

Diffusion Asymptotics for Sequential Experiments

Stefan Wager Kuang Xu
 swager@stanford.edu kuangxu@stanford.edu

Stanford Graduate School of Business

Abstract

We propose a new diffusion-asymptotic analysis for sequentially randomized experiments, including those that arise in solving multi-armed bandit problems. In an experiment with n time steps, we let the mean reward gaps between actions scale to the order $1/\sqrt{n}$ so as to preserve the difficulty of the learning task as n grows. In this regime, we show that the behavior of a class of sequentially randomized Markov experiments converges to a diffusion limit, given as the solution of a stochastic differential equation. The diffusion limit thus enables us to derive refined, instance-specific characterization of the stochastic dynamics of adaptive experiments. As an application of this framework, we use the diffusion limit to obtain several new insights on the regret and belief evolution of Thompson sampling. We show that a version of Thompson sampling with an asymptotically uninformative prior variance achieves nearly-optimal instance-specific regret scaling when the reward gaps are relatively large. We also demonstrate that, in this regime, the posterior beliefs underlying Thompson sampling are highly unstable over time.

Keywords: Multi-armed bandit, Thompson sampling, Stochastic differential equation.

1 Introduction

Sequential experiments, pioneered by Wald [1947] and Robbins [1952], involve collecting data over time using a design that adapts to past experience. The promise of sequential (or adaptive) experiments is that, relative to classical randomized trials, they can effectively concentrate power on studying the most promising alternatives and save on costs by helping us avoid repeatedly taking sub-optimal actions. Such experiments have now become widely adopted for automated decision making; for example, Hill et al. [2017] show how sequential experiments can be used to optimize the layout and content of a website, while Ferreira et al. [2018] discuss applications to pricing and online revenue management.

In automated decision-making applications, it is common to run multiple experiments in parallel while searching over a complex decision space, and so it is important to use robust methods that do not require frequent human oversight. As such, most existing in this area has focused on proving robust worst-case guarantees for adaptive learning that can

Draft version: June 2021. We are grateful for detailed comments from Susan Athey, Peter Glynn, David Goldberg, Kei Hirano, Mike Harrison, David Hirshberg, Max Kasy, Emilie Kaufmann, Tor Lattimore, Dan Russo and David Simchi-Levi on an earlier version of this manuscript, and for valuable feedback from a number of colleagues as well as seminar participants at Columbia, Cornell, Heidelberg, Northwestern and Stanford.

help give us confidence in letting these algorithms run fairly independently; see [Bubeck and Cesa-Bianchi \[2012\]](#) for a review and discussion.

More recently, however, there has also been growing interest in using sequential experiments for high-stakes applications where the behavior of any learning algorithm is likely to be constantly scrutinized by a large research team: [Athey et al. \[2021\]](#) discuss the use of sequential experiments for learning better public health interventions, [Caria et al. \[2020\]](#) deploy them for targeting job search assistance for refugees, and [Kasy and Teytelboym \[2020\]](#) consider whom to test for an infectious disease in a setting where testing capacity is limited. In these applications where humans are likely to be closely monitoring and fine-tuning an adaptive experiment, it is valuable to be able to move past worst-case guarantees, and develop a more refined, instance-specific understanding of the stochastic behavior of adaptive experiments. In particular, it may be of interest to understand:

- How can we use domain specific knowledge to sharpen our understanding about how various sequential experiments would perform in practice?
- Beyond just the mean, what is the distribution of resource consumption and errors in an adaptive experiment?
- How does the behavior of adaptive experiments evolve over time, and what do sample paths of actions taken by adaptive experiments look like?

Available worst-case-focused formal results, however, do not provide sharp answers to these questions that would apply broadly to popular algorithms for adaptive experimentation, including the ones used in the studies cited above. A central difficulty here is simply that sequential experiments induce intricate dependence patterns in the data that make sharp finite-sample analysis of their behavior exceedingly delicate.

In this paper, we introduce a new approach to studying sequential experiments based on diffusion approximation. Specifically, we consider the behavior of adaptive experiments in a sequence of problems where, as the number of time steps n grows to infinity, the reward gap between different actions decays as $1/\sqrt{n}$. In this regime, we show that sequentially randomized Markov experiments, a general class of sequential experiments that includes several popular algorithms including Thompson sampling [[Thompson, 1933](#), [Russo et al., 2018](#)], converge weakly to a diffusion limit characterized as the solution to a stochastic differential equation (SDE). We then show that this diffusion limit enables us to derive practical, instance-specific and distributional insights about the behavior of sequential experiments, such as Thompson sampling.

The $1/\sqrt{n}$ arm gap scaling considered here is an important one. It can be thought of as a “moderate data” regime of sequential experimentation, where the problem’s difficulty, expressed in terms of the gap, is not sufficiently large relative to sample size so as to make identifying the optimal action asymptotically trivial, but large enough so that sub-optimal policies can lead to large regret.

In using diffusion approximation to obtain insights about the behavior of large, dynamic processes, our approach builds on a long tradition in operations research, especially in the context of queuing networks [[Iglehart and Whitt, 1970](#), [Harrison and Reiman, 1981](#), [Reiman, 1984](#), [Harrison, 1988](#), [Glynn, 1990](#), [Kelly and Laws, 1993](#), [Gamarnik and Zeevi, 2006](#)]. A key insight from this line of work is that, by focusing on a heavy-traffic limit where server utilization approaches full capacity and wait times diverge, the behavior of a queuing network can be well approximated by a Brownian limit where, as argued by [Kelly and Laws \[1993\]](#), “important features of good control policies are displayed in sharpest relief.”

Likewise, in our setting, we find that focusing on a scaling regime with weak signals and long time horizons enables us to capture key aspects of sequential experiments in terms of a tractable diffusion approximation.

1.1 Overview of Main Results

Throughout this paper, we focus on the following K -armed setting, also known as a stochastic K -armed bandit. There is a sequence of decision points $i = 1, 2, \dots, n$ at which an agent chooses which action $A_i \in \{1, \dots, K\}$ to take and then observes a reward $Y_i \in \mathbb{R}$. Here, Y_i is assumed to be drawn from a distribution P_{A_i} , where the action A_i may depend on past observations, and is conditionally independent from all other aspects of the system given the realization of A_i . Our goal is to use diffusion approximation to characterize the behavior of K -armed bandits in terms of a number of metrics, including regret

$$R^n = n \sup_{1 \leq k \leq K} \{\mu_k\} - \mathbb{E} \left[\sum_{i=1}^n \mu_{A_i} \right], \quad \mu_k = \mathbb{E}_{P_k} [Y], \quad (1.1)$$

i.e., the shortfall in rewards incurred by the bandit algorithm relative to always taking the best action.

The first part of our paper, Section 2, is focused on establishing a diffusion limit for a broad class of sequential experimentation, which we refer to as sequentially randomized Markov experiments. This result applies to a wide variety of adaptive experimentation rules arising in statistics, machine learning and behavioral economics. Of particular interest is a characterization of this diffusion limit using random time-change applied to a driving Brownian motion, which enables much of our subsequent theoretical analysis.

In the second part of our paper, Section 3, we use our diffusion limit to carry out a detailed study of Thompson sampling [Thompson, 1933, Russo et al., 2018], a widely used Bayesian heuristic for adaptive experimentation. One of our key findings is that it is essential to use what we refer to as “asymptotically undersmoothed” Thompson sampling, i.e., to use a disproportionately large prior variance so that its regularizing effect becomes vanishingly small in the limit. Without undersmoothing, we show that the regret of Thompson sampling can become unbounded as the reward gap between the best and second-best arms, Δ , grows, a rather counter-intuitive finding given that one would expect the learning task to become easier as the signal strength increases. In contrast, we prove that the regret of undersmoothed Thompson sampling decays faster than $1/\Delta^\beta$ for any $\beta < 1$, which is a nearly optimal regret profile [Mannor and Tsitsiklis, 2004]. To the best of our knowledge, this is one of the first known results that demonstrate instance-dependent regret optimality for Thompson sampling. Finally, we leverage the diffusion limit to obtain distributional properties of sample paths under Thompson sampling. We prove that Thompson sampling can be highly unstable with large swings in its posterior beliefs, even when the magnitudes of μ_k are arbitrarily large.

On the methodological front, our work introduces new tools to the study of adaptive experiments. We show weak convergence to the diffusion limit for sequentially randomized Markov experiments using the martingale framework developed by Stroock and Varadhan [2007], which hinges on showing that an appropriately scaled generator of the discrete-time Markov process converges to the infinitesimal generator of the diffusion process. Our analysis of the regret profiles of Thompson sampling in the diffusion limit relies on novel proof arguments that heavily exploit properties of Brownian motion, such as the Law of Iterated Logarithm.

1.2 Related Work

The multi-armed bandit problem is a popular framework for studying sequential experimentation; see [Bubeck and Cesa-Bianchi \[2012\]](#) for a broad discussion focused on bounds on the regret (1.1). An early landmark result in this setting is due to [Lai and Robbins \[1985\]](#), who showed that given any fixed set of arms $\{P_k\}_{k=1}^K$, a well-designed sequential algorithm can achieve regret that scales logarithmically with the number n of time steps considered, i.e., $R^n = \mathcal{O}_P(\log(n))$. Meanwhile, given any fixed time horizon n , it is possible to choose probability distributions $\{P_k\}_{k=1}^K$ such that the expected regret $\mathbb{E}[R^n]$ of any sequential algorithm is lower-bounded to order \sqrt{Kn} [[Auer et al., 2002](#)]. It is worth-noting that the problem instance that achieves the \sqrt{Kn} regret lower bound in [Auer et al. \[2002\]](#) involves the same mean reward scaling as we use for our diffusion limit, suggesting that the diffusion scaling proposed here captures the behavior of the most challenging (and thus potentially most interesting) sub-family of learning tasks.

Thompson sampling [[Thompson, 1933](#), [Russo et al., 2018](#)] has gained considerable popularity in recent years thanks to its simplicity and impressive empirical performance [[Chapelle and Li, 2011](#)]. Regret bounds for Thompson sampling have been established in the frequentist [[Agrawal and Goyal, 2017](#)] and Bayesian [[Bubeck and Liu, 2014](#), [Lattimore and Szepesvári, 2019](#), [Russo and Van Roy, 2016](#)] settings; the setup here belongs to the first category. None of the existing instance-dependent regret bounds, however, appear to have sufficient precision to yield meaningful characterization in our regime. For example, the instance-dependent upper bound in [Agrawal and Goyal \[2017\]](#) contains a constant to the order of $1/\Delta^4$, where Δ is the gap in mean reward between optimal and sub-optimal arms. This would thus lead to a trivial bound of $\mathcal{O}(n^2)$ in our regime, where mean rewards scale as $1/\sqrt{n}$. Furthermore, many of the existing bounds also require delicate assumptions on the reward distributions (e.g., bounded support, exponential family). In contrast, the diffusion asymptotics adopted in this paper are universal in the sense that they automatically allow us to obtain approximations for a much wider range of reward distributions, requiring only a bounded fourth moment.

Our choice of the diffusion scaling and the ensuing functional limit are motivated by insights from both queueing theory and statistics. The scaling plays a prominent role in heavy-traffic diffusion approximation in queueing networks [[Gamarnik and Zeevi, 2006](#), [Harrison and Reiman, 1981](#), [Reiman, 1984](#)]; in particular, [Harrison \[1988\]](#) uses diffusion approximation to study a dynamic control problem. Here, one considers a sequence of queueing systems in which the excessive service capacity, defined as the difference between arrival rate and service capacity, decays as $1/\sqrt{T}$, where T is the time horizon. Under this asymptotic regime, it is shown that suitably scaled queue-length and workload processes converge to reflected Brownian motion. Like in our problem, the diffusion regime here is helpful because it captures the most challenging problem instances, where the system is at once stable and exhibiting non-trivial performance variability. See also [Glynn \[1990\]](#) for an excellent survey for the use of diffusion approximation in operations research.

The diffusion scaling is further inspired by a recurring insight from statistics that, in order for asymptotic analysis to yield a normal limit that can be used for finite-sample insight, we need to appropriately down-scale the signal strength as the sample size gets large. One concrete example of this phenomenon arises when we seek to learn optimal decision rules from (quasi-)experimental data. Here, in general, optimal behavior involves regret that decays as $1/\sqrt{n}$ with the sample size [[Athey and Wager, 2021](#), [Hirano and Porter, 2009](#), [Kitagawa and Tetenov, 2018](#)]; however, this worst-case regret is only achieved if we

let effect sizes decay as $1/\sqrt{n}$. For any fixed sampling design, it’s possible to achieve faster than $1/\sqrt{n}$ rates asymptotically [Luedtke and Chambaz, 2017].

Diffusion approximations have been also been used for optimal stopping in sequential experiments [Siegmund, 1985]. In this literature, the randomization is typically fixed throughout the horizon. In contrast, in our multi-armed bandit setting the probabilities in the randomization depend on the history which creates a qualitatively different limit object. Our work is also broadly related, in spirit, to recent work on models of learning and experimentation using diffusion processes in the operations research literature [Araman and Caldentey, 2019, Harrison and Sunar, 2015, Wang and Zenios, 2020].

Finally, after we posted a first version of this paper, a number of authors have disseminated work that intersects with our main result. In the special case of Thompson sampling, Fan and Glynn [2021] derive the diffusion limit given in Theorems 1 and 3 using a different proof technique from us: As discussed further in Section 6, we derive the diffusion limit by studying the asymptotic behavior of the Markov chain associated with our sequential experiments following Stroock and Varadhan [2007], whereas Fan and Glynn [2021] develop a direct argument based on the continuous mapping theorem. Meanwhile, Hirano and Porter [2021] extend the “limits of experiments” analysis pioneered by Le Cam [1972] to discrete-time, batched sequential experiments. The limits of experiments paradigm also involves a $1/\sqrt{n}$ -scale local parametrization, and their results can be applied to sequentially randomized Markov experiments whose the randomization probabilities only change at a finite set of pre-specified times. Both Fan and Glynn [2021] and Hirano and Porter [2021] are based on research developed independently from ours.

2 Asymptotics for K-Armed Sequential Experiments

As discussed above, the first goal of this paper is to establish a diffusion limit for a class of sequential experiments. To this end, we first introduce a broad class of sequential experimentation schemes introduced in Section 2.1, which we refer to as sequentially randomized Markov experiments. We describe a diffusion scaling for sequential experiments in Section 2.2; then, in Section 2.3, we establish conditions under which—in this limit—sample paths of sequentially randomized Markov experiments converge weakly to the solution of a stochastic differential equation.

Throughout our analysis, we work within the following K -armed model. This model captures a number of interesting problems and is widely used in the literature [e.g., Bubeck and Cesa-Bianchi, 2012, Lai and Robbins, 1985]. However, we note that it does rule out some phenomena that may arise in applications; for example, we do not allow for distribution shift in the reward distribution of a given arm over time, and we do not allow for long-term consequences of actions, i.e., an action taken in period i cannot affect arm-specific period- i' reward distributions for any $i' > i$. Extending our asymptotic analysis to allow for distribution shift or long-term effects would be of considerable interest, but we do not pursue this line of investigation in the present paper.

Definition 1 (Stochastic K -Armed Bandit). A stochastic K -armed bandit is characterized by time horizon n and a set of K reward distributions $k = 1, \dots, K$. At each decision points $i = 1, 2, \dots, n$, an agent chooses which action $A_i \in \{1, \dots, K\}$ to take and then observes a reward $Y_i \in \mathbb{R}$. The action A_i is a random variable that is measurable with respect to the observed history $\{A_{i'}, Y_{i'}\}_{i'=1}^{i-1}$. Then, conditionally on the chosen action A_i , the reward Y_i is drawn from the distribution P_{A_i} , independently from the observed history.

2.1 Sequentially Randomized Markov Experiments

In the interest of generality, we state our main results in the context of sequentially randomized experiments whose sampling probabilities depend on past observations only through the state variables

$$Q_{k,i} = \sum_{j=1}^i \mathbf{1}(\{A_j = k\}), \quad S_{k,i} = \sum_{j=1}^i \mathbf{1}(\{A_j = k\}) Y_j, \quad (2.1)$$

where $Q_{k,i}$ counts the cumulative number of times arm k has been chosen by the time we collect the i -th sample, and $S_{k,i}$ measures its cumulative reward. When useful, we use the convention $Q_{k,0} = S_{k,0} = 0$. Working with this class of algorithms, which we refer sequentially randomized Markov experiments, enables us to state results that cover many popular ways of running sequential experiments without needing to derive specialized analyses for each of them.

Definition 2 (Sequentially Randomized Markov Experiment). A K -armed sequentially randomized Markov experiment chooses the i -th action A_i by taking a draw from a distribution

$$A_i \mid \{A_1, Y_1, \dots, A_{i-1}, Y_{i-1}\} \sim \text{Multinomial}(\pi_i), \quad (2.2)$$

where the sampling probabilities are computed using a measurable sampling function ψ ,

$$\psi : \mathbb{R}^K \times \mathbb{R}^K \rightarrow \Delta^K, \quad \pi_i = \psi(Q_{i-1}, S_{i-1}), \quad (2.3)$$

where $(Q_{i-1}) = (Q_{k,i-1})_{k=1,\dots,K}$, $S_{i-1} = (S_{k,i-1})_{k=1,\dots,K}$, and Δ^K is the K -dimensional unit simplex.

Remark 1 (Capturing Time Dependence). It is useful to note that the family of experiments described in Definition 2 contains those algorithms whose sampling probabilities may depend on the time period, i . This is captured through the fact that $\sum_k Q_{i,k} = i$, that is, the time period i can be simply read off by calculating the L_1 norm of the vector Q_i .

We now examine several examples of popular algorithms that fit under the sequentially randomized Markov experiment framework.

Example 3. Thompson sampling is a popular Bayesian heuristic for running sequential experiments [Thompson, 1933]. In Thompson sampling an agent starts with a prior belief distribution G_0 on the reward distributions $\{P_k\}_{k=1}^K$. Then, at each step i , the agent draws the k -th arm with probability $\rho_{k,i}$ corresponding to their posterior belief G_{i-1} that P_k has the highest mean, and any so-gathered information to update the posterior G_i using Bayes' rule. The motivation behind Thompson sampling is that it quickly converges to pulling the best arm, and thus achieves low regret [Agrawal and Goyal, 2017, Chapelle and Li, 2011]. Thompson sampling does not always satisfy Definition 2. However, widely used modeling choices involving exponential families for the $\{P_k\}_{k=1}^K$ and conjugate priors for G_0 result in these posterior probabilities $\rho_{k,i}$ satisfying the Markov condition (2.3) [Russo et al., 2018], in which case Thompson sampling yields a sequentially randomized Markov experiment in the sense of Definition 2. See Sections 3 for further discussion.

Example 4. Exploration sampling is a variant of Thompson sampling where, using notation from the above example, we pulling each arm with probability $\pi_{k,i} = \rho_{k,i}(1-\rho_{k,i}) / \sum_{l=1}^K \rho_{l,i}(1-\rho_{l,i})$ instead of $\pi_{k,i} = \rho_{k,i}$ [Kasy and Sautmann, 2021]. Exploration sampling is preferred to

Thompson sampling when the analyst is more interested in identifying the best arm than simply achieving low regret [Kasy and Sautmann, 2021, Russo, 2020]. Exploration sampling satisfies Definition 2 under the same conditions as Thompson sampling.

Example 5. A greedy agent may be tempted to always pull the arm with the highest apparent mean, $S_{i,k}/Q_{i,k}$; however, this strategy may fail to experiment enough and prematurely discard good arms due to early unlucky draws. A tempered greedy algorithm instead chooses

$$\pi_{i,k} = \exp \left[\alpha \frac{S_{k,i}}{Q_{k,i} + c} \right] / \sum_{l=1}^K \exp \left[\alpha \frac{S_{l,i}}{Q_{l,i} + c} \right], \quad (2.4)$$

where $\alpha, c > 0$ are tuning parameters that serve to govern the strength of the extent to which the agent focuses on the greedy choice and to protect against division by zero respectively. The selection choices (2.4) satisfy (2.3) and thus Definition 2 by construction.

Example 6. Similar learning dynamics arise in human psychology and behavioral economics where an agent chooses future actions with a bias towards those that have accrued higher (un-normalized) cumulative reward [Erev and Roth, 1998, Xu and Yun, 2020]. A popular example, known as Luce’s rule [Luce, 1959], uses sampling probabilities

$$\pi_{i,k} = (S_{k,i} \vee \alpha) / \sum_{l=1}^K (S_{l,i} \vee \alpha), \quad (2.5)$$

where $\alpha \geq 0$ is a tuning parameter governing the amount of baseline exploration. More generally, the agent may ascribe to arm k a weight $f(S_{k,i})$, where f is a non-negative potential function, and sample actions with probabilities proportional to the weights. The decision rule in (2.5) only depends on S and thus satisfies (2.3).

Example 7. The Exp3 algorithm, proposed by Auer et al. [2002], uses sampling probabilities

$$\pi_{i,k} = \exp \left[\alpha \sum_{j=1}^{i-1} \frac{\mathbf{1}(\{A_j = k\}) Y_j}{\pi_{j,k}} \right] / \sum_{l=1}^K \exp \left[\alpha \sum_{j=1}^{i-1} \frac{\mathbf{1}(\{A_j = l\}) Y_j}{\pi_{j,l}} \right], \quad (2.6)$$

where again $\alpha > 0$ is a tuning parameter. The advantage of Exp3 is that it can be shown to achieve low regret even when the underlying distributions $\{P_k\}_{k=1}^K$ may be non-stationary and change arbitrarily across samples. The sampling probabilities (2.6) do not satisfy (2.3), and so the Exp3 algorithm is not covered by the results given in this paper; however, it is plausible that a natural extension of our approach to non-stationary problems could be made accommodate it. We leave a discussion of non-stationary problems to future work.

Remark 2 (Continuity of the Sampling Function). With some appropriate adjustments, one may also express the upper-confidence bound (UCB) and ϵ -greedy algorithms in a form that is consistent with Definition 2. Unfortunately, our main results on diffusion approximation do not currently cover these two algorithms. The main reason is that the sampling functions ψ for these algorithms are discontinuous with respect to the underlying state (Q, S) . This causes a problem because the convergence to a diffusion limit, as well as the well-posedness of the limit stochastic integral, requires ψ to be appropriately continuous (Assumption 1). Modifying the UCB and ϵ -greedy in such a manner as to ensure some smoothness in the sampling probability should resolve this issue. Whether a well defined diffusion limit exists even under a discontinuous sampling function, such as that of vanilla UCB or ϵ -greedy, remains an open question.

2.2 A Diffusion Scaling

Next, we specify a sequence of experiments, indexed by n , that admits a diffusion limit. In order for sequentially randomized Markov experiments to admit a limit distribution, we will require both the reward distributions P^n and sampling functions ψ^n used in the n -th experiment to converge in an appropriate sense. First, we will assume that reward distributions satisfy the following scaling. Unless otherwise stated, all reward distributions are assumed to be in the diffusion regime for the remainder of the paper.

Definition 8 (Diffusion Regime of Reward Distributions). Consider a sequence of K -armed stochastic bandit problems in the sense of Definition 1, with reward distributions $\{P_k^n\}_{k,n \in \mathbb{N}}$. We say that this sequence resides in the diffusion regime if there exist $\mu, \sigma \in \mathbb{R}_+^K$ such that

$$\lim_{n \rightarrow \infty} \sqrt{n} \mu_k^n = \mu_k, \quad \lim_{n \rightarrow \infty} (\sigma_k^n)^2 = \sigma_k^2, \quad (2.7)$$

where $\mu_k^n = \mathbb{E}_{P_k^n}[Y]$ and $(\sigma_k^n)^2 = \text{Var}_{P_k^n}[Y]$.

Next, we require the sequence of sampling functions to converge in an appropriate sense. As discussed further below, the natural scaling of the $Q_{k,i}$ and $S_{k,i}$ state variables defined in (2.1) is

$$Q_{k,i}^n = \frac{1}{n} \sum_{j=1}^i \mathbf{1}(\{A_j = k\}), \quad S_{k,i}^n = \frac{1}{\sqrt{n}} \sum_{j=1}^i \mathbf{1}(\{A_j = k\}) Y_j. \quad (2.8)$$

We then say that a sequence of sampling functions ψ^n is convergent if it respects this scaling.

Definition 9. Writing sampling functions in a scale-adapted way as follows,

$$\bar{\psi}^n(q, s) = \psi^n(nq, \sqrt{n}s), \quad q \in [0, 1]^K, s \in \mathbb{R}^K, \quad (2.9)$$

we say that a sequence sampling functions ψ^n satisfying (2.3) is convergent if, for all values of $q \in [0, 1]^K$ and $s \in \mathbb{R}^K$, we have

$$\lim_{n \rightarrow \infty} \bar{\psi}^n(q, s) = \psi(q, s) \quad (2.10)$$

for a limiting sampling function $\psi : [0, 1]^K \times \mathbb{R}^K \rightarrow \Delta^K$.

Our first main result is that, given a sequence of reward distributions satisfying (2.7) and under a number of regularity conditions discussed further below, the sample paths of the scaled statistics $Q_{k,i}^n$ and $S_{k,i}^n$ of a sequentially randomized Markov experiments with convergent sampling functions converge in distribution to the solution to a stochastic differential equation

$$\begin{aligned} dQ_{k,t} &= \psi_k(Q_t, S_t) dt, \\ dS_{k,t} &= \mu_k \psi_k(Q_t, S_t) dt + \sigma_k \sqrt{\psi_k(Q_t, S_t)} dB_{k,t}, \end{aligned} \quad (2.11)$$

where $B_{\cdot,t}$ is a standard K -dimensional Brownian motion, μ_k and σ_k and the mean and variance parameters given in (2.7), and the time variable $t \in [0, 1]$ approximates the ratio i/n . A formal statement is given in Theorem 1.

When using our results in applications, one key practical consideration is in understanding conditions under which it is natural to consider a sequence of sampling functions ψ^n that

is convergent in the sense of Definition 9. The tempered greedy method from Example 5 can immediately be seen to be convergent, provided we use a sequence of tuning parameters α_n and c_n satisfying $\lim_{n \rightarrow \infty} \sqrt{n}\alpha_n = \alpha$ and $\lim_{n \rightarrow \infty} nc_n = c$ for some $\alpha, c \in \mathbb{R}_+$, resulting in a limiting sampling function

$$\psi(q, s) = \exp \left[\alpha \frac{s_k}{q_k + c} \right] / \sum_{l=1}^K \exp \left[\alpha \frac{s_l}{q_l + c} \right]. \quad (2.12)$$

For tempered greedy sampling to be interesting, we in general want the limit α to be strictly positive, else the claimed diffusion limit (2.11) will be trivial. Conversely, for the second parameter, both the limits $c > 0$ and $c = 0$ may be interesting, but working in the $c = 0$ limit may lead to additional technical challenges due to us getting very close to dividing by 0.

Meanwhile, as discussed further in Section 3, variants of Thompson sampling as in Examples 3 and 4 can similarly be made convergent via appropriate choices of the prior G_0 ; and we will again encounter questions regarding whether a scaled parameter analogous to c_n in Example 5 converges to 0 or has a strictly positive limit. Finally, similar convergence should hold for Luce's rule in Example 6 provided that $\sqrt{n}\alpha_n$ converges to a positive limit.

Remark 3 (Non-Zero Mean Rewards). The scaling condition in Definition 8 implies that, in large samples, all arms have roughly zero rewards on average, i.e., $\lim_{n \rightarrow \infty} \mu_k^n = 0$. In some applications, however, it may be more natural to consider a local expansion around non-zero mean rewards, where

$$\lim_{n \rightarrow \infty} \sqrt{n}(\mu_k^n - \mu_0) = \delta_k \quad (2.13)$$

for some potentially non-zero $\mu_0 \neq 0$ and $\delta \in \mathbb{R}^K$. In our general results, we focus on the setting from Definition 8; however, we note that, when applied to any translation-invariant algorithm (i.e., a sequential experiment whose sampling function is invariant to adding a fixed offset to all rewards), any results proven under the setting of Definition 8 will also apply under (2.13). In Section 3, we will use this fact when studying a translation-invariant two-armed Thompson sampling algorithm.

2.3 Convergence to a Diffusion Limit

We are now ready to state our main result establishing a diffusion limit for sequentially randomized Markov experiments with convergent sampling functions. Given any Lipschitz sampling functions, we will show that a suitably scaled version of the process (Q_i^n, S_i^n) converges to an Itô diffusion process. In doing so, we make the following assumptions on the sampling functions.

Assumption 1. The following are true:

1. The limiting sampling function ψ is Lipschitz-continuous.
2. The convergence of $\bar{\psi}^n$ to ψ (Definition 9) occurs uniformly over compact sets.

Define \bar{Q}_t^n to be the linear interpolation of $Q_{[tn]}^n$,

$$\bar{Q}_{k,t}^n = (1 - tn + [tn])Q_{k,[tn]}^n + (tn - [tn])Q_{k,[tn]+1}^n, \quad t \in [0, 1], k = 1, \dots, K, \quad (2.14)$$

and define the process \bar{S}_t^n analogously. Let \mathcal{C} be the space of continuous functions $[0, 1] \mapsto \mathbb{R}^{2K}$ equipped with the uniform metric: $d(x, y) = \sup_{t \in [0, 1]} |x(t) - y(t)|$, $x, y \in \mathcal{C}$. We have the following result; the proof is given in Section 6.1.

Theorem 1. *Fix $K \in \mathbb{N}$, $\mu \in \mathbb{R}^K$ and $\sigma \in \mathbb{R}_+^K$. Suppose that we have a sequence of K -armed bandit problems as in Definition 1 whose reward distribution reside in the diffusion regime as per Definition 8, and have a convergent sequence of sequentially randomized Markov experiments following Definitions 2 and 9. Suppose furthermore that Assumption 1 holds, and $(\bar{Q}_0^n, \bar{S}_0^n) = 0$. Then, as $n \rightarrow \infty$, $(\bar{Q}_t^n, \bar{S}_t^n)_{t \in [0, 1]}$ converges weakly to $(Q_t, S_t)_{t \in [0, 1]} \in \mathcal{C}$, which is the unique solution to the following stochastic differential equation over $t \in [0, 1]$: for $k = 1, \dots, K$*

$$\begin{aligned} dQ_{k,t} &= \psi_k(Q_t, S_t) dt, \\ dS_{k,t} &= \psi_k(Q_t, S_t) \mu_k dt + \sqrt{\psi_k(Q_t, S_t)} \sigma_k dB_{k,t}, \end{aligned} \quad (2.15)$$

where $(Q_0, S_0) = 0$, and B_t is a standard Brownian motion in \mathbb{R}^K . Furthermore, for any bounded continuous function $f : \mathbb{R}^{2K} \mapsto \mathbb{R}$,

$$\lim_{n \rightarrow \infty} \mathbb{E} [f(\bar{Q}_t^n, \bar{S}_t^n)] = \mathbb{E} [f(Q_t, S_t)], \quad \forall t \in [0, 1]. \quad (2.16)$$

As an immediate corollary of Theorem 1, we obtain the following characterization of the finite-sample expected regret; the proof follows simply by setting $f(Q, S) := (\max_k \mu_k) - \langle Q, \mu \rangle$.

Corollary 2 (Convergence of Expected Regret). *Fix $K \in \mathbb{N}$, $\mu \in \mathbb{R}^K$ and $\sigma \in \mathbb{R}_+^K$. Suppose that Assumption 1 holds. Define*

$$R = (\max_k \mu_k) - \langle Q_1, \mu \rangle, \quad (2.17)$$

where $\{Q_t\}_{t \in [0, 1]}$ is given by the solution to (2.15). Then,

$$\mathbb{E} [R^n] = \mathbb{E} [R] \sqrt{n} + o(\sqrt{n}), \quad (2.18)$$

i.e., $\lim_{n \rightarrow \infty} \mathbb{E} [R^n] / \sqrt{n} = \mathbb{E} [R]$.

Finally, the following theorem gives a more compact representation of the stochastic differential equations in Theorem 1, showing that they can be written as a set of ordinary differential equations driven by a Brownian motion with a random time change, $t \Rightarrow Q_t$. The result will be useful, for instance, in our subsequent analysis of Thompson sampling in the super-diffusive regime. The proof is given in Section 6.2.

Theorem 3. *The limit stochastic differential equation in (2.15) can be equivalently written as*

$$dQ_{k,t} = \psi_k(Q_t \mu + \sigma W_{Q_t}, Q_t) dt, \quad k = 1, \dots, K, \quad (2.19)$$

with $Q_0 = 0$, where W is a K -dimensional standard Brownian motion. Here, $Q_t \mu$ and σW_{Q_t} are understood to be vectors of element-wise products, with $Q_t \mu = (Q_{k,t} \mu_k)_{k=1, \dots, K}$, and $\sigma W_{Q_t} = (\sigma_k W_{k, Q_{k,t}})_{k=1, \dots, K}$. In particular, we may also represent S_t explicitly as a function of Q and W :

$$S_{k,t} = Q_{k,t} \mu_k + \sigma_k W_{k, Q_{k,t}}, \quad k = 1, \dots, K, \quad t \in [0, 1]. \quad (2.20)$$

All proofs are deferred to Section 6. Our proof of Theorem 1 uses the Stroock-Varadhan program which is in turn based on the martingale characterization of diffusion [Durrett, 1996, Stroock and Varadhan, 2007]. The main technique is based on showing that the generator of the Markov chain associated with the sequentially randomized Markov experiment converges, in an appropriate sense, to the generator of the desired limit diffusion process. Meanwhile, Theorem 3 builds upon the convergence result in Theorem 1. The key additional step is to use the Skorohod’s representation theorem so as to allow us to couple all relevant sample paths, including, S , Q and a noise process U_t under a random time-change $U_t \rightarrow U_{Q_t}$, to a single driving Brownian motion, and show that they converge to the appropriate limits jointly.

3 Diffusion Analysis of Thompson Sampling

In the second part of this paper we use the diffusion limit derived above to give an in-depth analysis of the large-sample behavior of Thompson sampling, which is a successful and versatile approach to sequential experiments that is widely used in practice [e.g., Chapelle and Li, 2011, Ferreira et al., 2018, Hadad et al., 2021, Russo et al., 2018]. In doing so, we focus on Thompson sampling in the one- and two-armed bandit problems (i.e., with $K = 1$ or 2). In the one-armed bandit, an agent compares an arm with unknown mean reward to a deterministic outside option, while in the two-armed bandit, two arms with unknown mean rewards are compared to each other. The reason we focus on these settings is that they are simple to understand and have sampling functions that allow for closed-form expressions, yet still reveal a number of fundamental insights about Thompson sampling—notably around good choices of regularization and around stability.

We start our discussion in Sections 3.1 and 3.2 by specifying the variants of Thompson sampling studied, and by discussing prior choices that enable us to formally derive a diffusion limit. In some cases, the validity of this diffusion limit follows directly from Theorem 1; in others, however, the limiting sampling function is not Lipschitz and so a more delicate argument is required. We give formal results justifying these diffusion limits in Section 3.6.

Next, we carry out in Sections 3.3 and 3.4 an in-depth study of the regret of Thompson sampling using the diffusion limit. In particular, Section 3.4 contains our main theoretical results in this part of the paper, demonstrating a sharp performance separation between smooth and undersmoothed Thompson sampling, and showing that the latter achieves near-optimal, instance-dependent regret scaling when the arm gap is relative large. Finally, in Section 3.5 we use the diffusion limit to study distributional properties of undersmoothed Thompson sampling, and find striking instability properties.

Throughout the section we will focus on the notion of a limit regret, defined in (2.17):

$$R = (\max_k \mu_k) - \langle Q_1, \mu \rangle. \quad (3.1)$$

Theorem 1 immediately implies that the large-sample behavior of regret is captured by the limiting stochastic differential equation under our diffusion asymptotic setting, and that appropriately scaled regret converges in distribution to R . Thus, given access to the distribution of the final state Q_1 in our diffusion limit, we also get access to the distribution of regret.

3.1 One-Armed Thompson Sampling

Consider the following one-armed sequential experiment. In periods $i = 1, \dots, n$, an agent has an option to draw from a distribution P^n with (unknown) mean μ^n and (known) variance $(\sigma^n)^2$, or do nothing and receive zero reward. As discussed above, we are interested in the regime where $n \rightarrow \infty$, and $\lim_{n \rightarrow \infty} \sqrt{n} \mu^n = \mu$ for some fixed $\mu \in \mathbb{R}$ while $\lim_{n \rightarrow \infty} \sigma^n = \sigma$ remains constant. Following the paradigm of Thompson sampling, we study an algorithm where the agent starts with a prior belief distribution G_0^n on P^n . Then, at each step i , the agent draws a new sample with probability $\pi_i^n = \mathbb{P}_{G_{i-1}^n}(\mu^n > 0)$, and uses any so-gathered information to update the posterior G_i^n using Bayes' rule. Furthermore, we assume that the agent takes P^n to be a Gaussian distribution with (unknown) mean μ^n and (known) variance σ^2 , and sets G_0^n to be a Gaussian prior on μ^n with mean 0. Thus, writing I_i for the event that we draw a sample in the i -th period and Y_i for the observed outcome, we get

$$\begin{aligned} \mu^n | G_i^n &\sim \mathcal{N} \left(\frac{\sigma^{-2} \sum_{j=1}^i I_j Y_j}{\sigma^{-2} \sum_{j=1}^i I_j + (\nu^n)^{-2}}, \frac{1}{\sigma^{-2} \sum_{j=1}^i I_j + (\nu^n)^{-2}} \right), \\ \pi_i &= \Phi \left(\frac{\sigma^{-2} \sum_{j=1}^i I_j Y_j}{\sqrt{\sigma^{-2} \sum_{j=1}^i I_j + (\nu^n)^{-2}}} \right), \end{aligned} \quad (3.2)$$

where $(\nu^n)^2$ is the prior variance and Φ is the standard Gaussian cumulative distribution function. Qualitatively, one can motivate this sampling scheme by considering an agent gambling at a slot machine: Here, μ^n represents the expected reward from playing, and the agent's propensity to play depends on the strength of their belief that this expected reward is positive.

The one-armed Thompson sampling algorithm defined above is clearly a sequentially randomized Markov experiment; thus, given our general results, we expect it to admit a diffusion limit when given a sequence of problems in the diffusion regime—provided the underlying sampling functions are convergent. In our case, the key remaining question is how we should scale the prior variance ν^n used in the Thompson sampling heuristic. A first option that leads to a convergent sampling function is to choose ν^n such that

$$\lim_{n \rightarrow \infty} (\nu^n)^{-2}/n = c > 0, \quad (3.3)$$

in which case (2.11) suggests that we should expect scaled sample paths of S^n and Q^n to converge weakly to

$$dQ_t = \pi_t dt, \quad dS_t = \mu \pi_t dt + \sqrt{\pi_t} dB_t, \quad \pi_t = \Phi \left(\frac{S_t}{\sigma \sqrt{Q_t + \sigma^2 c}} \right), \quad (3.4)$$

with $S_0 = Q_0 = 0$. Here, the corresponding sampling function is Lipschitz-continuous, meaning that the above diffusion limit in fact follows immediately from Theorem 1. We refer to Thompson sampling with this scaling of the prior variance as *smoothed* Thompson sampling.

Alternatively, one could also consider a setting where $\nu^n = \nu > 0$ does not scale with n , or where ν^n decays slowly enough that:

$$c = \lim_{n \rightarrow \infty} (\nu^n)^{-2}/n = 0 \quad (3.5)$$

This is the scaling of Thompson sampling that is most commonly considered in practice; for example, ν^n is simply set to 1 in [Agrawal and Goyal, 2017]. We refer to Thompson sampling with this type of prior variance scaling as *undersmoothed*. As before, (2.11) suggests weak convergence to

$$dQ_t = \pi_t dt, \quad dS_t = \mu \pi_t dt + \sqrt{\pi_t} dB_t, \quad \pi_t = \Phi\left(\frac{S_t}{\sigma \sqrt{Q_t}}\right), \quad (3.6)$$

with $S_0 = Q_0 = 0$. In this case, however, the sampling function is no longer Lipschitz-continuous in Q , and so validity of the limit (3.6) no longer follows immediately from Theorem 1. Rather, we need to rely on some further analysis, which we defer to Section 3.6.1 (Theorem 8). For now we simply note, as seen in Figure 1, the diffusion approximations given above appear to be numerically accurate in finite samples, both in the smoothed and undersmoothed regimes.

3.2 Translation-Invariant Two-Armed Thompson Sampling

Next, we consider Thompson sampling in the two-armed setting, where both arms' mean rewards are unknown. In periods $i = 1, \dots, n$, an agent chooses which of two distributions P_1^n or P_2^n to draw from, each with (unknown) mean μ_k^n and (known) variance $(\sigma_k^n)^2$. To streamline notation, we will present the analysis for the case where $\sigma_1^n = \sigma_2^n = \sigma^n$, with the understanding that all results stated in this section will generalize to the case with heterogeneous reward variances in a straightforward manner.

In a finite-horizon, pre-limit system, the agent uses the following version of translation-invariant Thompson sampling based on reasoning about the posterior distribution of the arm difference $\delta^n = \mu_1^n - \mu_2^n$. The agent starts with one draw from each arm, and subsequently pulls arm 1 in period i with probability:

$$\pi_i = \Phi\left(\frac{\alpha_i^{-2} \Delta_i}{\sqrt{\alpha_i^{-2} + (\nu^n)^{-2}}}\right), \quad \text{with} \quad \Delta_i = \frac{S_{1,i}}{Q_{1,i}} - \frac{S_{2,i}}{Q_{2,i}}, \quad \alpha_i^2 = \frac{\sigma^2 i}{Q_{1,i} Q_{2,i}}, \quad (3.7)$$

where ν^n is interpreted as the prior standard deviation for δ^n , Δ_i the empirical mean of δ^n , and α_i^2 the variance associated with the noisy realizations of rewards. Here, we note that $Q_{1,i} + Q_{2,i} = i$.

As usual, we focus on the behavior of (3.7) in the diffusion regime as in Definition 8. However, because this algorithm is translation invariant, we can always without loss of generality assume that $\mu_2^n = 0$ when studying its behavior. In that case—as discussed in Remark 3—our results apply as long as $\lim_{n \rightarrow \infty} \sqrt{n}(\mu_1^n - \mu_2^n) = \delta$ and $(\sigma_k^n)^2 = \sigma^2$, even if the mean arm rewards μ_k^n themselves may not converge to 0. When stating results below, we assume that $\mu_2^n = 0$, in which case $\mu_1^n = \delta^n$.

Now, to obtain a diffusion limit, we again need to choose a scaling for the prior standard deviation, ν^n , that yields a convergent sequence of sampling functions. As above, one option would be to use non-vanishing smoothing, $\lim_{n \rightarrow \infty} n(\nu^n)^{-2} = c > 0$, in which case we would expect weak convergence to the solution of the following SDE:

$$\begin{aligned} dS_{1,t} &= \delta \pi_t dt + \sqrt{\pi_t} \sigma dB_{1,t}, & dS_{2,t} &= \sqrt{1 - \pi_t} \sigma dB_{2,t}, \\ dQ_{1,t} &= \pi_t dt, & \pi_t &= \Phi\left(\frac{\sigma^{-2} Q_{1,t}(t - Q_{1,t})(S_{1,t}/Q_{1,t} - S_{2,t}/(t - Q_{1,t}))}{\sqrt{\sigma^{-2} t Q_{1,t}(t - Q_{1,t}) + t^2 c}}\right) \\ Q_{2,t} &= t - Q_{1,t}, \end{aligned} \quad (3.8)$$

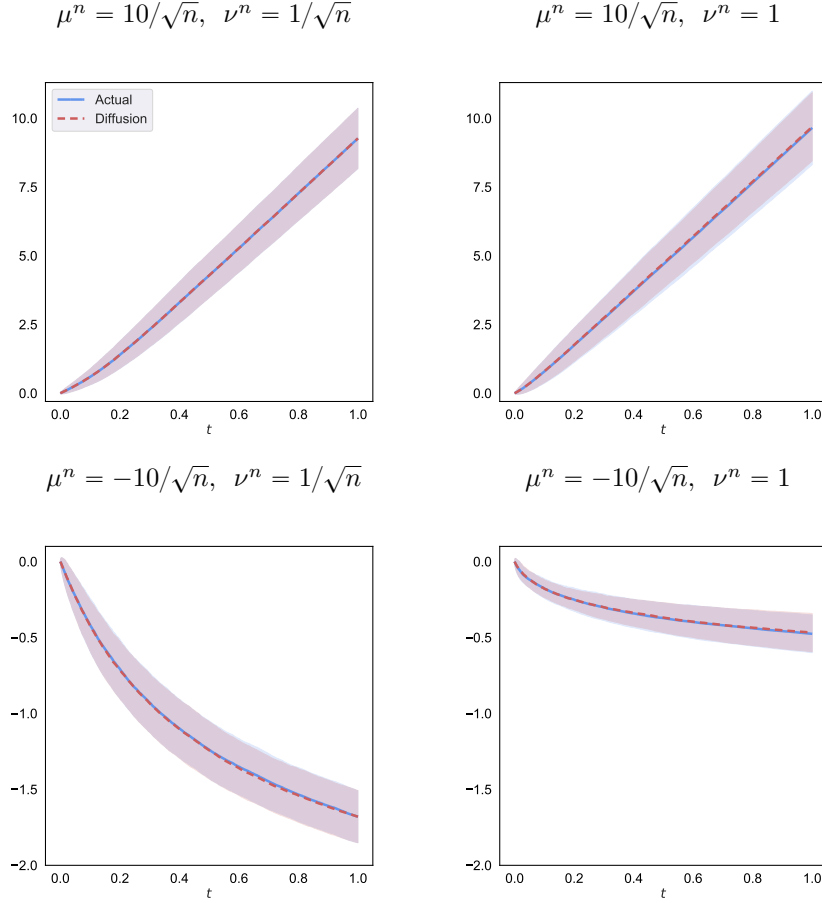


Figure 1: Convergence to the diffusion limit under one-armed Thompson sampling. The plots show the evolution of the scaled cumulative reward S_{nt}^n , and its limiting process S_t , with $n = 1500$ over 1000 simulation runs. The lines and shades represent the empirical mean and one empirical standard deviation from the mean, respectively. The two columns correspond to the smoothed (left) and undersmoothed (right) Thompson sampling, respectively.

with $S_{k,0} = Q_{k,0} = 0$. Meanwhile, an undersmoothed choice, $\lim_{n \rightarrow \infty} n(\nu^n)^{-2} = 0$, also yields convergent sampling functions, and suggests convergence to (3.8) with $c = 0$.

Justifying the diffusion limit for two-armed Thompson sampling is more delicate than in the one-armed setting. Even when $c > 0$, the sampling function associated with (3.8) is not Lipschitz-continuous with respect to (Q_t, S_t) near $t = 0$, so Theorem 1 cannot be evoked directly to show convergence. Fortunately, convergence to solutions of the SDE in (3.8) can be established using Theorem 1 by considering a version of (3.7) with an additional time-varying smoothing. Formal results justifying this limit are deferred to Section 3.6.2.

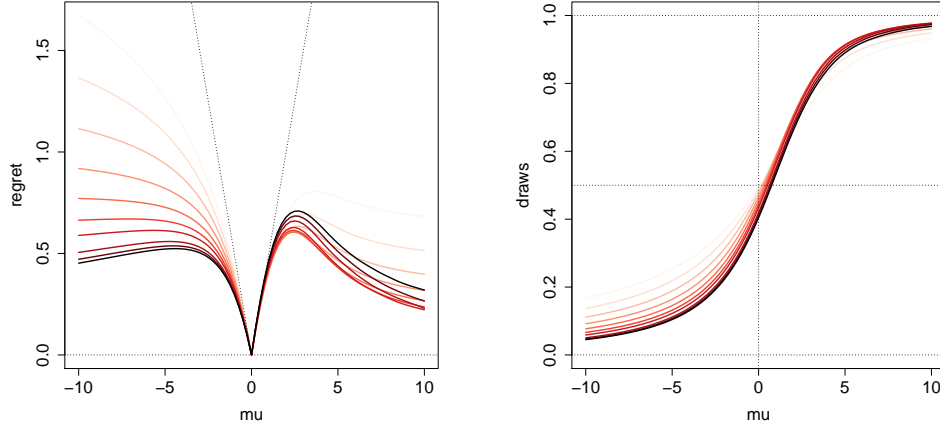


Figure 2: Regret profile under one-armed Thompson sampling, for $c = 1, 1/2, 1/4, 1/8, 1/16, 1/32, 1/64, 1/256, 1/1024$, and finally $c = 0$, with variance $\sigma^2 = 1$, as a function of the scaled mean reward of the unknown arm. The curves with positive values are shown in hues of red with darker-colored hues corresponding to smaller values of c , while $c = 0$ is shown in black.

3.3 Regret Profiles

Given our diffusion limits for one- and two-armed Thompson sampling derived above, we are now in a position to use them to study the large-sample behavior of the method. We start by using this limit to provide an exact, instance-specific characterization of the mean scaled regret of Thompson sampling, i.e., the limit of $\sqrt{n}\mathbb{E}[R^n]$. This exercise gives us a sharp picture of how the difficulty of Thompson sampling varies with signal strength, and also helps us understand the effect of the prior choice ν^n on performance.

In Figure 2, we plot both expected (scaled) regret and the mean (scaled) number of draws on the unknown arm $\mathbb{E}[Q_1]$ as a function of μ , across several choices of smoothing parameter c (throughout, we keep $\sigma^2 = 1$). This immediately highlights several properties of Thompson sampling; some well known, and others harder to see without our diffusion-based approach.

First, we see that when $\mu = 0$ we have $\mathbb{E}[Q_1] < 0.5$, meaning that bandits are biased towards being pessimistic about the value of an uncertain arm. This is in line with well established results about the bias of bandit algorithms [Nie et al., 2018, Shin et al., 2019]. Second, we see that these regret profiles are strikingly asymmetric: In general, getting low regret when $\mu > 0$ appears easier than when $\mu < 0$. This again matches what one might expect: When $\mu < 0$, there is a tension between learning more about the data-generating distribution (which requires pulling the arm), and controlling regret (which requires not pulling the arm), resulting in a tension between exploration and exploitation. In contrast, when $\mu > 0$, pulling the arm is optimal both for learning and for regret, and so as soon as the bandit acquires a belief that $\mu > 0$ they will pull the arm more and more frequently—thus reinforcing this belief.

Third, Figure 2 highlights an intriguing relationship between the regularization parameter c and regret. As predicted in Theorem 7, we see that regret in fact converges as $c \rightarrow 0$.

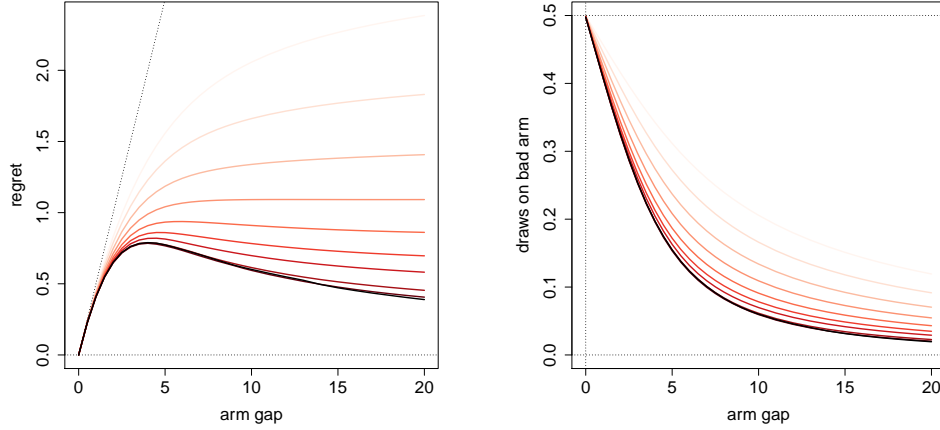


Figure 3: Regret profile for two-armed Thompson sampling, for $c = 1, 1/2, 1/4, 1/8, 1/16, 1/32, 1/64, 1/256, 1/1024$, and finally $c = 0$. We use $\sigma^2 = 1$ throughout. The left panel shows expected regret, while the right panel shows $\mathbb{E}[Q_1]$. The curves with positive values are shown in hues of red with darker-colored hues corresponding to smaller values of c , while $c = 0$ is shown in black. The algorithm (3.7) is translation-invariant and symmetric in its treatment of the arms, so regret only depends on the scaled arm gap $\delta = |\mu_1 - \mu_2|$.

What’s particularly interesting, and perhaps more surprising, is that setting $c = 0$ is very close to being optimal regardless of the true value of μ . When $\mu < 0$, any deviations from instance-wise optimality are not visible given the resolution of the plot, whereas for some values of $\mu > 0$ the choice $c = 0$ is sub-optimal but not by much.

Furthermore, the Bayesian heuristic behind Thompson sampling appears to be mostly uninformative about which choices of c will perform well in terms of regret. For example, when $\mu = -2$, one might expect a choice $c = 1/4$ to be well justified, as $c = 1/4$ arises by choosing a prior standard deviation of $\nu = 2$, in line with the effect size. But this is in fact a remarkably poor choice here: By setting $c = 0$ we achieve expected scaled regret of 0.44, but with $c = 1/4$ this number climbs by 41% up to 0.62. In other words, while Bayesian heuristics may be helpful in qualitatively motivating Thompson sampling, our diffusion-based analysis gives a much more precise understanding of the instance-based behavior of the method.

Figure 3 shows the limiting mean scaled regret of two-armed Thompson sampling in the diffusion limit, with $\sigma^2 = 1$. At a high level, the qualitative implications closely mirror those from the one-armed bandit as reported in Figure 2. The behavior of Thompson sampling converges as $c \rightarrow 0$, and its regret properties with $c = 0$ are in general very strong. If anything, the $c = 0$ choice is now even more desirable than before: With the one-armed case, this choice was modestly but perceptibly dominated by other choices of c for some values of $\mu > 0$, but here $c = 0$ is effectively optimal across all δ to within the resolution displayed in Figure 3.

Another interesting observation from Figure 3 is that, in the undersmoothed regime, the regret is maximized around $\delta = 4$. We note that, in a randomized trial with $\pi_t = 0.5$ throughout, $\delta = 4$ corresponds to an effect size that is twice the standard deviation of its

difference-in-means estimator. In other words, $\delta = 4$ is an effect size that can be reliably detected using a randomized trial run on all samples, but that would be difficult to detect using just a fraction of the data. The fact that regret is maximized around $\delta = 4$ is consistent with an intuition that the hardest problems for bandit algorithms are those with effects we can detect—but just barely.

3.4 The Super-Diffusive Limit

Above, we computed the mean scaled regret for Thompson sampling in the diffusion limit, for effect sizes $\mu^n \approx \mu/\sqrt{n}$ for moderate values of μ (e.g., $-10 \leq \mu \leq 10$), and found that undersmoothing (i.e., using $c = 0$) is overall a robust and versatile choice for getting good instance-specific regret. Given this observation, it is natural to ask: Does the good behavior of undersmoothing persist in problems with relatively strong effect sizes (i.e., with $\mu^n \approx \mu/\sqrt{n}$ with $|\mu| \gg 1$), or is it specific to the range of effect sizes considered in Figures 2 and 3?

To get a deeper understanding of this phenomenon, we now pursue a more formal analysis of the interplay between c and the limit regret when the arm gap is large. Specifically, we study the regret scaling of one- and two-armed Thompson sampling in what we refer to as the super-diffusive regime: We first take the diffusion limit as $n \rightarrow \infty$ for a fixed μ , and subsequently look at how the resulting limiting process behaves in the limit as the arm gap tends to infinity. In other words, this is a regime of diffusion processes where the magnitude of difference in mean rewards between the two arms is relatively large.

Overall, our results not only corroborate our above observations on the robustness of undersmoothing, but in fact suggest that undersmoothing is the only robust choice of regularization in problems where effect sizes may be relatively large. We find that there is a sharp separation in the regret performance of Thompson sampling in the super-diffusive regime whereby the regret of undersmoothed Thompson sampling ($c = 0$) gets lower and lower as effects get large, whereas the regret of smoothed Thompson sampling ($c > 0$) does not decay to zero in the super-diffusive limit, and may even diverge. As revealed in the proof, the reason smoothed Thompson sampling does poorly in the super-diffusive regime is that non-trivial regularization prevents the algorithm from being responsive enough to new information, and especially from shutting down a very bad arm fast enough.

We start with a result in the one-armed case. The proof of this result relies on a delicate study of the diffusion limit. The main difficulty arises from the dynamics of undersmoothed Thompson sampling near $t = 0$. In this regime, the sampling probability is highly sensitive to the empirical mean rewards of the arms, which can oscillate wildly. Our main approach here is rely on the law of iterated logarithm of Brownian motion in order to obtain some form of “regularity” in the behavior of the sampling probabilities near $t = 0$. This regularity property is further combined with a set of carefully chosen events, conditioning and stopping times to arrive at the desirable regret scaling. Below, we use the following notation: As x tends to a limit, we write $f(x) < g(x)$, if for any $\beta \in (0, 1)$, we have $f(x)/g(x)^\beta \rightarrow 0$.

Theorem 4. *Consider the diffusion limit associated with one-armed Thompson sampling given in (3.4), where $\lim_{n \rightarrow \infty} (\nu^n)^{-2}/n = c$.*

1. *If $c > 0$, then, almost surely,*

$$\lim_{\mu \rightarrow -\infty} R = \infty, \quad \text{and} \quad \liminf_{\mu \rightarrow +\infty} R > 0. \quad (3.9)$$

2. If $c = 0$, then, almost surely,

$$R < 1/|\mu|, \quad \text{as } |\mu| \rightarrow \infty. \quad (3.10)$$

Here, R is the limit regret defined in (2.17). The almost-sure statement is with respect to the probability measure associated with the driving Brownian motion $\{W_t\}_{t \in [0,1]}$ as in Theorem 3.

Meanwhile, the following result shows that the desirable regret scaling of undersmoothed Thompson sampling also extend to the two-armed setting. In this case, the analysis is substantially more complex due to the fact that both arms' mean rewards are uncertain. The key idea behind the proof of Theorem 5 is to leverage a novel approximate reduction from the two-armed dynamics into two separate one-armed problems; the proof is given in Section 6.4, which also includes a high-level description of the strategy.

Theorem 5. *Consider the limit regret R associated with undersmoothed two-armed Thompson sampling with $c = 0$. Then, almost surely,*

$$R < 1/\delta, \quad \text{as } \delta \rightarrow \infty, \quad (3.11)$$

where δ is the (scaled) arm gap.

Finally, we note that the above not only establish the robustness of undersmoothed Thompson sampling across a wide variety of effect sizes; they also provide a quantitative characterization of the regret profile when effect sizes are strong. In particular, in the two-armed case, we find that regret decays faster than $1/\delta^{1-\epsilon}$ for any $\epsilon > 0$ as $\delta \rightarrow \infty$ in the super-diffusive regime. Encouragingly, this scaling nearly matches a known instance-dependent regret lower bounds in the frequentist stochastic bandit literature. Mannor and Tsitsiklis [2004] show that, for any bandit algorithm, there exists an instance with arm gap Δ under which expected regret is at least

$$CK \frac{1}{\Delta} \log \left(\frac{\Delta^2 n}{K} \right), \quad (3.12)$$

where Δ the mean reward gap between best and second-best arms, and C is a universal constant. Applying this result to the two-armed setting, with $\Delta = \delta/\sqrt{n}$ and $K = 2$, (3.12) would suggest that

$$\frac{\mathbb{E}[R^n]}{\sqrt{n}} \geq 2C \frac{1}{\delta} \log \left(\frac{n(\delta/\sqrt{n})^2}{2} \right) = 2C \frac{\log(|\delta|/2)}{\delta}. \quad (3.13)$$

Comparing the above with (3.10) in Theorem 4 and (3.11) in Theorem 5 shows that the regret scaling of undersmoothed Thompson sampling matches this lower bound up to an arbitrarily small polynomial factor.

To the best of our knowledge, this is the first formal result suggesting that Thompson sampling achieves anything close to instance-optimal behavior as the effect size grows large; see also discussion in Section 4. Furthermore, it is interesting to note that algorithms known to attain regret upper bounds on the order of (3.12) tend to rely on substantially more sophisticated mechanisms, such as adaptive arm elimination and time-dependent confidence intervals [Auer and Ortner, 2010]. It is thus both surprising and encouraging that such a simple and easily implementable heuristic as Thompson sampling should achieve near-optimal instance-dependent regret. We are hopeful that similar insights can be generalized to Thompson sampling applied to general K -armed bandits for $K > 2$.

Remark 4 (Expected Regret in the Super Diffusive Regime). We note that the regret characterizations in Theorems 4 and 5 are given in an almost-sure sense (with respect to the measure associated with the driving Brownian motion), and it is natural to ask whether analogous statements can be established for expected regret $\mathbb{E}[R]$ as well. Since regret is always non-negative, almost-sure regret lower bounds immediately extend to expected regret. Specifically, it follows from part 1 of Theorem 4 that for any $c > 0$, expected limit regret under one-armed Thompson sampling satisfies

$$\lim_{\mu \rightarrow -\infty} \mathbb{E}[R] = \infty, \quad \text{and} \quad \liminf_{\mu \rightarrow +\infty} \mathbb{E}[R] > 0. \quad (3.14)$$

Unfortunately, almost-sure regret upper bounds, on the other hand, do not extend immediately. This is because R can be as large as μ in the worst case, which diverges as $|\mu| \rightarrow \infty$, and as such we do not have an easy tightness property to rely on in order to extend the almost-sure guarantees to expected regret. Showing the same super diffusive upper bound holds for expected regret is an important question for further work.

3.5 The (In)stability of Undersmoothed Thompson Sampling

The diffusion limit also allow us to conduct refined performance analysis that goes beyond mean rewards. In this section, we use the diffusion characterization to unearth some interesting distributional properties of Thompson sampling.

Both our numerical results and the theoretical analysis in the proceeding sections point to the fact that undersmoothed Thompson sampling ($c = 0$) yields far superior total regret than its smooth counterpart. This performance improvement from undersmoothing, however, does not come for free. Although undersmoothed Thompson sampling identifies and focuses on the correct best action often enough to achieve low average regret, it is also liable to fail completely and double down on a bad arm.

As a first lens on the instability of Thompson sampling, Figure 4 displays the distribution of regret for undersmoothed two-armed Thompson sampling for a variety of arm gaps δ . Interestingly, we see that for all considered values of δ , the distribution of regret is noticeably not unimodal. Rather, there is a primary mode corresponding to the bulk of realizations where Thompson sampling gets reasonably low regret, but there is also a second mode near $R \approx \delta$. Recall that, if $\mu_1 > \mu_2$, then regret measures the frequency of draws on the second arm $R = |\delta| Q_{2,1}$, and in particular $|R| \leq |\delta|$ almost surely. Thus realizations of Thompson sampling with $R \approx \delta$ correspond to cases where the algorithm almost immediately settled on the bad arm, and never really even gave the good arm a chance.

To get a more formal handle on the instability of undersmoothed Thompson sampling, it is helpful to consider the evolution of the sampling probabilities π_t over time. Qualitatively, these sampling probabilities π_t correspond to the subjective time- t beliefs about which arms are best, as held by the “agent” running the algorithm. In this case, we find that undersmoothed one-armed Thompson sampling will always lead the algorithm be convinced of the “wrong” reality with arbitrarily high confidence at some during the sampling process, no matter how large the magnitude of the actual mean reward. Formally, in the one-armed case, we show the following.

Theorem 6. *Consider the sampling probability $\pi_t = \Phi(S_t/(\sigma\sqrt{Q_t}))$ associated with the diffusion limit for one-arm undersmoothed Thompson sampling ($c = 0$). Then, fixing any*

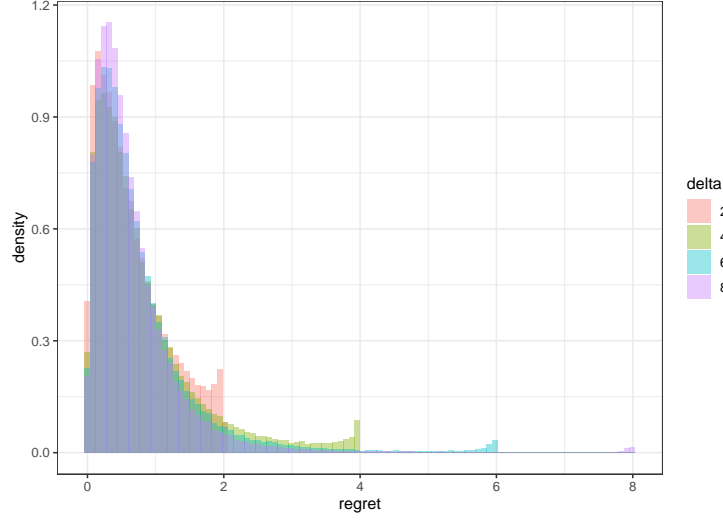


Figure 4: Distribution of the (scaled) regret for two-armed Thompson sampling in the undersmoothed regime (i.e., with $c = 0$), as a function of (scaled) arm gap δ . The histograms are aggregated over 100,000 realization of the limiting stochastic differential equation.

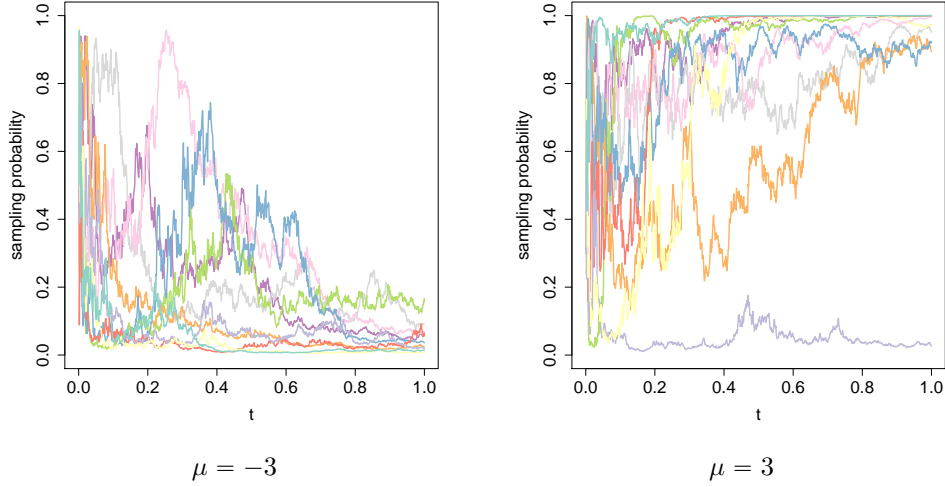


Figure 5: Sample paths of the sampling probability π_t in one-armed Thompson sampling as defined in (3.4), in the undersmoothed regime (i.e., with $c = 0$).

effect size $\mu \in \mathbb{R}$ and confidence level $\eta \in (0, 1)$, we have, for all $\epsilon \in (0, 1)$,

$$\mathbb{P} \left(\sup_{t \in [0, \epsilon]} \pi_t \geq 1 - \eta \right) = \mathbb{P} \left(\inf_{t \in [0, \epsilon]} \pi_t \leq \eta \right) = 1. \quad (3.15)$$

In other words, in the undersmoothed limit, Thompson sampling will almost always

at some early point in time be arbitrarily convinced about μ having the wrong sign; and this holds no matter how large $|\mu|$ really is. However, Thompson sampling will eventually recover, thus achieving low regret. We further illustrate sample paths in the case with $c = 0$ in Figure 5. At the very least, this finding again challenges a perspective that would take Thompson sampling literally as a principled Bayesian algorithm (since in this case we'd expect belief distributions to follow martingale updates), and instead highlights that Thompson sampling has subtle and unexpected behaviors that can only be elucidated via dedicated methods.

The key idea in the proof of Theorem 6 is to use the time-changed form of the diffusion limit given in Theorem 3, which gives us a characterization

$$Q_0 = 0, \quad dQ_t = \pi_t dt, \quad \pi_t = \Phi \left(\frac{Q_t \mu + W_{Q_t}}{\sigma \sqrt{Q_t}} \right), \quad (3.16)$$

where W_q is a standard Brownian motion. By the law of iterated logarithm, we know that W_q/\sqrt{q} will get arbitrarily close to $\pm\infty$ as $q \downarrow 0$, which we can use to show that π_t will in turn spend time arbitrarily close to 0 and 1 as $t \downarrow 0$. The proof is given in Section 6.5. A similar result also holds for the 2-armed case, although its proof is not as immediate so we do not provide it here.

3.6 Justifying the Diffusion Limits for Thompson Sampling

We now return to giving a formal justification to the diffusion limits employed in Sections 3.1 and 3.2. Recall that with the exception of smoothed one-armed Thompson sampling, the sampling function is not Lipschitz for two-armed sampling (smoothed or undersmoothed) and undersmoothed one-armed Thompson sampling. As a result, Theorem 1 cannot be used directly in showing convergence to a diffusion limit. Nevertheless, we will show that with additional analysis, we can still use Theorem 1 to rigorously justify the diffusion limits analyzed in Sections 3.1 and 3.2 as limits of suitably convergent pre-limit sample paths.

We begin with the equivalent ODE characterization of the diffusion limit for one- and two-armed Thompson sampling using the random-time change in Theorem 3. These ODEs will be used repeatedly in the sequel. Fix $c \geq 0$. For one-armed Thompson sampling, the following expresses the SDE in (3.4) in an ODE form:

$$dQ_t = \Phi \left(\frac{Q_t \mu + W_{Q_t}}{\sigma \sqrt{Q_t + \sigma^2 c}} \right) dt, \quad S_t = \mu Q_t + \sigma W_{Q_t}, \quad (3.17)$$

with $Q_0 = S_0 = 0$ and W is a standard one-dimensional Brownian motion. Likewise, for two-armed Thompson sampling, the following is the ODE corresponding to the SDE in (3.8):

$$\begin{aligned} dQ_{1,t} &= \Phi \left(\frac{\sigma^{-2} (Q_{1,t} Q_{2,t} \delta + Q_{2,t} W_{1,Q_{1,t}} - Q_{1,t} W_{2,Q_{2,t}})}{\sqrt{\sigma^{-2} t Q_{1,t} Q_{2,t} + t^2 c}} \right) dt, \\ dQ_{2,t} &= \Phi \left(\frac{-\sigma^{-2} (Q_{1,t} Q_{2,t} \delta - Q_{2,t} W_{2,Q_{1,t}} + Q_{1,t} W_{1,Q_{2,t}})}{\sqrt{\sigma^{-2} t Q_{1,t} Q_{2,t} + t^2 c}} \right) dt, \\ S_{k,t} &= \mu_k Q_{k,t} + \sigma_k W_{k,Q_{k,t}}, \quad k = 1, 2, \end{aligned} \quad (3.18)$$

with $Q_{\cdot,0} = S_{\cdot,0} = 0$, where $\delta = \mu_1 - \mu_2$ and W is a standard two-dimensional Brownian motion.

3.6.1 Diffusion Limits for Undersmoothed One-Armed Thompson Sampling

First, we look at the case of one-armed undersmoothed Thompson sampling. Note that if we were to directly set $c = 0$ in (3.4), the resulting drift ψ now violates the Lipschitz condition required by Theorem 1. To give the diffusion limit under $c = 0$ a meaningful interpretation, we will show next that, almost surely,

1. The stochastic differential equation (3.4) admits a unique solution at $c = 0$.
2. Pre-limit sample paths of Thompson sampling converge to this limit under a sequence of appropriately scaled prior variances.

For the first objective, we will construct the SDE solution at $c = 0$ by considering the sequence of solutions under a diminishing, but strictly positive, sequence of c . The following theorem shows that the sequence of diffusion processes almost surely admits a unique limit as $c \rightarrow 0$; we provide the proof in Section 6.6. The key step of the proof hinges on establishing that, as $c \rightarrow 0$, the drift term of the diffusion does not exhibit too wild of an oscillation near $t = 0$. This would further allow us to use the dominated convergence theorem and show the soundness of the limit solution.

Theorem 7. *The diffusion limit $(Q_t)_{t \in [0,1]}$ under Thompson sampling converges uniformly to a limit \tilde{Q} as $c \rightarrow 0$ almost surely. Furthermore, \tilde{Q} is a strong solution to the stochastic differential equation:*

$$d\tilde{Q}_t = \Phi \left(\frac{\mu\sqrt{\tilde{Q}_t}}{\sigma} + \frac{W_{\tilde{Q}_t}}{\sigma\sqrt{\tilde{Q}_t}} \right) dt, \quad \tilde{Q}_0 = 0. \quad (3.19)$$

With Theorem 7 in hand, the second objective is relatively simple. The next result combines Theorems 1 and 7 to show that if ν^n scales at an appropriate rate, then the pre-limit sample path of one-armed Thompson sampling converges to the diffusion limit (3.19) with $c = 0$. The claim is proved by taking a triangulation limit across ν^n and n ; the proof is included in Section 6.7.

Theorem 8. *There exists a sequence $(\nu^n)_{n \in \mathbb{N}}$ with $\lim_{n \rightarrow \infty} (\nu^n)^{-2} = 0$, such that almost surely $(\bar{Z}^n)_{t \in [0,1]}$ converges the solution to (3.19) as $n \rightarrow \infty$.*

3.6.2 Diffusion Limits for Two-Armed Thompson Sampling

We now turn to two-armed Thompson sampling as presented in Section 3.2. Recall that the algorithm begins by pulling once each of the two arms, and thereafter samples arm 1 in period i with probability:

$$\pi_i = \Phi \left(\frac{\alpha_i^{-2} \Delta_i}{\sqrt{\alpha_i^{-2} + (\nu^n)^{-2}}} \right), \quad \text{with} \quad \Delta_i = \frac{S_{1,i}}{Q_{1,i}} - \frac{S_{2,i}}{Q_{2,i}}, \quad \alpha_i^2 = \frac{\sigma^2 i}{Q_{1,i} Q_{2,i}}, \quad (3.20)$$

The sampling probability π_i yields, in the diffusion limit, the following drift for Q_1

$$\pi_t = \Phi \left(\frac{\sigma^{-2} Q_{1,t} (t - Q_{1,t}) (S_{1,t}/Q_{1,t} - S_{2,t}/(t - Q_{1,t}))}{\sqrt{\sigma^{-2} t Q_{1,t} (t - Q_{1,t}) + t^2 c}} \right). \quad (3.21)$$

Contrasting this with one-armed Thompson sampling (3.4), we notice that that regularization term c in the drift has become $t^2 c$. This is challenging to work with since the latter results in a drift function that is not Lipschitz-continuous. To remedy this issue and ultimately show convergence to a meaningful diffusion limit, we will introduce some additional tempering in the sampling probability, so that π_i becomes

$$\pi_i = \Phi \left(\frac{\alpha_i^{-2} \Delta_i}{\sqrt{\alpha_i^{-2} + (\nu^n)^{-2} + (\zeta^n)^{-2} (n^2/i^2)}} \right), \quad (3.22)$$

with

$$\lim_{n \rightarrow \infty} (\zeta^n)^{-2}/n = d > 0. \quad (3.23)$$

This version of Thompson sampling thus tempers the sampling probability so that there is more regularization in earlier experiments in a manner that compensates for the lack of data.

The benefit of introducing the tempering term is that now the Thompson sampling algorithm admits a Lipschitz-continuous sampling function, which allows us to evoke Theorem 1 to conclude that the sample paths of Q converges to the diffusion limit given in Theorem 1. The drift for arm 1 is given by:

$$\pi_t = \Phi \left(\frac{\sigma^{-2} Q_{1,t} (t - Q_{1,t}) (S_{1,t}/Q_{1,t} - S_{2,t}/(t - Q_{1,t}))}{\sqrt{\sigma^{-2} t Q_{1,t} (t - Q_{1,t}) + t^2 c + d}} \right). \quad (3.24)$$

We now ready to return to the original diffusion equations (3.7) without tempering. The following theorem shows that the ODE (3.18) admits a solution almost surely whenever $c \geq 0$, though the solution may not be unique; the proof, presented in Section 6.8, is accomplished by considering a sequence of solutions to the ODE with vanishing ζ and evoking the Arzela-Ascoli theorem. This result can be seen as analogous to Theorem 7 for the one-armed setting, albeit lacking the uniqueness. Nevertheless, since the regret characterization in Theorem 5 holds for any solution, it is sufficient for our purpose.

Theorem 9. *Suppose that $\lim_{n \rightarrow \infty} (\nu^n)^{-2}/n = c$ and $\lim_{n \rightarrow \infty} (\zeta^n)^{-2}/n = d$ for $c \geq 0$ and $d > 0$. Consider a sequence $\{d_j\}_{j \in \mathbb{N}}$ where $d_j \downarrow 0$ as $j \rightarrow \infty$. Denote by Q^j a solution with drift given by (3.24) and $d = d_j$. The following holds almost surely:*

1. $\{Q^j\}_{j \in \mathbb{N}}$ is tight in the sense that any subsequence of $\{Q^j\}$ admits a further subsequence that converges to a limit uniformly over $[0, 1]$. We say that Q is a limit function if it is a limit point for one of these convergent sub-sequences.
2. Fix Q to be a limit function. Then, for all $t \in (0, 1]$, Q is differentiable and satisfies the ODE (3.18).

Finally, similar to Theorem 8 we show that pre-limit sample paths converge to a solution of (3.18) under a sequence of appropriately vanishing ζ^n ; the proof is essentially identical to Theorem 8 and therefore omitted.

Theorem 10. *Fix $c \geq 0$. There exist sequences $\{\zeta^n\}_{n \in \mathbb{N}}$ and $\{\nu^n\}_{n \in \mathbb{N}}$, with $\lim_{n \rightarrow \infty} (\nu^n)^{-2}/n = c$, under which $(\bar{Z}^n)_{t \in [0, 1]}$ converges almost surely to one of the limit functions in Theorem 9.*

	Regret (original)	Scaled regret	Reference
UCB	$3\Delta_n + 16 \log(n)/\Delta_n$	$+\infty$	LS20, §7.1
Thompson	$\log(n)\Delta_n/D(\mu_1^n, \mu_2^n)$	$+\infty$	AG17
MOSS	$39\sqrt{2n} + \Delta_n$	55.2	LS20, §9.1
Impr. UCB	$\min \left\{ \Delta_n + 32 \frac{\log((\Delta^n)^2 n)}{\Delta^n} + \frac{96}{\Delta^n}, n\Delta^n \right\}$	$\min \left\{ 64 \frac{\log(\Delta)}{\Delta} + \frac{96}{\Delta}, \Delta \right\}$	AO10
Oracle ETC	$\min \left\{ \frac{4}{\Delta^n} \left(1 + \log \left(\frac{n(\Delta^n)^2}{4} \right)^+ \right), n\Delta^n \right\}$	$\min \left\{ \frac{4}{\Delta} (1 + \log(\Delta^2/4)^+), \Delta \right\}$	LS20, §6.1

Figure 6: Comparison with existing finite-time instance-dependent regret bounds. The column of scaled regret corresponds to what the bound would become under diffusion scaling, where $\mu_k^n = \mu_k/\sqrt{n}$ and $\Delta^n = \Delta/\sqrt{n}$, with $n \rightarrow \infty$. The algorithms under consideration are the upper confidence bound (UCB) algorithm of [Lai and Robbins \[1985\]](#), Thompson sampling, the Minimax Optimal Strategy in the Stochastic case (MOSS) from [Audibert et al. \[2009\]](#), improved UCB by [Auer and Ortner \[2010\]](#), and an oracle explore-then-commit (ETC) baseline that takes uniformly random actions up to a deterministic time chosen using a-priori knowledge of the effect size $|\delta|$, and then commits to the most promising arm for the rest of time. The specific bounds are as reported in [Agrawal and Goyal \[2017, AG17\]](#), [Auer and Ortner \[2010, AO10\]](#) and [Lattimore and Szepesvári \[2020, LS20\]](#).

4 Comparison to Finite-Sample Upper Bounds

One question for further discussion is in understanding how results derived from the diffusion-based approach developed in this paper should be compared to traditional finite-sample guarantees for bandit algorithms. Overall, existing worst-case finite-sample results are not powerful enough to yield sharp distributional results of the type given here, so we did not discuss them in our analysis of Thompson sampling given in Section 3. However, one point where a more direct comparison is possible is in the context of instance-specific mean regret guarantees as discussed in Section 3.3. Although the existing literature does not in general have exact finite-sample and instance-specific results on mean regret, there are a number of instance-specific *upper bounds* that can be compared to our exact diffusion-based results for Thompson sampling.

In Table 6, we collect a number of state-of-the-art finite-sample bounds for two-armed bandits. We report results both in the original finite-sample form of the bound, and the (scaled) limit of the bound under the diffusion scaling from Definition 8. A first interesting finding is that many available regret bounds are in fact vacuous in the diffusion limit, and do not provide any control on regret. For example, while we know from our diffusion-based analysis that Thompson sampling gets bounded (and in fact quite good) regret in the diffusion limit, the strongest available finite-sample regret bound for Thompson sampling, due to [Agrawal and Goyal \[2017\]](#), diverges in the diffusion limit. The only upper bound given in Table 6 that both remains finite in the diffusion limit and has meaningful instance-dependent behavior (i.e., that improves as the scaled arm gap δ gets large) is the bound of [Auer and Ortner \[2010\]](#) for improved UCB. We also note that the oracle explore-then-commit gets good instance-dependent behavior; however, this algorithm relies on a-priori knowledge of the effect size $|\delta|$, and so it is not a feasible baseline.

The next question is, in cases where the bounds from Table 6 are not vacuous in the

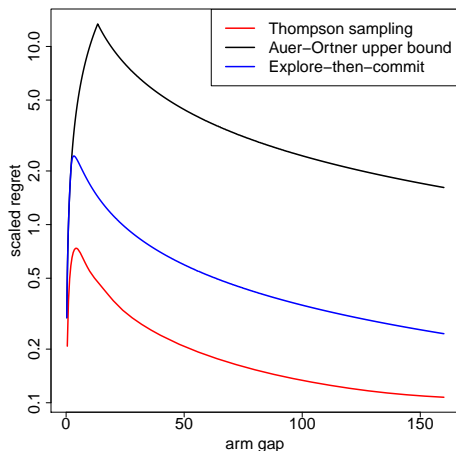


Figure 7: Comparison between the scaled regret in the diffusion limit for undersmoothed Thompson sampling and existing bounds in Table 6, as function of the arm gap. Here, we have a two-armed bandit with $\sigma = 1$. We plot all bounds with a finite scaled regret, with the exception of MOSS, which has a constant value of 55.2 in this case.

diffusion limit, how do they compare with the exact results for Thompson sampling we get in the diffusion limit? Figure 7 compares our regret limit to the relevant upper bounds and reveals that the limiting regret we obtain for Thompson sampling in the diffusion limit is much lower than any of the available upper bounds. Surprisingly, Thompson sampling even outperforms the oracle explore-then-commit algorithm where the duration of exploration is optimized with a-priori knowledge of the effect size and time horizon. These examples seem to suggest that existing finite-sample instance-dependent regret upper bounds could still be improved substantially, possibly by leveraging the diffusion asymptotics advanced in this work.

5 Discussion

In this paper, we introduced an asymptotic regime under which sequentially randomized experiments converge to a diffusion limit. In particular, the limit cumulative reward is obtained by applying a random time change to a constant-drift Brownian motion, where the time change is in turn given by cumulative sampling probabilities (Theorem 3). We then applied this result to derive sharp insights about the behavior of one- and two-armed Thompson sampling.

A first class of natural follow-up questions is in seeing whether diffusion limits hold for broader classes of sequential experiments. For example, can our results be extended to the case of contextual bandits, or to bandit problems with continuous action spaces that arise, e.g., with pricing? Another practical question is whether the approach used here can be used to build confidence intervals using data from sequential experiments, thus adding to the line of work pursued by Hadad et al. [2021], Howard et al. [2018], Zhang et al. [2020], and others.

Further along, a potentially interesting avenue for investigation is whether the diffusion

limit derived here is useful in understanding human—as opposed to algorithmic—learning. Throughout this paper, we have considered Thompson sampling and related algorithms as a class of sequential experiments designed by an investigator, and have discussed how different design choices (e.g., around smoothing) affect the performance of the learning algorithms. However, following [Erev and Roth \[1998\]](#) and [Xu and Yun \[2020\]](#), we could alternatively use sequentially randomized Markov experiments as models for how humans (or human communities) learn over time, and use our results to make qualitative predictions about their behavior.

For example, it may be of interest to examine Thompson sampling as a model for how a scientific community collects and assimilates knowledge about different medical treatments. Qualitatively, this would correspond to a hypothesis that the number of scientists investigating any specific treatment should be proportional to the consensus beliefs that the treatment is best given the available evidence at the time (i.e., that scientists at any given time prioritize investigating the treatments that appear most promising). In this case, our [Theorem 6](#) would provide conditions under which we predict consensus beliefs to first temporarily concentrate around sub-optimal treatments before eventually reaching the truth. More broadly, the many unintuitive phenomena arising from our diffusion limits could yield a number of valuable insights on how we collect and process information.

6 Proofs of Main Results

Terminology: We will use the term “almost all sample paths” throughout the proofs to refer to sample paths of the Brownian motion that belong to a set of probability measure one. We will also use the following asymptotic notation: $f(x) \ll g(x)$ indicates that $f(x)/g(x) \rightarrow 0$, similarly for $f(x) \gg g(x)$.

6.1 Proof of [Theorem 1](#)

Proof. The proof is based on the martingale framework of [Stroock and Varadhan \[2007\]](#). Let us first introduce some notation to streamline the presentation of the proof. Define $Z_t = (\bar{Q}_t, \bar{S}_t)$. Denote by \mathcal{I}_S and \mathcal{I}_Q the indices in Z_t corresponding to the coordinates of S and Q , respectively. Both sets are understood to be an ordered set of K elements, where the subscript is for distinguishing whether the ordering is applied to S versus Q . For $z \in \mathbb{R}_+^K \times \mathbb{R}^K$, define the functions $(b_k)_{k \in \mathcal{I}_Q \cup \mathcal{I}_S}$,

$$b_k(z) = \begin{cases} \psi_k(z), & k \in \mathcal{I}_Q, \\ \psi_k(z)\mu_k, & k \in \mathcal{I}_S. \end{cases} \quad (6.1)$$

For $1 \leq k, l \leq K$, define

$$\eta_{k,l}(z) = \begin{cases} \sqrt{\psi_k(z)}\sigma_k, & \text{if } k = l \in \mathcal{I}_S, \\ 0, & \text{otherwise.} \end{cases} \quad (6.2)$$

Then, the Itô diffusion SDE in [\(2.15\)](#) can be written more compactly as

$$dZ_t = b(Z_t)dt + \eta(Z_t)dB_t, \quad t \in [0, 1], \quad (6.3)$$

with $Z_0 = 0$.

Next, we briefly review the relevant results of the Stroock and Varadhan program. Fix $d \in \mathbb{N}$. Let $(Z_i^n)_{i \in \mathbb{N}}$ be a sequence of time-homogeneous Markov chains taking values in \mathbb{R}^d , indexed by $n \in \mathbb{N}$. Denote by Π^n the transition kernel of Z^n :

$$\Pi^n(z, A) = \mathbb{P}(Z_{i+1}^n \in A \mid Z_i^n = z), \quad z \in \mathbb{R}^d, A \subseteq \mathbb{R}^d. \quad (6.4)$$

Let \bar{Z}_t^n be the piece-wise linear interpolation of Z_{nt}^n :

$$\bar{Z}_t^n = (1 - tn + \lfloor tn \rfloor)Z_{\lfloor tn \rfloor}^n + (tn - \lfloor tn \rfloor)Z_{\lfloor tn \rfloor + 1}^n, \quad t \in [0, 1]. \quad (6.5)$$

Define $K^n(z, A)$ to be the scaled transition kernel:

$$K^n(z, A) = n\Pi^n(z, A). \quad (6.6)$$

Finally, define the functions

$$\begin{aligned} a_{k,l}^n(z) &= \int_{x: |z-x| \leq 1} (x_k - z_k)(x_l - z_l) K^n(z, dx), \\ b_k^n(z) &= \int_{x: |z-x| \leq 1} (x_k - z_k) K^n(z, dx), \\ \Delta_\epsilon^n(z) &= K^n(z, \{x : |x - z| > \epsilon\}). \end{aligned}$$

We will use the following result. A proof of the theorem can be found in [Stroock and Varadhan \[2007, Chapter 11\]](#) or [Durrett \[1996, Chapter 8\]](#). For conditions that ensure the uniqueness and existence of the Itô diffusion (6.10), see [Karatzas and Shreve \[2005, Chapter 5, Theorem 2.9\]](#).

Theorem 11. *Fix d . Let $\{a_{k,l}\}_{1 \leq k,l \leq d}$ and $\{b_k\}_{1 \leq k \leq d}$ be bounded Lipschitz-continuous functions from \mathbb{R}^d to \mathbb{R} . Suppose that for all $k, l \in \{1, \dots, d\}$ and $\epsilon, R > 0$*

$$\lim_{n \rightarrow \infty} \sup_{z: |z| < R} |a_{k,l}^n(z) - a_{k,l}(z)| = 0, \quad (6.7)$$

$$\lim_{n \rightarrow \infty} \sup_{z: |z| < R} |b_k^n(z) - b_k(z)| = 0, \quad (6.8)$$

$$\lim_{n \rightarrow \infty} \sup_{z: |z| < R} \Delta_\epsilon^n(z) = 0. \quad (6.9)$$

If $\bar{Z}_0^n \rightarrow z_0$ as $n \rightarrow \infty$, then $(\bar{Z}_t^n)_{t \in [0,1]}$ converges weakly in \mathcal{C} to the unique solution to the stochastic differential equation

$$dZ_t = b(Z_t)dt + \eta(Z_t)dB_t, \quad (6.10)$$

*where $Z_0 = z_0$ and $\{\eta_{k,l}\}_{1 \leq k,l \leq d}$ are dispersion functions such that $a(z) = \eta(z)\eta^\top(z)$.*¹

We are now ready to prove Theorem 1. We will use the compact representation of the SDE given in (6.3), with $\bar{Z}^n = (\bar{Q}^n, \bar{S}^n)$, $Z_t = (Q_t, S_t)$ and b and η defined as in (6.1) and (6.2), respectively. To prove the convergence of \bar{Z}^n to the suitable diffusion limit, we will

¹The decomposition from a to η is unique only up to rotation. However, the resulting stochastic differential equation is uniquely defined by a . This is because the distribution of the standard Brownian motion is invariant under rotation, and hence any valid decomposition would lead to the same stochastic differential equation.

evoke Theorem 11 (here $d \leftrightarrow 2K$). It suffices to verify the convergence of the corresponding generators in (6.7) through (6.8). To start, the following technical lemma [Durrett, 1996, Section 8.8] will simplify the task of proving convergence by removing the need of truncation in the integral; the proof is given in Appendix B.1. Define

$$\begin{aligned} m_p^n(z) &= \int |x - z|^p K^n(z, dx), & \tilde{a}_{k,l}^n(z) &= \int (x_k - z_k)(x_l - z_l) K^n(z, dx), \\ \tilde{b}_k^n(z) &= \int (x_k - z_k) K^n(z, dx). \end{aligned}$$

Lemma 12. *Fix $p \geq 2$ and suppose that for all $R < \infty$,*

$$\lim_{n \rightarrow \infty} \sup_{z: |z| < R} m_p^n(z) = 0. \quad (6.11)$$

$$\lim_{n \rightarrow \infty} \sup_{z: |z| < R} |\tilde{a}_{k,l}^n(z) - a_{k,l}(z)| = 0, \quad (6.12)$$

$$\lim_{n \rightarrow \infty} \sup_{z: |z| < R} |\tilde{b}_k^n(z) - b_k(z)| = 0, \quad (6.13)$$

Then, the convergence in (6.7) through (6.9) holds.

In what follows, we will use $z = (q, s)$ to denote a specific state of the Markov chain \bar{Z}^n . The transition kernel of the pre-limit chain \bar{Z}^n can be written as

$$\Pi^n((q, s), (q + e_k/n, s + e_k ds/\sqrt{n})) = \bar{\psi}_k^n(q, s) P_k^n(ds), \quad k = 1, \dots, K, \quad (6.14)$$

and zero elsewhere, where $e_k \in \{0, 1\}^K$ is the unit vector where the k th entry is equal to 1 and all other entries are 0, and $\{P_k^n\}_{k=1, \dots, K}$ are the reward probability measures. Define $K^n(z, A) = n\Pi^n(z, A)$.

We next define the limiting functions a and b . The function b is defined as in (6.1):

$$b_k(z) = \begin{cases} \psi_k(z), & k \in \mathcal{I}_Q, \\ \psi_k(z)\mu_k, & k \in \mathcal{I}_S, \end{cases} \quad (6.15)$$

and we let

$$a_{ij}(z) = (\eta\eta^\top)_{k,l}(z) \quad (6.16)$$

where η is defined in (6.2). That is,

$$a_{k,l}(z) = \begin{cases} \psi_k(z)\sigma_k^2, & \text{if } k = l \in \mathcal{I}_S, \\ 0, & \text{otherwise.} \end{cases} \quad (6.17)$$

Fix $R > 0$. We show that the corresponding a^n and b^n converge to the functions a and b defined above, uniformly over the compact set $\{z : |z| \leq R\}$. In light of Lemma 12, it suffices to verify the convergence in (6.11) through (6.13) for $p = 4$. Starting with (6.11),

we have that

$$\begin{aligned}
m_4^n(z) &= \int |z' - z|^4 n \Pi^n(z, dz') \\
&= \sum_{k=1}^K n \bar{\psi}_k^n(z) \int_{w \in \mathbb{R}} \left(\frac{1}{n^2} + \frac{w^2}{n} \right)^2 P_k^n(dw) \\
&\leq \sum_{k=1}^K n \bar{\psi}_k^n(z) \left(\frac{2}{n^4} + \frac{1}{n^2} \int_{w \in \mathbb{R}} w^4 P_k^n(dw) \right) \\
&= \frac{2}{n} + \frac{1}{n} \mathbb{E}_{Z \sim P_k^n} [Z^4] \\
&\xrightarrow{n \rightarrow \infty} 0,
\end{aligned} \tag{6.18}$$

as $n \rightarrow 0$, uniformly over all z , where the last step follows from the assumption that the reward distributions admit bounded fourth moments. This shows (6.11).

For the drift term b , we consider the following two cases; together, they prove (6.13).

Case 1, $k \in \mathcal{I}_Q$. For all $k \in \mathcal{I}_Q$, and $n \in \mathbb{N}$,

$$\begin{aligned}
\tilde{b}_k^n(z) &= \int (q'_k - q_k) K^n(z, dz') \\
&= \frac{1}{n} (n \bar{\psi}_k^n(z)) \\
&\xrightarrow{n \rightarrow \infty} \psi_k(z).
\end{aligned} \tag{6.19}$$

Case 2, $k \in \mathcal{I}_S$. For all $k \in \mathcal{I}_S$,

$$\begin{aligned}
\tilde{b}_k^n(z) &:= \int (s'_k - s_k) K^n(z, dz') \\
&= \bar{\psi}_k^n(z) n \int \frac{w}{\sqrt{n}} P_k^n(dw) \\
&= \bar{\psi}_k^n(z) \mu_k \\
&\xrightarrow{n \rightarrow \infty} b_k(z).
\end{aligned} \tag{6.20}$$

For the variance term a , we consider the following three cases:

Case 1, $k, l \in \mathcal{I}_Q$. Note that under the multi-armed bandit model, only one arm can be chosen at each time step. This means that only one coordinate of \bar{Q}^n can be updated at time, immediately implying that for all n and $k, l \in \mathcal{I}_Q$, $k \neq l$,

$$\tilde{a}_{k,l}^n(z) = \int (q'_k - q_k)(q'_l - q_l) K^n(z, dz') = 0. \tag{6.21}$$

For the case $k = l$, we note that for all $k \in \mathcal{I}_Q$, and all sufficiently large n

$$\tilde{a}_{k,k}^n(z) = \frac{1}{n^2} n \bar{\psi}_k^n(z) \xrightarrow{n \rightarrow \infty} 0. \tag{6.22}$$

Case 2: $k \in \mathcal{I}_Q, l \in \mathcal{I}_S$, or $k \in \mathcal{I}_S$ and $l \in \mathcal{I}_Q$.

$$\begin{aligned}
\tilde{a}_{k,l}^n(z) &= \int (q'_k - q_k)(q'_l - q_l) K^n(z, dz') \\
&= \frac{\bar{\psi}_k^n(z)}{n} \int w(nP_k^n(\sqrt{n}dw)) \\
&= \bar{\psi}_k^n(z) \mathbb{E}_{Z \sim P_k^n} [Z] \\
&= \bar{\psi}_k^n(z) \mu_k / \sqrt{n} \\
&\xrightarrow{n \rightarrow \infty} 0.
\end{aligned} \tag{6.23}$$

Case 3: $k, l \in \mathcal{I}_S$. This case divides into two further sub-cases. Suppose that $k \neq l$. Similar to the logic in Case 1, because only one coordinate of \bar{Q}^n can be updated at a given time step, we have

$$\tilde{a}_{k,l}^n(z) = 0, \quad k \neq l. \tag{6.24}$$

Suppose now that $k = l$. We have

$$\begin{aligned}
\tilde{a}_{k,l}^n(z) &= \int (q'_k - q_k)^2 K^n(z, dz') \\
&= \bar{\psi}_k^n(z) \int w^2(nP_k^n(\sqrt{n}dw)) \\
&= \bar{\psi}_k^n(z) \mathbb{E}_{Z \sim P_k^n} [Z^2] \\
&\xrightarrow{n \rightarrow \infty} \psi_k(z) \sigma_k^2 \\
&= a_{k,l}(z).
\end{aligned} \tag{6.25}$$

We note that due to Assumption 1, the convergence of \tilde{b}^n , \tilde{a}^n and m_p^n to their respective limits holds uniformly over compact sets. We have thus verified the conditions in Lemma 12, further implying (6.7) through (6.8). Note that because ψ_k is bounded and Lipschitz-continuous, so are a and b . This proves the convergence of \bar{Z}^n to the diffusion limit in \mathcal{C} .

Finally, to prove the convergence of $\mathbb{E}[f(\bar{Z}_1^n)]$ to $\mathbb{E}[f(Z_t)]$, note that the weak convergence of \bar{Z}^n in \mathcal{C} implies that the marginal distribution, \bar{Z}_t^n converges weakly to Z_t , as $n \rightarrow \infty$. The result then follows immediately from the continuous mapping theorem and the bounded convergence theorem. This completes the proof of Theorem 1. \square

6.2 Proof of Theorem 3

Proof. It suffices to show that (2.20) holds. We will begin with a slightly different, but equivalent, characterization of the pre-limit bandit dynamics. Consider the n th problem instance. Denote by $\tilde{Y}_{k,j}$ the reward obtained from the j th pull of arm k . Then, we have that for a fixed k , $\tilde{Y}_{k,\cdot}$ is an i.i.d. sequence, independent from all other aspects of the system, and

$$S_{k,i} = \sum_{j=1}^{Q_{k,i}} \tilde{Y}_{k,j}. \tag{6.26}$$

We can further write

$$\tilde{Y}_{k,j} = \mu_k / \sqrt{n} + U_{k,j}, \tag{6.27}$$

where $U_{k,j}$ is a zero-mean random variable with variance σ_k^2 . Define the scaled process:

$$U_{k,i}^n = \frac{1}{\sqrt{n}} \sum_{j=1}^i U_{k,j}. \quad (6.28)$$

We thus arrive at the following expression for the diffusion-scaled cumulative reward:

$$\bar{S}_{k,t}^n = \bar{Q}_{k,t}^n \mu_k + U_{k,n\bar{Q}_{k,t}^n}^n, \quad i = 1, \dots, n. \quad (6.29)$$

Denote by $\bar{U}_{k,t}^n$ to be the linear interpolation of $U_{k,[nt]}^n$ for $t \in [0, 1]$. By Donsker's theorem, there exists a K -dimensional standard Brownian motion W such that \bar{U}^n converges to $\{\sigma_k W_{k,\cdot}\}_{k=1,\dots,K}$ weakly in \mathcal{C} . Evoking Theorem 1 and the Skorohod's representation theorem [Billingsley, 1999, Theorem 6.7], we may construct a probability space on which the following convergences in \mathcal{C} occur almost surely:

$$\bar{S}_{k,t}^n \rightarrow S_{k,t}, \quad \bar{Q}_{k,t}^n \rightarrow Q_{k,t}, \quad \bar{U}_t^n k, t \rightarrow \sigma_k W_{k,t}, \quad (6.30)$$

as $n \rightarrow \infty$, where S and Q are diffusion processes satisfying the SDEs in Theorem 1.

We now combine (6.29) and (6.30), along with the fact that W is uniformly continuous in the compact interval $[0, 1]$, to conclude that almost surely

$$U_{k,n\bar{Q}_{k,i}^n}^n \rightarrow \sigma_k W_{k,Q_{k,t}}, \quad (6.31)$$

in \mathcal{C} . This further implies that S also satisfies

$$S_{k,t} = Q_{k,t} \mu_k + \sigma_k W_{k,Q_{k,t}}, \quad (6.32)$$

proving our claim. □

6.3 Proof of Theorem 4

Proof. Define the limit cumulative regret R_t as:

$$R_t = (\mu)_+ t - \mu Q_t, \quad t \in [0, 1], \quad (6.33)$$

where Q_t is the diffusion limit associated with Q_i^n . Note that R_1 corresponds to the scaled cumulative regret R in (2.17).

We will be working with the random-time-change version of the diffusion as per (3.17). It turns out that the diffusion limit exhibits distinct qualitative behavior in the super-diffusive regime, depending on whether $c > 0$ or $c = 0$, and whether μ tends to negative or positive infinity (some of these differences can be readily observed by comparing the sub-plots in Figure 1). As such, our proof will be divided into the same four cases. We will make repeated use of the following stopping time:

$$\tau(q) = \inf\{t : Q_t \geq q\}, \quad (6.34)$$

with $\tau(q) := 1$ if $Q_1 \leq q$.

Case 1: $c > 0, \mu \rightarrow -\infty$. Effectively, this portion of the proof will show that due to the smoothing, the algorithm does not shut down arm 1 quickly enough, thus leading to an unbounded regret. From (3.17), the drift of Q_t is given by

$$\Pi(c, Q_t) = \Phi \left(\frac{Q_t \mu + W_{Q_t}}{\sigma \sqrt{Q_t + \sigma^2 c}} \right). \quad (6.35)$$

For the sake of contradiction, suppose there exists a constant $C > 0$ such that for all sufficiently large μ :

$$\sup_{t \in [0, 1]} Q_t \leq C/|\mu|. \quad (6.36)$$

This would imply that there exists $B > 0$ such that for all sufficiently large μ ,

$$\frac{Q_t \mu + W_{Q_t}}{\sigma \sqrt{Q_t + \sigma^2 c}} \geq \frac{-C + B}{2\sigma^2 \sqrt{c}}, \quad \forall t \in [0, 1]. \quad (6.37)$$

That is, the drift of Q_t is positive and bounded away from zero for all t , independently of μ . This implies that

$$\liminf_{\mu \rightarrow -\infty} Q_1 > 0, \quad (6.38)$$

leading to a contradiction with (6.36). We conclude that $Q_1 \gg 1/|\mu|$ as $\mu \rightarrow -\infty$ and hence

$$\lim_{\mu \rightarrow -\infty} R_1 = \lim_{\mu \rightarrow -\infty} |\mu| Q_1 = \infty, \quad (6.39)$$

as claimed.

Case 2: $c > 0, \mu \rightarrow \infty$. Because $c > 0$, $|Q_t| \leq 1$, and W_t is bounded over $t \in [0, 1]$, we have that there exist constants $b_1, b_2 > 0$ such that for all $\mu > 0$

$$\Phi \left(\frac{Q_t \mu + W_{Q_t}}{\sigma \sqrt{Q_t + \sigma^2 c}} \right) \leq \Phi \left(\frac{Q_t \mu + b_1}{b_2} \right), \quad \forall t \in [0, 1]. \quad (6.40)$$

Note that since $Q_0 = 0$, and the derivative of Q_t is bounded from above by 1 for all t , we have that for all $\mu > 1$:

$$Q_t \leq 1/\mu, \quad \forall t \leq 1/\mu. \quad (6.41)$$

Combining the above two equations shows that $Q_{1/\mu}$ is strictly less than $1/\mu$:

$$Q_{1/\mu} \leq \int_0^{1/\mu} \Phi \left(\frac{Q_t \mu + W_{Q_t}}{\sigma \sqrt{Q_t + \sigma^2 c}} \right) ds \leq \Phi \left(\frac{1 + b_1}{b_2} \right) / \mu. \quad (6.42)$$

We deduce from the above that for all large μ the regret incurred by time $t = 1/\mu$ is bounded away from zero, and so

$$R \geq \mu(1/\mu - Q_{1/\mu}) \geq 1 - \Phi \left(\frac{1 + b_1}{b_2} \right) > 0. \quad (6.43)$$

This proves the claim.

Case 3: $c = 0, \mu \rightarrow -\infty$. Fix $\alpha \in (1, 2)$. Recall the stopping time from (6.34) with threshold q . We will be particularly interested in the case where

$$q = |\mu|^{-\alpha}.$$

Decompose R_1 into two components:

$$R_1 = R_{\tau(|\mu|^{-\alpha})} + (R_1 - R_{\tau(|\mu|^{-\alpha})}). \quad (6.44)$$

We next bound the two terms on the right-hand side of the equation above separately. Since Q_t is non-decreasing, we have

$$R_{\tau(|\mu|^{-\alpha})} = |\mu| \cdot Q_{\tau(|\mu|^{-\alpha})} = |\mu|^{-(\alpha-1)}. \quad (6.45)$$

For the second term, the intuition is that by the time Q_t reaches $|\mu|^{-\alpha}$, the drift in Q_t will have already become overwhelmingly small for the rest of the time horizon. To make this rigorous, note the following facts:

1. By the law of iterated logarithm of Brownian motion, along almost all sample paths of W , there exists constant C such that

$$\limsup_{\mu \rightarrow -\infty} \sup_{x \in [\tau(|\mu|^{-\alpha}), 1]} \left| \frac{W_x}{\sqrt{x}} \right| \leq C \sqrt{\log \log(|\mu|^\alpha)}. \quad (6.46)$$

2. $\mu \sqrt{Q_{\tau(|\mu|^{-\alpha})}} = -|\mu|^{1-\alpha/2}$, and therefore $\frac{\mu \sqrt{Q_{\tau(|\mu|^{-\alpha})}}}{2} \ll -\sqrt{\log \log(|\mu|^\alpha)}$ as $\mu \rightarrow -\infty$.

Combining these facts along with the normal cdf tail bounds from Lemma 18 in Appendix A, we have that along almost all sample paths of W , there exists constant $b > 0$, such that for all sufficiently small μ ,

$$\begin{aligned} R_1 - R_{\tau(|\mu|^{-\alpha})} &= |\mu| \int_{\tau(|\mu|^{-\alpha})}^1 \Pi(0, Q_t) dt \\ &\leq |\mu| \left(\sup_{t \in [\tau(|\mu|^{-\alpha}), 1]} \Phi \left(\frac{\mu \sqrt{Q_t}}{\sigma} + \frac{W_{Q_t}}{\sigma \sqrt{Q_t}} \right) \right) \\ &\stackrel{(a)}{\leq} |\mu| \Phi \left(-|\mu|^{1-\alpha/2} + C \sqrt{\log \log(|\mu|^\alpha)} \right) \\ &\leq |\mu| \exp(-|\mu|^b), \end{aligned} \quad (6.47)$$

where (a) follows from the aforementioned facts. Putting together (6.45) and (6.47) shows that

$$R_1 \leq \mu^{-(\alpha-1)} + |\mu| \exp(-|\mu|^b) \stackrel{\mu \rightarrow \infty}{\prec} \mu^{-(\alpha-1)}, \quad a.s. \quad (6.48)$$

This proves the claim by noting that the above holds for all $\alpha \in (1, 2)$.

Case 4: $c = 0, \mu \rightarrow \infty$. In this case, we would like to argue that Q_t will increase rapidly as μ grows. Let η be a function of the form:

$$\eta(x) = 1 - x^{-\alpha}, \quad (6.49)$$

where $\alpha \in (1, 2)$ is a constant; the value of α will be specified in a later part of the proof.

The remainder of the proof will be centered around the dynamics of Q before and after the following stopping time:

$$\tau(\eta(\mu)) = \inf\{t : Q_t \geq 1 - \mu^{-\alpha}\}. \quad (6.50)$$

It follows from the definition that if

$$\tau(\eta(\mu)) < 1, \quad (6.51)$$

then

$$Q_1 \geq \eta(\mu) = 1 - \mu^{-\alpha}. \quad (6.52)$$

Therefore, if we can show that almost surely (6.51) holds for all sufficiently large μ , then it follows that for all large μ , the desired inequality holds:

$$R_1 = \mu(1 - Q_1) \leq \mu^{-(\alpha-1)}, \quad a.s. \quad (6.53)$$

The remainder of the proof is devoted to showing (6.51). A main challenge in this part of the proof is that the dynamics of $W_{Q_t}/\sqrt{Q_t}$, and by consequence that of Q_t , is highly volatile near $t = 0$. To obtain a handle on the behavior of these quantities, we will use the following trick by performing a change of variable in time so that the integrand below changes from s to $1/s$. We have that

$$\begin{aligned} \tau(\eta(\mu)) &= \int_0^{\eta(\mu)} 1/\Phi\left(\frac{s\mu + W_s}{\sigma\sqrt{s}}\right) ds \\ &= \int_{\eta(\mu)^{-1}}^{\infty} u^{-2} \left(\Phi\left(\frac{\mu}{\sigma\sqrt{u}} + \frac{\tilde{W}_u}{\sigma\sqrt{u}}\right) \right)^{-1} du \\ &= \int_{\eta(\mu)^{-1}}^{\infty} u^{-2} \xi(\mu, u) du \end{aligned} \quad (6.54)$$

where $u = 1/s$,

$$\xi(\mu, u) := \left(\Phi\left(\frac{\mu}{\sigma\sqrt{u}} + \frac{\tilde{W}_u}{\sigma\sqrt{u}}\right) \right)^{-1}. \quad (6.55)$$

and

$$\tilde{W}_t = tW_{1/t}. \quad (6.56)$$

Importantly, it is well known that if W_t is a standard Brownian motion, then so is \tilde{W}_t .

We now bound the above integral using a truncation argument. For $K > \eta(\mu)^{-1}$, we write

$$\begin{aligned} \tau(\eta(\mu)) &= \int_{\eta(\mu)^{-1}}^{\infty} u^{-2} \xi(\mu, u) du \\ &= \int_{\eta(\mu)^{-1}}^K u^{-2} \xi(\mu, u) du + \int_K^{\infty} u^{-2} \xi(\mu, u) du \\ &\leq \left(\sup_{u \in [\eta(\mu)^{-1}, K]} \xi(\mu, u) \right) \int_{\eta(\mu)^{-1}}^{\infty} u^{-2} du + \int_K^{\infty} u^{-2} \xi(\mu, u) du \\ &= \left(\sup_{u \in [\eta(\mu)^{-1}, K]} \xi(\mu, u) \right) \eta(\mu) + \int_K^{\infty} u^{-2} \xi(\mu, u) du. \end{aligned} \quad (6.57)$$

The following lemma bounds the second term in the above equation; the proof is given in Appendix B.2.

Lemma 13. *For any $\delta \in (0, 1)$, there exists $C > 0$ such that, along almost all sample paths of W , for all large μ and K :*

$$\int_K^\infty u^{-2} \xi(\mu, u) du \leq CK^{-(1-\delta)}. \quad (6.58)$$

Bounding the first term in (6.57) is more delicate, and will involve taking μ to infinity in a manner that depends on K . Fix any $\gamma \in (0, 1)$ and consider a parameterization of μ where

$$\mu_K = K^{\frac{1}{2}+\gamma}, \quad K \in \mathbb{N}. \quad (6.59)$$

By law of iterated logarithm (Lemma 19 in Appendix A), and noting that $\eta(\mu) < 1$, we have that there exists $C > 0$ such that for all sufficiently large K

$$\inf_{u \in [\eta(\mu_K)^{-1}, K]} \frac{\tilde{W}_u}{\sqrt{u}} \geq -C\sqrt{\log \log K}, \quad a.s. \quad (6.60)$$

Combining this with the lower bound on the normal cdf (Lemma 18), we have, for all large K ,

$$\begin{aligned} \sup_{u \in [\eta(\mu_K)^{-1}, K]} \xi(\mu_K, u) &\leq 1 + \frac{\exp(-(\mu_K/\sqrt{K} - C\sqrt{\log \log K})^2)}{\mu_K/\sqrt{K} - C\sqrt{\log \log K}} \\ &= 1 + \frac{\exp(-(K^{1/2+\gamma}/\sqrt{K} - C\sqrt{\log \log K})^2)}{K^{1/2+\gamma}/\sqrt{K} - C\sqrt{\log \log K}} \\ &\leq 1 + \exp\left(-(K^\gamma - C\sqrt{\log \log K})^2 - \gamma \log K\right) \\ &\leq 1 + \exp(-K^\gamma). \end{aligned} \quad (6.61)$$

Fix $\nu \in (0, 1)$, and $\delta, \gamma \in (0, 1/4)$ such that

$$2 > \frac{1-\delta}{1/2+\gamma} > 2-\nu. \quad (6.62)$$

Note that such δ and γ exist for any ν , so long as we ensure that both δ and γ are sufficiently close to 0. Combining (6.57), (6.61) and Lemma 13, we have that there exist $c_1, c_2 > 0$ such that, along almost all sample paths of W , for all large K :

$$\begin{aligned} \tau_{\mu_K} &\leq \left(\sup_{u \in [\eta(\mu_K)^{-1}, K]} \xi(\mu_K, u) \right) \eta(\mu_K) + \int_K^\infty u^{-2} \xi(\mu_K, u) du \\ &\leq (1 + \exp(-K^\gamma)) \eta(\mu_K) + c_1 K^{-(1-\delta)} \\ &\leq \eta(K^{1/2+\gamma}) + c_2 K^{-(1-\delta)}, \quad a.s. \end{aligned} \quad (6.63)$$

Recall that $\eta(\mu) = 1 - \mu^{-\alpha}$. We now choose α to be such that

$$2 - \nu < \alpha < \frac{1-\delta}{(1/2+\gamma)} < 2. \quad (6.64)$$

Under this choice of α (which exists because of (6.62)), we have that for all sufficiently large K

$$\tau_{\mu_K} \leq 1 - K^{-\alpha(1/2+\gamma)} + c_2 K^{-(1-\delta)} < 1, \quad a.s., \quad (6.65)$$

where the last inequality follows from (6.64). Combining the above equation, (6.51), (6.53) and the fact that ν can be arbitrarily close to 0, we have thus shown that for all $\alpha \in (1, 2)$,

$$R_1 \leq \mu(1 - \eta(\mu)) = \mu^{-(\alpha-1)}, \quad a.s., \quad (6.66)$$

for all large μ . This proves our main claim in this case, that is, almost surely

$$R_1 < 1/\mu, \quad \text{as } \mu \rightarrow \infty. \quad (6.67)$$

□

6.4 Proof of Theorem 5

Proof. We now prove that the $1/\delta$ regret scaling holds under undersmoothed two-armed Thompson sampling, as the arm gap δ tends to infinity. Before delving into the details, let us first point out an intriguing connection between the two-armed bandit analyzed here and the one-armed version in Theorem 4, and its consequence for how the proof will be carried out. In the one-armed setting, a crucial simplifying feature is that there is no uncertainty associated with the second, default arm, whereas in the two-armed case, both arms' mean rewards are uncertain. This manifests in there being only one driving Brownian motion in the description of the diffusion limit for one-armed Thompson sampling, versus two independent Brownian motions in the two-armed case.

Fortunately, the aforementioned distinction also suggests a plan of attack for analyzing the two-armed bandit. Consider the diffusion process at a small positive time $t = v > 0$. At this point, we have little knowledge of the behavior of Q_t because its dynamics near $t = 0$ is highly volatile due to a lack of data during this period. However, we do know that the total arm pulls up to this point have to add up to v , and so at least one of the two arms has been pulled by the amount $v/2$. The key insight here is that, depending on which arm has been pulled more by this point, we obtain, from time $t = v$ onward, a version of diffusion that mirrors one of the two super-diffusive regimes under the *one-arm* Thompson sampling, i.e., $\mu \rightarrow -\infty$ or $\mu \rightarrow \infty$. Specifically:

1. If arm 1 (superior arm) has been pulled by at least $v/2$, then we can approximately treat arm 1 as the “certain” arm, and our problem can be approximately mapped to a one-armed bandit with $\mu \rightarrow -\infty$.
2. If, on the other hand, arm 2 (inferior arm) has been pulled by $v/2$, then arm 2 can be viewed as the certain arm, and the problem can be roughly reduced to a one-armed bandit with $\mu \rightarrow +\infty$.

The above heuristic argument sets the stage for how the proof will proceed: we will consider two separate cases depending on the realization of $Q_{1,t}$ at a carefully chosen, early point in time, and subsequently manipulate the drift equation to exploit the above-mentioned symmetry. That being said, the reduction from the two-armed bandit into two separate one-armed bandits is not exact, and the remainder of our proof is centered around using delicate estimates to make the above connection precise.

We now present the formal proof of Theorem 5. We will be working with the random-time-change version of the diffusion as per (3.18). For clarity of notation, we will fix $\sigma = 1$ throughout the proof. All results extend easily to the case of an arbitrary, fixed σ . Fix $\alpha \in (1, 2)$. Define

$$v = 2\delta^{-\alpha}, \quad (6.68)$$

and the events

$$\mathcal{E}_k = \{Q_{k,v} \geq \delta^{-\alpha}\}, \quad k = 1, 2. \quad (6.69)$$

Because $t = Q_{1,t} + Q_{2,t}$ for all t , we have that either \mathcal{E}_1 or \mathcal{E}_2 occurs almost surely. The proof will be completed by showing the claimed regret bound by conditioning upon each of these two events separately. Without loss of generality, the analysis of each case assumes the corresponding event occurs with strictly positive probability. Should one of the events occurs with probability zero, that portion of the proof can simply be ignored and should not impact the overall claim.

Case 1. First, suppose that \mathcal{E}_2 has occurred, and without stating otherwise, this conditioning will be assumed throughout this portion of the proof. From (3.18), we have that

$$dQ_{1,t} = \Phi \left(\frac{Q_{2,t}}{\sqrt{tQ_{2,t}}} \left(\frac{\delta Q_{1,t} + W_{1,Q_{1,t}}}{\sqrt{Q_{1,t}}} \right) - \sqrt{\frac{Q_{1,t}}{t}} \frac{W_{2,Q_{2,t}}}{\sqrt{Q_{2,t}}} \right) dt, \quad (6.70)$$

Define functions

$$g_1(t) = \frac{Q_{2,t}}{\sqrt{tQ_{2,t}}}, \quad g_2(t) = -\sqrt{\frac{Q_{1,t}}{t}} \frac{W_{2,Q_{2,t}}}{\sqrt{Q_{2,t}}}. \quad (6.71)$$

Then,

$$dQ_{1,t} = \Phi \left(g_1(t) \left(\frac{\delta Q_{1,t} + W_{1,Q_{1,t}}}{\sqrt{Q_{1,t}}} \right) + g_2(t) \right) dt. \quad (6.72)$$

The following lemma shows that the terms g_1 and g_2 are appropriately bounded, and will be crucial to our subsequent analysis; the proof is presented Appendix B.3.

Lemma 14. *The following is true almost surely:*

1. $1 \geq g_1(t) \geq \delta^{-\alpha/2}$ for all $t \geq v$.
2. There exists constant C , such that for all large δ and $t \geq v$,

$$|g_2(t)| \leq C\sqrt{\log \log \delta}. \quad (6.73)$$

Recall the stopping time defined in (6.34):

$$\tau(q) = \inf\{t : Q_{1,t} \geq q\}. \quad (6.74)$$

Using the fact that $Q_{1,\tau(t)} = t$, we have that

$$\begin{aligned} \tau(q) &= \int_0^q \tau'(s) ds \\ &= \int_0^q 1/\Phi \left(g_1(\tau(s)) \left(\frac{\delta s + W_{1,s}}{\sqrt{s}} \right) + g_2(\tau(s)) \right) ds \end{aligned} \quad (6.75)$$

For $q > Q_{1,v}$, we have, from the change of variables $s = u^{-1}$ (see e.g., the discussion preceding (6.54) where this transformation was first used),

$$\begin{aligned} \tau(q) &= v + \int_{Q_{1,v}}^q 1/\Phi \left(g_1(\tau(s)) \left(\frac{\delta s + W_{1,s}}{\sqrt{s}} \right) + g_2(\tau(s)) \right) ds \\ &= v + \int_{1/q}^{1/Q_{1,v}} u^{-2}/\Phi \left(g_1(\tau(u^{-1})) \left(\frac{\delta}{\sqrt{u}} + \frac{\tilde{W}_u}{\sqrt{u}} \right) + g_2(\tau(u^{-1})) \right) du. \end{aligned} \quad (6.76)$$

Note that $\tau(s) \geq v$ for all $s \geq Q_{1,v}$.

We now employ a truncation argument similar to that in the proof of Theorem 4, Case 4 (6.57). Define

$$\xi(\delta, u) = 1/\Phi \left(g_1(\tau(u^{-1})) \left(\frac{\delta}{\sqrt{u}} + \frac{\tilde{W}_u}{\sqrt{u}} \right) + g_2(\tau(u^{-1})) \right). \quad (6.77)$$

For $M \in \mathbb{R}$, $1/q < M < 1/Q_{1,v}$, we have, from (6.76), that

$$\begin{aligned} \tau(q) &\leq v + \int_{1/q}^M u^{-2} \xi(\delta, u) du + \int_M^{1/Q_{1,v}} u^{-2} \xi(\delta, u) du \\ &\leq v + \left(\sup_{u \in [1/q, M]} \xi(\delta, u) \right) q + \int_M^{1/Q_{1,v}} u^{-2} \xi(\delta, u) du. \end{aligned} \quad (6.78)$$

The last term in the above equation can be bounded by the following lemma. The result is analogous to Lemma 13, with the key difference being that now the instance-specific parameter δ also features in the bound; the proof is given in Appendix B.4.

Lemma 15. *For any $\beta \in (0, 1)$, there exist constants $B, C > 0$ such that for all large δ and M :*

$$\int_M^{1/Q_{1,v}} u^{-2} \xi(\delta, u) du \leq C(\log \delta)^B M^{-(1-\beta)}, \quad a.s. \quad (6.79)$$

Fix $\epsilon \in (1, 2)$ and define

$$q^*(\delta) = 1 - \delta^{-\epsilon}. \quad (6.80)$$

Applying Lemma 15 to (6.78), we have that for all $\beta \in (0, 1)$ and sufficiently large M and δ :

$$\tau(q^*(\delta)) \leq v + \left(\sup_{u \in [q^*(\delta)^{-1}, M]} \xi(\delta, u) \right) q^*(\delta) + C(\log \delta)^B M^{1-\beta}. \quad (6.81)$$

To bound the term in the middle, we again resort to a double limit, in which M and δ will tend to infinity simultaneously. Fix $\gamma > \frac{\alpha/2}{2-\alpha}$, and consider a parameterization of δ :

$$\delta_M = M^{1/2+\gamma}. \quad (6.82)$$

By the law of iterated logarithm (Lemma 19 in Appendix A), we have that there exists $C > 0$ such that for all large δ and M

$$\inf_{u \in [q^*(\delta_M)^{-1}, M]} \frac{\tilde{W}_u}{u} \geq -C\sqrt{\log \log M}. \quad (6.83)$$

We have that for all large M

$$\begin{aligned}
& \sup_{u \in [q^*(\delta_M)^{-1}, M]} \xi(\delta_M, u) \\
&= \sup_{u \in [q^*(\delta_M)^{-1}, M]} 1/\Phi \left(g_1(\tau(u^{-1})) \left(\frac{\delta}{\sqrt{u}} + \frac{\tilde{W}_u}{\sqrt{u}} \right) + g_2(\tau(u^{-1})) \right) \\
&\stackrel{(a)}{\leq} \sup_{u \in [q^*(\delta_M)^{-1}, M]} 1/\Phi \left(g_1(\tau(u^{-1})) \left(\frac{\delta}{\sqrt{u}} + \frac{\tilde{W}_u}{\sqrt{u}} \right) - c_1 \sqrt{\log \log \delta_M} \right) \\
&\stackrel{(b)}{\leq} \sup_{u \in [q^*(\delta_M)^{-1}, M]} 1/\Phi \left((\delta_M)^{1-\alpha/2}/\sqrt{u} - c_2 \left(\sqrt{\log \log u} + \sqrt{\log \log \delta_M} \right) \right) \\
&\leq 1/\Phi \left((\delta_M)^{1-\alpha/2}/M^{1/2} - c_2 \left(\sqrt{\log \log M} + \sqrt{\log \log \delta_M} \right) \right) \\
&= 1/\Phi \left(M^{\gamma(1-\alpha/2)-\alpha/4} - c_2 \left(\sqrt{\log \log M} + \sqrt{\log \log \delta_M} \right) \right) \\
&\stackrel{(c)}{\leq} 1 + \frac{\exp \left(- \left(M^{\gamma(1-\alpha/2)-\alpha/4} - c_2 \left(\sqrt{\log \log M} + \sqrt{\log \log \delta_M} \right) \right)^2 \right)}{M^{\gamma(1-\alpha/2)-\alpha/4} - c_2 \left(\sqrt{\log \log M} + \sqrt{\log \log \delta_M} \right)} \\
&\leq 1 + \exp(-M^{\gamma(1-\alpha/2)-\alpha/4}), \tag{6.84}
\end{aligned}$$

where we note that the exponent $(\gamma(1-\alpha/2)-\alpha/4)$ is strictly positive based on the definition of γ . The steps are based on:

- (a): $g_2(t) \leq C\sqrt{\log \log \delta}$ when $t \geq v$.
- (b): $1 \geq g_1(t) \geq \delta^{-\alpha/2}$ for $t \in [v, 1)$, and (6.83).
- (c): The lower bound on the normal cdf for $x > 0$ in Lemma 18, Appendix A.

We now substitute (6.84) into (6.81), and recall that $v = 2\delta^{-\alpha}$ and $q^*(\delta) = 1 - \delta^{-\epsilon}$. We have

$$\begin{aligned}
\tau(q^*(\delta_M)) &\leq v + \left(\sup_{u \in [q^*(\delta_M)^{-1}, M]} \xi(\delta_M, u) \right) q^*(\delta_M) + C(\log \delta_M)^B M^{1-\beta} \\
&\leq 2\delta_M^{-\alpha} + \left(1 + \exp(-M^{\gamma(1-\alpha/2)-\alpha/4}) \right) (1 - \delta_M^{-\epsilon}) + C(\log \delta)^B M^{1-\beta} \\
&\leq 1 - \delta_M^{-\epsilon} + 2\delta_M^{-\alpha} + C(\log \delta_M)^B M^{1-\beta} + \exp(-M^{\gamma(1-\alpha/2)-\alpha/4}). \tag{6.85}
\end{aligned}$$

Let us choose the parameters ϵ and β as follows:

1. Let $\epsilon \in (1, \alpha)$, so that $\delta^{-\epsilon} \gg \delta^{-\alpha}$ as $\delta \rightarrow \infty$.
2. Choose β to be sufficiently close to 1 such that $C(\log \delta_M)^B M^{1-\beta} < \delta_M^{-\alpha}$ for all large M .

Under these choices of parameters, we see that $\delta_M^{-\epsilon}$ is orders-of-magnitude larger than the sum of the last three terms in (6.85), and we have that for all large M :

$$\tau(q^*(\delta_M)) \leq 1 - \delta_M^{-\epsilon}/2 < 1. \tag{6.86}$$

We can now turn to the final objective, R_1 . (6.86) shows that $Q_{1,1} \geq 1 - \delta_M^{-\epsilon}$. We have that for all large M

$$R = \delta_M(1 - Q_{1,1}) \leq \delta_M(1 - q^*(\delta_M)) = \delta_M^{-(\epsilon-1)}. \quad (6.87)$$

Finally, notice that since α may take any value in $(1, 2)$, so can $\epsilon - 1$ take on any value in $(0, 1)$. Considering that δ_M is a continuous increasing function of M that tends to infinity as $M \rightarrow \infty$, we conclude that $R < 1/\delta$ as $\delta \rightarrow \infty$, as claimed.

Case 2. Next, we will look at the case where \mathcal{E}_1 occurred. The structure of this part of the proof loosely mirrors Case 3 in the proof of Theorem 4, where arm 1's reward is more certain compared to that of arm 2. With this in mind, we will instead look at the drift of $Q_{2,\cdot}$:

$$dQ_{2,t} = \Phi \left(\frac{Q_{1,t}}{\sqrt{tQ_{1,t}}} \left(\frac{-\delta Q_{2,t} - W_{2,Q_{2,t}}}{\sqrt{Q_{2,t}}} \right) + \sqrt{\frac{Q_{2,t}}{t}} \frac{W_{1,Q_{1,t}}}{\sqrt{Q_{1,t}}} \right) dt. \quad (6.88)$$

We now redefine g_1 and g_2 in a fashion that mirrors symmetrically the first case:

$$g_1(t) = \frac{Q_{1,t}}{\sqrt{tQ_{1,t}}}, \quad g_2(t) = \sqrt{\frac{Q_{2,t}}{t}} \frac{W_{1,Q_{1,t}}}{\sqrt{Q_{1,t}}}. \quad (6.89)$$

We obtain

$$dQ_{2,t} = \Phi \left(g_1(t) \left(\frac{-\delta Q_{2,t} - W_{2,Q_{2,t}}}{\sqrt{Q_{2,t}}} \right) + g_2(t) \right) dt. \quad (6.90)$$

Note that due to the symmetry, the bounds on g_1 and g_2 as laid out in Lemma 14 continue to hold.

We now define τ in terms of $Q_{2,\cdot}$:

$$\tau(q) = \inf\{t : Q_{2,t} = q\}, \quad (6.91)$$

and define

$$\tau_\delta^* = \tau(\delta^{-\alpha}). \quad (6.92)$$

Because \mathcal{E}_1 is assumed to have occurred, we have that

$$\tau_\delta^* \geq v. \quad (6.93)$$

The following property on the drift of $Q_{2,\cdot}$ will be used in the remainder of the proof; the proof is given in Appendix B.5.

Lemma 16. *Along almost all sample paths of W , there exists a constant C such that for all large δ ,*

$$\frac{d}{dt}Q_{2,t} \leq \Phi \left(-g_1(t)\sqrt{Q_{2,t}}\delta + C\sqrt{\log \log \delta} \right), \quad \forall t \in [\tau_\delta^*, 1). \quad (6.94)$$

Decompose the regret as follows:

$$R = R_{\tau(\delta^{-\alpha})} + (R - R_{\tau(\delta^{-\alpha})}). \quad (6.95)$$

For the first term, we have that

$$R_{\tau(\delta^{-\alpha})} \leq \delta \cdot \delta^{-\alpha} = \delta^{-(\alpha-1)}. \quad (6.96)$$

To bound the second term, we will again aim to show that once $Q_{2,\cdot}$ reaches $\delta^{-\alpha}$, its drift would become overwhelmingly small as δ gets large. Compared to the one-armed setting, a major obstacle in this case is that the $\delta^{-\alpha/2}$ uniform lower bound on $g_1(t)$ in Lemma 14 turns out to be too weak for our purpose. We will rely on the following stronger lower bound on $g_1(t)$; the proof is given in Appendix B.6.

Lemma 17. *Along almost all sample paths of W , we have that for all sufficiently large δ , under the conditioning of event \mathcal{E}_1 ,*

$$g_1(t) \geq \sqrt{1/3}, \quad \forall t \in [\tau_\delta^*, 1). \quad (6.97)$$

With the strengthened lower bound on $g_1(t)$ at hand, we are now ready to bound the second term in (6.95). Combining Lemmas 16 and 17, we have that for all large δ

$$\begin{aligned} \sup_{t \in [\tau_\delta^*, 1)} \frac{d}{dt} Q_{2,t} &\leq \sup_{t \in [\tau_\delta^*, 1)} \Phi \left(-\sqrt{1/3} \sqrt{Q_{2,t}} \delta + C \sqrt{\log \log \delta} \right) \\ &\leq \Phi \left(-\sqrt{1/3} \delta^{1-\alpha/2} + C \sqrt{\log \log \delta} \right) \\ &< \exp(-\delta^{1-\alpha/2}), \end{aligned} \quad (6.98)$$

where the second inequality follows from the definition of τ_δ^* . We can write the regret term as

$$R - R_{\tau_\delta^*} = \delta \int_{\tau_\delta^*}^1 \frac{d}{dt} Q_{2,t} < \delta \exp(-\delta^{1-\alpha/2}) \ll \delta^{-1}. \quad (6.99)$$

This shows that the regret R is dominated by $R_{\tau_\delta^*}$, so that

$$R < \delta^{-(\alpha-1)}, \quad (6.100)$$

and the claim follows from the fact that α can be arbitrarily close to 2.

6.5 Proof of Theorem 6

Proof. Using the random-time change characterization of the diffusion limit from Theorem 3, we have that

$$\pi_t = \Phi \left(\frac{\mu \sqrt{Q_t}}{\sigma} + \frac{W_{Q_t}}{\sigma \sqrt{Q_t}} \right). \quad (6.101)$$

By the law of iterated logarithm (Lemma 19, Appendix A) of Brownian motion, as well as the fact that $\lim_{t \downarrow 0} Q_t = 0$, we have that almost surely

$$\limsup_{t \downarrow 0} \frac{W_{Q_t}/\sqrt{Q_t}}{\sqrt{2 \log \log(1/Q_t)}} = \liminf_{t \downarrow 0} \frac{W_{Q_t}/\sqrt{Q_t}}{-\sqrt{2 \log \log(1/Q_t)}} = 1. \quad (6.102)$$

This further implies that, almost surely, the term $\frac{W_{Q_t}}{\sigma \sqrt{Q_t}}$ in the expression of π_t will oscillate between arbitrarily large positive and negative values as $t \rightarrow 0$. Since the term $\frac{\mu \sqrt{Q_t}}{\sigma}$ is bounded, this proves our claim. \square

6.6 Proof of Theorem 7

Proof. We will use the characterization of the diffusion process in (3.17). Although our main focus is on the process Z_t restricted to the $[0, 1]$ interval, the diffusion process itself is in fact well defined on $t \in [0, \infty)$. A useful observation we will make here is that, for any $c > 0$ and μ , we have that almost surely

$$Q_t \rightarrow \infty, \quad \text{as } t \rightarrow \infty. \quad (6.103)$$

This fact can be verified by noting that if Q_t were bounded over $[0, \infty)$, then its drift $\Pi(c, Q_t)$ would have been bounded from below by a strictly positive constant, leading to a contradiction. Recall the stopping time from (6.34), and here we emphasize the dependence on c :

$$\tau^c(q) = \inf\{t : Q_t \geq q\}. \quad (6.104)$$

Note that for all $c > 0$, Q_t and τ^c are increasing and continuous and $Q_t = \tau^{-1}(t)$.

As $c \rightarrow 0$, we face the same challenge of the volatile behavior of the drift function near $t = 0$, as was discussed in the paragraph preceding (6.54), and we will resort to the same change of variable technique as that in the proof of Theorem 4, (6.54). From (3.17), we have that for $q > 0$:

$$\begin{aligned} \tau^c(q) &= \int_0^q 1/\Phi\left(\frac{s\mu + W_s}{\sigma\sqrt{s + \sigma^2 c}}\right) ds \\ &= \int_{1/q}^\infty u^{-2}/\Phi\left(\frac{\mu + uW_{1/u}}{\sigma\sqrt{u + u^2\sigma^2 c}}\right) du \\ &= \int_{1/q}^\infty u^{-2}/\Phi\left(\frac{\mu + \tilde{W}_u}{\sigma\sqrt{u + u^2\sigma^2 c}}\right) du \\ &= \int_{1/q}^\infty h^c(u) du \end{aligned} \quad (6.105)$$

where $\tilde{W}_t = tW_{1/t}$ is a standard Brownian motion, and

$$h^c(u) = u^{-2}/\Phi\left(\frac{\mu + \tilde{W}_u}{\sigma\sqrt{u + u^2\sigma^2 c}}\right). \quad (6.106)$$

By the law of iterated logarithm of Brownian motion (Lemma 19 in Appendix A), we have that

$$\limsup_{t \rightarrow \infty} \frac{|\tilde{W}_t|}{\sqrt{2t \log \log t}} = 1, \quad a.s. \quad (6.107)$$

Fix a sample path of W such that the above is satisfied. Then, there exist $M \in (1/q, \infty)$ and $D > 0$, such that

$$\tilde{W}_t \geq -D\sqrt{t \log \log t}, \quad \forall t \geq M. \quad (6.108)$$

We now consider two cases depending on the sign of μ . First, suppose that $\mu \geq 0$. Define

$$g(u) = \begin{cases} u^{-2}/\Phi\left(\frac{\mu - |\tilde{W}_u|}{\sigma\sqrt{u}}\right), & 1/q \leq u < M, \\ u^{-2}/\Phi\left(-\frac{D}{\sigma}\sqrt{\log \log u}\right), & u \geq M. \end{cases} \quad (6.109)$$

where $M \in (1/q, \infty)$ is chosen such that (6.108) holds. It follows from the definitions that

$$h^c(u) \leq g(u), \quad \forall c \geq 0, u \geq 1/q. \quad (6.110)$$

We now show that g is integrable over $u \in [1/q, \infty)$. It suffices to show that

$$\int_M^\infty g(u) du < \infty. \quad (6.111)$$

Recall the following lower bound on the cdf of standard normal from Lemma 18 in Appendix A: for all sufficiently small x

$$\Phi(x) \geq \frac{1}{\sqrt{2\pi}} \frac{-x}{1+x^2} \exp(-x^2/2), \quad x < 0. \quad (6.112)$$

We thus have that, for all sufficiently large u ,

$$\begin{aligned} g(u) &= u^{-2} / \Phi\left(-\frac{D}{\sigma} \sqrt{\log \log u}\right) \\ &\leq u^{-2} \sqrt{2\pi} \frac{1 + \frac{D^2}{\sigma^2} \log \log u}{\frac{D}{\sigma} \sqrt{\log \log u}} \exp\left(\frac{D^2}{2\sigma^2} \log \log u\right) \\ &\leq 2u^{-2} \sqrt{2\pi} \frac{D}{\sigma} \sqrt{\log \log u} (\log u)^{D^2/2\sigma^2} \\ &\leq b_1 u^{-2} (\log u)^{b_2}, \end{aligned} \quad (6.113)$$

where b_1 and b_2 are positive constants. Noting that

$$(\log u)^\alpha \ll \sqrt{u} \quad (6.114)$$

as $u \rightarrow \infty$ for any constant $\alpha > 0$, we have that $b_1 u^{-2} (\log u)^{b_2}$ is integrable over (M, ∞) for all sufficiently large M . This proves the integrability of g in (6.111).

Using (6.110), (6.111) and the dominated convergence theorem, we thus conclude that, for all $q > 0$,

$$\begin{aligned} \lim_{c \downarrow 0} \tau^c(q) &= \int_{1/q}^\infty \lim_{c \rightarrow \infty} \left(u^{-2} / \Phi\left(\frac{\mu + \tilde{W}_u}{\sigma \sqrt{u + u^2 \sigma^2 c}}\right) \right) du \\ &= \int_0^q 1 / \Phi\left(\frac{s\mu + W_s}{\sigma \sqrt{s}}\right) ds := \tau^0(q), \quad a.s. \end{aligned} \quad (6.115)$$

Recall that $Q_t = (\tau^c)^{-1}(t)$, the above thus implies that, for all $t \in [0, 1]$,

$$Q_t \xrightarrow{c \downarrow 0} \tilde{Q}_t := (\tau^0)^{-1}(t), \quad a.s. \quad (6.116)$$

Finally, the point-wise convergence implies uniform convergence over the compact interval $[0, 1]$ since Q_t is 1-Lipschitz. This proves our claim in the case where $\mu \geq 0$. The case for $\mu < 0$ follows an essentially identical set of arguments after adjusting the constants M and D in (6.109), recognizing that the behavior of $\mu + \tilde{W}^u$ is largely dominated by that of \tilde{W}^u when u is large, which can be in turn bounded by the law of iterated logarithm. This completes the proof of Theorem 7. \square

6.7 Proof of Theorem 8

Proof. Fix a sequence $\{c_i\}$ such that $c_i \downarrow 0$ as $i \rightarrow \infty$. Consider a set $\{\nu_{i,n}\}_{i,n \in \mathbb{N}}$ where $\nu_{i,n} > 0$ and

$$\lim_{n \rightarrow \infty} (\nu_{i,n})^{-2}/n = c_i, \quad \forall i \in \mathbb{N}. \quad (6.117)$$

Denote by $\bar{Z}^{i,n}$ the pre-limit sample path with prior standard deviation $\nu_{i,n}$. Let Z^i be the limit diffusion for when $c = c_i$, with Z^∞ being the solution for when $c = 0$. We thus have that almost surely

$$\begin{aligned} \bar{Z}^{i,n} &\longrightarrow Z^i, & \text{as } n \rightarrow \infty, \text{ for all } i, \\ Z^i &\longrightarrow Z^\infty, & \text{as } i \rightarrow \infty, \end{aligned} \quad (6.118)$$

where these two limits follow from Theorems 1 and 7, respectively. It thus follows from (6.118) that there exists a sequence $\{i_n\}$, such that almost surely

$$\bar{Z}^{i_n,n} \longrightarrow Z^\infty, \quad \text{as } n \rightarrow \infty. \quad (6.119)$$

Setting $\nu^n = \nu_{i_n,n}$ thus proves the claim. \square

6.8 Proof of Theorem 9

Proof. The first claim follows directly from the fact that Q^j are 1-Lipschitz and the Arzela-Ascoli theorem. For the second claim, let $\{Q^{j_k}\}_{k \in \mathbb{N}}$ be the subsequence that converges uniformly to Q . Because $W_{1,\cdot}$ and $W_{2,\cdot}$ are uniformly bounded in $[0, 1]$ almost surely, it follows that $\frac{d}{dt}Q_t^{j_k}$ converges uniformly to $\frac{d}{dt}Q_t$ over all compact subsets of $(0, 1]$ as $k \rightarrow \infty$. This proves the claim. \square

References

- Shipra Agrawal and Navin Goyal. Near-optimal regret bounds for Thompson sampling. *Journal of the ACM*, 64(5):1–24, 2017.
- Victor F Araman and Rene Caldentey. Diffusion approximations for a class of sequential testing problems. *Available at SSRN 3479676*, 2019.
- Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1):133–161, 2021.
- Susan Athey, Sarah Baird, Vitor Hadad, Julian Jamison, Craig McIntosh, Berk Özler, Luca Parisotto, and Dohbit Sama. Increasing the take-up of long acting reversible contraceptives among adolescents and young women in Cameroon. 2021.
- Jean-Yves Audibert, Sébastien Bubeck, et al. Minimax policies for adversarial and stochastic bandits. In *COLT*, volume 7, pages 1–122, 2009.
- Peter Auer and Ronald Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.

- Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 1999.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1): 1–122, 2012.
- Sébastien Bubeck and Che-Yu Liu. Prior-free and prior-dependent regret bounds for thompson sampling. In *2014 48th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–9. IEEE, 2014.
- Stefano Caria, Maximilian Kasy, Simon Quinn, Soha Shami, and Alex Teytelboym. An adaptive targeted field experiment: Job search assistance for refugees in jordan. 2020.
- Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems*, pages 2249–2257, 2011.
- Richard Durrett. *Stochastic calculus: a practical introduction*, volume 6. CRC press, 1996.
- Ido Erev and Alvin E Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, pages 848–881, 1998.
- Lin Fan and Peter W Glynn. Diffusion approximations for thompson sampling. *arXiv preprint arXiv:2105.09232*, 2021.
- Kris Johnson Ferreira, David Simchi-Levi, and He Wang. Online network revenue management using thompson sampling. *Operations research*, 66(6):1586–1602, 2018.
- David Gamarnik and Assaf Zeevi. Validity of heavy traffic steady-state approximations in generalized Jackson networks. *The Annals of Applied Probability*, 16(1):56–90, 2006.
- Peter W Glynn. Diffusion approximations. *Handbooks in Operations research and management Science*, 2:145–198, 1990.
- Vitor Hadad, David A Hirshberg, Ruohan Zhan, Stefan Wager, and Susan Athey. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, 118(15), 2021.
- J Michael Harrison. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic differential systems, stochastic control theory and applications*, pages 147–186. Springer, 1988.
- J Michael Harrison and Martin I Reiman. Reflected Brownian motion on an orthant. *The Annals of Probability*, pages 302–308, 1981.
- J Michael Harrison and Nur Sunar. Investment timing with incomplete information and multiple means of learning. *Operations Research*, 63(2):442–457, 2015.
- Daniel N Hill, Houssam Nassif, Yi Liu, Anand Iyer, and SVN Vishwanathan. An efficient bandit algorithm for realtime multivariate optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1813–1821, 2017.

- Keisuke Hirano and Jack R Porter. Asymptotics for statistical treatment rules. *Econometrica*, 77(5):1683–1701, 2009.
- Keisuke Hirano and Jack R Porter. Asymptotic representations for sequential experiments. Cowles Foundation Conference on Econometrics, 2021.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Uniform, nonparametric, non-asymptotic confidence sequences. *arXiv preprint arXiv:1810.08240*, 2018.
- Donald L Iglehart and Ward Whitt. Multiple channel queues in heavy traffic. I. *Advances in Applied Probability*, 2(1):150–177, 1970.
- Ioannis Karatzas and Steven E Shreve. *Brownian Motion and Stochastic Calculus*. Springer, 2005.
- Maximilian Kasy and Anja Sautmann. Adaptive treatment assignment in experiments for policy choice. *Econometrica*, 89(1):113–132, 2021.
- Maximilian Kasy and Alexander Teytelboym. Adaptive targeted infectious disease testing. *Oxford Review of Economic Policy*, 36(Supplement_1):S77–S93, 2020.
- FP Kelly and CN Laws. Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *Queueing systems*, 13(1):47–86, 1993.
- Toru Kitagawa and Aleksey Tetenov. Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616, 2018.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. *arXiv preprint arXiv:1902.00470*, 2019.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Lucien M Le Cam. Limits of experiments. In Elizabeth L Scott Lucien M Le Cam, Jerzy Neyman, editor, *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability*, volume 6, pages 245–261. University of California Press Berkeley–Los Angeles, 1972.
- R Duncan Luce. *Individual choice behavior: A theoretical analysis*. John Wiley & Son, 1959.
- Alexander R Luedtke and Antoine Chambaz. Faster rates for policy learning. *arXiv preprint arXiv:1704.06431*, 2017.
- Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.
- Xinkun Nie, Xiaoying Tian, Jonathan Taylor, and James Zou. Why adaptively collected data have negative bias and how to correct for it. In *International Conference on Artificial Intelligence and Statistics*, pages 1261–1269. PMLR, 2018.
- Martin I Reiman. Open queueing networks in heavy traffic. *Mathematics of operations research*, 9(3):441–458, 1984.

- Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- Daniel Russo. Simple Bayesian algorithms for best-arm identification. *Operations Research*, 68(6), 2020.
- Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of Thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016.
- Daniel Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. *Foundations and Trends in Machine Learning*, 11(1):1–96, 2018.
- Jaehyeok Shin, Aaditya Ramdas, and Alessandro Rinaldo. On the bias, risk and consistency of sample means in multi-armed bandits. *arXiv preprint arXiv:1902.00746*, 2019.
- David Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer Science & Business Media, 1985.
- Daniel W Stroock and SR Srinivasa Varadhan. *Multidimensional diffusion processes*. Springer, 2007.
- William 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.
- Abraham Wald. *Sequential analysis*. John Wiley & Sons, New York, 1947.
- Zhengli Wang and Stefanos Zenios. Adaptive design of clinical trials: A sequential learning approach. *Available at SSRN*, 2020.
- Kuang Xu and Se-Young Yun. Reinforcement with fading memories. *Mathematics of Operations Research*, 45(4):1258–1288, 2020.
- Kelly Zhang, Lucas Janson, and Susan Murphy. Inference for batched bandits. In H Larochelle, M Ranzato, R Hadsell, M F Balcan, and H Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9818–9829. Curran Associates, Inc., 2020.

A Technical Lemmas

We will use the following technical lemmas repeatedly.

Lemma 18 (Gaussian Tail Bounds). *For all $x < -\sqrt{2\pi/(9-2\pi)}$:*

$$\Phi(x) \leq \frac{1}{|x|} \exp(-x^2/2), \quad \Phi(x) \geq \frac{1}{3|x|} \exp(-x^2/2). \quad (\text{A.1})$$

This immediately implies that for all $x > \sqrt{2\pi/(9-2\pi)}$:

$$\Phi(x) \leq 1 - \frac{1}{3x} \exp(-x^2/2), \quad \Phi(x) \geq 1 - \frac{1}{x} \exp(-x^2/2). \quad (\text{A.2})$$

Proof. For the lower bound, we have that for all $x < 0$:

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-x}^{\infty} \exp(-s^2/2) ds \stackrel{(a)}{\leq} \frac{1}{\sqrt{2\pi}} \int_{-x}^{\infty} \frac{-x}{s} \exp(-s^2/2) ds < \frac{1}{|x|} \exp(x^2/2),$$

where (a) follows from the fact that $-x/s \leq 1$ for all $s \geq -x$. For the upper bound, define $f(x) = \frac{x}{\sqrt{2\pi(x^2+1)}} \exp(-x^2/2) - \Phi(-x)$. We have that $f(0) = -\Phi(0) < 0$, $\lim_{x \rightarrow \infty} f(x) = 0$, and

$$f'(x) = \frac{1}{\sqrt{2\pi}(1+x^2)^2} \exp(-x^2/2) > 0, \quad \forall x > 0. \quad (\text{A.3})$$

This implies that $f(x) < 0$ for all $x > 0$, which further implies that

$$\Phi(-x) \geq \frac{|x|}{\sqrt{2\pi}(x^2+1)} \exp(-x^2/2), \quad \forall x < 0. \quad (\text{A.4})$$

The claim follows by noting that $\frac{|x|}{\sqrt{2\pi}(x^2+1)} \geq \frac{1}{3|x|}$ whenever $|x| \geq \sqrt{2\pi/(9-2\pi)}$. \square

The next result is a well-known, fundamental property of the Brownian motion. Proofs can be found in standard texts (e.g., [Karatzas and Shreve, 2005, Theorem 9.23]).

Lemma 19 (Law of Iterated Logarithm). *Let W_t be a standard Brownian motion. Then, almost surely,*

$$\limsup_{t \rightarrow \infty} \frac{W_t}{\sqrt{2t \log \log t}} = \limsup_{t \downarrow 0} \frac{W_t}{\sqrt{2t \log \log(1/t)}} = 1. \quad (\text{A.5})$$

$$\liminf_{t \rightarrow \infty} \frac{W_t}{\sqrt{2t \log \log t}} = \liminf_{t \downarrow 0} \frac{W_t}{\sqrt{2t \log \log(1/t)}} = -1. \quad (\text{A.6})$$

B Additional Proofs

B.1 Proof of Lemma 12

Proof. The proof is based on Durrett [1996], and we include it here for completeness. For Δ_ϵ^n , note that for all $\epsilon > 0$,

$$\Delta_\epsilon^n(z) \leq \frac{1}{\epsilon^p} m_p^n(z). \quad (\text{B.1})$$

The convergence of (6.11) thus implies that of (6.9). For b_k^n , note that

$$|\tilde{b}_k^n(z) - b_k^n(z)| = \int_{x: |x-z| > 1} |x-z| K^n(z, dx) \leq m_p^n(z), \quad (\text{B.2})$$

where the last step follows from the assumption $p \geq 2$. We have thus proven that (6.11) and (6.13) together imply (6.8). Finally, for $a_{k,l}^n$, we have

$$\begin{aligned} |\tilde{a}_{k,l}^n(z) - a_{k,l}^n(z)| &\leq \int_{x: |x-z| > 1} |(x_k - z_k)(x_l - z_l)| K^n(z, dx) \\ &\leq \int_{x: |x-z| > 1} |x-z|^2 K^n(z, dx) \leq m_p^n(z), \end{aligned} \quad (\text{B.3})$$

whenever $p \geq 2$, where the first step follows from the Cauchy-Schwartz inequality, and the second step from the observation that $(x_k - z_k)^2 \leq |x-z|^2$ for all k . This shows (6.11) and (6.12) together imply (6.7), completing our proof. \square

B.2 Proof of Lemma 13

By the law of iterated logarithm (Lemma 19 in Appendix A), we have that there exists constant $C > 0$ such that for all large u

$$\tilde{W}_u/\sqrt{u} \leq -C\sqrt{\log \log u}, \quad a.s. \quad (\text{B.4})$$

Using the lower bound on the normal cdf:

$$\Phi(x) \geq \frac{1}{\sqrt{2\pi}} \frac{-x}{1+x^2} \exp(-x^2/2), \quad x < 0,$$

We have that for all large u , almost surely:

$$\xi(\mu, u) = 1/\Phi\left(\frac{\mu}{\sigma\sqrt{u}} + \frac{\tilde{W}_u}{\sigma\sqrt{u}}\right) \leq 1/\Phi\left(\frac{\tilde{W}_u}{\sigma\sqrt{u}}\right) \leq b_1(\log u)^{b_2}, \quad (\text{B.5})$$

where b_1 and b_2 are positive constants that do not depend on u , which further implies that for any $\alpha \in (0, 1)$

$$\int_K^\infty u^{-2} \xi(\mu, u) \leq 1/K^{1-\alpha}, \quad (\text{B.6})$$

for all large K . □

B.3 Proof of Lemma 14

Proof. The first claim follows directly from the definitions of g_1 and \mathcal{E}_2 . In particular, for all $t \geq \nu$, $t \leq 1$, and under \mathcal{E}_2 ,

$$g_1(t) \geq \frac{\sqrt{Q_{2,t}}}{\sqrt{t}} \geq \sqrt{Q_{2,t}} \geq \delta^{-\alpha/2}. \quad (\text{B.7})$$

The second claim follows from the law of iterated logarithm (Lemma 19) applied to $W_{2,\cdot}$ and the fact that $Q_{1,t} \leq t$. □

B.4 Proof of Lemma 15

Proof. The proof of the result will use the law of iterated logarithm of Brownian motion, along with the bounds we have developed for the functions g_1 and g_2 in Lemma 14. We have that there exist constants b_1, \dots, b_4 and c_1, \dots, c_6 , such that for all sufficiently large

M :

$$\begin{aligned}
& \int_M^{1/Q_{1,v}} u^{-2} \xi(\delta, u) du \\
&= \int_M^{1/Q_{1,v}} u^{-2} / \Phi \left(g_1(\tau(u^{-1})) \left(\frac{\delta}{\sqrt{u}} + \frac{\tilde{W}_u}{\sqrt{u}} \right) + g_2(\tau(u^{-1})) \right) du \\
&\stackrel{(a)}{\leq} \int_M^\infty u^{-2} / \Phi \left(- \left| \frac{\tilde{W}_u}{\sqrt{u}} \right| - c_1 \sqrt{\log \log \delta} \right) du \\
&\stackrel{(b)}{\leq} \int_M^\infty u^{-2} / \Phi \left(-c_3 (\sqrt{\log \log u} + \sqrt{\log \log \delta}) \right) du \\
&\stackrel{(c)}{\leq} c_4 \int_M^\infty u^{-2} \left(\sqrt{\log \log u} + \sqrt{\log \log \delta} \right) \exp \left(c_3^2 \left(\sqrt{\log \log u} + \sqrt{\log \log \delta} \right)^2 \right) du \\
&\stackrel{(d)}{\leq} c_4 \int_M^\infty u^{-2} \left(\sqrt{\log \log u} + \sqrt{\log \log \delta} \right) (\log u)^{b_1} (\log \delta)^{b_2} du \\
&\leq c_5 (\log \delta)^{b_3} \int_M^\infty u^{-2} (\log u)^{b_4} du \\
&\leq c_6 (\log \delta)^{b_3} M^{1-\beta}.
\end{aligned} \tag{B.8}$$

where the various steps are based on the following facts:

- (a): $g_1(t) \leq 1$, $|g_2(t)| \leq c_1 \sqrt{\log \log \delta}$ for all $t \geq v$ and that $\tau(u^{-1}) \geq v$ when $u \leq 1/Q_{1,u}$.
- (b): Law of iterated logarithm applied to \tilde{W} .
- (c): Lower bound on the normal cdf (Lemma 18 in Appendix A): $\Phi(x) \geq \frac{1}{\sqrt{2\pi}} \frac{-x}{1+x^2} \exp(-x^2/2)$, for $x < 0$.
- (d): Cauchy-Schwartz.

This proves our claim. □

B.5 Proof of Lemma 16

Proof. From (6.90), we have that

$$\begin{aligned}
\frac{d}{dt} Q_{2,t} &= \Phi \left(g_1(t) \left(\frac{-\delta Q_{2,t} - W_{2,Q_{2,t}}}{\sqrt{Q_{2,t}}} \right) + g_2(t) \right) \\
&\leq \Phi \left(-g_1(t) \sqrt{Q_{2,t}} \delta + \left| \frac{W_{2,Q_{2,t}}}{\sqrt{Q_{2,t}}} \right| + g_2(t) \right) \\
&\leq \Phi \left(-g_1(t) \sqrt{Q_{2,t}} \delta + C \sqrt{\log \log \delta} \right),
\end{aligned} \tag{B.9}$$

where the first inequality follows from the fact that $g_1(t) \leq 1$. The second follows from applying the law of iterated logarithm to $\left| \frac{W_{2,Q_{2,t}}}{\sqrt{Q_{2,t}}} \right|$, the fact that $Q_{2,t} \geq \delta^{-\alpha}$ by the definition of τ_δ^* , and the upper bound on $|g_2(t)|$ from Lemma 14. □

B.6 Proof of Lemma 17

Proof. Define the function

$$h(t) = \frac{Q_{2,t}}{t}. \quad (\text{B.10})$$

We have

$$g_1(t) = \sqrt{\frac{Q_{1,t}}{t}} = \sqrt{1 - h(t)}. \quad (\text{B.11})$$

Our goal is to show that $h(t)$ stays below $2/3$ for all $t \in [\tau_\delta^*, 1)$. First, we verify that

$$h(\tau_\delta^*) \leq 1/2. \quad (\text{B.12})$$

This is true under the condition of \mathcal{E}_1 , by noting that $Q_{1,\tau_\delta^*} \geq Q_{2,\tau_\delta^*}$ and $Q_{1,t} + Q_{2,t} = t$, and hence

$$h(\tau_\delta^*) = \frac{Q_{2,\tau_\delta^*}}{Q_{1,\tau_\delta^*} + Q_{2,\tau_\delta^*}} \leq 1/2. \quad (\text{B.13})$$

Suppose, for the sake of contradiction that $h(t_0) = 1/2$ for some $t_0 \geq \tau_\delta^*$. In light of (B.12), it suffices to show that $h'(t_0) < 0$, which would imply that h will not increase beyond $1/2$ to ever reach $2/3$. We have that

$$h'(t) = \frac{\frac{d}{dt}Q_{2,t}}{t} - \frac{Q_{2,t}}{t} \frac{1}{t} = \frac{\frac{d}{dt}Q_{2,t} - h(t)}{t}. \quad (\text{B.14})$$

Evaluating this derivative at $t = t_0$ yields

$$h'(t_0) = \frac{\frac{d}{dt}Q_{2,t} - h(t_0)}{t_0} = \frac{\frac{d}{dt}Q_{2,t_0} - 1/2}{t_0}. \quad (\text{B.15})$$

It thus suffices to show that $\frac{d}{dt}Q_{2,t_0} < 1/2$. To this end, note that by the definition of t_0 ,

$$h(t) \leq 1/2, \quad \forall t \in [\tau_\delta^*, t_0]. \quad (\text{B.16})$$

By Lemma 16, we have that for all large δ and $t \in [\tau_\delta^*, t_0)$,

$$\begin{aligned} \frac{d}{dt}Q_{2,t} &\leq \Phi \left(-g_1(t)\sqrt{Q_{2,t}}\delta + C\sqrt{\log \log \delta} \right) \\ &= \Phi \left(-\sqrt{1 - h(t)}\sqrt{Q_{2,t}}\delta + C\sqrt{\log \log \delta} \right) \\ &\stackrel{(a)}{\leq} \Phi \left(-\sqrt{\frac{Q_{2,t}}{2}}\delta + C\sqrt{\log \log \delta} \right) \\ &\stackrel{(b)}{\leq} \Phi \left(-\delta^{1-\alpha/2}/\sqrt{2} + C\sqrt{\log \log \delta} \right) \\ &< 1/2, \end{aligned} \quad (\text{B.17})$$

where step (a) follows from (B.16) and (b) from the fact that $Q_{2,t} \geq \delta^{-\alpha}$ for all $t \geq \tau_\delta^*$. The final inequality follows from the fact that $\alpha < 2$ and hence $\delta^{1-\alpha/2} \gg \sqrt{\log \log \delta}$ for all large δ . This proves the claim. \square