Let $\tilde{q}^1 = \text{proj}(\hat{q}, Q^1)$. Applying Claim G.9 again, we have $\|\hat{q} - \tilde{q}^1\|_2 \leq C(2C\alpha_0) = 2C^2\alpha_0$. If
 1121 $\tilde{q}^1 \geq 0$, then STOP here.

1122 Otherwise, let j be the index which violates the non-negativity constraint found using the same
 1123 method as before; except this time, we draw a line between \tilde{q}^1 towards $p^0 \in Q$. We let p^1 be the first
 1124 point where $p^1 \geq 0$. Then, we repeat the above process. We define Q^2 to be the same polytope as Q^1 ,
 1125 with the additional constraint that $q_j = 0$. \hat{Q}^2 is defined as \hat{Q}^1 with the additional constraint $q_j = \hat{q}_j$.
 1126 Then, $\hat{q}_j \leq 2C^2\alpha_0$. Therefore, $d(Q^2, \hat{Q}^2) \leq d(Q^1, \hat{Q}^1) + 2C^2\alpha_0 \leq 2C\alpha_0 + 2C^2\alpha_0 \leq 4C^2\alpha_0$.
 1127 Applying Claim G.9, we get $\|\hat{q} - \tilde{q}^2\|_2 \leq C(4C^2\alpha_0) = 4C^3\alpha_0$. If $\tilde{q}^2 \geq 0$, then STOP here.

1128 **After stopping:** If this process stopped at iteration m , then $\tilde{q}^m \in Q$ and $\|\hat{q} - \tilde{q}^m\|_2 \leq 2^m C^{m-1} \alpha_0$.
 1129 It must be that $m \leq n$. If $\alpha_0 = \frac{\varepsilon}{2^m C^{n-1}}$, then $\|\hat{q} - \tilde{q}^m\|_2 \leq \varepsilon$. Then, $\|\text{proj}(\hat{q}, Q) - \hat{q}\|_2 \leq \varepsilon$. Let
 1130 $\delta > 0$ such that $H_t(\delta)$ implies $d(Q, \hat{Q}) \leq \alpha_0$. \square

1131 *Proof of Lemma G.4.* This proof follows essentially the same steps as the proof of Lemma G.3 by
 1132 swapping Q and \hat{Q} . The main difference is that we are projecting q onto $Q(\hat{s}_t, \hat{\theta}_t)$, but this must hold
 1133 for all possible values of $\hat{s}_t, \hat{\theta}_t$ (using a single δ). Due to this, the only thing we have to change from
 1134 the proof of Lemma G.3 is Claim G.9. We must show that there exists a constant C where Claim G.9
 1135 is satisfied for all possible values of $\hat{s}_t, \hat{\theta}_t$. The only place where C relies on a property of the polytope
 1136 $P_{\mathcal{I}}$ is in choosing λ . Therefore our goal is to uniformly upper bound $\max_{\mathcal{I}} \|\hat{A}_{\mathcal{I}}^{\top} (\hat{A}_{\mathcal{I}} \hat{A}_{\mathcal{I}}^{\top})^{-1}\|_2$ for
 1137 all possible $\hat{A}_{\mathcal{I}}$ that can be induced by all possible $\hat{s}_t, \hat{\theta}_t$.

1138 Note that since we assume that $H_t(\delta_0)$ holds, the possible matrices \hat{A} lie in a compact space (since
 1139 every element of the matrix \hat{A} can be at most δ_0 apart). Since $\|A^{\top} (AA^{\top})^{-1}\|_2$ is a continuous
 1140 function of the elements of the matrix A , $\lambda_1 = \max_{\hat{A}} \|\hat{A}^{\top} (\hat{A} \hat{A}^{\top})^{-1}\|_2$ exists. Moreover, for any
 1141 \mathcal{I} , $\|\hat{A}_{\mathcal{I}}^{\top} (\hat{A}_{\mathcal{I}} \hat{A}_{\mathcal{I}}^{\top})^{-1}\|_2 \leq C(n) \|\hat{A}^{\top} (\hat{A} \hat{A}^{\top})^{-1}\|_2$ for a constant $C(n)$. Therefore, by replacing λ with
 1142 $\lambda_1 C(n)$, Claim G.9 holds. \square