

An Information-Theoretic Analysis of Thompson Sampling

Daniel Russo

DJRUSO@STANFORD.EDU

*Department of Management Science and Engineering
Stanford University
Stanford, California 94305*

Benjamin Van Roy

BVR@STANFORD.EDU

*Departments of Management Science and Engineering and Electrical Engineering
Stanford University
Stanford, California 94305*

Editor: Kevin Murphy

Abstract

We provide an information-theoretic analysis of Thompson sampling that applies across a broad range of online optimization problems in which a decision-maker must learn from partial feedback. This analysis inherits the simplicity and elegance of information theory and leads to regret bounds that scale with the entropy of the optimal-action distribution. This strengthens preexisting results and yields new insight into how information improves performance.

Keywords: Thompson sampling, online optimization, multi-armed bandit, information theory, regret bounds

1. Introduction

This paper considers the problem of repeated decision making in the presence of model uncertainty. A decision-maker repeatedly chooses among a set of possible actions, observes an outcome, and receives a reward representing the utility derived from this outcome. The decision-maker is uncertain about the underlying system and is therefore initially unsure of which action is best. However, as outcomes are observed, she is able to learn over time to make increasingly effective decisions. Her objective is to choose actions sequentially so as to maximize the expected cumulative reward.

We focus on settings with *partial feedback*, under which the decision-maker does not generally observe what the reward would have been had she selected a different action. This feedback structure leads to an inherent tradeoff between *exploration and exploitation*: by experimenting with poorly understood actions one can learn to make more effective decisions in the future, but focusing on better understood actions may lead to higher rewards in the short term. The classical multi-armed bandit problem is an important special case of this formulation. In such problems, the decision-maker only observes rewards she receives, and rewards from one action provide no information about the reward that can be attained by selecting other actions. We are interested here primarily in models and algorithms that accommodate cases where the number of actions is very large and there is a richer

information structure relating actions and observations. This category of problems is often referred to as *online optimization with partial feedback*.

A large and growing literature treats the design and analysis of algorithms for such problems. An online optimization algorithm typically starts with two forms of prior knowledge. The first – *hard knowledge* – posits that the mapping from action to outcome distribution lies within a particular family of mappings. The second – *soft knowledge* – concerns which of these mappings are more or less likely to match reality. Soft knowledge evolves with observations and is typically represented in terms of a probability distribution or a confidence set.

Much recent work concerning online optimization algorithms focuses on establishing performance guarantees in the form of *regret bounds*. Surprisingly, virtually all of these regret bounds depend on hard knowledge but not soft knowledge, a notable exception being the bounds of Srinivas et al. (2012) which we discuss further in Section 1.2. Regret bounds that depend on hard knowledge yield insight into how an algorithm’s performance scales with the complexity of the family of mappings and are useful for delineating algorithms on that basis, but if a regret bound does not depend on soft knowledge, it does not have much to say about how future performance should improve as data is collected. The latter sort of insight should be valuable for designing better ways of trading off between exploration and exploitation, since it is soft knowledge that is refined as a decision-maker learns. Another important benefit to understanding how performance depends on soft knowledge arises in practical applications: when designing an online learning algorithm, one may have access to historical data and want to understand how this prior information benefits future performance.

In this paper, we establish regret bounds that depend on both hard and soft knowledge for a simple online learning algorithm alternately known as *Thompson sampling*, *posterior sampling*, or *probability matching*. The bounds strengthen results from prior work not only with respect to Thompson sampling but relative to regret bounds for *any* online optimization algorithm. Further, the bounds offer new insight into how regret depends on soft knowledge. Indeed, forthcoming work of ours leverages this to produce an algorithm that outperforms Thompson sampling.

We found information theory to provide tools ideally suited for deriving our new regret bounds, and our analysis inherits the simplicity and elegance enjoyed by work in that field. Our formulation encompasses a broad family of information structures, including as special cases multi-armed bandit problems with independent arms, online optimization problems with full information, linear bandit problems, and problems with combinatorial action sets and “semi-bandit” feedback. We leverage information theory to provide a unified analysis that applies to each of those special cases.

A novel feature of our bounds is their dependence on the entropy of the optimal-action distribution. To our knowledge, these are the first bounds on the expected regret of *any* algorithm that depend on the magnitude of the agent’s uncertainty about which action is optimal. The fact that our bounds only depend on uncertainties relevant to optimizing performance highlights the manner in which Thompson sampling naturally exploits complex information structures. Further, in practical contexts, a decision-maker may begin with an understanding that some actions are more likely to be optimal than others. For example, when dealing with a shortest path problem, one might expect paths that traverse fewer

edges to generally incur less cost. Our bounds are the first to formalize the performance benefits afforded by such an understanding.

1.1 Preview of Results

Our analysis is based on a general probabilistic, or Bayesian, formulation in which uncertain quantities are modeled as random variables. In principle, the optimal strategy for such a problem could be calculated via dynamic programming, but for problem classes of practical interest this would be computationally intractable. *Thompson sampling* serves as a simple and elegant heuristic strategy. In this section, we provide a somewhat informal problem statement, and a preview of our main results about Thompson sampling. In the next subsection we discuss how these results relate to the existing literature.

A decision maker repeatedly chooses among a finite set of possible actions \mathcal{A} and upon taking action $a \in \mathcal{A}$ at time $t \in \mathbb{N}$ she observes a random outcome $Y_{t,a} \in \mathcal{Y}$. She associates a reward with each outcome as specified by a reward function $R : \mathcal{Y} \rightarrow \mathbb{R}$. The outcomes $(Y_{t,a})_{t \in \mathbb{N}}$ are drawn independently over time from a fixed probability distribution p_a^* over \mathcal{Y} . The decision maker is uncertain about the distribution of outcomes $p^* = (p_a^*)_{a \in \mathcal{A}}$, which is itself distributed according to a prior distribution over a family \mathcal{P} of such distributions. However, she is able to learn about p^* as the outcomes of past decisions are observed, and this allows her to learn over time to attain improved performance.

We first present a special case of our main result that applies to online linear optimization problems under bandit feedback. This result applies when each action is associated with a d -dimensional feature vector and the mean reward of each action is the inner product between an unknown parameter vector and the action's known feature vector. More precisely, suppose $\mathcal{A} \subset \mathbb{R}^d$ and that for every $p \in \mathcal{P}$ there is a vector $\theta_p \in \mathbb{R}^d$ such that

$$\mathbb{E}_{y \sim p_a} [R(y)] = a^T \theta_p \tag{1}$$

for all $a \in \mathcal{A}$. When an action is sampled, a random reward in $[0, 1]$ is received, where the mean reward is given by (1). Then, our analysis establishes that the expected cumulative regret of Thompson sampling up to time T is bounded by

$$\sqrt{\frac{\text{Entropy}(A^*)dT}{2}}, \tag{2}$$

where $A^* \in \arg \max_{a \in \mathcal{A}} \mathbb{E} [R(Y_{t,a})|p^*]$ denotes the optimal action. This bound depends on the time horizon, the entropy of the of the prior distribution of the optimal action A^* , and the dimension d of the linear model.

Because the entropy of A^* is always less than $\log |\mathcal{A}|$, (2) yields a bound of order $\sqrt{\log(|\mathcal{A}|)dT}$, and this scaling cannot be improved in general (see Section 6.5). The bound (2) is stronger than this worst-case bound, since the entropy of A^* can be much smaller than $\log(|\mathcal{A}|)$.

Thompson sampling incorporates prior knowledge in a flexible and coherent way, and the benefits of this are reflected in two distinct ways by the above bound. First, as in past work (see e.g. Dani et al., 2008; Rusmevichientong and Tsitsiklis, 2010), the bound depends on the dimension of the linear model instead of the number of actions. This reflects that the

algorithm is able to learn more rapidly by exploiting the known model, since observations from selecting one action provide information about rewards that would have been generated by other actions. Second, the bound depends on the entropy of the prior distribution of A^* instead of a worst case measure like the logarithm of the number of actions. This highlights the benefit of prior knowledge that some actions are more likely to be optimal than others. In particular, this bound exhibits the natural property that as the entropy of the prior distribution of A^* goes to zero, expected regret does as well.

Our main result extends beyond the class of linear bandit problems. Instead of depending on the linear dimension of the model, it depends on a more general measure of the problem’s information complexity: what we call the problem’s *information ratio*. By bounding the information ratio in specific settings, we recover the bound (2) as a special case, along with bounds for problems with full feedback and problems with combinatorial action sets and “semi-bandit” feedback.

1.2 Related Work

Though Thompson sampling was first proposed in 1933 (Thompson, 1933), until recently it was largely ignored in the academic literature. Interest in the algorithm grew after empirical studies (Scott, 2010; Chapelle and Li, 2011) demonstrated performance exceeding state of the art. Over the past several years, it has also been adopted in industry.¹ This has prompted a surge of interest in providing theoretical guarantees for Thompson sampling.

One of the first theoretical guarantees for Thompson sampling was provided by May et al. (2012), but they showed only that the algorithm converges asymptotically to optimality. Agrawal and Goyal (2012); Kauffmann et al. (2012); Agrawal and Goyal (2013a) and Korda et al. (2013) studied on the classical multi-armed bandit problem, where sampling one action provides no information about other actions. They provided frequentist regret bounds for Thompson sampling that are asymptotically optimal in the sense defined by Lai and Robbins (1985). To attain these bounds, the authors fixed a specific uninformative prior distribution, and studied the algorithm’s performance assuming this prior is used.

Our interest in Thompson sampling is motivated by its ability to incorporate rich forms of prior knowledge about the actions and the relationship among them. Accordingly, we study the algorithm in a very general framework, allowing for an *arbitrary* prior distribution over the true outcome distributions $p^* = (p_a^*)_{a \in \mathcal{A}}$. To accommodate this level of generality while still focusing on finite-time performance, we study the algorithm’s *expected regret* under the prior distribution. This measure is sometimes called *Bayes risk* or *Bayesian regret*.

Our recent work (Russo and Van Roy, 2014) provided the first results in this setting. That work leverages a close connection between Thompson sampling and upper confidence bound (UCB) algorithms, as well as existing analyses of several UCB algorithms. This confidence bound analysis was then extended to a more general setting, leading to a general regret bound stated in terms of a new notion of model complexity – what we call the eluder dimension. While the connection with UCB algorithms may be of independent interest,

1. Microsoft (Graepel et al., 2010), Google analytics (Scott, 2014) and LinkedIn (Tang et al., 2013) have used Thompson sampling.

it’s desirable to have a simple, self-contained, analysis that does not rely on the—often complicated—construction of confidence bounds.

Agrawal and Goyal (2013a) provided the first “distribution independent” bound for Thompson sampling. They showed that when Thompson sampling is executed with an independent uniform prior and rewards are binary the algorithm satisfies a frequentist regret bound² of order $\sqrt{|\mathcal{A}|T \log(T)}$. Russo and Van Roy (2014) showed that, for an arbitrary prior over bounded reward distributions, the expected regret of Thompson sampling under this prior is bounded by a term of order $\sqrt{|\mathcal{A}|T \log(T)}$. Bubeck and Liu (2013) showed that this second bound can be improved to one of order $\sqrt{|\mathcal{A}|T}$ using more sophisticated confidence bound analysis, and also studied a problem setting where the regret of Thompson sampling is bounded uniformly over time. In this paper, we are interested mostly in results that replace the explicit $\sqrt{|\mathcal{A}|}$ dependence on the number of actions with a more general measure of the problem’s information complexity. For example, as discussed in the last section, for the problem of linear optimization under bandit feedback one can provide bounds that depend on the dimension of the linear model instead of the number of actions.

To our knowledge, Agrawal and Goyal (2013b) are the only other authors to have studied the application of Thompson sampling to linear bandit problems. They showed that, when the algorithm is applied with an uncorrelated Gaussian prior over θ_{p^*} , it satisfies frequentist regret bounds of order $\frac{d^2}{\epsilon} \sqrt{T^{1+\epsilon}} (\log(Td))$. Here ϵ is a parameter used by the algorithm to control how quickly the posterior concentrates. Russo and Van Roy (2014) allowed for an arbitrary prior distribution, and provided a bound on expected regret (under this prior) of order $d \log(T) \sqrt{T}$. Unlike the bound (2), these results hold whenever \mathcal{A} is a compact subset of \mathbb{R}^d , but we show in Appendix D.1 that through discretization argument the bound (2) also yields a similar bound whenever \mathcal{A} is compact. In the worst case, that bound is of order $d \sqrt{T \log(T)}$.

Other very recent work (Gopalan et al., 2014) provided a general asymptotic guarantee for Thompson sampling. They studied the growth rate of the cumulative number of times a suboptimal action is chosen as the time horizon T tends to infinity. One of their asymptotic results applies to the problem of online linear optimization under “semi-bandit” feedback, which we study in Section 6.6.

An important aspect of our regret bound is its dependence on soft knowledge through the entropy of the optimal-action distribution. One of the only other regret bounds that depends on soft-knowledge was provided very recently by Li (2013). Inspired by a connection between Thompson sampling and exponential weighting schemes, that paper introduced a family of Thompson sampling like algorithms and studied their application to contextual bandit problems. While our analysis does not currently treat contextual bandit problems, we improve upon their regret bound in several other respects. First, their bound depends on the entropy of the prior distribution of mean rewards, which is never smaller, and can be much larger, than the entropy of the distribution of the optimal action. In addition, their bound has an order $T^{2/3}$ dependence on the problem’s time horizon, and, in order to guarantee each action is explored sufficiently often, requires that actions are frequently selected uniformly at random. In contrast, our focus is on settings where the number of actions is large and the goal is to learning without sampling each one.

2. They bound regret conditional on the true reward distribution: $\mathbb{E}[\text{Regret}(T, \pi^{\text{TS}}) | p^*]$.

Another regret bound that to some extent captures dependence on soft knowledge is that of Srinivas et al. (2012). This excellent work focuses on extending algorithms and expanding theory to address multi-armed bandit problems with arbitrary reward functions and possibly an infinite number of actions. In a sense, there is no hard knowledge while soft knowledge is represented in terms of a Gaussian process over a possibly infinite dimensional family of functions. An upper-confidence-bound algorithm is proposed and analyzed. Our earlier work (Russo and Van Roy, 2014) showed similar bounds also apply to Thompson sampling. Though our results here do not treat infinite action spaces, it should be possible to extend the analysis in that direction. One substantive difference is that our results apply to a much broader class of models: distributions are not restricted to Gaussian and more complex information structures are allowed. Further, the results of Srinivas et al. (2012) do not capture the benefits of soft knowledge to the extent that ours do. For example, their regret bounds do not depend on the mean of the reward function distribution, even though mean rewards heavily influence the chances that each action is optimal. Our regret bounds, in contrast, establish that regret vanishes as the probability that a particular action is optimal grows.

Our review has discussed the recent literature on Thompson sampling as well as two papers that have established regret bounds that depend on soft knowledge. There is of course an immense body of work on alternative approaches to efficient exploration, including work on the Gittins index approach (Gittins et al., 2011), Knowledge Gradient strategies (Ryzhov et al., 2012), and upper-confidence bound strategies (Lai and Robbins, 1985; Auer et al., 2002). In an adversarial framework, authors often study exponential-weighting schemes or, more generally, strategies based on online stochastic mirror descent. Bubeck and Cesa-Bianchi (2012) provided a general review of upper-confidence strategies and algorithms for the adversarial multi-armed bandit problem.

2. Problem Formulation

The decision-maker sequentially chooses actions $(A_t)_{t \in \mathbb{N}}$ from the action set \mathcal{A} and observes the corresponding outcomes $(Y_{t,A_t})_{t \in \mathbb{N}}$. To avoid measure-theoretic subtleties, we assume the space of possible outcomes \mathcal{Y} is a subset of a finite dimensional Euclidean space. There is a random outcome $Y_{t,a} \in \mathcal{Y}$ associated with each $a \in \mathcal{A}$ and time $t \in \mathbb{N}$. Let $Y_t \equiv (Y_{t,a})_{a \in \mathcal{A}}$ be the vector of outcomes at time $t \in \mathbb{N}$. The “true outcome distribution” p^* is a distribution over $\mathcal{Y}^{|\mathcal{A}|}$ that is itself randomly drawn from the family of distributions \mathcal{P} . We assume that, conditioned on p^* , $(Y_t)_{t \in \mathbb{N}}$ is an iid sequence with each element Y_t distributed according to p^* . Let p_a^* be the marginal distribution corresponding to $Y_{t,a}$.

The agent associates a reward $R(y)$ with each outcome $y \in \mathcal{Y}$, where the reward function $R : \mathcal{Y} \rightarrow \mathbb{R}$ is fixed and known. Uncertainty about p^* induces uncertainty about the true optimal action, which we denote by $A^* \in \arg \max_{a \in \mathcal{A}} \mathbb{E}[R(Y_{t,a})|p^*]$. The T period *regret* of the sequence of actions A_1, \dots, A_T is the random variable,

$$\text{Regret}(T) := \sum_{t=1}^T [R(Y_{t,A^*}) - R(Y_{t,A_t})], \tag{3}$$

which measures the cumulative difference between the reward earned by an algorithm that always chooses the optimal action, and actual accumulated reward up to time T . In this paper we study expected regret

$$\mathbb{E}[\text{Regret}(T)] = \mathbb{E} \left[\sum_{t=1}^T [R(Y_{t,A^*}) - R(Y_{t,A_t})] \right], \quad (4)$$

where the expectation is taken over the randomness in the actions A_t and the outcomes Y_t , and over the prior distribution over p^* . This measure of performance is commonly called *Bayesian regret* or *Bayes risk*.

Filtrations and randomized policies: We define all random variables with respect to a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Actions are chosen based on the history of past observations and possibly some external source of randomness. To represent this external source of randomness more formally, we introduce a sequence of random variables $(U_t)_{t \in \mathbb{N}}$, where for each $i \in \mathbb{N}$, U_i is jointly independent of $\{U_t\}_{t \neq i}$, the outcomes $\{Y_{t,a}\}_{t \in \mathbb{N}, a \in \mathcal{A}}$, and p^* . We fix the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$ where $\mathcal{F}_t \subset \mathcal{F}$ is the sigma-algebra generated by $(A_1, Y_{1,A_1}, \dots, A_{t-1}, Y_{t-1,A_{t-1}})$. The action A_t is measurable with respect to the sigma-algebra generated by (\mathcal{F}_t, U_t) . That is, given the history of past observations, A_t is random only through its dependence on U_t .

The objective is to choose actions in a manner that minimizes expected regret. For this purpose, it's useful to think of the actions as being chosen by a *randomized policy* π , which is an \mathcal{F}_t -adapted sequence $(\pi_t)_{t \in \mathbb{N}}$. An action is chosen at time t by randomizing according to $\pi_t(\cdot) = \mathbb{P}(A_t \in \cdot | \mathcal{F}_t)$, which specifies a probability distribution over \mathcal{A} . We explicitly display the dependence of regret on the policy π , letting $\mathbb{E}[\text{Regret}(T, \pi)]$ denote the expected value given by (4) when the actions (A_1, \dots, A_T) are chosen according to π .

Further Assumptions: To simplify the exposition, our main results will be stated under two further assumptions. The first requires that rewards are uniformly bounded, effectively controlling the worst-case variance of the reward distribution. In particular, this assumption is used only in proving Fact 9. In the technical appendix, we show that Fact 9 can be extended to the case where reward distributions are sub-Gaussian, which yields results in that more general setting.

Assumption 1 $\sup_{\bar{y} \in \mathcal{Y}} R(\bar{y}) - \inf_{\underline{y} \in \mathcal{Y}} R(\underline{y}) \leq 1$.

Our next assumption requires that the action set is finite. In the technical appendix we show that some cases where \mathcal{A} is infinite can be addressed through discretization arguments.

Assumption 2 \mathcal{A} is finite.

Because the Thompson sampling algorithm only chooses actions from the support of A^* , all of our results hold when the finite set \mathcal{A} denotes only the actions in the support of A^* . This difference can be meaningful. For example, when \mathcal{A} is a polytope and the objective function is linear, the support of A^* contains only extreme points of the polytope: a finite set.

3. Basic Measures and Relations in Information Theory

Before proceeding, we will define several common information measures – entropy, mutual information, and Kullback-Leibler divergence — and state several facts about these measures that will be referenced in our analysis. When all random variables are discrete, each of these facts is shown in chapter 2 of Cover and Thomas (2012). A treatment that applies to general random variables is provided in chapter 5 of Gray (2011).

Before proceeding, we will define some simple shorthand notation. Let $P(X) = \mathbb{P}(X \in \cdot)$ denote the distribution function of random variable X . Similarly, define $P(X|Y) = \mathbb{P}(X \in \cdot|Y)$ and $P(X|Y = y) = \mathbb{P}(X \in \cdot|Y = y)$.

Throughout this section, we will fix random variables X , Y , and Z that are defined on a joint probability space. We will allow Y and Z to be general random variables, but will restrict X to be supported on a finite set \mathcal{X} . This is sufficient for our purposes, as we will typically apply these relations when X is A^* , and is useful, in part, because the entropy of a general random variable can be infinite.

The *Shannon entropy* of X is defined as

$$H(X) = - \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) \log \mathbb{P}(X = x).$$

The first fact establishes uniform bounds on the entropy of a probability distribution.

Fact 1 $0 \leq H(X) \leq \log(|\mathcal{X}|)$.

If Y is a discrete random variable, the entropy of X conditional on $Y = y$ is

$$H(X|Y = y) = \sum_{x \in \mathcal{X}} \mathbb{P}(X = x|Y = y) \log \mathbb{P}(X = x|Y = y)$$

and the conditional entropy is of X given Y is

$$H(X|Y) = \sum_y H(X|Y = y) \mathbb{P}(Y = y).$$

For a general random variable Y , the conditional entropy of X given Y is,

$$H(X|Y) = \mathbb{E}_Y \left[- \sum_{x \in \mathcal{X}} \mathbb{P}(X = x|Y) \log \mathbb{P}(X = x|Y) \right],$$

where this expectation is taken over the marginal distribution of Y . For two probability measures P and Q , if P is absolutely continuous with respect to Q , the *Kullback–Leibler divergence* between them is

$$D(P||Q) = \int \log \left(\frac{dP}{dQ} \right) dP \tag{5}$$

where $\frac{dP}{dQ}$ is the Radon–Nikodym derivative of P with respect to Q . This is the expected value under P of the log-likelihood ratio between P and Q , and is a measure of how different P and Q are. The next fact establishes the non-negativity of Kullback–Leibler divergence.

Fact 2 (*Gibbs’ inequality*) For any probability distributions P and Q such that P is absolutely continuous with respect to Q , $D(P||Q) \geq 0$ with equality if and only if $P = Q$ P -almost everywhere.

The *mutual information* between X and Y

$$I(X; Y) = D(P(X, Y) || P(X)P(Y)) \tag{6}$$

is the Kullback–Leibler divergence between the joint distribution of X and Y and the product of the marginal distributions. From the definition, it’s clear that $I(X; Y) = I(Y; X)$, and Gibbs’ inequality implies that $I(X; Y) \geq 0$ and $I(X; Y) = 0$ when X and Y are independent.

The next fact, which is Lemma 5.5.6 of Gray (2011), states that the mutual information between X and Y is the expected reduction in the entropy of the posterior distribution of X due to observing Y .

Fact 3 (*Entropy reduction form of mutual information*)

$$I(X; Y) = H(X) - H(X|Y)$$

The mutual information between X and Y , conditional on a third random variable Z is

$$I(X; Y|Z) = H(X|Z) - H(X|Y, Z),$$

the expected additional reduction in entropy due to observing Y given that Z is also observed. This definition is also a natural generalization of the one given in (6), since

$$I(X; Y|Z) = \mathbb{E}_Z [D(P((X, Y)|Z) || P(X|Z)P(Y|Z))].$$

The next fact shows that conditioning on a random variable Z that is independent of X and Y does not affect mutual information.

Fact 4 *If Z is jointly independent of X and Y , then $I(X; Y|Z) = I(X; Y)$.*

The mutual information between a random variable X and a collection of random variables (Z_1, \dots, Z_T) can be expressed elegantly using the following “chain rule.”

Fact 5 (*Chain Rule of Mutual Information*)

$$I(X; (Z_1, \dots, Z_T)) = I(X; Z_1) + I(X; Z_2 | Z_1) + \dots + I(X; Z_T | Z_1, \dots, Z_{T-1}).$$

We provide some details related to the derivation of Fact 6 in the appendix.

Fact 6 (*KL divergence form of mutual information*)

$$\begin{aligned} I(X; Y) &= \mathbb{E}_X [D(P(Y|X) || P(Y))] \\ &= \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) D(P(Y|X = x) || P(Y)) \end{aligned}$$

While Facts 1 and 6, are standard properties of mutual information, it's worth highlighting their surprising power. It's useful to think of X as being A^* , the optimal action, and Y as being $Y_{t,a}$, the observation when selecting some action a . Then, combining these properties, we see that the next observation Y is expected to greatly reduce uncertainty about the optimal action A^* if and only if the distribution of Y varies greatly depending on the realization of A^* , in the sense that $D(P(Y|A^* = a^*) || P(Y))$ is large on average. This fact is crucial to our analysis.

One implication of the next fact is that the expected reduction in entropy from observing the outcome $Y_{t,a}$ is always at least as large as that from observing the reward $R(Y_{t,a})$.

Fact 7 (*Weak Version of the Data Processing Inequality*) *If $Z = f(Y)$ for a deterministic function f , then $I(X; Y) \geq I(X; Z)$. If f is invertible, so Y is also a deterministic function of Z , then $I(X; Y) = I(X; Z)$.*

We close this section by stating a fact that guarantees $D(P(Y|X = x) || P(Y))$ is well defined. It follows from a general property of conditional probability: for any random variable Z and event $E \subset \Omega$, if $\mathbb{P}(E) = 0$ then $\mathbb{P}(E|Z) = 0$ almost surely.

Fact 8 *For any $x \in \mathcal{X}$ with $\mathbb{P}(X = x) > 0$, $P(Y|X = x)$ is absolutely continuous with respect to $P(Y)$.*

3.1 Notation Under Posterior Distributions

As shorthand, we let

$$\mathbb{P}_t(\cdot) = \mathbb{P}(\cdot | \mathcal{F}_t) = \mathbb{P}(\cdot | A_1, Y_{1,A_1}, \dots, A_{t-1}, Y_{t-1,A_{t-1}})$$

and $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_t]$. As before, we will define some simple shorthand notation for the distribution function of a random variable under \mathbb{P}_t . Let $P_t(X) = \mathbb{P}_t(X \in \cdot)$, $P_t(X|Y) = \mathbb{P}_t(X \in \cdot | Y)$ and $P_t(X|Y = y) = \mathbb{P}_t(X \in \cdot | Y = y)$.

The definitions of entropy and mutual information implicitly depend on some base measure over the sample space Ω . We will define special notation to denote entropy and mutual information under the posterior measure $\mathbb{P}_t(\cdot)$. Define

$$\begin{aligned} H_t(X) &= - \sum_{x \in \mathcal{X}} \mathbb{P}_t(X = x) \log \mathbb{P}_t(X = x) \\ H_t(X|Y) &= \mathbb{E}_t \left[- \sum_{x \in \mathcal{X}} \mathbb{P}_t(X = x|Y) \log \mathbb{P}_t(X = x|Y) \right] \\ I_t(X; Y) &= H_t(X) - H_t(X|Y). \end{aligned}$$

Because these quantities depend on the realizations of $A_1, Y_{1,A_1}, \dots, A_{t-1}, Y_{t-1,A_{t-1}}$, they are random variables. By taking their expectation, we recover the standard definition of conditional entropy and conditional mutual information:

$$\begin{aligned} \mathbb{E}[H_t(X)] &= H(X|A_1, Y_{1,A_1}, \dots, A_{t-1}, Y_{t-1,A_{t-1}}) \\ \mathbb{E}[I_t(X; Y)] &= I(X; Y|A_1, Y_{1,A_1}, \dots, A_{t-1}, Y_{t-1,A_{t-1}}). \end{aligned}$$

4. Thompson Sampling

The Thompson sampling algorithm simply samples actions according to the posterior probability they are optimal. In particular, actions are chosen randomly at time t according to the sampling distribution $\pi_t^{\text{TS}} = \mathbb{P}(A^* = \cdot | \mathcal{F}_t)$. By definition, this means that for each $a \in \mathcal{A}$, $\mathbb{P}(A_t = a | \mathcal{F}_t) = \mathbb{P}(A^* = a | \mathcal{F}_t)$. This algorithm is sometimes called *probability matching* because the action selection distribution is *matched* to the posterior distribution of the optimal action.

This conceptually elegant probability matching scheme often admits a surprisingly simple and efficient implementation. Consider the case where $\mathcal{P} = \{p_\theta\}_{\theta \in \Theta}$ is some parametric family of distributions. The true outcome distribution p^* corresponds to a particular random index $\theta^* \in \Theta$ in the sense that $p^* = p_{\theta^*}$ almost surely. Practical implementations of Thompson sampling typically use two simple steps at each time t to randomly generate an action from the distribution α_t . First, an index $\hat{\theta}_t \sim \mathbb{P}(\theta^* \in \cdot | \mathcal{F}_t)$ is sampled from the posterior distribution of the true index θ^* . Then, the algorithm selects the action $A_t \in \arg \max_{a \in \mathcal{A}} \mathbb{E} \left[R(Y_{t,a}) | \theta^* = \hat{\theta}_t \right]$ that would be optimal if the sampled parameter were actually the true parameter. We next provide an example of a Thompson sampling algorithm designed to address the problem of online linear optimization under bandit feedback.

4.1 Example of Thompson Sampling

Suppose each action $a \in \mathcal{A} \subset \mathbb{R}^d$ is defined by a d -dimensional feature vector, and almost surely there exists $\theta^* \in \mathbb{R}^d$ such that for each $a \in \mathcal{A}$, $\mathbb{E}[R(Y_{t,a}) | p^*] = a^T \theta^*$. Assume θ^* is drawn from a normal distribution $N(\mu_0, \Sigma_0)$. When a is selected, only the realized reward $Y_{t,a} = R(Y_{t,a}) \in \mathbb{R}$ is observed. For each action a , reward noise $R(Y_{t,a}) - \mathbb{E}[R(Y_{t,a}) | p^*]$ follows a Gaussian distribution with known variance. One can show that, conditioned on the history of observed data \mathcal{F}_t , θ^* remains normally distributed. Algorithm 1 provides an implementation of Thompson sampling for this problem. The expectations in step 3 can be computed efficiently via Kalman filtering.

Algorithm 1 Linear–Gaussian Thompson Sampling

- 1: **Sample Model:**
 $\hat{\theta}_t \sim N(\mu_{t-1}, \Sigma_{t-1})$
 - 2: **Select Action:**
 $A_t \in \arg \max_{a \in \mathcal{A}} \langle a, \hat{\theta}_t \rangle$
 - 3: **Update Statistics:**
 $\mu_t \leftarrow \mathbb{E}[\theta^* | \mathcal{F}_t]$
 $\Sigma_t \leftarrow \mathbb{E}[(\theta^* - \mu_t)(\theta^* - \mu_t)^\top | \mathcal{F}_t]$
 - 4: **Increment t and Goto Step 1**
-

Algorithm 1 is efficient as long as the linear objective $\langle a, \hat{\theta}_t \rangle$ can be maximized efficiently over the action set \mathcal{A} . For this reason, the algorithm is implementable in important cases where other popular approaches, like the ConfidenceBall₂ algorithm of Dani et al. (2008),

are computationally intractable. Because the posterior distribution of θ^* has a closed form, Algorithm 1 is particularly efficient. Even when the posterior distribution is complex, however, one can often generate samples from this distribution using Markov chain Monte Carlo algorithms, enabling efficient implementations of Thompson sampling. A more detailed discussion of the strengths and potential advantages of Thompson sampling can be found in earlier work (Scott, 2010; Chapelle and Li, 2011; Russo and Van Roy, 2014; Gopalan et al., 2014).

5. The Information Ratio and a General Regret Bound

Our analysis will relate the expected regret of Thompson sampling to its expected information gain: the expected reduction in the entropy of the posterior distribution of A^* . The relationship between these quantities is characterized by what we call the *information ratio*,

$$\Gamma_t := \frac{\mathbb{E}_t [R(Y_{t,A^*}) - R(Y_{t,A_t})]^2}{I_t(A^*; (A_t, Y_{t,A_t}))}$$

which is the ratio between the square of expected regret and information gain in period t . Recall that, as described in Subsection 3.1, the subscript t on \mathbb{E}_t and I_t indicates that these quantities are evaluated under the posterior measure $\mathbb{P}_t(\cdot) = \mathbb{P}(\cdot|\mathcal{F}_t)$.

Notice that if the information ratio is small, Thompson sampling can only incur large regret when it is expected to gain a lot of information about which action is optimal. This suggests its expected regret is bounded in terms of the maximum amount of information any algorithm could expect to acquire, which is at most the entropy of the prior distribution of the optimal action. Our next result shows this formally. We provide a general upper bound on the expected regret of Thompson sampling that depends on the time horizon T , the entropy of the prior distribution of A^* , $H(A^*)$, and any worst-case upper bound on the information ratio Γ_t . In the next section, we will provide bounds on Γ_t for some of the most widely studied classes of online optimization problems.

Proposition 1 *For any $T \in \mathbb{N}$, if $\Gamma_t \leq \bar{\Gamma}$ almost surely for each $t \in \{1, \dots, T\}$,*

$$\mathbb{E} [\text{Regret}(T, \pi^{\text{TS}})] \leq \sqrt{\bar{\Gamma} H(A^*) T}.$$

Proof Recall that $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot|\mathcal{F}_t]$ and we use I_t to denote mutual information evaluated under the base measure $\mathbb{P}_t(\cdot) = \mathbb{P}(\cdot|\mathcal{F}_t)$. Then,

$$\begin{aligned} \mathbb{E} [\text{Regret}(T, \pi^{\text{TS}})] &\stackrel{(a)}{=} \mathbb{E} \sum_{t=1}^T \mathbb{E}_t [R(Y_{t,A^*}) - R(Y_{t,A_t})] = \mathbb{E} \sum_{t=1}^T \sqrt{\Gamma_t I_t(A^*; (A_t, Y_{t,A_t}))} \\ &\leq \sqrt{\bar{\Gamma}} \left(\mathbb{E} \sum_{t=1}^T \sqrt{I_t(A^*; (A_t, Y_{t,A_t}))} \right) \\ &\stackrel{(b)}{\leq} \sqrt{\bar{\Gamma} T \mathbb{E} \sum_{t=1}^T I_t(A^*; (A_t, Y_{t,A_t}))}, \end{aligned}$$

where (a) follows from the tower property of conditional expectation, and (b) follows from the Cauchy-Schwartz inequality. We complete the proof by showing that expected information gain cannot exceed the entropy of the prior distribution. For the remainder of this proof, let $Z_t = (A_t, Y_{t,A_t})$. Then, as discussed in Subsection 3.1,

$$\mathbb{E}[I_t(A^*; Z_t)] = I(A^*; Z_t | Z_1, \dots, Z_{t-1}),$$

and therefore

$$\begin{aligned} \mathbb{E} \sum_{t=1}^T I_t(A^*; Z_t) &= \sum_{t=1}^T I(A^*; Z_t | Z_1, \dots, Z_{t-1}) \stackrel{(c)}{=} I(A^*; (Z_1, \dots, Z_T)) \\ &= H(A^*) - H(A^* | Z_1, \dots, Z_T) \\ &\stackrel{(d)}{\leq} H(A^*), \end{aligned}$$

where (c) follows from the chain rule for mutual information (Fact 5), and (d) follows from the non-negativity of entropy (Fact 1). \blacksquare

6. Bounding the Information Ratio

This section establishes upper bounds on the information ratio in several important settings. This yields explicit regret bounds when combined with Proposition 1, and also helps to clarify the role the information ratio plays in our results: it roughly captures the extent to which sampling some actions allows the decision maker to make inferences about *different* actions. In the worst case, the information ratio depends on the number of actions, reflecting the fact that actions could provide no information about others. For problems with full information, the information ratio is bounded by a numerical constant, reflecting that sampling one action perfectly reveals the rewards that would have been earned by selecting any other action. The problems of online linear optimization under “bandit feedback” and under “semi-bandit feedback” lie between these two extremes, and the information ratio provides a natural measure of each problem’s information structure. In each case, our bounds reflect that Thompson sampling is able to automatically exploit this structure.

We will compare our upper bounds on expected regret with known lower bounds. Some of these lower bounds were developed and stated in an adversarial framework, but were proved using the *probabilistic method*; authors fixed a family of distributions \mathcal{P} and an initial distribution over p^* and lower bounded the expected regret under this environment of any algorithm. This provides lower bounds on $\inf_{\pi} \mathbb{E}[\text{Regret}(T, \pi)]$ in our framework at least for particular problem instances. Unfortunately, we are not aware of any general prior-dependent lower bounds, and this remains an important direction for the field.

Our bounds on the information ratio Γ_t will hold at any time t and under any posterior measure \mathbb{P}_t . To simplify notation, in our proofs we will omit the subscript t from $\mathbb{E}_t, \mathbb{P}_t, P_t, A_t, Y_t, H_t$, and I_t .

6.1 An Alternative Representation of the Information Ratio

Recall that the information ratio is defined to be

$$\frac{\mathbb{E}[R(Y_{A^*}) - R(Y_A)]^2}{I(A^*; (A, Y_A))}.$$

The following proposition expresses the information ratio of Thompson sampling in a form that facilitates further analysis. The proof uses that Thompson sampling matches the action selection distribution to the posterior distribution of the optimal action, in the sense that $\mathbb{P}(A^* = a) = \mathbb{P}(A = a)$ for all $a \in \mathcal{A}$.

Proposition 2

$$\begin{aligned} I(A^*; (A, Y_A)) &= \sum_{a \in \mathcal{A}} \mathbb{P}(A = a) I(A^*; Y_a) \\ &= \sum_{a, a^* \in \mathcal{A}} \mathbb{P}(A^* = a) \mathbb{P}(A^* = a^*) [D(P(Y_a | A^* = a^*) || P(Y_a))]. \end{aligned}$$

and

$$\mathbb{E}[R(Y_{A^*}) - R(Y_A)] = \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) (\mathbb{E}[R(Y_a) | A^* = a] - \mathbb{E}[R(Y_a)])$$

The numerator of the information ratio measures the average difference between rewards generated from $P(Y_a)$, the posterior predictive distribution at a , and $P(Y_a | A^* = a)$, the posterior predictive distribution at a conditioned on a being the optimal action. It roughly captures how much knowing that the *selected action is optimal* influences the expected reward observed. The denominator measures how much, on average, knowing *which action is optimal* changes the observations at the selected action. Intuitively, the information ratio tends to be small when knowing which action is optimal significantly influences the anticipated observations at many other actions.

It's worth pointing out that this form of the information ratio bears a superficial resemblance to fundamental complexity terms in the multi-armed bandit literature. The results of Lai and Robbins (1985) and Agrawal et al. (1989) show the optimal asymptotic growth rate of regret is characterized by a ratio where the numerator depends on the difference between means of the reward distributions and the denominator depends on Kullback–Leibler divergences.

Proof Both proofs will use that the action A is selected based on past observations and independent random noise. Therefore, conditioned on the history, A is jointly independent

of A^* and the outcome vector $Y \equiv (Y_a)_{a \in \mathcal{A}}$.

$$\begin{aligned}
 I(A^*; (A, Y_A)) &\stackrel{(a)}{=} I(A^*; A) + I(A^*; Y_A | A) \\
 &\stackrel{(b)}{=} I(A^*; Y_A | A) \\
 &= \sum_{a \in \mathcal{A}} \mathbb{P}(A = a) I(A^*; Y_A | A = a) \\
 &\stackrel{(c)}{=} \sum_{a \in \mathcal{A}} \mathbb{P}(A = a) I(A^*; Y_a) \\
 &\stackrel{(d)}{=} \sum_{a \in \mathcal{A}} \mathbb{P}(A = a) \left(\sum_{a^* \in \mathcal{A}} \mathbb{P}(A^* = a^*) D(P(Y_a | A^* = a^*) || P(Y_a)) \right) \\
 &= \sum_{a, a^* \in \mathcal{A}} \mathbb{P}(A^* = a) \mathbb{P}(A^* = a^*) [D(P(Y_a | A^* = a^*) || P(Y_a))],
 \end{aligned}$$

where (a) follows from the chain rule for mutual information (Fact 5), (b) uses that A and A^* are independent and the mutual information between independent random variables is zero (Fact 4), (c) uses Fact 4 and that A is jointly independent of Y and A^* , and equality (d) uses Fact 6. Now, the numerator can be rewritten as,

$$\begin{aligned}
 \mathbb{E}[R(Y_{A^*}) - R(Y_A)] &= \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) \mathbb{E}[R(Y_a) | A^* = a] - \sum_{a \in \mathcal{A}} \mathbb{P}(A = a) \mathbb{E}[R(Y_a) | A = a] \\
 &= \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) (\mathbb{E}[R(Y_a) | A^* = a] - \mathbb{E}[R(Y_a)]),
 \end{aligned}$$

where the second equality uses that $\mathbb{P}(A = a) = \mathbb{P}(A^* = a)$ by the definition of Thompson sampling, and that Y is independent of the chosen action A . \blacksquare

6.2 Preliminaries

Here we state two basic facts that are used in bounding the information ratio. Proofs of both results are provided in the appendix for completeness.

The first fact lower bounds the Kullback–Leibler divergence between two bounded random variables in terms of the difference between their means. It follows trivially from an application of Pinsker’s inequality.

Fact 9 *For any distributions P and Q such that that P is absolutely continuous with respect to Q , any random variable $X : \Omega \rightarrow \mathcal{X}$ and any $g : \mathcal{X} \rightarrow \mathbb{R}$ such that $\sup g - \inf g \leq 1$,*

$$\mathbb{E}_P[g(X)] - \mathbb{E}_Q[g(X)] \leq \sqrt{\frac{1}{2} D(P || Q)},$$

where \mathbb{E}_P and \mathbb{E}_Q denote the expectation operators under P and Q .

Because of Assumption 1, this fact shows

$$\mathbb{E}[R(Y_a) | A^* = a^*] - \mathbb{E}[R(Y_a)] \leq \sqrt{\frac{1}{2} D(P(Y_a | A^* = a^*) || P(Y_a))}.$$

By the Cauchy–Schwartz inequality, for any vector $x \in \mathbb{R}^n$, $\sum_i x_i \leq \sqrt{n}\|x\|_2$. The next fact provides an analogous result for matrices. For any rank r matrix $M \in \mathbb{R}^{n \times n}$ with singular values $\sigma_1, \dots, \sigma_r$, let

$$\|M\|_* := \sum_{i=1}^r \sigma_i, \quad \|M\|_F := \sqrt{\sum_{k=1}^m \sum_{j=1}^n M_{i,j}^2} = \sqrt{\sum_{i=1}^r \sigma_i^2}, \quad \text{Trace}(M) := \sum_{i=1}^n M_{ii},$$

denote respectively the Nuclear norm, Frobenius norm and trace of M .

Fact 10 For any matrix $M \in \mathbb{R}^{k \times k}$,

$$\text{Trace}(M) \leq \sqrt{\text{Rank}(M)} \|M\|_F.$$

6.3 Worst Case Bound

The next proposition provides a bound on the information ratio that holds whenever rewards are bounded, but that has an explicit dependence on the number of actions. This scaling cannot be improved in general, but we go on to show tighter bounds are possible for problems with different *information structures*.

Proposition 3 For any $t \in \mathbb{N}$, $\Gamma_t \leq |\mathcal{A}|/2$ almost surely.

Proof We bound the numerator of the information ratio by $|\mathcal{A}|/2$ times its denominator:

$$\begin{aligned} \mathbb{E} [R(Y_{A^*}) - R(Y_A)]^2 &\stackrel{(a)}{=} \left(\sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) (\mathbb{E} [R(Y_a) | A^* = a] - \mathbb{E}[R(Y_a)]) \right)^2 \\ &\stackrel{(b)}{\leq} |\mathcal{A}| \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a)^2 (\mathbb{E} [R(Y_a) | A^* = a] - \mathbb{E}[R(Y_a)])^2 \\ &\leq |\mathcal{A}| \sum_{a, a^* \in \mathcal{A}} \mathbb{P}(A^* = a) \mathbb{P}(A^* = a^*) (\mathbb{E} [R(Y_a) | A^* = a^*] - \mathbb{E}[R(Y_a)])^2 \\ &\stackrel{(c)}{\leq} \frac{|\mathcal{A}|}{2} \sum_{a, a^* \in \mathcal{A}} \mathbb{P}(A^* = a) \mathbb{P}(A^* = a^*) D(P(Y_a | A^* = a^*) || P(Y_a)) \\ &\stackrel{(d)}{=} \frac{|\mathcal{A}| I(A^*; (A, Y))}{2} \end{aligned}$$

where (b) follows from the Cauchy–Schwarz inequality, (c) follows from Fact 9, and (a) and (d) follow from Proposition 2. \blacksquare

Combining Proposition 3 with Proposition 1 shows that $\mathbb{E} [\text{Regret}(T, \pi^{\text{TS}})] \leq \sqrt{\frac{1}{2} |\mathcal{A}| H(A^*) T}$. Bubeck and Liu (2013) show $\mathbb{E} [\text{Regret}(T, \pi^{\text{TS}})] \leq 14\sqrt{|\mathcal{A}|T}$ and that this bound is order optimal, in the sense that for any time horizon T and number of actions $|\mathcal{A}|$ there exists a prior distribution over p^* such that $\inf_{\pi} \mathbb{E} [\text{Regret}(T, \pi)] \geq \frac{1}{20} \sqrt{|\mathcal{A}|T}$.

6.4 Full Information

Our focus in this paper is on problems with *partial feedback*. For such problems, what the decision maker observes depends on the actions selected, which leads to a tension between exploration and exploitation. Problems with full information arise as an extreme point of our formulation where the outcome $Y_{t,a}$ is perfectly revealed by observing $Y_{t,\tilde{a}}$ for any $\tilde{a} \neq a$; what is learned does not depend on the selected action. The next proposition shows that under full information, the information ratio is bounded by $1/2$.

Proposition 4 *Suppose for each $t \in \mathbb{N}$ there is a random variable $Z_t : \Omega \rightarrow \mathcal{Z}$ such that for each $a \in \mathcal{A}$, $Y_{t,a} = (a, Z_t)$. Then for all $t \in \mathbb{N}$, $\Gamma_t \leq 1/2$ almost surely.*

Proof As before we bound the numerator of Γ_t by $\sqrt{1/2}$ times the denominator:

$$\begin{aligned}
 \mathbb{E}[R(Y_{A^*}) - R(Y_A)] &\stackrel{(a)}{=} \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) (\mathbb{E}[R(Y_a)|A^* = a] - \mathbb{E}[R(Y_a)]) \\
 &\stackrel{(b)}{\leq} \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) \sqrt{\frac{1}{2} D(P(Y_a|A^* = a) || P(Y_a))} \\
 &\stackrel{(c)}{\leq} \sqrt{\frac{1}{2} \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) D(P(Y_a|A^* = a) || P(Y_a))} \\
 &\stackrel{(d)}{=} \sqrt{\frac{1}{2} \sum_{a, a^* \in \mathcal{A}} \mathbb{P}(A^* = a) \mathbb{P}(A^* = a^*) D(P(Y_a|A^* = a^*) || P(Y_a))} \\
 &\stackrel{(e)}{=} \sqrt{\frac{I(A^*; (A, Y))}{2}},
 \end{aligned}$$

where again (a) and (e) follow from Proposition 2, (b) follows from Fact 9 and (c) follows from Jensen’s inequality. Equality (d) follows because $D(P(Y_a|A^* = a^*) || P(Y_a)) = D(P(Z|A^* = a^*) || P(Z))$ does not depend on the sampled action a under full information. ■

Combining this result with Proposition 1 shows $\mathbb{E}[\text{Regret}(T, \pi^{\text{TS}})] \leq \sqrt{\frac{1}{2} H(A^*) T}$. Further, a worst-case bound on the entropy of A^* shows that $\mathbb{E}[\text{Regret}(T, \pi^{\text{TS}})] \leq \sqrt{\frac{1}{2} \log(|\mathcal{A}|) T}$. The bound here improves upon this worst case bound since $H(A^*)$ can be much smaller than $\log(|\mathcal{A}|)$ when the prior distribution is informative.

6.5 Linear Optimization Under Bandit Feedback

The stochastic linear bandit problem has been widely studied (e.g Dani et al., 2008; Rusmevichientong and Tsitsiklis, 2010; Abbasi-Yadkori et al., 2011) and is one of the most important examples of a multi-armed bandit problem with “correlated arms.” In this setting, each action is associated with a finite dimensional feature vector, and the mean reward generated by an action is the inner product between its known feature vector and some unknown parameter vector. Because of this structure, observations from taking one action allow the decision-maker to make inferences about other actions. The next proposition

bounds the information ratio for such problems. Its proof is essentially a generalization of the proof of Proposition 3.

Proposition 5 *If $\mathcal{A} \subset \mathbb{R}^d$ and for each $p \in \mathcal{P}$ there exists $\theta_p \in \mathbb{R}^d$ such that for all $a \in \mathcal{A}$*

$$\mathbb{E}_{y \sim p_a} [R(y)] = a^T \theta_p,$$

then for all $t \in \mathbb{N}$, $\Gamma_t \leq d/2$ almost surely.

Proof Write $\mathcal{A} = \{a_1, \dots, a_K\}$ and, to reduce notation, for the remainder of this proof let $\alpha_i = \mathbb{P}(A^* = a_i)$. Define $M \in \mathbb{R}^{K \times K}$ by

$$M_{i,j} = \sqrt{\alpha_i \alpha_j} (\mathbb{E}[R(Y_{a_i}) | A^* = a_j] - \mathbb{E}[R(Y_{a_i})]),$$

for all $i, j \in \{1, \dots, K\}$. Then, by Proposition 2,

$$\mathbb{E}[R(Y_{A^*}) - R(Y_A)] = \sum_{i=1}^K \alpha_i (\mathbb{E}[R(Y_{a_i}) | A^* = a_i] - \mathbb{E}[R(Y_{a_i})]) = \text{Trace}(M).$$

Similarly, by Proposition 2,

$$\begin{aligned} I(A^*; (A, Y_A)) &= \sum_{i,j} \alpha_i \alpha_j D(P(Y_{a_i} | A^* = a_j) || P(Y_{a_i})) \\ &\stackrel{(a)}{\geq} 2 \sum_{i,j} \alpha_i \alpha_j (\mathbb{E}[R(Y_{a_i}) | A^* = a_j] - \mathbb{E}[R(Y_{a_i})])^2 \\ &= 2 \|M\|_{\mathbb{F}}^2, \end{aligned}$$

where inequality (a) uses Fact 9. This shows, by Fact 10, that

$$\frac{\mathbb{E}[R(Y_{A^*}) - R(Y_A)]^2}{I(A^*; (A, Y_A))} \leq \frac{\text{Trace}(M)^2}{2 \|M\|_{\mathbb{F}}^2} \leq \frac{\text{Rank}(M)}{2}.$$

We now show $\text{Rank}(M) \leq d$. Define

$$\mu = \mathbb{E}[\theta_{p^*}] \quad \mu^j = \mathbb{E}[\theta_{p^*} | A^* = a_j].$$

Then, by the linearity of the expectation operator,

$$M_{i,j} = \sqrt{\alpha_i \alpha_j} ((\mu^j - \mu)^T a_i)$$

and therefore

$$M = \begin{bmatrix} \sqrt{\alpha_1} (\mu^1 - \mu)^T \\ \vdots \\ \sqrt{\alpha_K} (\mu^K - \mu)^T \end{bmatrix} \begin{bmatrix} \sqrt{\alpha_1} a_1 & \cdots & \cdots & \sqrt{\alpha_K} a_K \end{bmatrix}.$$

This shows M is the product of a K by d matrix and a d by K matrix, and hence has rank at most d . \blacksquare

This result shows that $\mathbb{E} [\text{Regret}(T, \pi^{\text{TS}})] \leq \sqrt{\frac{1}{2}H(A^*)dT} \leq \sqrt{\frac{1}{2}\log(|\mathcal{A}|)dT}$ for linear bandit problems. Dani et al. (2007) show this bound is order optimal, in the sense that for any time horizon T and dimension d if the actions set is $\mathcal{A} = \{0, 1\}^d$, there exists a prior distribution over p^* such that $\inf_{\pi} \mathbb{E} [\text{Regret}(T, \pi)] \geq c_0 \sqrt{\log(|\mathcal{A}|)dT}$ where c_0 is a constant that is independent of d and T . The bound here improves upon this worst case bound since $H(A^*)$ can be much smaller than $\log(|\mathcal{A}|)$ when the prior distribution is informative.

6.6 Combinatorial Action Sets and “Semi-Bandit” Feedback

To motivate the information structure studied here, consider a simple resource allocation problem. There are d possible projects, but the decision-maker can allocate resources to at most $m \leq d$ of them at a time. At time t , project $i \in \{1, \dots, d\}$ yields a random reward $\theta_{t,i}$, and the reward from selecting a subset of projects $a \in \mathcal{A} \subset \{a' \subset \{0, 1, \dots, d\} : |a'| \leq m\}$ is $m^{-1} \sum_{i \in a} \theta_{t,i}$. In the linear bandit formulation of this problem, upon choosing a subset of projects a the agent would only observe the overall reward $m^{-1} \sum_{i \in a} \theta_{t,i}$. It may be natural instead to assume that the outcome of each selected project ($\theta_{t,i} : i \in a$) is observed. This type of observation structure is sometimes called “semi-bandit” feedback (Audibert et al., 2013).

A naive application of Proposition 5 to address this problem would show $\Gamma_t \leq d/2$. The next proposition shows that since the entire parameter vector ($\theta_{t,i} : i \in a$) is observed upon selecting action a , we can provide an improved bound on the information ratio. The proof of the proposition is provided in the appendix.

Proposition 6 *Suppose $\mathcal{A} \subset \{a \subset \{0, 1, \dots, d\} : |a| \leq m\}$, and that there are random variables ($\theta_{t,i} : t \in \mathbb{N}, i \in \{1, \dots, d\}$) such that*

$$Y_{t,a} = (\theta_{t,i} : i \in a) \quad \text{and} \quad R(Y_{t,a}) = \frac{1}{m} \sum_{i \in a} \theta_{t,i}.$$

Assume that the random variables $\{\theta_{t,i} : i \in \{1, \dots, d\}\}$ are independent conditioned on \mathcal{F}_t and $\theta_{t,i} \in [-\frac{1}{2}, \frac{1}{2}]$ almost surely for each (t, i) . Then for all $t \in \mathbb{N}$, $\Gamma_t \leq \frac{d}{2m^2}$ almost surely.

In this problem, there are as many as $\binom{d}{m}$ actions, but because Thompson sampling exploits the structure relating actions to one another, its regret is only polynomial in m and d . In particular, combining Proposition 6 with Proposition 1 shows $\mathbb{E} [\text{Regret}(T, \pi^{\text{TS}})] \leq \frac{1}{m} \sqrt{dH(A^*)T}$. Since $H(A^*) \leq \log |\mathcal{A}| = O(m \log(\frac{d}{m}))$ this also yields a bound of order $\sqrt{\frac{d}{m} \log(\frac{d}{m}) T}$. As shown by Audibert et al. (2013), the lower bound for this problem is of order $\sqrt{\frac{d}{m} T}$, so our bound is order optimal up to a $\sqrt{\log(\frac{d}{m})}$ factor.³ It’s also worth

3. In their formulation, the reward from selecting action a is $\sum_{i \in a} \theta_{t,i}$, which is m times larger than in our formulation. The lower bound stated in their paper is therefore of order \sqrt{mdT} . They don’t provide a complete proof of their result, but note that it follows from standard lower bounds in the bandit literature. In the proof of Theorem 5 in that paper, they construct an example in which the decision

pointing that, although there may be as many as $\binom{d}{m}$ actions, the action selection step of Thompson sampling is computationally efficient whenever the offline decision problem $\max_{a \in \mathcal{A}} \theta^T a$ can be solved efficiently.

7. Conclusion

This paper has provided a new analysis of Thompson sampling based on tools from information theory. As such, our analysis inherits the simplicity and elegance enjoyed by work in that field. Further, our results apply to a much broader range of information structures than those studied in prior work on Thompson sampling. Our analysis leads to regret bounds that highlight the benefits of soft knowledge, quantified in terms of the entropy of the optimal-action distribution. Such regret bounds yield insight into how future performance depends on past observations. This is key to assessing the benefits of exploration, and as such, to the design of more effective schemes that trade off between exploration and exploitation. In forthcoming work, we leverage this insight to produce an algorithm that outperforms Thompson sampling.

While our focus has been on providing theoretical guarantees for Thompson sampling, we believe the techniques and quantities used in the analysis may be of broader interest. Our formulation and notation may be complex, but the proofs themselves essentially follow from combining known relations in information theory with the tower property of conditional expectation, Jensen’s inequality, and the Cauchy–Schwartz inequality. In addition, the information theoretic view taken in this paper may provide a fresh perspective on this class of problems.

Acknowledgments

Daniel Russo is supported by a Burt and Deedee McMurty Stanford Graduate Fellowship. This work was supported, in part, by the National Science Foundation under Grant CMMI-0968707. The authors would like to thank the anonymous reviewers for their helpful comments.

Appendix A. Proof of Basic Facts

A.1 Proof of Fact 9

This result is a consequence of *Pinsker’s inequality*, (see Lemma 5.2.8 of Gray (2011)) which states that

$$\sqrt{\frac{1}{2}D(P||Q)} \geq \|P - Q\|_{\text{TV}} \tag{7}$$

maker plays m bandit games in parallel, each with d/m actions. Using that example, and the standard bandit lower bound (see Theorem 3.5 of Bubeck and Cesa-Bianchi (2012)), the agent’s regret from each component must be at least $\sqrt{\frac{d}{m}T}$, and hence her overall expected regret is lower bounded by a term of order $m\sqrt{\frac{d}{m}T} = \sqrt{mdT}$.

where $\|P - Q\|_{\text{TV}}$ is the total variation distance between P and Q . When Ω is countable, $\|P - Q\|_{\text{TV}} = \frac{1}{2} \sum_{\omega} |P(\omega) - Q(\omega)|$. More generally, if P and Q are both absolutely continuous with respect to some base measure μ , then $\|P - Q\|_{\text{TV}} = \frac{1}{2} \int_{\Omega} \left| \frac{dP}{d\mu} - \frac{dQ}{d\mu} \right| d\mu$, where $\frac{dP}{d\mu}$ and $\frac{dQ}{d\mu}$ are the Radon–Nikodym derivatives of P and Q with respect to μ . We now prove Fact 9.

Proof Choose a base measure μ so that P and Q are absolutely continuous with respect to μ . This is always possible: since P is absolutely continuous with respect to Q by hypothesis, one could always choose this base measure to be Q . Let $f(\omega) = g(X(\omega)) - \inf_{\omega'} g(X(\omega')) - 1/2$ so that $f : \Omega \rightarrow [-1/2, 1/2]$ and f and $g(X)$ differ only by a constant. Then,

$$\begin{aligned} \sqrt{\frac{1}{2} D(P||Q)} &\geq \frac{1}{2} \int \left| \frac{dP}{d\mu} - \frac{dQ}{d\mu} \right| d\mu \geq \frac{1}{2} \int \left| 2 \left(\frac{dP}{d\mu} - \frac{dQ}{d\mu} \right) f \right| d\mu \\ &\geq \int \left(\frac{dP}{d\mu} - \frac{dQ}{d\mu} \right) f d\mu \\ &= \int f dP - \int f dQ \\ &= \mathbb{E}_P[g(X)] - \mathbb{E}_Q[g(X)], \end{aligned}$$

where the first inequality is Pinsker's inequality. ■

A.2 Proof of Fact 10

Proof Fix a rank r matrix $M \in \mathbb{R}^{k \times k}$ with singular values $\sigma_1 \geq \dots \geq \sigma_r$. By the Cauchy Shwartz inequality,

$$\|M\|_* \stackrel{\text{Def}}{=} \sum_{i=1}^r \sigma_i \leq \sqrt{r} \sqrt{\sum_{i=1}^r \sigma_i^2} \stackrel{\text{Def}}{=} \sqrt{r} \|M\|_F. \quad (8)$$

Now, we show

$$\begin{aligned} \text{Trace}(M) = \text{Trace} \left(\frac{1}{2} M + \frac{1}{2} M^T \right) &\stackrel{(a)}{\leq} \left\| \frac{1}{2} M + \frac{1}{2} M^T \right\|_* \stackrel{(b)}{\leq} \frac{1}{2} \|M\|_* + \frac{1}{2} \|M^T\|_* \\ &\stackrel{(c)}{=} \|M\|_* \stackrel{(d)}{\leq} \sqrt{r} \|M\|_F. \end{aligned}$$

Here (b) follows from the triangle inequality and the fact that norms are homogeneous of degree one. Inequality (c) uses that the singular values of M and M^T are the same, and inequality (d) follows from equation (8).

To justify inequality (a), we show that for any symmetric matrix W , $\text{Trace}(W) \leq \|W\|_*$. To see this, let W be a rank r matrix and let $\tilde{\sigma}_1 \geq \dots \geq \tilde{\sigma}_r$ denote its singular values. Since W is symmetric, it has r nonzero real valued eigenvalues $\lambda_1, \dots, \lambda_r$. If these are sorted so that $|\lambda_1| \geq \dots \geq |\lambda_r|$ then $\tilde{\sigma}_i = |\lambda_i|$ for each $i \in \{1, \dots, r\}$.⁴ This shows $\text{Trace}(M) = \sum_{i=1}^r \lambda_i \leq \sum_{i=1}^r \tilde{\sigma}_i = \|M\|_*$. ■

4. This fact is stated, for example, in Appendix A.5 of Boyd and Vandenberghe (2004).

Appendix B. Proof of Proposition 6

The proof relies on the following lemma, which lower bounds the information gain due to selecting an action a . The proof is provided below, and relies on the chain–rule of Kullback–Leibler divergence.

Lemma 1 *Under the conditions of Proposition 6, for any $a \in \mathcal{A}$,*

$$I(A^*; Y_a) \geq 2 \sum_{i \in a} \mathbb{P}(i \in A^*) (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i])^2$$

Proof In the proof of this lemma, for any $i \in a$ we let $\theta_{<i} = (\theta_j : j < i, j \in a)$. Since $a \in \mathcal{A}$ is fixed throughout this proof, we do not display the dependence of $\theta_{<i}$ on a . Recall that by Fact 6,

$$I(A^*; Y_a) = \sum_{a^* \in \mathcal{A}} \mathbb{P}(A^* = a^*) D(P(Y_a | A^* = a^*) || P(Y_a)).$$

where here,

$$\begin{aligned} D(P(Y_a | A^* = a^*) || P(Y_a)) &= D(P((\theta_i : i \in a) | A^* = a^*) || P((\theta_i : i \in a))) \\ &\stackrel{(a)}{=} \sum_{i \in a} \mathbb{E} \left[D(P(\theta_i | A^* = a^*, \theta_{<i}) || P(\theta_i | \theta_{<i})) \Big| A^* = a^* \right] \\ &\stackrel{(b)}{\geq} 2 \sum_{i \in a} \mathbb{E} \left[(\mathbb{E}[\theta_i | A^* = a^*, \theta_{<i}] - \mathbb{E}[\theta_i | \theta_{<i}])^2 \Big| A^* = a^* \right] \\ &\stackrel{(c)}{=} 2 \sum_{i \in a} \mathbb{E} \left[(\mathbb{E}[\theta_i | A^* = a^*, \theta_{<i}] - \mathbb{E}[\theta_i])^2 \Big| A^* = a^* \right] \\ &\stackrel{(d)}{\geq} 2 \sum_{i \in a} (\mathbb{E}[\theta_i | A^* = a^*] - \mathbb{E}[\theta_i])^2. \end{aligned}$$

Equality (a) follows from the chain rule of Kullback–Leibler divergence, (b) follows from Fact 9, (c) follows from the assumption that $(\theta_i : 1 \leq i \leq d)$ are independent conditioned on any history of observations, and (d) follows from Jensen’s inequality and the tower property

of conditional expectation. Now, we can show

$$\begin{aligned}
 I(A^*; Y_a) &\geq 2 \sum_{i \in a} \sum_{a^* \in \mathcal{A}} \mathbb{P}(A^* = a^*) (\mathbb{E}[\theta_i | A^* = a^*] - \mathbb{E}[\theta_i])^2 \\
 &= 2 \sum_{i \in a} \mathbb{P}(i \in A^*) \sum_{a^* \in \mathcal{A}} \frac{\mathbb{P}(A^* = a^*)}{\mathbb{P}(i \in A^*)} (\mathbb{E}[\theta_i | A^* = a^*] - \mathbb{E}[\theta_i])^2 \\
 &\geq 2 \sum_{i \in a} \mathbb{P}(i \in A^*) \sum_{a^*: i \in a^*} \frac{\mathbb{P}(A^* = a^*)}{\mathbb{P}(i \in A^*)} (\mathbb{E}[\theta_i | A^* = a^*] - \mathbb{E}[\theta_i])^2 \\
 &= 2 \sum_{i \in a} \mathbb{P}(i \in A^*) \sum_{a^*: i \in a^*} \mathbb{P}(A^* = a^* | i \in A^*) (\mathbb{E}[\theta_i | A^* = a^*] - \mathbb{E}[\theta_i])^2 \\
 &\stackrel{(e)}{\geq} 2 \sum_{i \in a} \mathbb{P}(i \in A^*) \left(\sum_{a^*: i \in a^*} \mathbb{P}(A^* = a^* | i \in A^*) (\mathbb{E}[\theta_i | A^* = a^*] - \mathbb{E}[\theta_i]) \right)^2 \\
 &\stackrel{(f)}{=} 2 \sum_{i \in a} \mathbb{P}(i \in A^*) (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i])^2,
 \end{aligned}$$

where (e) and (f) follow from Jensen's inequality and the tower property of conditional expectation, respectively. \blacksquare

Lemma 2 *Under the conditions of Proposition 6,*

$$I(A^*; (A, Y_A)) \geq 2 \sum_{i=1}^d \mathbb{P}(i \in A^*)^2 (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i])^2.$$

Proof By Lemma 1 and the tower property of conditional expectation,

$$\begin{aligned}
 I(A^*; (A, Y_A)) &= \mathbb{P}(A^* = a) I(A^*; Y_a) \\
 &\geq 2 \sum_{a \in \mathcal{A}} \mathbb{P}(A^* = a) \left[\sum_{i \in a} \mathbb{P}(i \in A^*) (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i])^2 \right] \\
 &= 2 \sum_{i=1}^d \sum_{a: i \in a} \mathbb{P}(A^* = a) (\mathbb{P}(i \in A^*) (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i])^2) \\
 &= 2 \sum_{i=1}^d \mathbb{P}(i \in A^*)^2 (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i])^2.
 \end{aligned}$$

Now, we complete the proof of Proposition 6. \blacksquare

Proof The proof establishes that the numerator of the information ratio is less than $d/2m^2$ times its denominator:

$$\begin{aligned}
 m\mathbb{E}[R(Y_{A^*}) - R(Y_A)] &= \mathbb{E} \sum_{i \in A^*} \theta_i - \mathbb{E} \sum_{i \in A} \theta_i \\
 &= \sum_{i=1}^d \mathbb{P}(i \in A^*) \mathbb{E}[\theta_i | i \in A^*] - \sum_{i=1}^d \mathbb{P}(i \in A) \mathbb{E}[\theta_i | i \in A] \\
 &= \sum_{i=1}^d \mathbb{P}(i \in A^*) (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i]) \\
 &\leq \sqrt{d} \sqrt{\sum_{i=1}^d \mathbb{P}(i \in A^*)^2 (\mathbb{E}[\theta_i | i \in A^*] - \mathbb{E}[\theta_i])^2} \\
 &\leq \sqrt{\frac{dI(A^*; (A, Y_A))}{2}}.
 \end{aligned}$$

■

Appendix C. Proof of Fact 6

We could not find a reference that proves Fact 6 in a general setting, and will therefore provide a proof here.

Consider random variables $X : \Omega \rightarrow \mathcal{X}$ and $Y : \Omega \rightarrow \mathcal{Y}$ where \mathcal{X} is assumed to be a finite set but Y is a general random variable. We show

$$I(X; Y) = \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) D(\mathbb{P}(Y \in \cdot | X = x) \| \mathbb{P}(Y \in \cdot)).$$

When Y has finite support this result follows easily by using that $\mathbb{P}(X = x, Y = y) = \mathbb{P}(X = x)\mathbb{P}(Y = y | X = x)$. For an extension to general random variables Y we follow chapter 5 of Gray (2011).

The extension follows by considering quantized versions of the random variable Y . For a finite partition $\mathcal{Q} = \{Q_i \subset \mathcal{Y}\}_{i \in I}$, we denote by $Y_{\mathcal{Q}}$ the quantization of Y defined so that $Y_{\mathcal{Q}}(\omega) = i$ if and only if $Y(\omega) \in Q_i$. As shown in Gray (2011), for two probability measures P_1 and P_2 on (Ω, \mathcal{F}) ,

$$D(P_1(Y \in \cdot) \| P_2(Y \in \cdot)) = \sup_{\mathcal{Q}} D(P_1(Y_{\mathcal{Q}} \in \cdot) \| P_2(Y_{\mathcal{Q}} \in \cdot)), \quad (9)$$

where the supremum is taken over all finite partitions of \mathcal{Y} , and whenever $\overline{\mathcal{Q}}$ is a refinement of \mathcal{Q} ,

$$D(P_1(Y_{\overline{\mathcal{Q}}} \in \cdot) \| P_2(Y_{\overline{\mathcal{Q}}} \in \cdot)) \geq D(P_1(Y_{\mathcal{Q}} \in \cdot) \| P_2(Y_{\mathcal{Q}} \in \cdot)). \quad (10)$$

Since $I(X; Y) = D(\mathbb{P}(X \in \cdot, Y \in \cdot) \| \mathbb{P}(X \in \cdot)\mathbb{P}(Y \in \cdot))$ these properties also apply to $I(X; Y_{\mathcal{Q}})$. Now, for each $x \in \mathcal{X}$ we can introduce a sequence of successively refined partitions $(\mathcal{Q}_n^x)_{n \in \mathbb{N}}$

so that

$$\begin{aligned} D(\mathbb{P}(Y \in \cdot | X = x) \parallel \mathbb{P}(Y \in \cdot)) &= \lim_{n \rightarrow \infty} D(\mathbb{P}(Y_{\mathcal{Q}_n^x} \in \cdot | X = x) \parallel \mathbb{P}(Y_{\mathcal{Q}_n^x} \in \cdot)) \\ &= \sup_{\mathcal{Q}} D(\mathbb{P}(Y_{\mathcal{Q}} \in \cdot | X = x) \parallel \mathbb{P}(Y_{\mathcal{Q}} \in \cdot)). \end{aligned}$$

Let the finite partition \mathcal{Q}_n denote the refinement of the partitions $\{\mathcal{Q}_n^x\}_{x \in \mathcal{X}}$. That is, for each x , every set in \mathcal{Q}_n^x can be written as the union of sets in \mathcal{Q}_n . Then,

$$\begin{aligned} I(X; Y) &\geq \lim_{n \rightarrow \infty} I(X; Y_{\mathcal{Q}_n}) \\ &= \lim_{n \rightarrow \infty} \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) D(\mathbb{P}(Y_{\mathcal{Q}_n} \in \cdot | X = x) \parallel \mathbb{P}(Y_{\mathcal{Q}_n} \in \cdot)) \\ &\geq \lim_{n \rightarrow \infty} \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) D(\mathbb{P}(Y_{\mathcal{Q}_n^x} \in \cdot | X = x) \parallel \mathbb{P}(Y_{\mathcal{Q}_n^x} \in \cdot)) \\ &= \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) D(\mathbb{P}(Y \in \cdot | X = x) \parallel \mathbb{P}(Y \in \cdot)) \\ &= \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) \sup_{\mathcal{Q}^x} D(\mathbb{P}(Y_{\mathcal{Q}^x} \in \cdot | X = x) \parallel \mathbb{P}(Y_{\mathcal{Q}^x} \in \cdot)) \\ &\geq \sup_{\mathcal{Q}} \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) D(\mathbb{P}(Y_{\mathcal{Q}} \in \cdot | X = x) \parallel \mathbb{P}(Y_{\mathcal{Q}} \in \cdot)) \\ &= \sup_{\mathcal{Q}} I(X; Y_{\mathcal{Q}}) = I(X; Y), \end{aligned}$$

hence the inequalities above are equalities and our claim is established.

Appendix D. Technical Extensions

D.1 Infinite Action Spaces

In this section, we discuss how to extend our results to problems where the action space is infinite. For concreteness, we will focus on linear optimization problems. Specifically, throughout this section our focus is on problems where $\mathcal{A} \subset \mathbb{R}^d$ and for each $p \in \mathcal{P}$ there is some $\theta_p \in \Theta \subset \mathbb{R}^d$ such that for all $a \in \mathcal{A}$

$$\mathbb{E}_{y \sim p_a} [R(y)] = a^T \theta_p.$$

We assume further that there is a fixed constant $c \in \mathbb{R}$ such that $\sup_{a \in \mathcal{A}} \|a\|_2 \leq c$ and $\sup_{\theta \in \Theta} \|\theta\|_2 \leq c$. Because \mathcal{A} is compact, it can be extremely well approximated by a finite set. In particular, we can choose a finite cover $\mathcal{A}_\epsilon \subset \mathcal{A}$ with $\log |\mathcal{A}_\epsilon| = O(d \log(dc/\epsilon))$ such that for every $a \in \mathcal{A}$, $\min_{\tilde{a} \in \mathcal{A}_\epsilon} \|a - \tilde{a}\|_2 \leq \epsilon$. By the Cauchy-Schwartz inequality, this implies that $|(a - \tilde{a})^T \theta_p| \leq c\epsilon$ for every θ_p .

We introduce a quantization $A_{\mathcal{Q}}^*$ of the random variable A^* . $A_{\mathcal{Q}}^*$ is supported only on the set \mathcal{A}_ϵ , but satisfies $\|A^*(\omega) - A_{\mathcal{Q}}^*(\omega)\|_2 \leq \epsilon$ for each $\omega \in \Omega$.

Consider a variant of Thompson sampling that randomizes according to the distribution of A_Q^* . That is, for each $a \in \mathcal{A}_\epsilon$, $\mathbb{P}(A_t = a | \mathcal{F}_t) = \mathbb{P}(A_Q^* = a | \mathcal{F}_t)$. Then, our analysis shows,

$$\begin{aligned} \mathbb{E} \sum_{t=1}^T [R(Y_{t,A^*}) - R(Y_{t,A_t})] &= \mathbb{E} \sum_{t=1}^T [R(Y_{t,A^*}) - R(Y_{t,A_Q^*})] + \mathbb{E} \sum_{t=1}^T [R(Y_{t,A_Q^*}) - R(Y_{t,A_t})] \\ &\leq \epsilon T + \sqrt{dT H(A_Q^*)}. \end{aligned}$$

This new bound depends on the time horizon T , the dimension d of the linear model, the entropy of the prior distribution of the quantized random variable A_Q^* and the discretization error ϵT . Since $H(A_Q^*) \leq \log(|\mathcal{A}_\epsilon|)$, by choosing $\epsilon = (cT)^{-1}$ and choosing a finite cover with $\log(|\mathcal{A}_\epsilon|) = O(d \log(dcT))$ we attain a regret bound of order $d\sqrt{T \log(dcT)}$. Note that for $d \leq T$, this bound is of order $d\sqrt{T \log(cT)}$, but for $d > T \geq 1$, we have a trivial regret bound of $T < \sqrt{dT}$.

Here, for simplicity, we have considered a variant of Thompson sampling that randomizes according to the posterior distribution of the quantized random variable A_Q^* . However, with a somewhat more careful analysis, one can provide a similar result for an algorithm that randomizes according to the posterior distribution of A^* .

D.2 Unbounded Noise

In this section we relax Assumption 1, which requires that reward distributions are uniformly bounded. This assumption is required for Fact 9 to hold, but otherwise was never used in our analysis. Here, we show that an analogue to Fact 9 holds in the more general setting where reward distributions are *sub-Gaussian*. This yields more general regret bounds.

A random variable is sub-Gaussian if its moment generating function is dominated by that of a Gaussian random variable. Gaussian random variables are sub-Gaussian, as are real-valued random variables with bounded support.

Definition 1 Fix a deterministic constant $\sigma \in \mathbb{R}$. A real-valued random variable X is *sub-Gaussian* if for all $\lambda \in \mathbb{R}$

$$\mathbb{E}[\exp\{\lambda X\}] \leq \exp\{\lambda^2 \sigma^2 / 2\}.$$

More generally X is σ sub-Gaussian conditional on a sigma-algebra $\mathcal{G} \subset \mathcal{F}$ if for all $\lambda \in \mathbb{R}$,

$$\mathbb{E}[\exp\{\lambda X\} | \mathcal{G}] \leq \exp\{\lambda^2 \sigma^2 / 2\}$$

almost surely.

The next lemma extends Fact 9 to sub-Gaussian random variables.

Lemma 3 Suppose there is a \mathcal{F}_t measurable random variable σ such that for each $a \in \mathcal{A}$, $R(Y_{t,a})$ is σ sub-Gaussian conditional on \mathcal{F}_t . Then for every $a, a^* \in \mathcal{A}$

$$\mathbb{E}_t[R(Y_{t,a}) | A^* = a^*] - \mathbb{E}_t[R(Y_{t,a})] \leq \sigma \sqrt{2D (P_t(Y_{t,a} | A^* = a^*) || P_t(Y_{t,a}))},$$

almost surely.

Using this Lemma in place of Fact 9 in our analysis leads to the following corollary.

Corollary 1 *Fix a deterministic constant $\sigma \in \mathbb{R}$. Suppose that conditioned on \mathcal{F}_t , $R(Y_{t,a}) - \mathbb{E}[R(Y_{t,a})|\mathcal{F}_t]$ is σ sub-Gaussian. Then*

$$\Gamma_t \leq 2|\mathcal{A}|\sigma^2$$

almost surely for each $t \in \mathbb{N}$. Furthermore, $\Gamma_t \leq 2\sigma^2$ under the conditions of Proposition 4, $\Gamma_t \leq 2d\sigma^2$ under the conditions of Proposition 5, and $\Gamma_t \leq \frac{2d\sigma^2}{m^2}$ under the conditions of Proposition 6.

It's worth highlighting two cases under which the conditions of Corollary 1 are satisfied. The first is a setting with a Gaussian prior and Gaussian reward noise. The second case is when under any model $p \in \mathcal{P}$ expected rewards lie in a bounded interval and reward noise is sub-Gaussian.

1. Suppose that for each $a \in \mathcal{A}$, $R(Y_{t,a})$ follows a Gaussian distribution with variance less than σ^2 , $Y_{t,a} = R(Y_{t,a})$ and, reward noise $R(Y_{t,a}) - \mathbb{E}[R(Y_{t,a})|p^*]$ is Gaussian with known variance. Then, the posterior predictive distribution of $R(Y_{t,a})$, $\mathbb{P}(R(Y_{t,a}) \in \cdot | \mathcal{F}_t)$ is Gaussian with variance less than σ^2 for each $a \in \mathcal{A}$ and $t \in \mathbb{N}$.
2. Fix constants $C \in \mathbb{R}$ and $\sigma \in \mathbb{R}$. Suppose that almost surely $\mathbb{E}[R(Y_{t,a})|p^*] \in [-\frac{C}{2}, \frac{C}{2}]$ and reward noise is $R(Y_{t,a}) - \mathbb{E}[R(Y_{t,a})|p^*]$ is σ sub-Gaussian. Then, conditioned on the history, $R(Y_{t,a}) - \mathbb{E}[R(Y_{t,a})|\mathcal{F}_t]$ is $\sqrt{C^2 + \sigma^2}$ sub-Gaussian almost surely.

The first case relies on the fact that $\mathbb{E}[R(Y_{t,a})|p^*]$ is C -sub-Gaussian, as well as the following fact about sub-Gaussian random variables.

Fact 11 *Consider random variables (X_1, \dots, X_K) . If for each X_i there is a deterministic σ_i such that X_i is σ_i sub-Gaussian conditional on X_1, \dots, X_{i-1} , then $\sum_{i=1}^K X_i$ is $\sqrt{\sum_{i=1}^K \sigma_i^2}$ sub-Gaussian.*

Proof For any $\lambda \in \mathbb{R}$,

$$\begin{aligned} \mathbb{E} \left[\exp \left\{ \lambda \sum_{i=1}^K X_i \right\} \right] &= \mathbb{E} \left[\prod_{i=1}^K \exp \{ \lambda X_i \} \right] = \mathbb{E} \left[\prod_{i=1}^K \mathbb{E} \left[\exp \{ \lambda X_i \} \middle| X_1, \dots, X_{i-1} \right] \right] \\ &\leq \prod_{i=1}^K \exp \left\{ \frac{\lambda^2 \sigma_i^2}{2} \right\} = \exp \left\{ \frac{\lambda^2 (\sum_{i=1}^K \sigma_i^2)}{2} \right\}. \end{aligned}$$

■

The proof of Lemma 3 relies on the following property of Kullback–Leibler divergence.

Fact 12 *(Variational form of KL-Divergence given in Theorem 5.2.1 of Gray (2011)) Fix two probability distributions P and Q such that P is absolutely continuous with respect to Q . Then,*

$$D(P||Q) = \sup_X \{ \mathbb{E}_P[X] - \log \mathbb{E}_Q[\exp\{X\}] \},$$

where \mathbb{E}_P and \mathbb{E}_Q denote the expectation operator under P and Q respectively, and the supremum is taken over all real valued random variables X such that $\mathbb{E}_P[X]$ is well defined and $\mathbb{E}_Q[\exp\{X\}] < \infty$.

We now show that Lemma 3 follows from the variational form of KL–Divergence.

Proof Define the random variable $X = R(Y_{t,a}) - \mathbb{E}_t[R(Y_{t,a})]$. Then, for arbitrary $\lambda \in \mathbb{R}$, applying Fact 12 to λX yields

$$\begin{aligned} D(P_t(Y_{t,a}|A^* = a^*) || P_t(Y_{t,a})) &\geq \lambda \mathbb{E}_t[X|A^* = a^*] - \log \mathbb{E}_t[\exp\{\lambda X\}] \\ &= \lambda (\mathbb{E}_t[R(Y_{t,a})|A^* = a^*] - \mathbb{E}_t[R(Y_{t,a})]) - \log \mathbb{E}[\exp\{\lambda X\}|\mathcal{F}_t] \\ &\geq \lambda (\mathbb{E}_t[R(Y_{t,a})|A^* = a^*] - \mathbb{E}_t[R(Y_{t,a})]) - (\lambda^2 \sigma^2 / 2). \end{aligned}$$

Maximizing over λ yields the result. ■

References

- Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24, 2011.
- R. Agrawal, D. Teneketzis, and V. Anantharam. Asymptotically efficient adaptive allocation schemes for controlled iid processes: Finite parameter space. *IEEE Transactions on Automatic Control*, 34(3), 1989.
- S. Agrawal and N. Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proceedings of the 21st Annual Conference on Learning Theory*, 2012.
- S. Agrawal and N. Goyal. Further optimal regret bounds for thompson sampling. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, pages 99–107, 2013a.
- S. Agrawal and N. Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 127–135, 2013b.
- J.-Y. Audibert, S. Bubeck, and G. Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2013.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- S.P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and trends in machine learning*, 5(1):1–122, 2012.
- S. Bubeck and C.-Y. Liu. Prior-free and prior-dependent regret bounds for Thompson sampling. In *Advances in Neural Information Processing Systems*, 2013.
- O. Chapelle and L. Li. An empirical evaluation of Thompson sampling. In *Neural Information Processing Systems (NIPS)*, 2011.
- T.M. Cover and J.A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

- V. Dani, S.M. Kakade, and T.P. Hayes. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, pages 345–352, 2007.
- V. Dani, T.P. Hayes, and S.M. Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Annual Conference on Learning Theory*, pages 355–366, 2008.
- J. Gittins, K. Glazebrook, and R. Weber. *Multi-armed bandit allocation indices*. John Wiley & Sons, Ltd, 2011. ISBN 9780470980033.
- A. Gopalan, S. Mannor, and Y. Mansour. Thompson sampling for complex online problems. In *Proceedings of the 31st International Conference on Machine Learning*, pages 100–108, 2014.
- T. Graepel, J.Q. Candela, T. Borchert, and R. Herbrich. Web-scale Bayesian click-through rate prediction for sponsored search advertising in Microsoft’s Bing search engine. In *Proceedings of the 27th International Conference on Machine Learning*, pages 13–20, 2010.
- R.M. Gray. *Entropy and information theory*. Springer, 2011.
- E. Kauffmann, N. Korda, and R. Munos. Thompson sampling: an asymptotically optimal finite time analysis. In *International Conference on Algorithmic Learning Theory*, 2012.
- N. Korda, E. Kaufmann, and R. Munos. Thompson sampling for one-dimensional exponential family bandits. In *Advances in Neural Information Processing Systems*, 2013.
- T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- L. Li. Generalized Thompson sampling for contextual bandits. *arXiv preprint arXiv:1310.7163*, 2013.
- B.C. May, N. Korda, A. Lee, and D.S. Leslie. Optimistic Bayesian sampling in contextual-bandit problems. *The Journal of Machine Learning Research*, 98888:2069–2106, 2012.
- P. Rusmevichientong and J.N. Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.
- D. Russo and B. Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 2014. URL <http://dx.doi.org/10.1287/moor.2014.0650>.
- I.O. Ryzhov, W.B. Powell, and P.I. Frazier. The knowledge gradient algorithm for a general class of online learning problems. *Operations Research*, 60(1):180–195, 2012.
- S.L. Scott. A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26(6):639–658, 2010.
- S.L. Scott. Overview of content experiments: Multi-armed bandit experiments, 2014. URL <https://support.google.com/analytics/answer/2844870?hl=en>. [Online; accessed 9-February-2014].

- N. Srinivas, A. Krause, S.M. Kakade, and M. Seeger. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.
- L. Tang, R. Rosales, A. Singh, and D. Agarwal. Automatic ad format selection via contextual bandits. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, pages 1587–1594. ACM, 2013.
- W.R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.