



Satisficing in Multi-Armed Bandit Problems

Paul Reverdy, *Member, IEEE*, Vaibhav Srivastava, and Naomi Ehrich Leonard, *Fellow, IEEE*

Abstract—Satisficing is a relaxation of maximizing and allows for less risky decision making in the face of uncertainty. We propose two sets of satisficing objectives for the multi-armed bandit problem, where the objective is to achieve reward-based decision-making performance above a given threshold. We show that these new problems are equivalent to various standard multi-armed bandit problems with maximizing objectives and use the equivalence to find bounds on performance. The different objectives can result in qualitatively different behavior; for example, agents explore their options continually in one case and only a finite number of times in another. For the case of Gaussian rewards we show an additional equivalence between the two sets of satisficing objectives that allows algorithms developed for one set to be applied to the other. We then develop variants of the Upper Credible Limit (UCL) algorithm that solve the problems with satisficing objectives and show that these modified UCL algorithms achieve efficient satisficing performance.

Index Terms—Multi-armed bandit, upper credible limit (UCL).

I. INTRODUCTION

ENGINEERING solutions to decision-making problems are often designed to maximize an objective function. However, in many contexts maximization of an objective function is an unreasonable goal, either because the objective itself is poorly defined or because solving the resulting optimization problem is intractable or costly. In these contexts, it is valuable to consider alternative decision-making frameworks.

Herbert Simon considered alternative models of rational decision-making [30] with the goal of making them “compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist.” A major feature of the models he considered is what he called “satisficing”. In [30], he discussed in very broad terms a variety of

simplifications to the classical economic concept of rationality, most importantly the idea that payoffs should be simple, defined by doing well relative to some threshold value. In [31], he introduced the word “satisficing”, a combination of the words “satisfy” and “suffice”, to refer to this thresholding concept and illustrated it using a mathematical model of foraging. He also briefly discussed how satisficing relates to problems in inventory control and more complicated decision processes like playing chess.

Since Simon’s pioneering work, satisficing has been studied in many fields such as psychology [29], economics [6], management science [23], [37], and ecology [36], [8]. In engineering, satisficing is of interest for the same reasons that motivated its introduction in the social science literature, specifically that it can simplify decision-making problems: as compared to maximizing it allows for less risky decision making in the face of uncertainty. Furthermore, many engineering problems are naturally posed using a satisficing objective, such as choosing a design that meets given specifications, but where the designers may be indifferent among any such designs. Satisficing is well defined even if there are several competing performance measures that trade off in complicated ways, whereas maximizing may be poorly defined without additional information about preferences.

Satisficing has been studied in the engineering literature in several contexts. In [25], the authors studied design optimization using a satisficing objective and found that it is effective in many practical fields. In [14], the authors studied control theory using a satisficing objective function, and in [38], the authors used satisficing to study optimal software design. In [10], the authors used a multi-armed bandit algorithm to construct robots that actively adapt their control policies to mitigate damage, such as actuator failures. In order to speed the convergence of their algorithm, they only sought to identify control policies with performance above a set threshold, rather than to identify an optimal policy. The theory that we develop in this paper formalizes their notion of thresholding and provides bounds on performance.

In this paper, we consider satisficing in the stochastic multi-armed bandit problem [28], for which a decision maker sequentially chooses one of a set of alternative options, called arms, and earns a reward drawn from a stationary probability distribution associated with that arm. The standard multi-armed bandit problem uses a maximizing objective on accumulated reward. For this objective there is a known performance bound in terms of expected *regret*, which is the expected difference between the reward received by the decision maker and the maximum reward possible.

Manuscript received July 20, 2016; accepted November 27, 2016. Date of publication December 22, 2016; date of current version July 26, 2017. This work was supported in part by ONR grant N00014-14-1-0635 and ARO grant W911NF-14-1-0431. P. Reverdy was supported through a NDSEG Fellowship. Recommended by Associate Editor Q.-S. Jia.

P. Reverdy is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: preverdy@seas.upenn.edu).

V. Srivastava is with the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI 48824 USA (e-mail: vaibhav@egr.msu.edu).

N. E. Leonard is with the Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: naomi@princeton.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2016.2644380

Since the standard notion of regret is defined relative to the unknown optimum, it can only be computed by an omniscient agent; this notion of regret is not computable by a decision maker faced with a multi-armed bandit problem. Nevertheless, it is a useful theoretical concept, which facilitates the analysis of algorithms designed to solve bandit problems. We extend the notion of regret to satisficing objectives and use it to analyze new algorithms.

In contrast to the standard stochastic multi-armed bandit problem in which the agent seeks to determine, with certainty, the option with maximum mean reward, the satisficing multi-armed bandit problem seeks to determine, with a desired confidence, a satisfying option. We characterize satisficing in multi-armed bandit problems using three separate features of the satisficing objective.

The first feature selects the quantity on which the satisficing objective is defined. We consider two such quantities: (i) the unknown mean reward of the selected option, and (ii) the instantaneous observed reward.

The second feature treats the satisfaction aspect of the satisficing problem. In particular, it selects if the objective function should be optimizing, or if it should be satisfying.

The third feature treats the sufficing aspect of the satisficing problem. In particular, it selects if the decision-making algorithm should be certain that the optimizing/satisfying criterion is met, or if it is sufficient for the algorithm to meet a desired threshold in confidence about the criterion. Different combinations of the above three features of satisficing lead to eight satisficing objectives that we discuss in this paper.

We begin by defining the four objectives for the case where the satisficing quantity is the unknown mean reward. We show that the bandit problem with each of these four objectives is equivalent to a previously studied bandit problem and use the equivalence to derive a performance bound for the satisficing problems. These four objectives seek an arm with satisfyingly high mean reward without regard to that reward's dispersion. To develop objectives with improved robustness properties, we then consider the case where the satisficing quantity is the instantaneous observed reward. We extend the first four objectives to this case by adding an additional layer of thresholding, which defines four more objectives. When the reward distributions belong to location-scale families, there is an equivalence between the objectives defined in terms of mean reward and the robust objectives defined in terms of instantaneous reward, which we prove for Gaussian rewards.

For simplicity of exposition, we then specialize to Gaussian multi-armed bandit problems, where the reward distributions are Gaussian with unknown mean and known variance. For such problems, we develop several modifications of the UCL algorithm that we developed in previous work [27]. These algorithms solve the problem with the satisficing mean reward objectives (and thus also with the robust objectives). We show that these algorithms achieve efficient performance. These results extend our previous work [26] by incorporating the concept of sufficiency into the satisficing objective, as well as by adding several new algorithms and their associated analysis.

The assumption of Gaussian rewards with known variance is not required, but it allows us to focus on the different notions of regret, which is the main contribution of this paper. We later show how the known variance assumption can be relaxed. Our methods also extend immediately to many other important classes of reward distributions, including distributions with bounded support and sub-Gaussian distributions. We show how to extend our methods in these cases and provide references to the relevant literature for other extensions.

The remainder of the paper is structured as follows. In Section II we review the standard stochastic multi-armed bandit problem and the associated performance bounds. In Section III we propose the satisficing objectives and bound performance in terms of these objectives. In Section IV we specialize to the case of Gaussian rewards and show the equivalence between the satisficing in mean reward objectives and the satisficing in instantaneous observed reward objectives. In Section V we review the UCL algorithm, and in Section VI we design modified versions of the UCL algorithm for the satisficing objectives. We show that these modified algorithms achieve efficient performance for Gaussian rewards. We show the results of numerical simulations in Section VII and in Section VIII we conclude.

II. THE STOCHASTIC MULTI-ARMED BANDIT PROBLEM

In the stochastic multi-armed bandit problem a decision-making agent sequentially chooses one among a set of N options called *arms* in analogy with the lever of a slot machine. A single-levered slot machine is called a *one-armed bandit*, so the case of $N \geq 2$ options is called a *multi-armed bandit*.

The decision-making agent collects reward $r_t \in \mathbb{R}$ by choosing arm i_t at each time $t \in \{1, \dots, T\}$, where $T \in \mathbb{N}$ is the horizon length for the sequential decision process. The reward from option $i \in \{1, \dots, N\}$ is sampled from a stationary probability distribution ν_i which has an unknown mean $m_i \in \mathbb{R}$. The decision-maker's objective is to maximize some function of the sequence of rewards $\{r_t\}$ by sequentially picking arms i_t using only the information available at time t .

A. Maximization Objective

In the standard multi-armed bandit problem, the agent's objective is to maximize the expected cumulative reward

$$J = \mathbb{E} \left[\sum_{t=1}^T r_t \right] = \sum_{t=1}^T m_{i_t}. \quad (1)$$

Equivalently, by defining $m_{i^*} = \max_i m_i$ and $R_t = m_{i^*} - m_{i_t}$, the expected *regret* at time t , minimizing (1) can be formulated as minimizing the cumulative expected regret defined by

$$\sum_{t=1}^T R_t = T m_{i^*} - \sum_{i=1}^N m_i \mathbb{E} [n_i^T] = \sum_{i=1}^N \Delta_i \mathbb{E} [n_i^T], \quad (2)$$

where n_i^T is the number of times arm i has been chosen up to time T , $\Delta_i = m_{i^*} - m_i$ is the expected regret due to picking arm i instead of arm i^* , and the expectation is over the possible rewards and decisions made by the agent.

The interpretation of (2) is that suboptimal arms $i \neq i^*$ should be chosen as rarely as possible. This is a non-trivial task since the mean rewards m_i are initially unknown to the decision-maker, who must try arms to learn about their rewards while preferentially picking arms that appear more rewarding. The tension between these requirements is known as the *explore-exploit* tradeoff and is common to many problems in machine learning and adaptive control.

B. Bound on Optimal Performance

Optimal performance in a bandit problem corresponds to picking suboptimal arms as rarely as possible, as shown by (2). Lai and Robbins [20] studied the standard stochastic multi-armed bandit problem and showed that any policy solving the problem must pick each suboptimal arm $i \neq i^*$ a number of times that is at least logarithmic in the time horizon T , i.e.,

$$\mathbb{E} [n_i^T] \geq \left(\frac{1}{D(\nu_i || \nu_{i^*})} + o(1) \right) \log T, \quad (3)$$

where $o(1) \rightarrow 0$ as $T \rightarrow +\infty$. The quantity $D(\nu_i || \nu_{i^*}) := \int \nu_i(r) \log \frac{\nu_i(r)}{\nu_{i^*}(r)} dr$ is the Kullback-Leibler divergence between the reward density ν_i of any suboptimal arm and the reward density ν_{i^*} of the optimal arm. Equation (3) implies that cumulative expected regret must grow at least logarithmically in time.

The bound (3) is asymptotic in time, but researchers (e.g., [4], [13], [27]) have developed algorithms that achieve cumulative expected regret that is bounded by a logarithmic term uniformly in time, sometimes with the same constant as in (3). Cumulative expected regret that is uniformly bounded in time by a logarithmic term is often called *logarithmic regret* for short. In the literature, algorithms that achieve logarithmic regret with a leading term that is within a constant factor of that in (3) are considered to have optimal performance.

C. Multiple Plays

Anantharam *et al.* [2] studied a generalization of the multi-armed bandit problem in which the agent picks $k \geq 1$ arms at each time t , which they called the multi-armed bandit problem with multiple plays. The case $k = 1$ corresponds to the standard multi-armed bandit problem defined above.

In the spirit of [2], let σ be a permutation of $\{1, \dots, N\}$ such that $m_{\sigma(1)} \geq m_{\sigma(2)} \geq \dots \geq m_{\sigma(N)}$. For the multi-armed bandit problem with k plays, the optimal policy with full information corresponds to picking the arms $\sigma(1), \dots, \sigma(k)$, called the *k-best* arms [2]. In the case $k = 1$, $\sigma(1) = i^*$, the optimal arm defined above. For the case of general $k \geq 1$, the cumulative expected regret for the multi-armed bandit problem with multiple plays is defined as follows [2]:

$$T \sum_{i=1}^k m_{\sigma(i)} - \sum_{i=1}^N m_i \mathbb{E} [n_i^T], \quad (4)$$

which is a straightforward generalization of the regret (2). The suboptimal arms $\sigma(k+1), \dots, \sigma(N)$ are called the *k-worst* arms [2]. Define $\Delta_i^{(k)} = m_{\sigma(k)} - m_i$ for each *k-worst* arm i . The quantity $\Delta_i^{(k)}$ is the generalization of the expected regret Δ_i

for the problem with multiple plays, where the expected value of the optimal policy is that of the *k* best arms.

As in the case of a single play, optimal performance corresponds to picking suboptimal (i.e., *k-worst*) arms as rarely as possible. By [2] each *k-worst* arm i must be picked a number of times that is at least logarithmic in the time horizon T , i.e.,

$$\mathbb{E} [n_i^T] \geq \left(\frac{1}{D(\nu_i || \nu_{\sigma(k)})} + o(1) \right) \log T. \quad (5)$$

This bound can be interpreted as a generalization of the Lai-Robbins bound (3) where the Kullback-Leibler divergence is taken with respect to the *kth* best arm $\sigma(k)$ rather than the first best arm $\sigma(1)$ (i.e., i^* in the case $k = 1$).

D. PAC Bounds

In the standard multi-armed bandit problem and the multi-armed bandit problem with multiple plays, regret is defined in terms of the unknown mean reward values m_i . These regret definitions imply that avoiding regret requires identifying optimal arms with certainty. The requirement to identify optimal arms with certainty is characteristic of a maximizing decision-making strategy. In contrast, a satisficing decision-making agent should seek arms that are “good enough”. In this context, satisficing corresponds to finding arms that are optimal with high probability rather than with certainty.

The Probably Approximately Correct (PAC) model for learning introduced by Valiant [34] provides a natural way to capture this aspect of satisficing. Even-Dar *et al.* [11], [12] and Mannor and Tsitsiklis [22] studied the multi-armed bandit problem using the PAC model and defined an ϵ -optimal arm i as one for which $m_i > m_{i^*} - \epsilon$, i.e., the mean reward is within ϵ of the optimum value. Equivalently, an ϵ -optimal arm is an arm i for which the expected regret Δ_i is at most ϵ . Under the PAC model one wishes to find an ϵ -optimal arm with probability of at least $1 - \delta$. With probability one, this can be achieved in a finite number of samples, so performance guarantees take the form of bounds on the number of samples required, which is referred to as *sample complexity*. In our notation, we denote sample complexity by T^* , as it is the value of the horizon length at which sampling terminates.

When the rewards are Bernoulli distributed with unknown success probabilities p_i , the following lower bound holds [22]:

$$\mathbb{E} [T^*] \geq \mathcal{O} \left(\frac{1}{\epsilon^2} \log(1/\delta) \right). \quad (6)$$

A similar result was reported in [11] for T^* , rather than its expected value. In other words, one must sample an arm at least $\log(1/\delta)/\epsilon^2$ times to be able to declare that it is ϵ -optimal with probability at least $1 - \delta$.

Similar to the work of [2] extending Lai and Robbins’ bounds [20] to the case of multiple plays, Kalyanakrishnan *et al.* [15] extended the work of [12] from finding the ϵ -optimal arm to finding the m ϵ -best arms with probability at least $1 - \delta$. In [15] this problem is called *Explore- m* , and an algorithm that solves it (ϵ, m, δ) -optimal. Note that the problem in [12] is the special case *Explore-1*. The *Explore- m* problem is studied in [15] for

rewards that are Bernoulli distributed. It is proved that, for every (ϵ, m, δ) -optimal algorithm, there exists a bandit problem on which that algorithm has worst-case sample complexity of at least $\log(m/8\delta)$. Specifically, it is shown that there exists a bandit problem such that the number of samples T^* required to identify m ϵ -best arms obeys

$$T^* \geq \frac{1}{18375} \frac{N}{\epsilon^2} \log\left(\frac{m}{8\delta}\right). \quad (7)$$

This gives a worst-case bound on the number of times all arms need to be sampled to achieve (ϵ, m, δ) -optimality.

The bounds (6) and (7) were both formulated for the case of Bernoulli rewards, but it is straightforward to extend them to the case where the rewards are Gaussian distributed with unknown mean and known variance.

E. Gaussian Rewards

In this paper we focus on the case of Gaussian reward distributions, where the distribution ν_i of rewards associated with arm i is Gaussian with mean m_i , which is unknown to the decision maker, and variance $\sigma_{s,i}^2$, which is known to the decision maker from, e.g., previous observations or known measurement characteristics. Relaxation of the assumption of known variance is discussed in Remark 12. For the given case, the Kullback-Leibler divergence in (3) takes the value

$$D(\nu_i || \nu_{i^*}) = \frac{1}{2} \left(\frac{\Delta_i^2}{\sigma_{s,i^*}^2} + \frac{\sigma_{s,i}^2}{\sigma_{s,i^*}^2} - 1 - \log \frac{\sigma_{s,i}^2}{\sigma_{s,i^*}^2} \right). \quad (8)$$

This equation is more easily interpreted when the reward variances are uniform, i.e., $\sigma_{s,i}^2 = \sigma_s^2$ for each i . In some cases we assume uniform variance for simplicity of exposition, but the relevant results are readily generalized to the case of non-uniform variance. Assuming uniform variance, $D(\nu_i || \nu_{i^*}) = \Delta_i^2 / 2\sigma_s^2$, so the bound (3) is

$$\mathbb{E} [n_i^T] \geq \left(\frac{2\sigma_s^2}{\Delta_i^2} + o(1) \right) \log T. \quad (9)$$

This result can be interpreted as follows. For a given value of Δ_i , a larger variance σ_s^2 makes the rewards more variable and therefore it is more difficult to distinguish between the arms. For a given value of σ_s^2 , a larger value of Δ_i makes it easier to distinguish it from the optimal arm. The expressions for the problem with multiple plays (i.e., (5)) are identical except for substituting $\sigma(k)$ for i^* and $\Delta_i^{(k)}$ for Δ_i .

III. THE MULTI-ARMED BANDIT PROBLEM WITH SATISFICING OBJECTIVES

We now define the multi-armed bandit problem with satisficing objectives. We propose several new satisficing notions of regret and find associated bounds on optimal performance. These notions capture two dimensions of the satisficing problem: *satisfaction*, i.e., the agent's desire to obtain a reward that is above a certain threshold, and *sufficiency*, i.e., the agent's desire to attain a level of confidence that its choice of a given arm will bring them satisfaction. We define these notions first for satisficing in mean reward and then extend them to satisficing in instantaneous reward, which we refer to as *robust satisficing*.

A. Satisficing in Mean Reward

We define *satisfaction* in mean reward as having an expected reward m_{i_t} that is above a specified threshold value \mathcal{M} . Formally, we represent satisfaction in mean reward at time t by the variable s_t , defined as

$$s_t = \mathbb{1}(m_{i_t} > \mathcal{M}), \quad (10)$$

where $\mathbb{1}(\cdot)$ is the indicator function, equal to one if the argument is true and zero otherwise. The threshold \mathcal{M} is a free parameter that must be specified by the decision-making agent. Let $m_{i^*} = \max_i m_i$ be the maximum expected reward from any arm. The agent can never be satisfied if \mathcal{M} is greater than m_{i^*} , so we assume that $\mathcal{M} \leq m_{i^*}$ to make the problem feasible. If $\mathcal{M} > m_{\sigma(2)}$, i.e., greater than the mean reward of the second-best arm, then arm $\sigma(1) = i^*$ is the only one that is satisfying in mean reward.

As in the multi-armed bandit problem with multiple plays, let σ be a permutation of $\{1, \dots, N\}$ such that $m_{\sigma(1)} \geq m_{\sigma(2)} \geq \dots \geq m_{\sigma(N)}$. Let k be the largest integer such that $m_{\sigma(k)} \geq \mathcal{M}$. The arms $\{\sigma(1), \dots, \sigma(k)\}$ are the k -best arms defined by the satisfaction threshold \mathcal{M} . For each arm i , define the thresholded expected regret $\Delta_i^{\mathcal{M}} = \max\{\mathcal{M} - m_i, 0\}$. For each k -best arm, the thresholded regret is zero, and for each k -worst arm $i \in \{\sigma(k+1), \dots, \sigma(N)\}$, the value $\Delta_i^{\mathcal{M}} > 0$ quantifies the extent to which the arm is unsatisfying in mean rewards. Note that if $\mathcal{M} = m_{i^*}$, $\Delta_i^{\mathcal{M}} = \Delta_i$, which is the standard measure of expected regret. We refer to the k -best and k -worst arms as *satisfying* and *non-satisfying* arms, respectively.

The satisfaction variable s_t defined in (10) can be written as a function of the sign of $\Delta_{i_t}^{\mathcal{M}}$:

$$s_t = \mathbb{1}(\Delta_{i_t}^{\mathcal{M}} = 0).$$

The quantity s_t is deterministic. However, since the agent does not know the value of $\Delta_i^{\mathcal{M}}$ associated with any given arm, they must learn it by sampling rewards from the various arms and updating their beliefs accordingly. Adopting a Bayesian framework, we assume s_t is a realization of a binary random variable S_t . Due to the stochastic nature of the rewards the agent will have less than perfect confidence in their beliefs about the value of s_t .

We distinguish satisficing objectives in mean reward according to the degree $\delta \in [0, 1]$ of confidence the agent seeks in their beliefs, which we call *sufficiency in mean reward*. We define an arm i to be (δ) -sufficing in mean reward if

$$\Pr [S_t = 1] \geq 1 - \delta,$$

where the probability is evaluated based on the agent's current beliefs. For non-zero values of δ , the agent finds it sufficient to have finite confidence that they are satisfied, while for $\delta = 0$, the agent wants certainty that they are satisfied. The agent cannot achieve certainty in finite time, so these two cases result in qualitatively different behavior: $\delta = 0$ means the agent will never stop exploring, while $\delta > 0$ means the agent will settle on a set of acceptable options after finite time.

TABLE I
TABLE OF THE FOUR DIFFERENT REGRET CONCEPTS, AND RESULTING PROBLEMS, ASSOCIATED WITH THE SATISFICING-IN-MEAN-REWARD MULTI-ARMED BANDIT PROBLEM

Threshold level	Seek certainty ($\delta = 0$)	Suffice ($\delta > 0$)
$\mathcal{M} > m_{\sigma(2)}$	1) Standard bandit	3) δ -sufficing
$\mathcal{M} \leq m_{\sigma(2)}$	2) Satisfaction-in-mean-rwd	4) (\mathcal{M}, δ) -satisficing

The *satisficing-in-mean-reward objective* is

$$\sum_{t=1}^T \mathbb{1}((s_t = 1) \text{ or } \Pr[S_t = 1] > 1 - \delta). \quad (11)$$

The objective (11) is maximized if, at each time, a satisfying option is selected, or the probability that the option is satisfying is sufficiently high. The event that an option is satisfying is not known a priori and must be learned by exploration. This results in an explore-exploit tradeoff as in the standard multi-armed bandit problem.

To quantify the optimal explore-exploit tradeoff in the spirit of the Lai-Robbins bound we introduce the following notion of the *expected satisficing regret* at time t , R_t , defined by

$$R_t = \Delta_{i_t}^{\mathcal{M}} \mathbb{1}(\Pr[S_t = 1] < 1 - \delta). \quad (12)$$

If the agent is insufficiently certain of being satisfied by the choice of i_t , they incur expected regret of $\Delta_{i_t}^{\mathcal{M}}$. Otherwise, they incur no regret.

We define the satisficing-in-mean-reward multi-armed bandit problem in terms of minimizing cumulative expected satisficing regret.

Definition 1 (Satisficing-In-Mean-Reward Multi-Armed Bandit Problem): The *satisficing-in-mean-reward multi-armed bandit problem* is to minimize the cumulative sum of the expected satisficing regret (12):

$$J_R = \mathbb{E} \left[\sum_{t=1}^T R_t \right]. \quad (13)$$

The satisficing-in-mean-reward bandit problem has two parameters: \mathcal{M} and δ . These parameters characterize the agent's thresholds for satisfaction and sufficiency, respectively. For purposes of analysis we distinguish four cases as a function of the parameter values. For the satisfaction threshold $\mathcal{M} \in \mathbb{R}$, the first case is setting $\mathcal{M} > m_{\sigma(2)}$, while the second case is setting $\mathcal{M} \leq m_{\sigma(2)}$. For the sufficiency threshold $\delta \in [0, 1]$, the first case is the certainty value $\delta = 0$, while the second case is $\delta \in (0, 1]$.

Table I summarizes the four problems that result from the interaction of the two dimensions of satisfaction and sufficiency. Problem 1 sets the satisfaction threshold $\mathcal{M} > m_{\sigma(2)}$ and the sufficiency threshold $\delta = 0$, which results in a standard bandit problem. We call Problem 2 with $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta = 0$ *satisfaction-in-mean-reward*. We call Problem 3 with $\mathcal{M} > m_{\sigma(2)}$ and $\delta \in (0, 1]$ *δ -sufficing*. Finally, we call Problem 4 with $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta \in (0, 1]$, *(\mathcal{M}, δ) -satisficing*.

Remark 1: We note that the distinction between Problems 1 and 2 and between Problems 3 and 4 is only due to the range of values \mathcal{M} can take. These problems can be thought of as a

single problem in which the choice of \mathcal{M} dictates the cardinality of the set of satisfying arms. However, the two ranges of thresholds $\mathcal{M} > m_{\sigma(2)}$ and $\mathcal{M} \leq m_{\sigma(2)}$ allow us to clearly contrast the satisficing problem with the standard problem. Assuming $\mathcal{M} > m_{\sigma(2)}$ in Problems 1 and 3 is equivalent to assuming that the agent seeks the (unknown) highest mean reward, which is consistent with the standard problem. The policies we define for Problems 1 and 3 do not rely on a known threshold \mathcal{M} . Assuming $\mathcal{M} \leq m_{\sigma(2)}$ is equivalent to assuming that the agent seeks to meet a (known) desired mean reward threshold. The policies we define for Problems 2 and 4 do rely on the threshold \mathcal{M} . These same assumptions analogously distinguish Problems 5 and 7 from Problems 6 and 8 defined in Section III-B. However, unlike the policies for Problems 1 and 3, the policies defined for Problems 5 and 7 do rely on $\mathcal{M} > m_{\sigma(2)}$ being known. We do not assume in any of the problems that the agent knows the permutation σ , so no policies depend on σ .

We develop performance bounds for each of these problems in terms of corollaries of the performance bounds presented in Section II. For the problems with $\delta = 0$, these bounds show that cumulative expected regret must grow at least at a logarithmic rate, while for the problems with $\delta > 0$, finite regret is possible.

Problem 1 Standard Bandit: The satisficing-in-mean-reward multi-armed bandit problem with $\mathcal{M} > m_{\sigma(2)}$ and $\delta = 0$ is a standard multi-armed bandit problem. Therefore, for this problem, the Lai-Robbins bound (3) holds, and the expected number of times a suboptimal arm i is chosen obeys

$$\mathbb{E}[n_i^T] \geq \left(\frac{1}{D(\nu_i || \nu_{i^*})} + o(1) \right) \log T.$$

As a direct consequence, the cumulative expected satisficing regret (13) grows at least logarithmically with time horizon T :

$$J_R \geq \left(\sum_{i=1}^N \frac{\Delta_i^{\mathcal{M}}}{D(\nu_i || \nu_{i^*})} + o(1) \right) \log T.$$

Problem 2 Satisfaction-in-Mean-Reward: The satisfaction-in-mean-reward problem, defined as the satisficing-in-mean-reward multi-armed bandit problem where $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta = 0$, also has a logarithmic lower bound on the cumulative expected satisficing regret:

Corollary 2 (Satisfaction-In-Mean-Reward Regret Bound): The satisfaction-in-mean-reward problem is a satisficing-in-mean-reward multi-armed bandit problem where the objective (13) is defined with $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta = 0$. Any policy solving the satisfaction-in-mean-reward problem obeys

$$\mathbb{E}[n_i^T] \geq \left(\frac{1}{D(\nu_i || \nu_{\sigma(k)})} + o(1) \right) \log T \quad (14)$$

for each non-satisfying arm i , where σ is a permutation of $\{1, \dots, N\}$ such that $m_{\sigma(1)} \geq m_{\sigma(2)} \geq \dots \geq m_{\sigma(N)}$ and k is the largest integer such that $m_{\sigma(k)} \geq \mathcal{M}$.

Proof: The definition of satisfaction (10) implies that performance bounds for the satisfaction-in-mean-reward problem and the multi-armed bandit problem with multiple plays are equivalent. Given a problem instance, the threshold \mathcal{M} induces the number k of satisfying arms, so performance can be analyzed as in the problem with multiple plays. The bound (5) applies to

the problem with multiple plays and the equivalence implies the result. ■

Problem 3 δ -Sufficing: The δ -sufficing problem, defined as the satisficing-in-mean-reward multi-armed bandit problem where $\mathcal{M} > m_{\sigma(2)}$ and $\delta \in (0, 1]$, admits policies that achieve cumulative expected regret that is a bounded function of T :

Corollary 3 (δ -Sufficing Regret Bound): The δ -sufficing problem is a satisficing-in-mean-reward multi-armed bandit problem where the objective (13) is defined with $\mathcal{M} > m_{\sigma(2)}$ and $\delta \in (0, 1]$. Any policy solving the δ -sufficing problem obeys

$$n_i^T \geq \mathcal{O}\left(\frac{1}{\epsilon^2} \log(1/\delta)\right) \quad (15)$$

for each suboptimal arm i , where $\epsilon = \Delta_i = \mathcal{M} - m_i$.

Proof: The definition of satisfaction (10) in the δ -sufficing problem implies that the agent incurs regret if the arm selected is not $(\epsilon = 0, \delta)$ -optimal. The bound (6) thus provides a lower bound on the number of times the agent must incur regret. ■

Problem 4 (\mathcal{M}, δ) -Satisficing: The (\mathcal{M}, δ) -satisficing problem, defined as the satisficing-in-mean-reward multi-armed bandit problem where $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta \in (0, 1]$, admits policies that achieve cumulative expected regret that is a bounded function of T :

Corollary 4 ((\mathcal{M}, δ) -Satisficing Regret Bound): The (\mathcal{M}, δ) -satisficing problem is a satisficing-in-mean-reward multi-armed bandit problem where the objective (13) is defined with $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta \in (0, 1]$. Any policy solving the (\mathcal{M}, δ) -satisficing multi-armed bandit problem obeys

$$T^* = \sum_{i=1}^N n_i^{T^*} \geq \frac{1}{18375} \frac{N}{\epsilon^2} \log\left(\frac{k}{8\delta}\right) \quad (16)$$

where σ is a permutation of $\{1, \dots, N\}$ such that $m_{\sigma(1)} \geq m_{\sigma(2)} \geq \dots \geq m_{\sigma(N)}$, k is the largest integer such that $m_{\sigma(k)} \geq \mathcal{M}$, and $\epsilon = \mathcal{M} - m_{\sigma(k)}$. Since only arms in $\{\sigma(k+1), \dots, \sigma(N)\}$ result in regret, the left hand side of (16) is an upper bound on the expected satisficing regret (13).

Proof: The definition of satisfaction (10) in the (\mathcal{M}, δ) -sufficing problem implies that an algorithm that minimizes satisficing regret is equivalent to an $(\epsilon = m_{\sigma(k)} - \mathcal{M}, k, \delta)$ -optimal algorithm in the sense of [15]. Therefore, the bound (7) applies to the (\mathcal{M}, δ) -sufficing problem. ■

Recall that T^* is the number of times all arms (including the optimal one) should be cumulatively sampled such that following T^* an (\mathcal{M}, ϵ) -optimal decision can be made. The lower bounds on both T^* and $n_i^{T^*}$ are independent of T , suggesting that for (\mathcal{M}, ϵ) -satisficing, a bounded regret can be achieved.

Corollaries 3 and 4 show that the worst-case regret is a bounded function of T for the sufficing problems, where $\delta > 0$. Therefore we can conclude that the expected regret for such problems can also be a bounded function of T . This is an important distinction from the maximizing problems, where $\delta = 0$: in such problems, the Lai-Robbins bound (3) implies that the expected regret must grow logarithmically with T . As is standard in the bandit literature, we say an algorithm has efficient performance if its regret matches, up to constant factors, the relevant growth rates: $\log T$ for maximizing problems and $\log(k/\delta)/\epsilon^2$ for sufficing problems.

B. Robust Satisficing in Instantaneous Reward

The four objectives defined in Section III-A above define satisfaction (10) in terms of the mean reward m_i from an arm i . This captures situations where the time scale for satisfaction spans numerous decision times. For example, consider foraging, where an animal must consume a minimum amount of food each day. If each decision time represents a small portion of the day, the total food consumed during the day represents the sum of numerous small rewards from each decision time. As long as the mean reward at each decision time is sufficiently high, the animal will meet its daily food requirement.

If, instead, the decision time scale is the same as the satisfaction time scale, it is more appropriate to define satisfaction at time t in terms of the reward r_t received at that time. This requires more robust algorithms, in the sense that they must ensure that each reward, rather than simply the mean reward, is satisfying with high probability. In this context we define satisfaction in two stages. First, we define *happiness* as receiving a reward r_t that is at least a threshold value $M \in \mathbb{R}$. We represent happiness at time t as the Bernoulli random variable h_t , defined as

$$h_t = \mathbb{1}(r_t > M). \quad (17)$$

We define the *success probability* of the happiness random variable h_t as

$$p_i = \Pr[h_t = 1 | i_t = i]. \quad (18)$$

The success probability p_i is the expected rate of happiness due to picking arm i . This defines a Bernoulli multi-armed bandit problem where the mean reward (i.e., happiness rate) is p_i . We then define satisfaction in terms of a threshold Π for this Bernoulli multi-armed bandit problem as we did in (10):

$$s_t = \mathbb{1}(p_{i_t} > \Pi). \quad (19)$$

Given the happiness threshold M , this definition is identical to the definition (10) of satisfaction where $m_i = p_i$, $p_{i^*} = \max_i p_i$, and $\mathcal{M} = \Pi$. Therefore the four satisficing multi-armed bandit problems defined in Table I can be used to define four additional problems in this context, which we call *robust satisficing*.

Definition 2 (Robust Satisficing Multi-Armed Bandit Problem): The *robust satisficing multi-armed bandit problem* is to minimize the cumulative sum of the expected satisficing regret (12):

$$J_R = \mathbb{E} \left[\sum_{t=1}^T R_t \right],$$

where the regret R_t is defined using the notion of satisfaction defined by (17)–(19).

A robust satisficing multi-armed bandit problem has three parameters: M , Π , and δ . We assume that M and Π are chosen such that there is at least one satisfying arm; otherwise, the expected regret must grow indefinitely. Table II summarizes the four robust satisficing multi-armed bandit problems that result from the interaction of the two dimensions of satisfaction and sufficiency, which we list below. We assume that ς is a permutation of $\{1, \dots, N\}$ such that $p_{\varsigma(1)} \geq p_{\varsigma(2)} \geq \dots \geq p_{\varsigma(N)}$.

TABLE II
TABLE OF THE FOUR DIFFERENT REGRET CONCEPTS, AND RESULTING PROBLEMS, ASSOCIATED WITH THE ROBUST SATISFICING MULTI-ARMED BANDIT PROBLEM

Threshold level	Seek certainty ($\delta = 0$)	Suffice ($\delta > 0$)
$\Pi > p_{\zeta(2)}$	5) Robust bandit	7) δ -robust sufficing
$\Pi \leq p_{\zeta(2)}$	6) Robust satisfaction	8) (Π, δ) -robust satisfying

The quantity p_i represents the probability of happiness (i.e., receiving a reward of at least M) due to choosing arm i

Problem 5. Robust Bandit: The robust bandit problem is defined as the robust satisfying multi-armed bandit problem where $\Pi > p_{\zeta(2)}$ and $\delta = 0$.

Problem 6. Robust Satisfaction: The robust satisfaction problem is defined as the robust satisfying multi-armed bandit problem where $\Pi \leq p_{\zeta(2)}$ and $\delta = 0$.

Problem 7. δ -Robust Sufficing: The δ -robust sufficing problem is defined as the robust satisfying multi-armed bandit problem where $\Pi > p_{\zeta(2)}$ and $\delta \in (0, 1]$.

Problem 8. (Π, δ) -Robust Satisficing: The (Π, δ) -robust satisfying problem is defined as the robust satisfying multi-armed bandit problem where $\Pi \leq p_{\zeta(2)}$ and $\delta \in (0, 1]$.

For a large class of reward distributions, there is an equivalence between Problems 5–8 defined in terms of r_t and Problems 1–4 defined in terms of m_i . By Lemma 5 below, when the rewards r_t follow a Gaussian distribution with unknown mean m_i and known variance $\sigma_{s,i}^2$, each problem in Table II is equivalent to the analogous problem in Table I.

IV. SATISFICING WITH GAUSSIAN REWARDS

In this section we study the Gaussian satisfying multi-armed bandit problem. This is the satisfying multi-armed bandit problem where the reward r_t due to selecting arm i_t is $r_t \sim \mathcal{N}(m_{i_t}, \sigma_{s,i_t}^2)$ and σ_{s,i_t}^2 is the known variance of arm i_t . In this case, we show a formal equivalence between the satisfying-in-mean-reward multi-armed bandit problems and the robust satisfying multi-armed bandit problems. The choice of Gaussian rewards facilitates modeling correlation dependencies among arms, which can be useful in applications.

A. Equivalence Lemma for Gaussian Rewards

For the Gaussian robust satisfying multi-armed bandit problem, define the quantity

$$x_i = \frac{m_i - M}{\sigma_{s,i}}, \quad (20)$$

which we call the *standardized mean reward*, for each arm i . The following lemma states that each Gaussian robust satisfying multi-armed bandit problem where satisfaction is defined by (19) is equivalent to a Gaussian satisfying-in-mean-reward multi-armed bandit problem where satisfaction is defined by (10) with standardized reward distributions.

Lemma 5 (Equivalence for Gaussian Rewards): Each Gaussian robust satisfying multi-armed bandit problem is equivalent to a Gaussian satisfying-in-mean-reward multi-armed bandit

problem with rewards $\tilde{r}_t \sim \mathcal{N}(x_{i_t}, 1)$ with x_i given by (20). That is, the ordering of the arms in terms of x_i is identical to the ordering in terms of p_i , and, in particular, the arm with maximal x_i is the arm with maximal p_i .

Proof: With Gaussian rewards, the probability (18) of happiness due to choosing arm i is

$$\begin{aligned} p_i &= \Pr[m_i + \sigma_{s,i}z \geq M] \\ &= \Phi\left(\frac{m_i - M}{\sigma_{s,i}}\right) = \Phi(x_i), \end{aligned} \quad (21)$$

where $z \sim \mathcal{N}(0, 1)$ is a standard normal random variable and $\Phi(z)$ is its cumulative distribution function. Let $i^* = \arg \max_i p_i$. The key insight is that $\Phi(\cdot)$ is a monotonically increasing function, which implies that the ordering of arms in terms of p_i is identical to the ordering in terms of x_i . In particular, arm i^* is the arm with maximal x_i . Therefore, satisfaction in terms of r_t is equivalent to satisfaction in terms of the mean reward x_i .

This is again a Gaussian bandit problem: consider the standardized reward

$$\tilde{r}_t = \frac{r_t - M}{\sigma_{s,i_t}}, \quad (22)$$

which is a Gaussian random variable $\tilde{r}_t \sim \mathcal{N}(x_{i_t}, 1)$. The quantity x_i plays the role of the mean reward m_i and the transformed rewards have uniform variance $\tilde{\sigma}_s^2 = 1$. Minimizing the robust satisfying regret in terms of r_t is equivalent to minimizing the satisfying regret in terms of x_i . ■

Lemma 5 has two implications for the relationship between Problems 5–8 and Problems 1–4 when rewards are Gaussian distributed. First, each Problem 5–8 inherits a regret bound from the corresponding Problem 1–4. Second, each Problem 5–8 can be solved by applying the algorithm developed for Problem 1–4 by first applying the standardization transformation (22) to the observed rewards.

Remark 6 (Location-Scale Families): Lemma 5 is easily generalized to reward distributions belonging to location-scale families. A location-scale family is a set of probability distributions closed under affine transformations, i.e., if the random variable X is in the family, so is the variable $Y = a + bX$, where $a, b \in \mathbb{R}$. Any random variable X in such a family with mean μ and standard deviation σ can be written as $X = \mu + \sigma Z$, where Z is a zero-mean, unit-variance member of the family. Examples include the uniform distribution and Student's t -distribution.

B. Application to the Gaussian Robust Satisficing Problems

In this section we show how to use the equivalence result of Lemma 5 for the full set of robust satisfying problems in the case of Gaussian rewards.

Recall from Lemma 5 that the probability of happiness (18) due to picking an arm i is p_i . In the proof of the lemma, we show that maximizing the probability of happiness is equivalent to maximizing the mean reward in a Gaussian multi-armed bandit problem with mean rewards $x_i = \Phi^{-1}(p_i)$, where x_i is the standardized mean reward $(m_i - M)/\sigma_{s,i}$.

Given an algorithm developed for one of the Problems 1–4 defined in Table I, it can be applied to the corresponding Problem 5–8 defined in Table II as follows. Standardize the observed rewards r_t and run the algorithm using the standardized rewards $\tilde{r}_t = (r_t - M)/\sigma_{s,i_t}$ as input. For example, Problem 5, the robust multi-armed bandit problem, can be solved by an algorithm designed to solve Problem 1, the standard bandit problem, where rewards are transformed according to (22) before being input to the algorithm. The same procedure allows one to apply algorithms developed for Problem 3, δ -sufficing, to Problem 7, δ -robust sufficing.

For Problem 6, robust satisfaction, and Problem 8, (Π, δ) -robust satisficing, we need a threshold X that is analogous to the threshold \mathcal{M} defined for Problem 2, satisfaction-in-mean-reward, and Problem 4, (\mathcal{M}, δ) -satisficing. We use the relationship between x_i and p_i to derive the threshold. In particular, for a robust satisficing problem with probability of happiness threshold Π , define the threshold X by

$$X = \Phi^{-1}(\Pi). \quad (23)$$

When the rewards are Gaussian distributed, we can apply algorithms developed for Problems 2 and 4 to the corresponding robust satisficing Problems 6 and 8 by standardizing rewards and using the threshold X defined in (23) in place of the threshold \mathcal{M} .

Lemma 5 implies that the efficient performance guarantees for algorithms designed for Problems 1–4 also hold when they are used to solve the robust satisficing Problems 5–8.

V. THE UCL ALGORITHM FOR GAUSSIAN MULTI-ARMED BANDIT PROBLEMS

In this section we review the UCL algorithm, a Bayesian algorithm we developed and analyzed in [27] to solve the standard Gaussian multi-armed bandit problem. The UCL algorithm was developed by applying the Bayesian upper confidence bound approach of [16] to the case of Gaussian rewards; the choice of Gaussian rewards facilitated the modeling of human decision-making behavior.

The UCL algorithm maintains a belief about the mean rewards \mathbf{m} by starting with a prior and updating it using Bayesian inference as new rewards are received. At each time t the algorithm chooses arm i_t using a heuristic that is a simple function of the current belief state. For uninformative priors, the UCL algorithm achieves logarithmic regret, i.e., optimal performance.

Uninformative priors correspond to having no information about the mean rewards. A major advantage of the UCL algorithm is its ability to incorporate information about the mean rewards through the use of a so-called informative prior. In [27], we showed that an appropriately-chosen prior can significantly increase the performance of the UCL algorithm. Several different UCL algorithms were developed in [27], including a stochastic decision rule to model human behavior; here we cover only the deterministic UCL algorithm, which, for brevity, we refer to as the UCL algorithm.

A. Prior

The prior distribution captures the agent’s knowledge about the vector of mean rewards \mathbf{m} before beginning the task. We assume that the prior distribution is multivariate Gaussian with mean $\boldsymbol{\mu}_0 \in \mathbb{R}^N$ and covariance $\Sigma_0 \in \mathbb{R}^{N \times N}$:

$$\mathbf{m} \sim \mathcal{N}(\boldsymbol{\mu}_0, \Sigma_0). \quad (24)$$

The i^{th} element of $\boldsymbol{\mu}_0$, denoted by μ_i^0 , represents the agent’s mean belief of the reward m_i associated with arm i . The (i, i) element of Σ_0 , denoted by $(\sigma_i^0)^2$, represents the agent’s uncertainty associated with that belief. Off-diagonal elements of Σ_0 , e.g., σ_{ij}^0 , represent the agent’s perceived relationship between m_i and m_j : if σ_{ij}^0 is positive, high values of m_i are correlated with high values of m_j , while if it is negative, high values of m_i correlate with low values of m_j . Any positive-definite matrix can be used as Σ_0 , but it is often useful to consider a structured parametrization, such as $\Sigma_0 = \sigma_0^2 \Sigma$, where $\sigma_0^2 > 0$ encodes the agent’s uncertainty. One important special case is an uncorrelated prior, where Σ is diagonal, which corresponds to the agent perceiving the rewards associated with different arms to be independent. Another important special case is an uninformative prior, which corresponds to complete uncertainty, i.e., the limit $\sigma_0^2 \rightarrow +\infty$; an uninformative prior can be thought of as a special case of an uncorrelated prior.

B. Inference Update

At each time t the agent picks an arm i_t and receives a reward r_t that is Gaussian distributed: $r_t \sim \mathcal{N}(m_{i_t}, \sigma_{s,i_t}^2)$. Bayesian inference provides an optimal solution to the problem of updating the belief state $(\boldsymbol{\mu}_t, \Sigma_t)$ (i.e., the sufficient statistics for estimating \mathbf{m}) to incorporate this new information. Let $\Lambda_t = \Sigma_t^{-1}$, and let $\boldsymbol{\phi}_t \in \mathbb{R}^N$ be the vector with element i_t equal to 1 and all other elements equal to zero. Then given the Gaussian prior (24), the Bayesian update equations are linear [17]:

$$\begin{aligned} \mathbf{q} &= \frac{r_t \boldsymbol{\phi}_t}{\sigma_{s,i_t}^2} + \Lambda_{t-1} \boldsymbol{\mu}_{t-1}, \\ \Lambda_t &= \frac{\boldsymbol{\phi}_t \boldsymbol{\phi}_t^T}{\sigma_{s,i_t}^2} + \Lambda_{t-1}, \quad \boldsymbol{\mu}_t = \Sigma_t \mathbf{q}. \end{aligned} \quad (25)$$

C. Decision Heuristic

At each time t the UCL algorithm computes a value Q_i^t for each arm i . The algorithm then picks the arm i_t that maximizes Q_i^t . That is, it picks

$$i_t = \arg \max_i Q_i^t. \quad (26)$$

The heuristic value Q_i^t is

$$Q_i^t = \mu_i^t + \sigma_i^t \Phi^{-1}(1 - \alpha_t), \quad (27)$$

where $\mu_i^t = (\boldsymbol{\mu}_t)_i$, $(\sigma_i^t)^2 = (\Sigma_t)_{ii}$, $\alpha_t = 1/(Kt)$, $K > 0$ is a tunable parameter, and $\Phi^{-1}(\cdot)$ is the quantile function of the standard normal random variable. The heuristic Q_i^t is a Bayesian upper limit for the value of m_i based on the information available at time t . It represents an optimistic assessment of the value of

m_i . The decision made can be thought of as the most optimistic one consistent with the current information.

D. Performance

In [27], we studied the case of homogeneous sampling noise (i.e., $\sigma_{s,i}^2 = \sigma_s^2$ for each i) and showed that the UCL algorithm achieves logarithmic cumulative expected regret uniformly in time. In particular, we proved that the following theorem holds. We define $\{R_t^{\text{UCL}}\}_{t \in \{1, \dots, T\}}$ as the sequence of expected regret for the deterministic UCL algorithm.

Theorem 7 (Regret of the Deterministic UCL Algorithm [27]): The following statements hold for the Gaussian multi-armed bandit problem and the deterministic UCL algorithm with uncorrelated uninformative prior and $K = 1$:

- 1) the expected number of times a suboptimal arm i is chosen until time T satisfies

$$\mathbb{E} [n_i^T] \leq \left(\frac{8\sigma_s^2}{\Delta_i^2} + 2 \right) \log T + 3 ;$$

- 2) the cumulative expected regret until time T satisfies

$$J_R = \sum_{t=1}^T R_t \leq \sum_{i=1}^N \Delta_i \left(\left(\frac{8\sigma_s^2}{\Delta_i^2} + 2 \right) \log T + 3 \right).$$

The implication of this theorem can be seen by comparing 1) with the Lai-Robbins bound (9): the UCL algorithm achieves logarithmic regret uniformly in time with a constant that differs from the optimal asymptotic one by a constant factor, and thus is considered to have optimal performance.

VI. ALGORITHMS FOR SATISFYING GAUSSIAN MULTI-ARMED BANDIT PROBLEMS

In this section we develop algorithms for solving Gaussian multi-armed bandit problems with the satisfying objectives proposed in Section III. All the algorithms consist of modified versions of the UCL algorithm. We analyze the algorithms and show that they achieve efficient performance. The UCL algorithm solves the standard Gaussian multi-armed bandit problem, i.e., the satisfying Gaussian multi-armed bandit problem with $\mathcal{M} > m_{\sigma(2)}$ and $\delta = 0$ (Problem 1). We develop three new UCL variants for Problems 2–4 in Table I. These algorithms can then be applied to Problems 5–8 in Table II. At the end of the section, we consider extensions to reward distributions other than the Gaussian with known variance.

A. Problem 2: Satisfaction-In-Mean-Reward UCL Algorithm

A simple modification of the UCL algorithm achieves logarithmic regret for the Gaussian satisfaction-in-mean-reward problem, which is the satisfying-in-mean-reward multi-armed bandit problem with $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta = 0$ (Problem 2). We define this algorithm, which we refer to as the *satisfaction-in-mean-reward UCL algorithm*, as follows.

As in (27), define the heuristic value Q_i^t as

$$Q_i^t = \mu_i^t + \sigma_i^t \Phi^{-1}(1 - \alpha_t),$$

where $\alpha_t = 1/(Kt)$ and $K > 0$ is again a tunable parameter.

Let $\mathcal{M} \in \mathbb{R}$ be the satisfaction threshold, so the agent is satisfied if it picks an arm with $m_i \geq \mathcal{M}$. Let the eligible set at time t be $\{i \mid Q_i^t \geq \mathcal{M}\}$. In contrast to the UCL selection scheme (26) that picks the arm with maximal Q_i^t , satisfaction-in-mean-reward UCL picks any arm in the eligible set. That is, if the eligible set is non-empty, then

$$i_t \in \{i \mid Q_i^t \geq \mathcal{M}\}, \quad (28)$$

or if the eligible set is empty, then satisfaction-in-mean-reward UCL picks the arm with maximal Q_i^t . Thus, if the most recently selected arm is in the eligible set, it may be selected again even if it does not have the maximal Q_i^t .

The satisfaction-in-mean-reward UCL algorithm achieves logarithmic cumulative expected satisfaction-in-mean-reward regret, as guaranteed by the following theorem.

Theorem 8 (Regret of the Satisfaction-In-Mean-Reward UCL Algorithm): Let a Gaussian multi-armed bandit problem with the satisfaction-in-mean-reward objective have at least one arm i that obeys $m_i > \mathcal{M}$, and, without loss of generality, assume $\sigma_{s,i}^2 = 1$ for each arm i . Then, the following statements hold for the satisfaction-in-mean-reward UCL algorithm with uncorrelated uninformative prior and $K = 1$:

- 1) the expected number of times a non-satisfying arm i is chosen until time T satisfies

$$\mathbb{E} [n_i^T] \leq \left(\frac{8}{(\Delta_i^{\mathcal{M}})^2} + 3 \right) \log T + 4;$$

- 2) the cumulative expected satisfaction-in-mean-reward regret until time T satisfies

$$J_{SM} \leq \sum_{i=1}^N \Delta_i^{\mathcal{M}} \left(\left(\frac{8}{(\Delta_i^{\mathcal{M}})^2} + 3 \right) \log T + 4 \right).$$

To prove Theorem 8 we use the following bound from [1].

Lemma 9 (Bounds on the Inverse Gaussian cdf): For the standard normal (i.e., Gaussian) random variable z and a constant $w \in \mathbb{R}_{\geq 0}$,

$$\Pr [z \geq w] \leq \frac{2e^{-w^2/2}}{\sqrt{2\pi}(w + \sqrt{w^2 + 8/\pi})} \leq \frac{1}{2} e^{-w^2/2}. \quad (29)$$

It follows from (29) that for any $\alpha \in [0.5, 1]$,

$$\Phi^{-1}(1 - \alpha) \leq \sqrt{-2 \log \alpha}. \quad (30)$$

Proof of Theorem 8: The proof proceeds as in the proof of Theorem 7 in [27], which itself follows the proofs in [4]. Let i be a non-satisfying arm, i.e., $m_i < \mathcal{M}$, and recall that i^* designates

the maximum mean reward. Then

$$\begin{aligned} \mathbb{E} [n_i^T] &= \sum_{t=1}^T \Pr [i_t = i] \\ &\leq \sum_{t=1}^T \Pr [Q_i^t \geq \mathcal{M}] \\ &\quad + \Pr \left[Q_i^t \geq Q_{i^*}^t \ \& \ \max_j Q_j^t < \mathcal{M} \right] \\ &\leq \eta + \sum_{t=1}^T (\Pr [Q_i^t \geq \mathcal{M}, n_i^t \geq \eta] \\ &\quad + \Pr [Q_i^t \geq Q_{i^*}^t, n_i^t \geq \eta]). \end{aligned}$$

The first term in the summand corresponds to the probability that the non-satisfying arm i is in the eligible set, while the second term corresponds to the probability that the eligible set is empty and that a non-satisfying arm appears better than an optimal arm.

The statement $Q_i^t \geq Q_{i^*}^t$ implies that at least one of the following inequalities holds:

$$\mu_i^t \geq m_i + C_i^t \quad (31)$$

$$\mu_{i^*}^t \leq m_{i^*} - C_{i^*}^t \quad (32)$$

$$m_{i^*} < m_i + 2C_i^t, \quad (33)$$

where $C_i^t = \sigma_i^t \Phi^{-1}(1 - \alpha_t)$ and $\alpha_t = 1/(Kt)$. Otherwise, if none of (31)–(33) holds, then

$$Q_{i^*}^t = \mu_{i^*}^t + C_{i^*}^t > m_{i^*} \geq m_i + 2C_i^t > \mu_i^t + C_i^t = Q_i^t.$$

We first analyze the probability that (31) holds. For an uncorrelated uninformative prior, μ_i^t is equal to \bar{m}_i^t , the empirical mean reward observed at arm i until time t , and $\sigma_i^t = 1/\sqrt{n_i^t}$. Therefore, for an uncorrelated uninformative prior,

$$Q_i^t = \bar{m}_i^t + \frac{1}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t).$$

Conditional on n_i^t , the empirical mean reward \bar{m}_i^t is itself a Gaussian random variable with mean m_i and standard deviation $1/\sqrt{n_i^t}$, so (31) holds if

$$\begin{aligned} \bar{m}_i^t &\geq m_i + \frac{1}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t) \\ \Leftrightarrow m_i + \frac{z}{\sqrt{n_i^t}} &\geq m_i + \frac{1}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t) \\ \Leftrightarrow z &\geq \Phi^{-1}(1 - \alpha_t), \end{aligned}$$

where $z \sim \mathcal{N}(0, 1)$ is a standard normal random variable. Thus, for an uninformative prior, $\Pr [(31) \text{ holds}] = \alpha_t = \frac{1}{Kt}$.

Similarly, (32) holds if

$$\begin{aligned} \bar{m}_{i^*}^t &\leq m_{i^*} - C_{i^*}^t \\ \Leftrightarrow m_{i^*} + \frac{z}{\sqrt{n_{i^*}^t}} &\leq m_{i^*} - \frac{1}{\sqrt{n_{i^*}^t}} \Phi^{-1}(1 - \alpha_t) \\ \Leftrightarrow z &\leq -\Phi^{-1}(1 - \alpha_t), \end{aligned}$$

where $z \sim \mathcal{N}(0, 1)$ is a standard normal random variable. Thus, for an uninformative prior, $\Pr [(32) \text{ holds}] = \alpha_t = \frac{1}{Kt}$.

Inequality (33) holds if

$$\begin{aligned} m_{i^*} &< m_i + \frac{2}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t) \\ \Leftrightarrow \Delta_i &< \frac{2}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t) \Leftrightarrow \frac{\Delta_i^2 n_i^t}{4} < -2 \log \alpha_t \quad (34) \\ \Rightarrow \frac{\Delta_i^2 n_i^t}{4} &< 2 \log t \Rightarrow \frac{\Delta_i^2 n_i^t}{4} < 2 \log T \end{aligned}$$

where $\Delta_i = m_{i^*} - m_i$ and inequality (34) follows from bound (30). Thus, for an uninformative prior, (33) never holds if

$$n_i^t \geq \frac{8}{\Delta_i^2} \log T. \quad (35)$$

Thus, for n_i^t sufficiently large, $\Pr [Q_i^t \geq Q_{i^*}^t] = 2/(Kt)$.

We now bound the probability $\Pr [Q_i^t \geq \mathcal{M}]$ that a non-satisfying arm i is in the eligible set. Note that $Q_i^t \geq \mathcal{M}$ implies that at least one of the following inequalities holds:

$$\mu_i^t \geq m_i + C_i^t \quad (36)$$

$$\mathcal{M} < m_i + 2C_i^t. \quad (37)$$

Otherwise, if neither (36) nor (37) holds, $\mathcal{M} \geq m_i + 2C_i^t > \mu_i^t + C_i^t = Q_i^t$ and arm i is not in the eligible set.

(36) is identical to (31) and (37) to (33). For an uninformative prior, $\Pr [(36) \text{ holds}] = \alpha_t = \frac{1}{Kt}$. And (37) holds if

$$\begin{aligned} \mathcal{M} &< m_i + \frac{2}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t) \\ \Leftrightarrow \Delta_i^{\mathcal{M}} &< \frac{2}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t) \\ \Leftrightarrow \frac{(\Delta_i^{\mathcal{M}})^2 n_i^t}{4} &< -2 \log(\alpha_t) \\ \Rightarrow \frac{(\Delta_i^{\mathcal{M}})^2 n_i^t}{4} &< 2 \log t \Rightarrow \frac{(\Delta_i^{\mathcal{M}})^2 n_i^t}{4} < 2 \log T. \end{aligned}$$

Thus, for an uninformative prior, (37) never holds if

$$n_i^t \geq \frac{8}{(\Delta_i^{\mathcal{M}})^2} \log T. \quad (38)$$

Since $m_{i^*} \geq \mathcal{M}$, for each non-satisfying arm i , $\Delta_i^{\mathcal{M}} \leq \Delta_i$. Thus, $1/(\Delta_i^{\mathcal{M}})^2 \geq 1/\Delta_i^2$ and (38) implies (35). So setting

$$\eta = \left\lceil \frac{8}{(\Delta_i^{\mathcal{M}})^2} \log T \right\rceil \quad (39)$$

yields the bound

$$\begin{aligned} \mathbb{E} [n_i^T] &\leq \eta + \sum_{t=1}^T (\Pr [Q_i^t \geq \mathcal{M}, n_i^t \geq \eta] \\ &\quad + \Pr [Q_i^t \geq Q_{i^*}^t, n_i^t \geq \eta]) \\ &< \left\lceil \frac{8}{(\Delta_i^{\mathcal{M}})^2} \log T \right\rceil + 3 \sum_{t=1}^T \frac{1}{t}. \end{aligned}$$

The sum can be bounded by the integral

$$\sum_{t=1}^T \frac{1}{t} \leq 1 + \int_1^T \frac{1}{t} dt = 1 + \log T,$$

yielding the bound in the first statement of the theorem:

$$\mathbb{E} [n_i^T] \leq \left(\frac{8}{(\Delta_i^{\mathcal{M}})^2} + 3 \right) \log T + 4.$$

The second statement of the theorem follows from the definition (12) of expected satisficing regret. ■

B. Problem 3: δ -Sufficing UCL Algorithm

An alternative modification of the UCL algorithm achieves finite satisficing regret in the Gaussian δ -sufficing problem, which is the satisficing-in-mean-reward multi-armed bandit with $\mathcal{M} > m_{\sigma(2)}$ and $\delta \in (0, 1]$ (Problem 3). For the agent, this can be thought of as wanting to have finite confidence that it has found the unknown optimal arm $\sigma(1)$. For the δ -sufficing problem, define the heuristic function

$$Q_i^t = \mu_i^t + \sigma_i^t \Phi^{-1} \left(1 - \frac{\delta}{2} \right).$$

We define the δ -sufficing UCL algorithm as the algorithm that selects arm $i_t = \arg \max_i Q_i^t$ at each decision time t . The δ -sufficing UCL algorithm achieves finite cumulative satisficing regret, as guaranteed by the following theorem.

Theorem 10: Consider the δ -sufficing UCL algorithm with an uninformative prior. The number of times the picked arm i_t is non-satisficing with probability greater than δ is upper bounded as

$$n_i^T < \frac{4\sigma_s^2}{\Delta_i^2} \left(\Phi^{-1} \left(1 - \frac{\delta}{2} \right) \right)^2 + 1.$$

Proof: We bound n_i^T by noting that a non-satisficing arm i is picked only if $Q_i^t \geq Q_{i^*}^t$, which can be decomposed as in the proof of Theorem 8 into the three conditions

$$\mu_i^t \geq m_i + C_i^t \quad (40)$$

$$\mu_{i^*}^t \leq m_{i^*} - C_{i^*}^t \quad (41)$$

$$m_{i^*} < m_i + 2C_i^t. \quad (42)$$

(42) is equivalent to

$$\Delta_i = m_{i^*} - m_i < 2C_i^t = \frac{2\sigma_s}{\sqrt{n_i^t}} \Phi^{-1}(1 - \delta/2).$$

Squaring and rearranging, we see that this never holds if

$$n_i^t > \frac{4\sigma_s^2}{\Delta_i^2} \left(\frac{\log(1/\delta)}{\log 2} + 1 \right) > \frac{4\sigma_s^2}{\Delta_i^2} (\Phi^{-1}(1 - \delta/2))^2 = \eta.$$

The same argument as in the proof of Theorem 8 shows that for $n_i^t \geq 1$, (40) and (41) each hold with probability at most $\delta/2$. Therefore, for $n_i^t > \eta + 1$, a non-satisficing arm is selected with probability at most δ . ■

Theorem 10 guarantees that the δ -sufficing UCL algorithm achieves finite regret. Furthermore, the algorithm is efficient

in that the regret matches the dependence on ϵ and δ in the bound (15). To see this, note that a non-satisficing arm i with Δ_i is an $\epsilon = \Delta_i$ -suboptimal arm, so Corollary 3 implies that n_i^T is lower bounded by $\mathcal{O}(\log(1/\delta)/\epsilon^2)$. The statement of Theorem 10 combined with the bound (30) on the inverse Gaussian cdf implies that n_i^T is upper bounded by $8\sigma_s^2 \log(2/\delta)/\Delta_i^2 + 1 = 8\sigma_s^2 \log(2/\delta)/\epsilon^2 + 1$, which matches the lower bound (15) up to constant factors.

C. Problem 4: (\mathcal{M}, δ) -Satisficing UCL Algorithm

A third modification of the UCL algorithm achieves finite satisficing regret in the Gaussian (\mathcal{M}, δ) -satisficing problem, which is the satisficing-in-mean-reward multi-armed bandit with $\mathcal{M} \leq m_{\sigma(2)}$ and $\delta \in (0, 1]$ (Problem 4). For the agent, this can be thought of as wanting to have finite confidence that it has found an arm whose mean reward is above a known threshold. For the (\mathcal{M}, δ) -satisficing problem, define the heuristic function

$$Q_i^t = \mu_i^t + \sigma_i^t \Phi^{-1} \left(1 - \frac{\delta}{3} \right).$$

Let the eligible set at time t be $\{i \mid Q_i^t \geq \mathcal{M}\}$. We define the (\mathcal{M}, δ) -satisficing UCL algorithm as the algorithm that selects arm $i_t \in \{i \mid Q_i^t \geq \mathcal{M}\}$, if the eligible set at time t is non-empty. Otherwise, if the eligible set is empty, the algorithm picks the arm with maximal Q_i^t .

The (\mathcal{M}, δ) -satisficing UCL algorithm achieves efficient performance as guaranteed by the following theorem.

Theorem 11: Consider the (\mathcal{M}, δ) -satisficing UCL algorithm with an uninformative prior. The number of times the picked arm i_t is non-satisficing with probability greater than δ is upper bounded as

$$n_i^T < \frac{4\sigma_s^2}{(\Delta_i^{\mathcal{M}})^2} (\Phi^{-1}(1 - \delta/3))^2 + 1.$$

Proof: The proof is very similar to the proofs of Theorems 8 and 10. As in Theorem 8, we bound n_i^T by

$$\begin{aligned} n_i^T &= \sum_{t=1}^T \mathbf{1}(i_t = i) \\ &\leq \eta + \sum_{t=1}^T (\mathbf{1}(Q_i^t \geq \mathcal{M}, n_i^t \geq \eta) + \mathbf{1}(Q_i^t \geq Q_{i^*}^t, n_i^t \geq \eta)). \end{aligned}$$

The condition $Q_i^t \geq \mathcal{M}$, which means arm i is in the eligible set, can be decomposed into the two conditions

$$\mu_i^t \geq m_i + C_i^t \quad (43)$$

$$\mathcal{M} < m_i + 2C_i^t. \quad (44)$$

Equation (44) is equivalent to

$$\Delta_i^{\mathcal{M}} = \mathcal{M} - m_i < 2C_i^t = \frac{2\sigma_s}{\sqrt{n_i^t}} \Phi^{-1}(1 - \delta/3).$$

Squaring and rearranging, we see that (44) never holds if

$$n_i^t > \frac{4\sigma_s^2}{(\Delta_i^{\mathcal{M}})^2} (\Phi^{-1}(1 - \delta/3))^2 = \eta.$$

The same argument as in the proof of Theorem 10 shows that for $n_i^t \geq 1$, (43) holds with probability at most $\delta/3$, so $n_i^t > \eta$ implies that a non-satisfying arm is in the eligible set with probability at most $\delta/3$.

As in the proof of Theorem 10, a non-satisfying arm i is picked due to the eligible set being empty only if $Q_i^t \geq Q_{i^*}^t$, where i^* is the arm with maximal mean reward. This condition can again be decomposed into the three conditions (40)–(42). Equation (42) does not hold if $n_i^t > \eta$, so the probability that $Q_i^t \geq Q_{i^*}^t$ is bounded by the probability that either (40) or (41) holds. For $n_i^t > 1$, each of these holds with probability $\delta/3$, so the probability of a non-satisfying arm being chosen due to the eligible set being empty is at most $2\delta/3$. Thus, for $n_i^t > \eta + 1$, a non-satisfying arm is selected with probability at most δ . ■

Theorem 11 guarantees that the (\mathcal{M}, δ) -satisficing UCL algorithm achieves finite regret. Furthermore, the algorithm is efficient in that the regret matches the dependence on ϵ and δ in the bound (16). Applying the bound (30) on the inverse Gaussian cdf to the statement in the theorem, we see that n_i^T is upper bounded by $8\sigma_s^2 \log(3/\delta) / (\Delta_i^{\mathcal{M}})^2$. Summing this bound over non-satisfying arms i shows that the total number of times the algorithm incurs regret is at most $8\sigma_s^2 \log(3/\delta) \sum_{\{i | \Delta_i^{\mathcal{M}} > 0\}} 1 / (\Delta_i^{\mathcal{M}})^2$. This matches the dependence on ϵ and δ in the bound (16) up to constant factors.

Note that lower bound (16) counts the number of selections of all arms including the optimal arm, while the upper bound counts only the suboptimal arms. Hence, we can only claim that we achieve cumulative regret bounded in T . With a better lower bound on n_i^T , we may be able to claim that, similar to δ -sufficing UCL, (\mathcal{M}, δ) -sufficing UCL achieves the optimal dependence on ϵ and δ . However, this remains an open problem to pursue.

D. Robust Satisficing UCL Algorithms

The UCL algorithm solves Problem 1, the Gaussian standard problem. The modified versions of the UCL algorithm in Sections VI-A, VI-B, and VI-C solve the other three Gaussian satisficing-in-mean-reward Problems 2–4. All four UCL algorithms achieve efficient performance in solving their respective problems, as guaranteed by Theorems 8, 10, and 11.

The equivalence result of Lemma 5 shows for Gaussian distributed rewards that we can modify the four UCL algorithms developed for Problems 1–4 to solve Problems 5–8 as follows. The modified UCL algorithms make decisions based on the standardized mean reward (20) using priors on the standardized mean rewards. A prior belief $\mathbf{m} \sim \mathcal{N}(\boldsymbol{\mu}_0, \Sigma_0)$ on the mean rewards \mathbf{m} is transformed into a prior belief on the standardized mean rewards $\mathbf{x} \sim \mathcal{N}(\tilde{\boldsymbol{\mu}}_0, \tilde{\Sigma}_0)$ by

$$(\tilde{\boldsymbol{\mu}}_0)_i = ((\boldsymbol{\mu}_0)_i - M) / \sigma_{s,i}, \quad (\tilde{\Sigma}_0)_{ij} = (\Sigma_0)_{ij} / (\sigma_{s,i} \sigma_{s,j}).$$

Problem 5: Robust UCL Algorithm: The *robust UCL algorithm* is the UCL algorithm where the prior is given in terms of the standardized mean rewards, and the observed reward r_t is standardized according to the transformation (22) before being input to the inference equations (25).

Problem 6: Robust Satisfaction UCL Algorithm: The *robust satisfaction UCL algorithm* is the satisfaction-in-

mean-reward UCL algorithm where the prior is given in terms of the standardized mean rewards, the observed reward r_t is standardized according to the transformation (22) before being input to the inference equations (25), and the parameter \mathcal{M} is set equal to $X = \Phi^{-1}(\Pi)$ defined in (23).

Problem 7: δ -Robust Sufficing Algorithm: The *δ -robust sufficing UCL algorithm* is the δ -sufficing UCL algorithm where the prior is given in terms of the standardized mean rewards, and the observed reward r_t is standardized according to the transformation (22) before being input to the inference equations (25).

Problem 8: (Π, δ) -Robust Sufficing Algorithm: The *(Π, δ) -robust sufficing UCL algorithm* is the (\mathcal{M}, δ) -satisficing UCL algorithm where the prior is given in terms of the standardized mean rewards, the observed reward r_t is standardized according to the transformation (22) before being input to the inference equations (25), and the parameter \mathcal{M} is set equal to $X = \Phi^{-1}(\Pi)$ defined in (23).

Lemma 5 implies that the performance guarantees that hold for the UCL algorithms developed for Problems 1–4 also hold for the four new UCL algorithms defined above when applied to Problems 5–8.

E. Relaxations of Gaussian and Known Variance Assumptions

The algorithms presented so far have been developed assuming that the reward distribution associated with each arm i is Gaussian with unknown mean m_i and known variance $\sigma_{s,i}^2$. The reward variance may be known, e.g., estimated from known sensor characteristics or prior data. When the reward variance is not known, a simple modification to the heuristic (27) yields an algorithm that achieves efficient performance. Similar simple modifications extend our results to the case where the reward distribution is sub-Gaussian, which includes distributions with bounded support. We state modifications for the case of an uninformative prior. Prior information can be incorporated using a conjugate prior, as discussed in [24].

Remark 12 (Gaussian Rewards With Unknown Variance): When the reward distribution is Gaussian with unknown variance, the heuristic developed by Auer *et al.* [4] for their algorithm UCB1-NORMAL results in algorithms that achieve efficient performance. Recall that n_i^t is the number of times arm i has been selected up to time t , and \bar{m}_i^t is the empirical mean reward observed at arm i up to time t . Define

$$q_i^t = \sum_{t | i_t = i} r_t^2$$

as the sum of the squared rewards obtained from arm i .

The UCB1-NORMAL algorithm is composed of two rules: if there is an arm that has been played less than $\lceil 8 \log t \rceil$ times, it selects that arm. Otherwise it selects the arm i that maximizes the heuristic

$$Q_i^{t, \text{UCB1-NORMAL}} = \bar{m}_i^t + \sqrt{16 \frac{q_i^t - n_i^t (\bar{m}_i^t)^2 \log t}{n_i^t - 1} \frac{1}{n_i^t}}.$$

This heuristic can be used directly in the standard and satisfaction-in-mean-reward UCL algorithms. For the δ -sufficing and (\mathcal{M}, δ) -satisficing UCL algorithms, use $\tilde{k} = 2$ and $\tilde{k} = 3$, respectively, in the heuristic

$$Q_i^t = \bar{m}_i^t + \sqrt{4 \frac{q_i^t - n_i^t (\bar{m}_i^t)^2 \log(\tilde{k}/\delta)}{n_i^t - 1} \frac{1}{n_i^t}}.$$

The Gaussian distribution with unknown mean and variance is again a location-scale family, so Lemma 5 implies that these modified algorithms can be used to solve the robust satisficing problems as well.

Prior information can be incorporated by means of a conjugate prior, as discussed in [24]. For the following remarks, define the generalized heuristic $\bar{Q}_i^t(\beta) = \bar{m}_i^t + \sqrt{\frac{\beta \log t}{n_i^t}}$.

Remark 13 (Sub-Gaussian Rewards): Another generalization of Gaussian rewards with known variance is the case where the reward distribution is sub-Gaussian, also known as light-tailed. The distribution of a random variable X is called sub-Gaussian if its moment generating function $M(u) = \mathbb{E}[\exp(uX)]$ is finite for all $u \in \mathbb{R}$. Then, one can find a constant ζ such that $M(u) \leq \exp(\zeta u^2/2)$ [9].

In this case, a heuristic function due to Liu and Zhou [21] $Q_i^{t,SG} = \bar{Q}_i^t(8\zeta)$ can be used to achieve efficient performance.

Remark 14 (Reward Distributions With Bounded Support): Another common assumption in the bandit literature is that the reward distributions are arbitrary but have a known bounded support $[a, b] \subset \mathbb{R}$. Without loss of generality, we assume that the support is contained in the unit interval $[0, 1]$. In this case the UCB1 heuristic due to Auer *et al.* [4] $Q_i^{t,UCB1} = \bar{Q}_i^t(2)$ can be used in the standard and satisfaction-in-mean-reward UCL algorithms.

For the δ -sufficing and (\mathcal{M}, δ) -satisficing UCL algorithms, $\tilde{k} = 2$ and $\tilde{k} = 3$, respectively, in the heuristic $Q_i^t = \bar{Q}_i^{(\tilde{k}/\delta)}(1/2)$ can be used to achieve efficient performance.

For the robust satisficing problems the relevant reward, happiness h_t (17), is a Bernoulli random variable which is supported on $[0, 1]$. Therefore, each robust satisficing problem can be solved by the appropriate variant of UCB1. However, if additional information is available about the distribution of the raw rewards r_t , e.g., that they are Gaussian with known variance, then the robust UCL algorithms can achieve improved performance relative to UCB1, for example if the Kullback-Leibler divergence between the r_t distributions is larger than the Bernoulli distributions associated with h_t .

Additional extensions to heavy-tailed distributions may be possible using the techniques of [7].

VII. NUMERICAL EXAMPLES

In this section, we present the results of numerical simulations of the modified UCL algorithms solving multi-armed bandit problems with Gaussian rewards and satisficing objectives. We consider both thresholding in the mean rewards m_i , as in Problems 1–4 (Table I), and thresholding in the instantaneous rewards, as in Problems 5–8 (Table II). In all of the cases presented, the algorithms used an uninformative prior. We use the

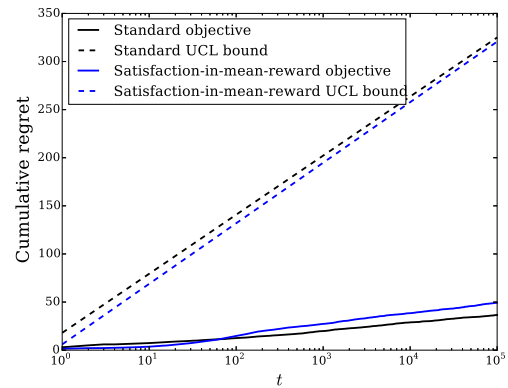


Fig. 1. Comparison of regret incurred by the UCL algorithms when solving the standard problem (Problem 1) and satisfaction-in-mean-reward problem (Problem 2). Both problems define regret by thresholding mean reward values; the standard bandit objective incurs regret when the mean reward of the chosen option is less than the maximum reward m_{i^*} , while the satisfaction-in-mean-reward problem incurs regret when the mean reward is less than $\mathcal{M} \leq m_{\sigma(2)}$, here set equal to 2.5. For both problems, the cumulative expected regret and its upper bound increase at a logarithmic rate since the agent seeks certainty that its threshold is met, which it cannot achieve in finite time.




Fig. 2. Comparison of regret incurred by the UCL algorithms when solving the δ - and (\mathcal{M}, δ) -satisficing problems, (Problems 3 and 4, respectively). As in Fig. 1, the problems define regret by thresholding the mean reward values; the δ -sufficing objective incurs regret when the mean reward of the chosen option is less than the maximum reward m_{i^*} , while the (\mathcal{M}, δ) -sufficing problem incurs regret when the mean reward is less than $\mathcal{M} \leq m_{\sigma(2)}$, here set equal to 2.5. In contrast to Fig. 1, the agent only seeks to have $1 - \delta = 95\%$ confidence that its threshold is met, which it can achieve in finite time. Thus, the upper bounds on cumulative expected regret are constant functions of horizon length and the mean regret plateaus at a finite value.

simulations to illustrate performance of the algorithms relative to the bounds proved in the theorems of Section IV. We also use the simulations to compare how the different algorithms trade off accumulation of reward with reduction in exploration cost as measured by number of switches among arms. As shown in the figures, satisficing can significantly decrease the exploration cost while incurring little cost in terms of the rewards received by the agent.

We first consider the satisficing objectives with thresholding in the mean rewards. We illustrate how the objectives of Problems 1 and 2 yield logarithmic regret (Fig. 1) whereas the objectives of Problems 3 and 4 yield finite regret (Fig. 2), as predicted by the bounds proved in Theorems 7, 8, 10 and 11.

For the simulations presented in Figs. 1 and 2, we set $N = 4$. The mean rewards \mathbf{m} were set equal to $[1\ 2\ 3\ 4]$ and the standard deviations $\sigma_{s,i}$ were each set equal to 1.

In Fig. 1, the agent’s regret is defined by comparing the mean rewards m_i with the satisfaction threshold \mathcal{M} . For the standard objective (Problem 1) the satisfaction level \mathcal{M} was set equal to $m_{i^*} = 4$, so the agent incurred regret if it selected any other than $i^* = 4$. For the satisfaction-in-mean-reward objective (Problem 2) the satisfaction level \mathcal{M} was set equal to 2.5, so the agent incurred regret if it selected arms 1 or 2. Fig. 1 plots mean cumulative regret from 100 simulations (solid lines) and the bounds on regret from Theorems 7 and 8 (dashed lines) for the standard and satisfaction-in-mean-reward UCL algorithms, respectively. We observe from Fig. 1 that the algorithms’ regret is significantly below the bounds, indicating that both bounds are conservative. Because both objectives set the sufficiency threshold $\delta = 0$ and define regret in terms of the unknown mean rewards m_i , the agent must achieve certainty about the mean reward values to stop incurring regret. It is impossible for the agent to achieve this certainty in finite time, so the mean regret and its bound both increase indefinitely at a logarithmic rate for both objectives.

Fig. 2, in contrast, shows that by setting the sufficiency threshold δ to a non-zero value, one can achieve finite regret. All the parameters for the simulations shown in Fig. 2 were identical to those for the simulations in Fig. 1, except that the sufficiency threshold δ was set equal to 0.05. Setting δ to a non-zero value transformed the standard and satisfaction-in-mean-reward objectives into the δ -sufficing and (\mathcal{M}, δ) -satisficing objectives, Problems 3 and 4, respectively. For these objectives, regret is again defined by thresholding the mean rewards m_i . However, rather than seeking certainty that the threshold is met, the agent only seeks to ensure that its threshold is met with a probability of at least $1 - \delta$. Because of the allowance of uncertainty, the agent only needs to perform a finite amount of exploration before settling on an arm that appears satisfying. The regret bounds in the figure follow from Theorems 10 and 11 for the δ -sufficing and (\mathcal{M}, δ) -satisficing objectives, respectively. As in Fig. 1, we observe from Fig. 2 that both bounds are conservative.

The reduction in exploring that comes with sufficing can be advantageous when exploring is costly, e.g., when there is a cost associated with making a switch from one arm to another. Fig. 3 suggests that this reduction in cost may require little sacrifice in reward. The upper panel plots mean cumulative reward from 100 simulations for both the standard bandit problem 1 and the δ -sufficing problem 3. The curve for δ -sufficing is slightly below that for the standard bandit problem, showing that it results in slightly lower cumulative rewards, but the difference is insignificant in comparison to the overall magnitude of the cumulative rewards. The lower panel plots the mean cumulative number of switches between arms for both algorithms and shows that δ -sufficing requires roughly half as much exploration.

We next consider the satisficing objectives with thresholding in the instantaneous rewards. Fig. 4 presents a simulation that demonstrates the equivalence result of Lemma 5 for the robust bandit (Problem 5) using the robust UCL algorithm. The happiness threshold M was set equal to 2. As in Figs. 1 and 2,




Fig. 3. Upper: Mean cumulative reward accrued by standard (Problem 1) and δ -sufficing (Problem 3) UCL algorithms. Lower: Mean cumulative number of switches between arms, quantifying the algorithms’ exploration costs. The δ -sufficing UCL algorithm achieves nearly the same cumulative rewards as the standard UCL algorithm, but with roughly half the exploration cost.

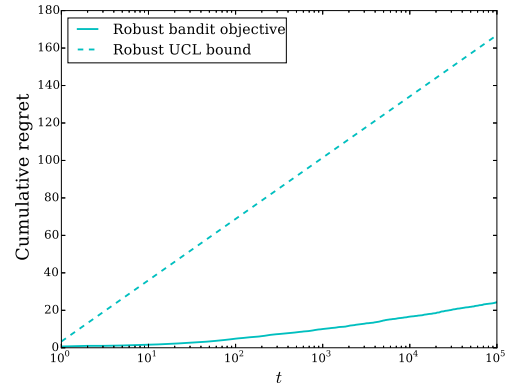


Fig. 4. Regret incurred by the robust UCL algorithm (Problem 5). The agent is happy if it receives a reward r_t that is at least equal to M , here set equal to 2. The agent seeks to maximize its probability of being happy at each time, so it incurs regret if it chooses an option with less than maximal happiness probability. As in Fig. 1, the agent seeks certainty that it maximizes its happiness probability, which it cannot achieve in finite time. So it incurs regret that increases indefinitely at a logarithmic rate.

the mean rewards \mathbf{m} were set equal to $[1\ 2\ 3\ 4]$, but for this simulation the standard deviations were set equal to $[1\ 1\ 1\ 3]$. So the standardized mean rewards were $\mathbf{x} = [-1\ 0\ 1\ \frac{2}{3}]$ and $i^* = 3$ the optimal arm, i.e., the arm with maximal happiness probability. Fig. 4 shows mean cumulative regret from 100 simulations (solid line) and the regret bound (dashed line) implied by Theorem 7 and the definition (19) of satisfaction. Because the objective requires identifying the arm with highest probability of satisfaction with certainty, both the regret and its upper bound increase indefinitely at a logarithmic rate, as in the problems illustrated in Fig. 1.

Fig. 5 shows the benefit of combining the robust objective with sufficing in the δ -robust bandit objective (Problem 7). By doing so, it is possible to retain the robustness benefit of the robust bandit objective (Problem 5) relative to the standard bandit objective (Problem 1) while reducing the exploration cost. The parameter values used in the δ -robust UCL algorithm (Problem 7) are the same as in Fig. 4, while the sufficiency parameter δ was set equal to 0.05. Fig. 5 shows how robust




Fig. 5. Upper: Mean cumulative number of times the agent is happy (i.e., has reward $r_t \geq M = 2$, (17)) when using the standard (Problem 1), robust UCL (Problem 5), and δ -robust (Problem 7) UCL algorithms. Lower: Mean cumulative number of switches between arms, quantifying the algorithms' exploration costs. The robust bandit objective is more robust in the sense that it is more likely to achieve a reward that is above the threshold M , but in this case incurs a greater exploration cost. The robust δ -sufficing objective combines the positive aspects of both the δ -sufficing and robust bandit objectives: exhibiting high robustness in terms of agent happiness and minimizing exploration costs.

satisficing algorithms (Problems 5 and 7) outperform the standard algorithm (Problem 1) for performance with respect to instantaneous reward as measured by happiness (17). There is a switching cost associated with achieving the higher rate of cumulative happiness. However, this cost is significantly reduced for the δ -robust algorithm (Problem 7), where sufficing is included, as compared to the robust algorithm (Problem 5) and it approaches the switching cost incurred by the standard UCL algorithm.

VIII. CONCLUSION

Satisficing, the concept of doing well relative to a reference value, is a useful alternative to maximizing that can be applied to a variety of decision-making scenarios. In this paper, we considered the multi-armed bandit problem using satisficing objectives. The multi-armed bandit problem is a canonical decision-making problem that is widely studied in machine learning and adaptive control using a maximization objective.

We proposed a system of eight objectives for stochastic multi-armed bandit problems that generalize the standard multi-armed bandit problem by capturing aspects of satisficing, notably thresholding effects that we termed *satisfaction* and *sufficiency*. We showed that each of the four objectives of Problems 1–4, defined by thresholding the unknown mean reward m_i associated with each arm, is equivalent to a related problem studied in the existing literature. We used these equivalences to derive bounds on efficient performance. For the four objectives of Problems 5–8 defined by thresholding the observed rewards r_t , we showed that, when the reward distributions belong to a location-scale family, each objective is equivalent to one of the first four objectives.

We then specialized to the case of Gaussian rewards (a particular location-scale family) and developed four variants of the UCL algorithm [27] to solve Problems 1–4 defined by threshold-

ing the mean rewards. We analyzed each algorithm and showed that it achieved efficient satisficing performance. We used the equivalency result (Lemma 5) to show how to apply the four variants of the UCL algorithm to Problems 5–8 defined by thresholding the instantaneous rewards and again achieve efficient satisficing performance.

Satisficing objectives that threshold the mean can reduce exploration costs (and thus risk) as compared to the standard problem where the objective is to maximize expected reward with certainty. Satisficing objectives that threshold observed rewards can result in more risk-averse and robust algorithms than objectives that account only for mean rewards [26]. Risk aversion and robustness are important for engineering applications (where standard bandit algorithms are known to have poor risk-aversion characteristics [3]). Thus, our proposed algorithms can be usefully applied to a range of engineering problems, notably those involving design of control policies.

Risk aversion and robustness are also important in the field of optimal foraging theory [8]. Foraging has been studied using the multi-armed bandit framework with a maximizing objective [19], [18], [32]. The satisficing objectives and algorithms that we have proposed in the present paper may provide an even more biologically plausible framework.

For any satisficing problem, selecting the appropriate satisfaction threshold remains an open problem. The results presented here provide an efficient policy once the satisfaction threshold has been chosen but leave the selection of the threshold up to the end user. The problem of choosing a threshold also arises when applying the Sequential Probability Ratio Test (SPRT) in hypothesis testing [35]. To apply the SPRT, one must select desired probabilities of type I and type II errors. Once these probabilities are selected the SPRT provides the optimal policy, but the SPRT itself does not provide optimal values for the error probabilities.

In many decision-making scenarios, maximizing a reward rate is used to optimize error probabilities [5]. A reward rate criterion may provide an optimal threshold for the satisficing algorithms developed here, but the specific criterion will depend on the decision-making scenario. Natural extensions also include considering cases where the mean rewards m_i are allowed to evolve over time, for example according to a jump process [33]; such evolution will likely encourage satisficing policies that incorporate adaptive thresholds. Finally, it is well understood that satisficing is an important feature of human decision making [29] and that the UCL algorithm can model many features of human decision making in bandit tasks [27]. New empirical work should be undertaken to compare the satisficing UCL algorithms with human behavior. It is an open empirical question to determine which of our notions of regret best explains human behavior.

ACKNOWLEDGEMENT

The authors wish to thank S. A. Levin for helpful discussions, as well as the three anonymous referees and the editor for their feedback which greatly strengthened the paper.

REFERENCES

- [1] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*. Dover Publications: Mineola, NY, 1964.
- [2] V. Anantharam, P. Varaiya, and J. Walrand, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays—Part I: I.I.D. rewards," *IEEE Trans. Autom. Control*, vol. 32, no. 11, pp. 968–976, 1987.
- [3] J.-Y. Audibert, R. Munos, and C. Szepesvári, "Exploration-exploitation tradeoff using variance estimates in multi-armed bandits," *Theoretical Comput. Sci.*, vol. 410, no. 19, pp. 1876–1902, 2009.
- [4] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Mach. Learning*, vol. 47, no. 2, pp. 235–256, 2002.
- [5] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen, "The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks," *Psychological Review*, vol. 113, no. 4, pp. 700–765, 2006.
- [6] R. Bordley and M. LiCalzi, "Decision analysis using targets instead of utility functions," *Dec. Economics and Finance*, vol. 23, no. 1, pp. 53–74, 2000.
- [7] S. Bubeck, N. Cesa-Bianchi, and G. Lugosi, "Bandits with heavy tail," *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7711–7717, 2013.
- [8] Y. Carmel and Y. Ben-Haim, "Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice?," *The Amer. Naturalist*, vol. 166, no. 5, pp. 633–641, 2005.
- [9] P. Chareka, O. Chareka, and S. Kennedy, "Locally sub-Gaussian random variables and the strong law of large numbers," *Atlantic Electron. J. Math.*, vol. 1, no. 1, pp. 75–81, 2006.
- [10] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, Apr. 2015.
- [11] E. Even-Dar, S. Mannor, and Y. Mansour. PAC bounds for multi-armed bandit and Markov decision processes. In *Computational Learning Theory*, pp. 255–270. Springer, 2002.
- [12] E. Even-Dar, S. Mannor, and Y. Mansour, "Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems," *J. Mach. Learning Research*, vol. 7, pp. 1079–1105, 2006.
- [13] A. Garivier and O. Cappé, "The KL-UCB algorithm for bounded stochastic bandits and beyond," In *JMLR: Workshop and Conference Proceedings*, vol. 19, COLT 2011, pp. 359–376, 2011.
- [14] M. Goodrich, W. Stirling, and R. Frost, "A theory of satisficing decisions and control," *IEEE Trans. Syst., Man and Cybern. A: Syst. Humans*, vol. 28, no. 6, pp. 763–779, 1998.
- [15] S. Kalyanakrishnan, A. Tewari, P. Auer, and P. Stone, "PAC subset selection in stochastic multi-armed bandits," In *Proc. 29th Int. Conf. on Mach. Learning (ICML-12)*, pp. 655–662, 2012.
- [16] E. Kaufmann, O. Cappé, and A. Garivier, "On Bayesian upper confidence bounds for bandit problems," In *Int. Conf. Artificial Intelligence and Statistics*, pp. 592–600. La Palma, Canary Islands, Spain, Apr. 2012.
- [17] S. M. Kay. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Prentice Hall: Englewood Cliffs, NJ, 1993.
- [18] T. Keasar, E. Rashkovich, D. Cohen, and A. Shmida, "Bees in two-armed bandit situations: Foraging choices and possible decision mechanisms," *Behavioral Ecology*, vol. 13, no. 6, pp. 757–765, 2002.
- [19] J. R. Krebs, A. Kacelnik, and P. Taylor, "Test of optimal sampling by foraging great tits," *Nature*, vol. 275, no. 5675, pp. 27–31, 1978.
- [20] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in Appl. Math.*, vol. 6, no. 1, pp. 4–22, 1985.
- [21] K. Liu and Q. Zhao, "Extended UCB policy for multi-armed bandit with light-tailed reward distributions," arXiv:1112.1768, Dec. 2011.
- [22] S. Mannor and J. N. Tsitsiklis, "The sample complexity of exploration in the multi-armed bandit problem," *J. Mach. Learning Research*, vol. 5, pp. 623–648, 2004.
- [23] T. M. Moe, "The new economics of organization," *Amer. J. of Political Sci.*, vol. 28, no. 4, pp. 739–777, 1984.
- [24] K. P. Murphy, "Conjugate Bayesian analysis of the Gaussian distribution," <https://www.cs.ubc.ca/~murphyk/Papers/bayesGauss.pdf>, 2007.
- [25] H. Nakayama and Y. Sawaragi. Satisficing trade-off method for multi-objective programming. In *Interactive Decision Analysis*, pp. 113–122. Springer: New York, 1984.
- [26] P. Reverdy and N. E. Leonard, "Satisficing in Gaussian bandit problems," In *Proc. IEEE Conf. Decision and Control*, pp. 5718–5723, 2014.
- [27] P. Reverdy, V. Srivastava, and N. E. Leonard, "Modeling human decision-making in generalized Gaussian multi-armed bandits," *Proc. IEEE*, vol. 102, no. 4, pp. 544–571, 2014.
- [28] H. Robbins, "Some aspects of the sequential design of experiments," *Bulletin of the Amer. Math. Soc.*, vol. 58, pp. 527–535, 1952.
- [29] B. Schwartz, A. Ward, J. Monterosso, S. Lyubomirsky, K. White, and D. R. Lehman, "Maximizing versus satisficing: Happiness is a matter of choice," *J. Personality and Social Psychology*, vol. 83, no. 5, pp. 1178, 2002.
- [30] H. A. Simon, "A behavioral model of rational choice," *The Quarterly J. of Econ.*, vol. 69, no. 1, pp. 99–118, 1955.
- [31] H. A. Simon, "Rational choice and the structure of the environment," *Psychological Review*, vol. 63, no. 2, pp. 129, 1956.
- [32] V. Srivastava, P. Reverdy, and N. E. Leonard, "On optimal foraging and multi-armed bandits," In *Proc. of the 51st Annu. Allerton Conf. on Commun., Control, and Computing*, pp. 494–499, 2013.
- [33] V. Srivastava, P. Reverdy, and N. E. Leonard, "Surveillance in an abruptly changing world via multiarmed bandits," In *Proc. IEEE Conf. Decision and Control*, pp. 692–697, 2014.
- [34] L. G. Valiant, "A theory of the learnable," *Commun. of the ACM*, vol. 27, no. 11, pp. 1134–1142, Nov. 1984.
- [35] A. Wald, "Sequential tests of statistical hypotheses," *Ann. of Math. Stat.*, vol. 16, no. 2, pp. 117–186, 1945.
- [36] D. Ward, "The role of satisficing in foraging theory," *Oikos*, pp. 312–317, 1992.
- [37] S. G. Winter, "The satisficing principle in capability learning," *Strategic Management Journal*, vol. 21, no. 10–11, pp. 981–996, 2000.
- [38] B. Yin *et al.*, "Finding optimal solution for satisficing non-functional requirements via 0–1 programming" In *Proc. IEEE 37th Annu. Computer Software and Applications Conf.*, pp. 415–424, 2013.

Paul Reverdy (M'14) received the B.S. degree in engineering physics and the B.A. degree in applied mathematics from the University of California, Berkeley, Berkeley, CA, USA, in 2007 and the M.A. and Ph.D. degrees in mechanical and aerospace engineering from Princeton University, Princeton, NJ, USA, in 2011 and 2014, respectively.

From 2007 to 2009, he worked as a Research Assistant at the Federal Reserve Board of Governors, Washington, DC, USA. He is currently a Postdoctoral Fellow with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA. His research interests are in the areas of control and robotics with current interests in human and automated decision making, engineering design, and navigation.

Vaibhav Srivastava received the B.Tech. degree in mechanical engineering from the Indian Institute of Technology Bombay, Mumbai, India, in 2007, and the M.S. degree in mechanical engineering, the M.A. degree in statistics, and the Ph.D. degree in mechanical engineering from the University of California at Santa Barbara, Santa Barbara, CA, in 2011 and 2012, respectively.

He served as a Lecturer and Associate Research Scholar with the Mechanical and Aerospace Engineering Department, Princeton University, Princeton, NJ from 2013–2016. He is an Assistant Professor of Electrical and Computer Engineering at Michigan State University. His research interests include modeling and analysis of human cognition; shared autonomous systems; socio-cognitive networks; computational networks; and robotic search and surveillance problems.

Naomi Ehrich Leonard (F'07) received the B.S.E. degree in mechanical engineering from Princeton University, Princeton, NJ, in 1985 and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1991 and 1994, respectively.

From 1985 to 1989, she was an Engineer in the electric power industry. She is the Edwin S. Wilsey Professor of Mechanical and Aerospace Engineering and Director of the Council on Science and Technology at Princeton University. She is also an associated faculty member of Princeton University's Program in Applied and Computational Mathematics. Her research and teaching are in control and dynamical systems with current interests in coordinated control for multi-agent systems, mobile sensor networks, collective animal behavior, and human decision-making dynamics.

Correction to “Satisficing in Multi-Armed Bandit Problems”

Paul Reverdy, Vaibhav Srivastava, and Naomi Ehrich Leonard

Abstract—An unfortunate mistake in the proof of Theorem 8 of the above paper is corrected.

We correct an error in the published proof of Theorem 8 of [2]. The error arises from an incorrect application of concentration inequalities. The correction follows the same structure as that published in [3, Appendix G], which corrects the proofs of performance bounds for UCL algorithms in [4] and thus Theorems 7 and 8 of [2]. For simplicity of presentation, we first state the correction and then provide the associated proof.

The heuristic value Q_i^t in [2, (27)] is

$$Q_i^t = \mu_i^t + \sigma_i^t \Phi^{-1}(1 - \alpha_t). \quad (\text{C1})$$

To correct Theorem 8 of [2], set $\alpha_t = 1/(Kt^a)$ with $a > 4/(3(1 - \epsilon^2/16))$, $\epsilon \in (0, 4)$, and $K = \sqrt{2\pi}e$. The last part of the statement of [2, Theorem 8] should be replaced by

“Then, the following statements hold for the satisfaction-in-mean-reward UCL algorithm with uncorrelated uninformative prior and $K = \sqrt{2\pi}e$:

- 1) the expected number of times a non-satisfying arm i is chosen until time T satisfies

$$\mathbb{E}[n_i^T] \leq \left(\frac{8a}{(\Delta_i^{\mathcal{M}})^2} \right) \log T + o(\log T);$$

- 2) the cumulative expected satisfaction-in-mean-reward regret until time T satisfies

$$J_{SM} \leq \sum_{i=1}^N \Delta_i^{\mathcal{M}} \left(\frac{8a}{(\Delta_i^{\mathcal{M}})^2} \right) \log T + o(\log T).”$$

For the δ -sufficing and (\mathcal{M}, δ) -satisficing UCL algorithms of [2], similar corrections also hold with Q_i^t defined by (C1) and a modification to α_t . For these algorithms, the modification to α_t and its consequences can be succinctly stated by referring to the following Lemma which is a straightforward application of Theorem 2 below.

Lemma 1. Let $\epsilon \in (0, 4)$ and define

$$\alpha_t = 1 - \Phi \left(\sqrt{\frac{2}{1 - \epsilon^2/16} \log \frac{\log((1 + \epsilon)t)}{\delta \log(1 + \epsilon)}} \right). \quad (\text{C2})$$

Then, at time t

$$\Pr[[2, (40)] \text{ holds}] = \Pr \left[\frac{\mu_i^t - m_i}{\sigma_i^t} \geq \Phi^{-1}(1 - \alpha_t) \right] \leq \delta.$$

The corrections to the four algorithms published in [2] and the corresponding corrected expressions for the performance bounds are summarized in Table I. For δ -Sufficing and (\mathcal{M}, δ) -Satisficing UCL, the bounds take the form

$$f(\delta, \Delta) := \frac{8\sigma_s^2}{\Delta^2(1 - \epsilon^2/16)} \log \frac{\log((1 + \epsilon)T)}{\delta \log(1 + \epsilon)} + 1. \quad (\text{C3})$$

TABLE I

SUMMARY OF THE CORRECTIONS FOR THE SATISFICING UCL ALGORITHMS. DEFINE Q_i^t BY (C1) WITH $\epsilon \in (0, 4)$, $K = \sqrt{2\pi}e$ AND SET α_t AS FOLLOWS. THE CORRECTED PERFORMANCE BOUNDS WITH f IS DEFINED BY (C3).

Algorithm	α_t	Bound
Deterministic UCL	$\alpha_t = 1/Kt^a, a > \frac{4}{3(1 - \epsilon^2/16)}$	$\mathbb{E}[n_i^T] \leq \frac{8a\sigma_s^2}{\Delta_i^2} \log T + o(\log T)$
Satisfaction-in-mean-reward UCL	$\alpha_t = 1/Kt^a, a > \frac{4}{3(1 - \epsilon^2/16)}$	$\mathbb{E}[n_i^T] \leq \frac{8a}{(\Delta_i^{\mathcal{M}})^2} \log T + o(\log T)$
δ -Sufficing UCL	α_t from Equation (C2), $\delta \mapsto \delta/2$	$n_i^T \leq f(\delta/2, \Delta_i)$
(\mathcal{M}, δ) -Satisficing UCL	α_t from Equation (C2), $\delta \mapsto \delta/3$	$n_i^T \leq f(\delta/3, \Delta_i^{\mathcal{M}})$

Note that with the correction, which accounts for the dependence of n_i^t on rewards accrued, the upper bound functional form (C3) is no longer independent of T . However, the dependence on T is of the form $\log \log T$, which is a very slowly increasing function of T . Therefore, in any realistic application the upper bound will effectively be constant and the qualitative result of [2] does not change.

REVISED PROOF

We employ the following concentration inequality from Garivier and Moulines [1] to fix the proof. Let $(X_t)_{t \geq 1}$ be a sequence of independent sub-Gaussian random variables with $\mathbb{E}[X_t] = \mu_t$, i. e., $\mathbb{E}[\exp(\lambda(X_t - \mu_t))] \leq \exp(\lambda^2 \sigma^2/2)$ for some variance parameter $\sigma > 0$. Consider a previsible sequence $(\epsilon_t)_{t \geq 1}$ of Bernoulli variables, i.e., for all $t > 0$, ϵ_t is deterministically known given $\{X_\tau\}_{0 < \tau < t}$. Let

$$s^t = \sum_{s=1}^t X_s \epsilon_s, m^t = \sum_{s=1}^t \mu_s \epsilon_s, n^t = \sum_{s=1}^t \epsilon_s.$$

Theorem 2 ([1, Theorem 22], [3, Theorem 11]). Let $(X_t)_{t \geq 1}$ be a sequence of sub-Gaussian¹ independent random variables with common variance parameter σ and let $(\epsilon_t)_{t \geq 1}$ be a pre-visible sequence of Bernoulli variables. Then, for all integers t and all $\delta, \epsilon > 0$,

$$\Pr \left[\frac{s^t - m^t}{\sqrt{n^t}} > \delta \right] \leq \left[\frac{\log t}{\log(1 + \epsilon)} \right] \exp \left(-\frac{\delta^2}{2\sigma^2} \left(1 - \frac{\epsilon^2}{16} \right) \right). \quad (\text{C4})$$

We will also use the following lower bound for $\Phi^{-1}(1 - \alpha)$, the quantile function of the normal distribution.

¹The result in [1, Theorem 22] is stated for bounded rewards, but it extends immediately to sub-Gaussian rewards by noting that the upper bound on the moment generating function for a bounded random variable obtained using a Hoeffding inequality has the same functional form as the sub-Gaussian random variable.

Proposition 3. For any $t \in \mathbb{N}$ and $a > 1$, the following holds:

$$\Phi^{-1} \left(1 - \frac{1}{\sqrt{2\pi e t^a}} \right) \geq \sqrt{\nu \log t^a}, \quad (\text{C5})$$

for any $0 < \nu \leq 1.59$.

Proof. We begin with the inequality $\Phi^{-1}(1 - \alpha) > \sqrt{-\log(2\pi\alpha^2(1 - \log(2\pi\alpha^2)))}$ established in [4]. It suffices to show that

$$-\log \left(\frac{1}{e t^2} \left(1 - \log \left(\frac{1}{e t^2} \right) \right) \right) - \nu \log t \geq 0,$$

for $\nu \in (0, 1.59]$. The left hand side of the above inequality is

$$g(t) := 1 - \log 2 + (2 - \nu) \log t - \log(1 + \log t).$$

It can be verified that g admits a unique minimum at $t = e^{(\nu-1)/(2-\nu)}$ and the minimum value is $\nu - \log 2 + \log(2 - \nu)$, which is positive for $0 < \nu \leq 1.59$. \square

In the following, we choose $\nu = 3/2$.

Correction to the proof of [2, Theorem 8]. The structure of the published proof carries through. Let i be a non-satisfying arm, i.e., $m_i < \mathcal{M}$, and recall that i^* denotes the arm with maximum mean reward. Let η be a positive integer and let $\epsilon \in (0, 4)$ and $a > 4/(3(1 - \epsilon^2/16))$.

We first analyze the probability that [2, Eq. (31)] holds by applying Theorem 2. Let $\{X_\tau\}_{0 < \tau < t}$ be the sequence of rewards associated with arm i , and let $(\epsilon_t)_{t \geq 1}$ equal 1 if the algorithm chooses arm i at time t . Note that, for an uncorrelated uninformative prior, $\mu_i^t = \bar{m}_i^t = s^t/n^t$, $\sigma_i^t = 1/\sqrt{n_i^t}$, $m_i = m^t/n^t$, and $n_i^t = n^t$. [2, Eq. (31)] is thus equivalent to

$$\frac{s^t}{n^t} - \frac{m^t}{n^t} \geq \frac{1}{\sqrt{n^t}} \Phi^{-1}(1 - \alpha_t) \Rightarrow \frac{s^t - m^t}{\sqrt{n^t}} \geq \Phi^{-1}(1 - \alpha_t).$$

Letting $\delta = \Phi^{-1}(1 - \alpha_t)$ and applying (C4) yields

$$\begin{aligned} \Pr[\text{[2, Eq. (31)] holds}] &= \Pr \left[\frac{s^t - m^t}{\sqrt{n^t}} \geq \delta \right] \\ &\leq \left[\frac{\log t}{\log(1 + \epsilon)} \right] \exp \left(-\frac{3 \log t^a}{4} \left(1 - \frac{\epsilon^2}{16} \right) \right) \\ &= \left[\frac{\log t}{\log(1 + \epsilon)} \right] t^{-\frac{3a(1-\epsilon^2)/16}{4}}, \end{aligned}$$

where the second inequality follows from (C5). The same bound holds for [2, Eq. (32)].

It can be verified that for the corrected Q_i^t in equation (C1), the constant ‘‘8’’ in [2, Eqns. (35, 38 and 39)] will be replaced by $8a$. Following the proof in [2] with the above corrections,

$$\mathbb{E}[n_i^T] \leq \left[\frac{8a}{(\Delta_i^{\mathcal{M}})^2} \log T \right] + \sum_{t=1}^T 3 \left[\frac{\log t}{\log(1 + \epsilon)} \right] t^{-\frac{3a(1-\epsilon^2)/16}{4}}.$$

The sum can be bounded by the integral

$$\int_1^T \left(\frac{\log t}{\log(1 + \epsilon)} + 1 \right) t^{-\frac{3a(1-\epsilon^2)/16}{4}} dt + 1. \quad (\text{C6})$$

It can be verified that the integral (C6) is of class $o(\log T)$ as long as the exponent $3a(1 - \epsilon^2/16)/4 > 1$. Putting everything together, we have

$$\mathbb{E}[n_i^T] \leq \frac{8a\sigma_s^2}{\Delta_i^2} \log T + o(\log T).$$

The second statement follows from the definition of the cumulative expected regret. \square

The corrections to the proofs of [2, Theorem 10] (δ -Sufficing UCL) and [2, Theorem 11] ((\mathcal{M}, δ) -Satisficing UCL) follow the same structure.

Correction to proof of [2, Theorem 10]. For the corrected α_t defined in equation (C2) with $\delta \mapsto \frac{\delta}{2}$, [2, Eq. (42)] is equivalent to

$$\Delta_i = m_{i^*} - m_i < 2C_i^t = \frac{2\sigma_s}{\sqrt{n_i^t}} \Phi^{-1}(1 - \alpha_t).$$

Squaring, rearranging, and applying Equation (C2), we see that this never holds if

$$n_i^t > \frac{8\sigma_s^2}{\Delta_i^2(1 - \epsilon^2/16)} \log \frac{2 \log((1 + \epsilon)t)}{\delta \log(1 + \epsilon)} = \eta.$$

Then, Lemma 1 implies that [2, Eqns. (40, 41)] each hold with probability at most $\delta/2$. Therefore, for $n_i^t > \eta + 1 = f(\delta/2, \Delta_i)$, a non-satisfying arm is selected with probability at most δ . \square

Correction to proof of [2, Theorem 11]. For the corrected α_t defined in equation (C2) with $\delta \mapsto \frac{\delta}{2}$, an argument analogous to that for [2, Eq. (42)] above shows that [2, Eq. (44)] never holds for $n_i^t > \eta = f(\delta/3, \Delta_i^{\mathcal{M}}) - 1$.

Applying Lemma 1 implies that [2, Eq. (43)] holds with probability at most $\delta/3$. Similarly to the corrected proof for [2, Theorem 10] above, for $n_i^t > \eta + 1 = f(\delta/3, \Delta_i^{\mathcal{M}})$, $Q_i^t \geq Q_{i^*}^t$ with probability at most $\frac{2\delta}{3}$. Thus, a non-satisfying arm is selected with probability at most δ . \square

REFERENCES

- [1] A. Garivier and E. Moulines. On upper-confidence bound policies for non-stationary bandit problems. *arXiv preprint arXiv:0805.3415*, 2008.
- [2] P. Reverdy, V. Srivastava, and N. E. Leonard. Satisficing in multi-armed bandit problems. *IEEE Transactions on Automatic Control*, 62(8):3788–3803, August 2017.
- [3] P. Reverdy, V. Srivastava, and N. E. Leonard. Modeling human decision-making in generalized Gaussian multi-armed bandits. *arXiv preprint 1307.6134v5*, 2019. Available <https://arxiv.org/abs/1307.6134v5>.
- [4] P. B. Reverdy, V. Srivastava, and N. E. Leonard. Modeling human decision making in generalized Gaussian multiarmed bandits. *Proceedings of the IEEE*, 102(4):544–571, 2014.