

Regret Analysis of the Finite-Horizon Gittins Index Strategy for Multi-Armed Bandits

Tor Lattimore

*Department of Computing Science
University of Alberta
Edmonton, Canada*

TOR.LATTIMORE@GMAIL.COM

Editors: Vitaly Feldman, Alexander Rakhlin and Ohad Shamir

Abstract

I prove near-optimal frequentist regret guarantees for the finite-horizon Gittins index strategy for multi-armed bandits with Gaussian noise and prior. Along the way I derive finite-time bounds on the Gittins index that are asymptotically exact and may be of independent interest. I also discuss computational issues and present experimental results suggesting that a particular version of the Gittins index strategy is an improvement on existing algorithms with finite-time regret guarantees such as UCB and Thompson sampling.

1. Introduction

The stochastic multi-armed bandit is a classical problem in sequential optimisation that captures a particularly interesting aspect of the dilemma faced by learning agents. How to explore an uncertain world, while simultaneously exploiting available information? Since [Robbins \(1952\)](#) popularised the problem there have been two main solution concepts. The first being the Bayesian approach developed by [Bradt et al. \(1956\)](#), [Gittins \(1979\)](#) and others, where research has primarily focussed on characterising optimal solutions. The second approach is frequentist, with the objective of designing policies that minimise various forms of regret ([Lai and Robbins, 1985](#)). The purpose of this article is to prove frequentist regret guarantees for popular Bayesian or near-Bayesian algorithms, which explicates the strong empirical performance of these approaches observed by [Kaufmann et al. \(2012a\)](#) and others.

In each round the learner chooses an arm $I_t \in \{1, \dots, d\}$ based on past observations and receives a Gaussian reward $X_t \sim \mathcal{N}(\mu_{I_t}, 1)$ where $\mu \in \mathbb{R}^d$ is the vector of unknown means. A strategy is a method for choosing I_t and is denoted by π . The performance of a particular strategy π will be measured in terms of the expected regret, which measures the difference between the expected cumulative reward of the optimal strategy that knows μ and the expected rewards of π . Let $\mu^* = \max_i \mu_i$ be the mean reward for the optimal arm, then the expected regret up to horizon n is defined by

$$R_\mu^\pi(n) = \mathbb{E} \left[\sum_{t=1}^n (\mu^* - \mu_{I_t}) \right], \quad (1)$$

where the expectation is taken with respect to the uncertainty in the rewards and any randomness in the strategy. If Q is a probability measure on \mathbb{R}^d , then the Bayesian regret is the expectation of

Eq. (1) with respect to the prior Q .

$$BR_Q^\pi(n) = \mathbb{E}_{\theta \sim Q} \left[\mathbb{E} \left[\sum_{t=1}^n (\mu^* - \mu_{I_t}) \middle| \mu = \theta \right] \right]. \quad (2)$$

I assume that Q is Gaussian with diagonal covariance matrix $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$.

A famous non-Bayesian strategy is UCB (Katehakis and Robbins, 1995; Agrawal, 1995; Auer et al., 2002), which chooses $I_t = t$ for rounds $t \in \{1, \dots, d\}$ and subsequently maximises an upper confidence bound.

$$I_t = \arg \max_i \hat{\mu}_i(t-1) + \sqrt{\frac{\alpha \log t}{T_i(t-1)}},$$

where $\hat{\mu}_i(t-1)$ is the empirical estimate of the return of arm i based on samples from the first $t-1$ rounds and $T_i(t-1)$ is the number of times that arm has been chosen. For $\alpha > 2$ and any choice of μ it can be shown that

$$R_\mu^{\text{UCB}}(n) = O \left(\sum_{i: \Delta_i > 0} \frac{\log(n)}{\Delta_i} + \Delta_i \right), \quad (3)$$

where $\Delta_i = \mu^* - \mu_i$ is the regret incurred by choosing arm i rather than the optimal arm. No strategy enjoying sub-polynomial regret for all mean vectors can achieve smaller asymptotic regret than Eq. (3), so in this sense the UCB strategy is optimal (Lai and Robbins, 1985).

The Bayesian strategy minimises Eq. (2), which appears to be a hopeless optimisation problem. A special case where it can be solved efficiently is called the one-armed bandit, which occurs when there are two arms ($d = 2$) and the expected return of the second arm is known ($\sigma_2^2 = 0$). Bradt et al. (1956) showed that the Bayesian optimal strategy involves choosing the first arm as long as its “index” is larger than the return of the second arm and thereafter choosing the second arm. The index depends only on the number of rounds remaining and the posterior distribution of the first arm and is computed by solving an optimal stopping problem. Another situation when Eq. (2) can be solved efficiently is when the horizon is infinite ($n = \infty$) and the rewards are discounted geometrically. Then Gittins (1979) was able to show that the Bayesian strategy chooses in each round the arm with maximal index. Gittins’ index is defined in the same fashion as the index of Bradt et al. (1956) but with obvious modifications to incorporate the discounting. The index has a variety of interpretations. For example, it is equal to the price per round that a rational learner should be willing to pay in order to play the arm until either the horizon or their choice of stopping time (Weber, 1992).

Gittins’ result is remarkable because it reduces the seemingly intractable problem of finding the Bayesian optimal strategy to solving an optimal stopping problem for each arm separately. In our setting, however, the horizon is finite and the rewards are not discounted, which means that Gittins’ result does not apply and the Bayesian optimality of Gittins index strategy is not preserved.¹

1. This contradicts the opposite claim made without proof by Kaufmann et al. (2012a) but counter-examples for Bernoulli noise have been known for some time and are given here for Gaussian noise in Appendix A. In fact, the Gittins index strategy is only Bayesian optimal in all generality for geometric discounting (Berry and Fristedt, 1985, §6).

Nevertheless, the finite-horizon Gittins index strategy has been suggested as a tractable heuristic for the Bayesian optimal policy by Niño-Mora (2011) and Kaufmann (2016), a claim that I support with empirical and theoretical results in the special case that $d = 2$ (with both means unknown). A brief chronological summary of the literature on multi-armed bandits as relevant for this article is given in Appendix K.

I make a number of contributions.

Contribution 1 (Section 2) Upper and lower bounds on the finite-horizon Gittins index for the Gaussian prior/noise model that match asymptotically and are near-optimal in finite time. Asymptotic approximations are known for the discounted case via an elegant embedding of the discrete stopping problem into continuous time and solving the heat equation as a model for Brownian motion (Bather, 1983; Yao, 2006). In the finite-horizon setting there are also known approximations, but again they are asymptotic in the horizon and are not suitable for regret analysis (Chernoff and Ray, 1965; Burnetas and Katehakis, 1997).

Contribution 2 (Section 3) A proof that the Gittins index strategy with the improper flat Gaussian prior enjoys finite-time regret guarantees comparable to those of UCB. There exists a universal constant $c > 0$ such that

$$R_{\mu}^{\text{Gittins}}(n) \leq c \cdot \left(\sum_{i:\Delta_i>0} \frac{\log(n)}{\Delta_i} + \Delta_i \right). \quad (4)$$

I also show that the Gittins index strategy is asymptotically order-optimal for arbitrary Gaussian priors. There exists a universal constant $c > 0$ such that for all non-degenerate Gaussian priors the Bayesian strategy satisfies

$$\limsup_{n \rightarrow \infty} \frac{R_{\mu}^{\text{Gittins}}(n)}{\log(n)} \leq c \cdot \left(\sum_{i:\Delta_i>0} \frac{1}{\Delta_i} \right).$$

While Eq. (4) depends on a particular choice of prior, the above result is asymptotic, but independent of the prior. This is yet another example of a common property of (near) Bayesian methods, which is that the effect of the prior is diluted in the limit of infinite data. A unique aspect of the bandit setting is that the strategy is collecting its own data, so although the result is asymptotic, one has to prove that the strategy is choosing the optimal arm sufficiently often to get started.

Contribution 3 (Appendix A) For the special case that there are two arms ($d = 2$) I show by a comparison to the Gittins index strategy that the fully Bayesian strategy (minimising Eq. (2)) also enjoys the two regret guarantees above.

Contribution 4 (Appendices B and C) I present a method of computing the index in the Gaussian case to arbitrary precision and with sufficient efficiency to simulate the Gittins index policy for horizons of order 10^4 . This is used to demonstrate empirically that the Gittins index strategy with a flat prior is competitive with state-of-the-art algorithms including UCB, Thompson sampling (Thompson, 1933), OCUCB (Lattimore, 2015a) and the fully Bayesian strategy. I propose an efficient approximation of the Gittins index that (a) comes with theoretical guarantees and (b) is empirically superior to Thompson sampling and UCB, and only marginally worse than OCUCB. The approximation may also lead to improvements for index-based algorithms in other settings.

2. Bounds on the Gittins Index

As previously remarked, the Gittins index depends only on the posterior mean and variance for the relevant arm and also on the number of rounds remaining in round t , which is denoted by $m = n - t + 1$. Let $\nu \in \mathbb{R}$ and $\sigma^2 \geq 0$ be the posterior mean and variance for some arm in a given round. Let $\mu \sim \mathcal{N}(\nu, \sigma^2)$ and Y_1, Y_2, \dots, Y_m be a sequence of random variables with $Y_t \sim \mathcal{N}(\mu, 1)$. Assume $\{Y_t\}_{t=1}^m$ are independent after conditioning on μ . The Gittins index for an arm with posterior mean ν and variance σ^2 and horizon $m = n - t + 1$ is given by

$$\gamma(\nu, \sigma^2, m) = \max \left\{ \gamma : \sup_{1 \leq \tau \leq m} \mathbb{E}[(\mu - \gamma)\tau] \geq 0 \right\}, \quad (5)$$

where the supremum is taken over stopping times with respect to the filtration generated by random variables Y_1, \dots, Y_m . Equivalently, $\mathbb{1}\{\tau = t\}$ must be measurable with respect to σ -algebra generated by Y_1, \dots, Y_t . I denote $\tau(\nu, \sigma^2, m)$ to be the maximising stopping time in Eq. (5) for $\gamma = \gamma(\nu, \sigma^2, m)$. The following upper and lower bound on γ is the main theorem of this section.

Theorem 1 *If $\beta \geq 1$ is chosen such that $\gamma(\nu, \sigma^2, m) = \nu + \sqrt{2\sigma^2 \log \beta}$. Then there exists a universal constant $c > 0$ such that*

$$c \min \left\{ \frac{m}{\log_+^{\frac{3}{2}}(m)}, \frac{m\sigma^2}{\log_+^{\frac{1}{2}}(m\sigma^2)} \right\} \leq \beta \leq \frac{m}{\log_+^{\frac{3}{2}}(m)},$$

where $\log_+(x) = \max\{1, \log(x)\}$.

Important Remark I use c, c' and c'' as *temporary* positive constants that change from theorem to theorem and proof to proof. A table of notation may be found in Appendix L.

The upper bound in Theorem 1 is rather trivial while the lower bound relies on a carefully constructed stopping time. As far as the asymptotics are concerned, as m tends to infinity the upper and lower bounds on β converge up to constant factors, which leads to

$$\exp \left(\left(\frac{\gamma(\nu, \sigma^2, m) - \nu}{\sqrt{2\sigma^2}} \right)^2 \right) = \Theta \left(\frac{m}{\log^{\frac{3}{2}} m} \right). \quad (6)$$

The $\log^{\frac{3}{2}}(m)$ looks bizarre, but is surprisingly well-justified as shall be seen in the next section. In order to obtain logarithmic regret guarantees it is necessary that the approximation be accurate in the exponential scale as in Eq. (6). Merely showing that $\log \beta / \log(m)$ tends to unity as m grows would not be sufficient. Empirically choosing $c = 1/4$ in the lower bound leads to an excellent approximation of the Gittins index (see Fig. 2 in Appendix C). The proof of Theorem 1 relies on a carefully chosen stopping time.

$$\tau = \min \left\{ m, \min \left\{ t \geq \frac{1}{\nu^2} : \hat{\mu}_t + \sqrt{\frac{4}{t} \log(4t\nu^2)} \leq 0 \right\} \right\}, \quad (7)$$

where $\hat{\mu}_t = \frac{1}{t} \sum_{s=1}^t Y_s$. Note that $\mathbb{1}\{\tau = t\}$ depends only on Y_1, Y_2, \dots, Y_t so this is a stopping time with respect to the right filtration.

Lemma 2 *If $\nu < 0$, then there exists a universal constant $c' \geq 2e$ such that:*

1. *If $\theta \in (-\infty, \nu]$, then $\mathbb{E}[\tau|\mu = \theta] \leq 1 + \frac{c'}{\nu^2}$.*
2. *If $\theta \in (\nu, 0)$, then $\mathbb{E}[\tau|\mu = \theta] \leq 1 + \frac{c'}{\theta^2} \log\left(\frac{e\nu^2}{\theta^2}\right)$.*
3. *If $\theta \in [0, \infty)$, then $\mathbb{E}[\tau|\mu = \theta] \geq m/2$.*

The proof of Lemma 2 may be found in Appendix D. Before the proof of Theorem 1 we need two more results. The proofs follow directly from the definition of the Gittins index and are omitted.

Lemma 3 *For all $m \geq 1$ and $\sigma^2 \geq 0$ and $\nu, \nu' \in \mathbb{R}$ we have $\gamma(\nu, \sigma^2, m) - \nu = \gamma(\nu', \sigma^2, m) - \nu'$.*

Lemma 4 *$\gamma(\nu, \sigma^2, m) \geq \nu$ for all ν, σ^2 and m .*

Proof of Theorem 1 Starting with the lower bound, by Lemma 3 it is enough to bound the Gittins index for any choice of ν . Let $\nu = -\sqrt{2\sigma^2 \log \beta}$ where $\beta \geq 1$ will be defined subsequently. I will shortly show that there exists a stopping time τ adapted to the filtration generated by Y_1, \dots, Y_m for which

$$\mathbb{E}[\mu\tau] \geq 0. \quad (8)$$

Therefore by the definition of the Gittins index we have $\gamma(\nu, \sigma^2, m) \geq 0$ and so by Lemma 3

$$(\forall \nu') \quad \gamma(\nu', \sigma^2, m) \geq \nu' + \sqrt{2\sigma^2 \log \beta}.$$

I now prove Eq. (8) holds for the stopping time given in Eq. (7). First note that if $\beta = 1$, then $\nu = 0$ and $\mathbb{E}[\mu\tau] = 0$ for the stopping time with $\tau = 1$ surely. Now assume $\beta > 1$ and define $P = \mathcal{N}(\nu, \sigma^2)$ to be a Gaussian measure on \mathbb{R} . Then

$$\mathbb{E}[\mu\tau] = \mathbb{E}_{\theta \sim P}[\theta \mathbb{E}[\tau|\mu = \theta]] = \int_0^\infty \theta \mathbb{E}[\tau|\mu = \theta] dP(\theta) + \int_{-\infty}^0 \theta \mathbb{E}[\tau|\mu = \theta] dP(\theta). \quad (9)$$

The integrals will be bounded separately. Define

$$\text{Erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty \exp(-y^2) dy \quad f(\beta) = \frac{1}{\beta} \sqrt{\frac{1}{2\pi}} - \sqrt{\frac{\log \beta}{2}} \text{Erfc}\left(\sqrt{\log \beta}\right).$$

A straightforward computation combined with Lemma 2 shows that

$$\int_0^\infty \theta \mathbb{E}[\tau|\mu = \theta] dP(\theta) \geq \frac{m}{2} \int_0^\infty \theta dP(\theta) = \frac{m\sigma}{2} f(\beta). \quad (10)$$

The function f satisfies $f(\beta)\beta \log_+(\beta) = \Theta(1)$ with the precise statement given in Lemma 22 in Appendix J. Let $W : [0, \infty) \rightarrow \mathbb{R}$ be the product logarithm, defined implicitly such that $x = W(x) \exp(W(x))$. For the second integral in Eq. (9) we can use Lemma 2 again, which for $\theta \in (\nu, 0)$ gives

$$\theta (\mathbb{E}[\tau|\mu = \theta] - 1) \geq \theta \min \left\{ m, \frac{c'}{\theta^2} \log \left(\frac{e\nu^2}{\theta^2} \right) \right\} \geq -\sqrt{c'm W \left(\frac{em\nu^2}{c'} \right)},$$

where in the first inequality we have exploited the fact that $\tau \leq m$ occurs surely and the second follows by noting that θ is negative and that θm is increasing in θ and $1/\theta \log(e\nu^2\theta^{-2})$ is decreasing. Let $\varepsilon = 1 - \sqrt{7/8}$, which is chosen such that by tail bounds on the Gaussian integral (Lemma 14):

$$\int_{\varepsilon\nu}^0 dP(\theta) \leq \exp\left(-\frac{(1-\varepsilon)^2\nu^2}{2\sigma^2}\right) = \left(\frac{1}{\beta}\right)^{\frac{7}{8}} \quad \text{and} \quad \int_{-\infty}^0 \theta dP(\theta) \geq \nu - \frac{\sigma}{\sqrt{2\pi}}.$$

Combining the above two displays with Lemma 2 we have for some universal constant $c'' > 0$

$$\begin{aligned} \int_{-\infty}^0 \theta \mathbb{E}[\tau | \mu = \theta] dP(\theta) &\geq \left(1 + \frac{c'}{\varepsilon^2\nu^2} \log\left(\frac{e}{\varepsilon^2}\right)\right) \int_{-\infty}^0 \theta dP(\theta) - \int_{\varepsilon\nu}^0 \sqrt{c'm W\left(\frac{em\nu^2}{c'}\right)} dP(\theta) \\ &\stackrel{(a)}{\geq} c'' \left(\left(1 + \frac{1}{\nu^2}\right) (\nu - \sigma) - \sqrt{m W\left(\frac{em\nu^2}{c'}\right)} \left(\frac{1}{\beta}\right)^{\frac{7}{8}} \right) \\ &\stackrel{(b)}{\geq} -c'' \left(\left(1 + \frac{1}{2\sigma^2 \log(\beta)}\right) \left(\sqrt{2\sigma^2 \log(\beta)} + \sigma\right) + \sqrt{m W(m\sigma^2 \log(\beta))} \left(\frac{1}{\beta}\right)^{\frac{7}{8}} \right), \end{aligned}$$

where in (a) I have substituted the bounds on the integrals in the previous two displays and naively chosen a large constant c'' to simplify the expression and in (b) I substitute $\nu = -\sqrt{2\sigma^2 \log(\beta)}$ and exploited the assumption that $c' \geq 2e$. Therefore by Eq. (10) and Eq. (9) we have

$$\mathbb{E}[\mu\tau] \geq \frac{m\sigma}{2} f(\beta) - c'' \left(\left(1 + \frac{1}{2\sigma^2 \log(\beta)}\right) \left(\sqrt{2\sigma^2 \log(\beta)} + \sigma\right) + \sqrt{m W(m\sigma^2 \log(\beta))} \beta^{-\frac{7}{8}} \right).$$

Therefore by expanding the bracketed product, dropping the non-dominant $1/(\sigma \log(\beta))$ term and upper bounding the sum by the max there exists a universal constant $c''' \geq 1$ such that the following implication holds:

$$m\sigma f(\beta) \geq c''' \max \left\{ \sigma, \frac{1}{\sqrt{\sigma^2 \log(\beta)}}, \sqrt{\sigma^2 \log(\beta)}, \sqrt{m W(m\sigma^2 \log(\beta))} \beta^{-\frac{7}{8}} \right\} \implies \mathbb{E}[\mu\tau] \geq 0.$$

Let $c \geq 1$ be a sufficiently large universal constant and define β by

$$\beta_1 = \frac{m}{c \log_+^{\frac{3}{2}}(m)} \quad \beta_2 = \frac{m\sigma^2}{c \log_+^{\frac{1}{2}}(m\sigma^2)} \quad \beta = \begin{cases} \min\{\beta_1, \beta_2\} & \text{if } \min\{\beta_1, \beta_2\} \geq 3 \\ 1 & \text{otherwise} \end{cases}$$

If $\beta \geq 3$, then the inequality leading to the implication is shown tediously by applying Lemma 22 to approximate $f(\beta)$ (see Appendix J for the gory algebraic details). If $\beta = 1$, then $\nu = 0$ and $\mathbb{E}[\mu\tau] \geq 0$ is trivial for the stopping time $\tau = 1$ surely. Finally we have the desired lower bound by Lemma 3,

$$\gamma(\nu', \sigma^2, m) \geq \nu' + \sqrt{2\sigma^2 \log(\beta)} \geq \nu' + \sqrt{2\sigma^2 \log \left(\frac{1}{3c} \min \left\{ \frac{m}{\log_+^{\frac{3}{2}}(m)}, \frac{m\sigma^2}{\log_+^{\frac{1}{2}}(m\sigma^2)} \right\} \right)}.$$

For the upper bound we can proceed naively and exploit the fact that the stopping time must be chosen such that $1 \leq \tau \leq m$ surely. This time choose ν such that $\gamma(\nu, \sigma^2, m) = 0$ and let $\beta \geq 1$ be such that $\nu = -\sqrt{2\sigma^2 \log(\beta)}$, which is always possible by Lemma 4. Then

$$0 = \sup_{\tau} \int_{-\infty}^{\infty} \theta \mathbb{E}[\tau | \mu = \theta] dP(\theta) \leq m \int_0^{\infty} \theta dP(\theta) + \int_{-\infty}^{\nu} \theta dP(\theta) = m\sigma f(\beta) - \frac{\sigma}{\sqrt{2\pi}} + \frac{\nu}{2}.$$

Rearranging and substituting the definition of ν and applying Lemma 22 in the appendix leads to

$$\frac{m}{\sqrt{8\pi\beta \log \beta}} \geq mf(\beta) \geq \frac{1}{\sqrt{2\pi}} + \sqrt{\frac{1}{2} \log \beta} \geq \sqrt{\frac{1}{2} \log \beta}.$$

Naive simplification leads to $\beta \leq m / \log^{\frac{3}{2}}(m)$ as required. ■

3. Regret Bounds for the Gittins Index Strategy

I start with the finite-time guarantees. Before presenting the algorithm and analysis we need some additional notation. The empirical mean of arm i after round t is denoted by

$$\hat{\mu}_i(t) = \frac{1}{T_i(t)} \sum_{s=1}^t \mathbb{1}\{I_s = i\} X_s \quad T_i(t) = \sum_{s \leq t} \mathbb{1}\{I_s = i\}.$$

I will avoid using $\hat{\mu}_i(t)$ for rounds t when $T_i(t) = 0$, so this quantity will always be well-defined. Let $\Delta_{\max} = \max_i \Delta_i$ and $\Delta_{\min} = \min \{\Delta_i : \Delta_i > 0\}$. The Gittins index strategy for a flat Gaussian prior is summarised in Algorithm 1. Since the flat prior is improper, the Gittins index is not initially defined. For this reason the algorithm chooses $I_t = t$ in rounds $t \in \{1, \dots, d\}$ after which the posterior has unit variance for all arms and the posterior mean is $\hat{\mu}_i(d)$. An alternative interpretation of this strategy is the limiting strategy as the prior variance tends to infinity for all arms.

Input: d, n
 Choose each arm once
for $t \in d + 1, \dots, n$ **do**
 | Choose $I_t = \arg \max_i \gamma(\hat{\mu}_i(t-1), T_i(t-1)^{-1}, n-t+1)$
end

Algorithm 1: Gittins strategy with flat Gaussian prior

Theorem 5 *Let π be the strategy given in Algorithm 1. Then there exist universal constants $c, c' > 0$ such that*

$$R_{\mu}^{\pi}(n) \leq c \sum_{i: \Delta_i > 0} \left(\frac{\log(n)}{\Delta_i} + \Delta_i \right) \quad R_{\mu}^{\pi}(n) \leq c' \left(\sqrt{dn \log(n)} + \sum_{i=1}^d \Delta_i \right)$$

While the problem dependent bound is asymptotically optimal up to constant factors, the problem independent bound is sub-optimal by a factor of $\sqrt{\log(n)}$ with both the MOSS algorithm by [Audibert and Bubeck \(2009\)](#) and OCUCB by [Lattimore \(2015a\)](#) matching the lower bound of $\Omega(\sqrt{dn})$

given by [Auer et al. \(1995\)](#). The main difficulty in proving [Theorem 5](#) comes from the fact that the Gittins index is smaller than the upper confidence bound used by UCB. This is especially true as t approaches the horizon when the Gittins index tends towards the empirical mean, while the UCB index actually grows. The solution is (a) to use very refined concentration guarantees and (b) to show that the Gittins strategy chooses near-optimal arms sufficiently often for $t \leq n/2$, ensuring that the empirical mean of these arms is large enough that bad arms are not chosen too often in the second half. Before the proof we require some additional definitions and lemmas. Assume for the remainder of this section that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_d$, which is non-restrictive, since if this is not the case, then the arms can simply be re-ordered. Let F be the event that there exists a $t \in \{1, \dots, n\}$ and $i \in \{1, \dots, d\}$ such that

$$|\hat{\mu}_i(t) - \mu_i| \geq \sqrt{\frac{2}{T_i(t)} \log(dn^2)}.$$

The index of the i th arm in round t is abbreviated to $\gamma_i(t) = \gamma(\hat{\mu}_i(t-1), T_i(t-1)^{-1}, n-t+1)$, which means that for rounds $t > d$ [Algorithm 1](#) is choosing $I_t = \arg \max_i \gamma_i(t)$. Define random variable $Z = \mu_1 - \min_{1 \leq t \leq n/2} \gamma_1(t)$, which measures how far below μ_1 the index of the first arm may fall some time in the first $n/2$ rounds. For each $i \in \{1, \dots, d\}$ define

$$u_i = \left\lceil \frac{32}{\Delta_i^2} \log(2dn^2) \right\rceil. \quad (11)$$

Now we are ready for the lemmas. First is the key concentration inequality that controls the probability that the Gittins index of the optimal arm drops far below the true mean. It is in the proof of [Theorem 6](#) that the odd-looking powers of the logarithmic terms in the bounds on the Gittins index are justified. Any higher power would lead to super-logarithmic regret, while a lower power could potentially lead to a sub-optimal trade-off between failure probability and exploration.

Lemma 6 *Let $c > 0$ be a universal constant, and $\Delta > 0$ and Y_1, Y_2, \dots be a sequence of i.i.d. random variables with $Y_t \sim \mathcal{N}(0, 1)$ and $S_t = \sum_{s=1}^t Y_s$. Then there exist universal constants $c' > 0$ and $n_0 \in \mathbb{N}$ such that whenever $n \geq n_0$.*

$$\mathbb{P} \left\{ \exists t : S_t \geq t\Delta + \max \left\{ 0, \sqrt{2t \log \left(\frac{cn}{2 \log_+^{\frac{3}{2}}(n/2)} \right)} \right\} \right\} \leq c' \cdot \frac{\log(n)}{n\Delta^2}$$

$$\mathbb{P} \left\{ \exists t : S_t \geq t\Delta + \max \left\{ 0, \sqrt{2t \log \left(\frac{cn}{2t \log_+^{\frac{1}{2}}(\frac{n}{2t})} \right)} \right\} \right\} \leq c' \cdot \frac{\log_+(n\Delta^2)}{n\Delta^2}.$$

The proof of [Lemma 6](#) follows by applying a peeling device and may be found in [Appendix F](#). Note that similar results exist in the literature. For example, by [Audibert and Bubeck \(2009\)](#); [Perchet and Rigollet \(2013\)](#) and presumably others. What is unique here is that the improved concentration guarantees for Gaussians are also being exploited.

Assumption 7 *Assume that $n \geq n_0$, which is non-restrictive since n_0 is a universal constant and [Theorem 5](#) holds trivially with $c = n_0$ for $n \leq n_0$.*

Lemma 8 *There exists a universal constant $c > 0$ such that for all $\Delta > 0$ we have*

$$\mathbb{P}\{Z \geq \Delta\} \leq c \cdot \frac{\log(n) + \log_+(n\Delta^2)}{n\Delta^2}.$$

Proof Apply Theorem 1 and Theorem 6 and the fact that $m = n - t + 1 \geq n/2$ for $t \leq n/2$. ■

Lemma 9 $\mathbb{E}[T_i(n)] \leq n/i$.

Proof The result follows from the assumption that $\mu_1 \geq \dots \geq \mu_d$, the definition of the algorithm, the symmetry of the Gaussian density and because the exploration bonus due to Gittins index is shift invariant (Lemma 3). ■

Lemma 10 $\mathbb{P}\{F\} \leq 1/n$.

Proof For fixed $T_i(t) = u$ apply the standard tail inequalities for the Gaussian (Lemma 14) to bound $\mathbb{P}\{|\hat{\mu}_i(t) - \mu_i| \geq \sqrt{2 \log(2dn^2)/u}\} \leq 1/(dn^2)$. The result is completed by applying the union bound over all arms and values of $u \in \{1, \dots, n\}$. ■

Lemma 11 *If F does not hold and i is an arm such that $Z < \Delta_i/2$, then $T_i(n/2) \leq u_i$.*

Proof Let $t \leq n/2$ be some round such that $T_i(t-1) = u_i$. Then by Theorem 1 and the definitions of F and u_i we have

$$\begin{aligned} \gamma_i(t) &\leq \hat{\mu}_i(t) + \sqrt{\frac{2}{u_i} \log(n)} \leq \mu_i + \sqrt{\frac{2}{u_i} \log(2dn^2)} + \sqrt{\frac{2}{u_i} \log(n)} \\ &\leq \mu_i + \frac{\Delta_i}{2} = \mu_1 - \frac{\Delta_i}{2} < \mu_1 - Z \leq \gamma_1(t). \end{aligned}$$

Therefore $I_t \neq i$ and so $T_i(n/2) \leq u_i$. ■

Proof of Theorem 5 The regret may be re-written as

$$R_\mu^\pi(n) = \sum_{i=1}^d \Delta_i \mathbb{E}[T_i(n)]. \quad (12)$$

To begin, let us naively bound the regret for nearly-optimal arms. Define a set $C \subset \{1, \dots, d\}$ by

$$C = \left\{ i : 0 < \Delta_i \leq 4\sqrt{\frac{8i}{n} \log(2dn^2)} \right\} = \left\{ i : 0 < \Delta_i \leq \frac{128i}{n\Delta_i} \log(2dn^2) \right\}.$$

Therefore by Lemma 9 we have

$$\sum_{i \in C} \Delta_i \mathbb{E}[T_i(n)] \leq \sum_{i \in C} \frac{128}{\Delta_i} \log(2dn^2). \quad (13)$$

For arms not in C we need to do some extra work. By Lemma 10

$$\begin{aligned} \sum_{i \notin C} \Delta_i \mathbb{E}[T_i(n)] &\leq \mathbb{P}\{F\} n \Delta_{\max} + \mathbb{E} \left[\mathbb{1}\{\neg F\} \sum_{i \notin C} \Delta_i T_i(n) \right] \\ &\leq \Delta_{\max} + \mathbb{E} \left[\mathbb{1}\{\neg F\} \sum_{i \notin C} \Delta_i T_i(n) \right]. \end{aligned} \quad (14)$$

From now on assume F does not hold while bounding the second term. By Lemma 11, if $Z < \Delta_i/2$, then $T_i(n/2) \leq u_i$. Define disjoint (random) sets $A, B \subseteq \{1, \dots, d\}$ by

$$B = \left\{ i : Z < \Delta_i/2 \text{ and } \sum_{j \geq i} u_j \leq \frac{n}{4} \right\} \quad \text{and} \quad A = \{1, \dots, d\} - B.$$

The set A is non-empty because $u_1 = \infty$, which implies that $1 \in A$. Let $i = \max A$, which satisfies either $Z \geq \Delta_i/2$ or $n \leq 4 \sum_{j \geq i} u_j$. Therefore

$$\begin{aligned} \sum_{k \in A} \Delta_k T_k(n) &\leq n \Delta_i \leq \max \left\{ 2nZ \mathbb{1}\{\Delta_{\min}/2 \leq Z\}, 4\Delta_i \sum_{j \geq i} u_j \right\} \\ &\leq 2nZ \mathbb{1}\{\Delta_{\min}/2 \leq Z\} + 4 \sum_{j: \Delta_j > 0} \Delta_j u_j. \end{aligned} \quad (15)$$

The next step is to show that there exists an arm in A that is both nearly optimal and has been chosen sufficiently often that its empirical estimate is reasonably accurate. This arm can then be used to show that bad arms are not chosen too often. From the definition of B and because F does not hold we have $\sum_{i \in A} T_i(n/2) = n/2 - \sum_{i \in B} T_i(n/2) \geq n/2 - \sum_{i \in B} u_i \geq n/4$. Therefore there exists an $i \in A$ such that $T_i(n/2) \geq n/(4|A|)$. Suppose $j \notin A \cup C$ and $\Delta_j \geq 4\Delta_i$ and $T_j(t) = u_j$. Then

$$\begin{aligned} \gamma_j(t) &\leq \mu_j + \frac{\Delta_j}{2} = \mu_i - \frac{\Delta_j}{2} + \Delta_i \leq \mu_i - \frac{\Delta_j}{4} < \mu_i - \sqrt{\frac{8j}{n} \log(dn^2)} \\ &\leq \mu_i - \sqrt{\frac{8|A|}{n} \log(dn^2)} \leq \hat{\mu}_i(t) \leq \gamma_i(t). \end{aligned} \quad (16)$$

Therefore $I_t \neq j$ and so $T_j(n) \leq u_j$. Now we consider arms with $\Delta_j < 4\Delta_i$. By the same reasoning as in Eq. (15) we have

$$\sum_{j: \Delta_j < 4\Delta_i} \Delta_j T_j(n) \leq 4n \Delta_i \leq 8nZ \mathbb{1}\{\Delta_{\min}/2 \leq Z\} + 16 \sum_{j: \Delta_j > 0} \Delta_j u_j.$$

Finally we have done enough to bound the regret due to arms not in C . By the above display and Eq. (15) and the sentence after Eq. (16) we have

$$\sum_{i \notin C} \mathbb{1}\{\neg F\} \Delta_i T_i(n) \leq 10nZ \mathbb{1}\{\Delta_{\min}/2 \leq Z\} + 20 \sum_{j: \Delta_j > 0} \Delta_j u_j + \sum_{j \notin C} \Delta_j u_j.$$

Combining this with Eq. (13) and the regret decomposition Eq. (12) and Eq. (14) we have

$$R_\mu^\pi(n) \leq \Delta_{\max} + \sum_{j \in C} \frac{128}{\Delta_j} \log(dn^2) + 21 \sum_{j: \Delta_j > 0} \Delta_j u_j + 10n \mathbb{E}[Z \mathbb{1}\{\Delta_{\min}/2 \leq Z\}]. \quad (17)$$

From Lemma 8 there exists a universal constant $c'' > 0$ such that

$$\begin{aligned} \mathbb{E}Z \mathbb{1}\{\Delta_{\min}/2 \leq Z\} &\leq \int_{\Delta_{\min}/2}^{\infty} \mathbb{P}\{Z \geq z\} dz + \frac{\Delta_{\min}}{2} \mathbb{P}\left\{Z \geq \frac{\Delta_{\min}}{2}\right\} \\ &\leq \frac{c''(\log(n) + \log_+(n\Delta_{\min}^2))}{n\Delta_{\min}}. \end{aligned}$$

Substituting into Eq. (17) and inserting the definition of u_i and naively simplifying completes the proof of the problem dependent regret bound. To prove the second bound in Theorem 5 it suffices to note that the total regret due to arms with $\Delta_i \leq \sqrt{d/n \log(n)}$ is at most $\sqrt{nd \log(n)}$. ■

Asymptotic regret for non-flat prior The previous results relied heavily on the choice of a flat prior. For other (Gaussian) priors it is still possible to prove regret guarantees, but now with a dependence on the prior. For the sake of simplicity I switch to asymptotic analysis and show that any negative effects of a poorly chosen prior wash out for sufficiently large horizons. Let ν_i and σ_i^2 be the prior mean and variance for arm i and let $\nu_i(t)$ and $\sigma_i^2(t)$ denote the posterior mean and variance at the end of round t . A simple computation shows that

$$\nu_i(t) = \left(\frac{\nu_i}{\sigma_i^2} + \sum_{s=1}^t \mathbb{1}\{I_s = i\} X_s \right) / \left(\frac{1}{\sigma_i^2} + T_i(t) \right) \quad \sigma_i^2(t) = \left(\frac{1}{\sigma_i^2} + T_i(t) \right)^{-1}. \quad (18)$$

The strategy chooses $I_t = \arg \max_i \gamma(\nu_i(t-1), \sigma_i^2(t-1), n-t+1)$. (19)

Theorem 12 *Assume that $\sigma_i^2 > 0$ for all i . Let π_n be the strategy given in Eq. (19), then there exists a universal $c > 0$ such that for $\mu \in \mathbb{R}^d$, $\limsup_{n \rightarrow \infty} R_\mu^{\pi_n}(n) / \log n \leq c \sum_{i: \Delta_i > 0} \Delta_i^{-1}$.*

The proof may be found in Appendix G. It should be emphasised that the limit in Theorem 12 is taken over a sequence of strategies and an increasing horizon. This is in contrast to similar results for UCB where the strategy is fixed and only the horizon is increasing.

4. Discussion

I have shown that the Gittins strategy enjoys finite-time regret guarantees, which explicates the excellent practical performance. The proof relies on developing tight bounds on the index, which are asymptotically exact. If the prior is mis-specified, then the resulting Gittins index strategy is order-optimal asymptotically, but its finite-time regret may be significantly worse. Experimental results show the Gittins strategy with a flat improper Gaussian prior is never much worse and often much better than the best frequentist algorithms (please see Appendix C). The index is non-trivial to compute, but reasonable accuracy is possible for horizons of $O(10^4)$. For larger horizons I propose a new index inspired by the theoretical results that is efficient and competitive with the state-of-the-art, at least in the worst case regime for which experiments were performed. There are a variety of open problems some of which are described below.

Alternative prior and noise models The Gaussian noise/prior was chosen for its simplicity, but bounded or Bernoulli rewards are also interesting. In the latter, the Gittins index can be computed by using a Beta prior and dynamic programming. Asymptotic approximations by [Burnetas and Katehakis \(1997\)](#) suggest that one might expect to derive the KL-UCB algorithm in this manner ([Cappé et al., 2013](#)), but finite-time bounds would be required to obtain regret guarantees. Many of the concentration inequalities used in the Gaussian case have information-theoretic analogues, so in principle I expect that substituting these results into the current technique should lead to good results ([Garivier, 2013](#)).

Finite-time impact of mis-specified prior For non-flat priors I only gave asymptotic bounds on the regret, showing that the Gittins index strategy will eventually recover from even the most poorly mis-specified prior. Of course it would be nice to fully understand how long the algorithms takes to recover by analysing its finite-time regret. Some preliminary work has been done on this question in a simplified setting for Thompson sampling by [Liu and Li \(2015\)](#). Unfortunately the results will necessarily be quite negative. If a non-flat prior is chosen in such a way that the resulting algorithm achieves unusually small regret with respect to a particular arm, then the regret it incurs on the remaining arms must be significantly larger ([Lattimore, 2015b](#)). In short, there is a large price to pay for favouring one arm over another and the Bayesian algorithm cannot save you. This is in contrast to predictive settings where a poorly chosen prior is quickly washed away by data.

Asymptotic optimality Theorem 12 shows that the Gittins index strategy is eventually *order-optimal* for any choice of prior, but the leading constant does not match optimal rate. The reason is that the failure probability for which the algorithm suffers linear regret appears to be $O(\log(n)/n)$ and so this term is not asymptotically insignificant. In contrast, the UCB confidence level is chosen such that the failure probability is $O(1/n)$, which is insignificant for large n .

Horizon effects Recent bandit algorithms including MOSS ([Audibert and Bubeck, 2009](#)) and OCUCB ([Lattimore, 2015a](#)) exploit the knowledge of the horizon. The Gittins strategy also exploits this knowledge, and it is not clear how it could be defined without a horizon (the index tends to infinity as n increases). The most natural approach would be to choose a prior on the unknown horizon, but what prior should you choose and is there hope to compute the index in that case?² You can also ask what is the benefit of knowing the horizon? The exploration bonus of the Gittins strategy (and MOSS and OCUCB) tend to zero as the horizon approaches, which makes these algorithms more aggressive than anytime algorithms such as UCB and Thompson sampling and improves practical performance.

Extending the model and computation issues For finite-armed bandits and simple priors the index may be approximated in polynomial time, with practical computation possible to horizons of $\sim 10^4$. It is natural to ask whether or not the computation techniques can be improved to compute the index for larger horizons. I am also curious to know if the classical results can be extended to more complicated settings such as linear bandits or partial monitoring. This has already been done to some extent (eg., restless bandits), but there are many open problems. The recent book by [Gittins et al. \(2011\)](#) is a broad reference for existing extensions.

2. An exponential prior leads to the discounted setting for which anytime regret guarantees are unlikely to exist, but where computation is efficient. A power-law would be a more natural choice, but the analysis becomes very non-trivial in that case.

Acknowledgements

My thanks to Marcus Hutter for several useful suggestions and to Dan Russo for pointing out that the finite-horizon Gittins strategy is not generally Bayesian optimal. The experimental component of this research was enabled by the support provided by WestGrid (www.westgrid.ca) and Compute Canada (www.computecanada.ca).

References

- Rajeev Agrawal. Sample mean based index policies with $O(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability*, pages 1054–1078, 1995.
- Shipra Agrawal and Navin Goyal. Further optimal regret bounds for Thompson sampling. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.
- Jean-Yves Audibert and Sébastien Bubeck. Minimax policies for adversarial and stochastic bandits. In *Proceedings of Conference on Learning Theory (COLT)*, pages 217–226, 2009.
- Peter Auer, Nicoló Cesa-Bianchi, Yoav Freund, and Robert E Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*, pages 322–331. IEEE, 1995.
- Peter Auer, Nicoló Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- John Bather. Optimal stopping of Brownian motion: A comparison technique. *Recent Advances in Statistics*, pages 19–50, 1983.
- Donald Berry and Bert Fristedt. *Bandit problems : sequential allocation of experiments*. Chapman and Hall, 1985.
- Dimitri P Bertsekas and John N Tsitsiklis. Neuro-dynamic programming: an overview. In *Decision and Control, 1995., Proceedings of the 34th IEEE Conference on*, volume 1, pages 560–564. IEEE, 1995.
- Russell N Bradt, SM Johnson, and Samuel Karlin. On sequential designs for maximizing the sum of n observations. *The Annals of Mathematical Statistics*, pages 1060–1074, 1956.
- Apostolos N Burnetas and Michael N Katehakis. On the finite horizon one-armed bandit problem. *Stochastic Analysis and Applications*, 16(1):845–859, 1997.
- Olivier Cappé, Aurélien Garivier, Odalric-Ambrym Maillard, Rémi Munos, and Gilles Stoltz. Kullback–Leibler upper confidence bounds for optimal sequential allocation. *The Annals of Statistics*, 41(3):1516–1541, 2013.
- Jhelum Chakravorty and Aditya Mahajan. Multi-armed bandits, Gittins index, and its calculation. *Methods and Applications of Statistics in Clinical Trials: Planning, Analysis, and Inferential Methods, Volume 2*, pages 416–435, 2013.
- Herman Chernoff and SN Ray. A bayes sequential sampling inspection plan. *The Annals of Mathematical Statistics*, pages 1387–1407, 1965.

- Aurélien Garivier. Informational confidence bounds for self-normalized averages and applications. *arXiv preprint arXiv:1309.3376*, 2013.
- John Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 41(2):148–177, 1979.
- John Gittins, Kevin Glazebrook, and Richard Weber. *Multi-armed bandit allocation indices*. John Wiley & Sons, 2011.
- Michael N Katehakis and Herbert Robbins. Sequential choice from several populations. *Proceedings of the National Academy of Sciences of the United States of America*, 92(19):8584, 1995.
- Emilie Kaufmann. On bayesian index policies for sequential resource allocation. *arXiv preprint arXiv:1601.01190*, 2016.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On Bayesian upper confidence bounds for bandit problems. In *International Conference on Artificial Intelligence and Statistics*, pages 592–600, 2012a.
- Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In NaderH. Bshouty, Gilles Stoltz, Nicolas Vayatis, and Thomas Zeugmann, editors, *Algorithmic Learning Theory*, volume 7568 of *Lecture Notes in Computer Science*, pages 199–213. Springer Berlin Heidelberg, 2012b. ISBN 978-3-642-34105-2. doi: 10.1007/978-3-642-34106-9_18.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Tor Lattimore. Optimally confident UCB : Improved regret for finite-armed bandits. Technical report, 2015a. URL <http://arxiv.org/abs/1507.07880>.
- Tor Lattimore. The pareto regret frontier for bandits. In *Proceedings of the 28th Conference on Neural Information Processing Systems (NIPS)*, 2015b.
- Hans R Lerche. *Boundary crossing of Brownian motion: Its relation to the law of the iterated logarithm and to sequential analysis*. Springer, 1986.
- Che-Yu Liu and Lihong Li. On the prior sensitivity of Thompson sampling. *arXiv preprint arXiv:1506.03378*, 2015.
- José Niño-Mora. Computing a classic index for finite-horizon bandits. *INFORMS Journal on Computing*, 23(2):254–267, 2011.
- Vianney Perchet and Philippe Rigollet. The multi-armed bandit problem with covariates. *The Annals of Statistics*, 41(2):693–721, 2013.
- Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- William 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.

John N Tsitsiklis. A short proof of the Gittins index theorem. *The Annals of Applied Probability*, pages 194–199, 1994.

Richard Weber. On the Gittins index for multiarmed bandits. *The Annals of Applied Probability*, 2(4):1024–1033, 1992.

Peter Whittle. Multi-armed bandits and the Gittins index. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 143–149, 1980.

Yi-Ching Yao. Some results on the Gittins index for a normal reward process. In *Time Series and Related Topics*, pages 284–294. Institute of Mathematical Statistics, 2006.

Appendix A. Comparison to Bayes

As remarked in the introduction, it turns out that both geometric discounting and an infinite horizon are crucial for the interchange argument used in all proofs of the Gittins index theorem and indeed the result is known to be false in general for non-geometric discounting as shown by [Berry and Fristedt \(1985\)](#). Thus the observation that the Gittins index strategy is not Bayesian optimal in the setting considered here should not come as a surprise, but is included for completeness. Let $n = 2$ and $\nu_1 = 0$ and $\sigma_1^2 = 1$ and $\sigma_2^2 = 1/2$. Then the Gittins indices are

$$\gamma(\nu_1, \sigma_1^2, 2) \approx 0.195183 \qquad \gamma(\nu_2, \sigma_2^2, 2) \approx \nu_2 + 0.112689.$$

Therefore the strategy based on Gittins index will choose the second action if $\nu_2 \gtrsim 0.082494$. Computing the Bayesian value of each choice is possible analytically

$$\begin{aligned} \sup_{\pi} \mathbb{E} \left[\sum_{t=1}^n X_t \middle| I_1 = 1 \right] &= \int_{-\infty}^{\infty} \max \{ \nu_2, \delta \} \frac{1}{\sqrt{\pi}} \exp(-\delta^2) d\delta \\ &= \frac{\exp(-\nu_2^2)}{2\sqrt{\pi}} + \frac{\nu_2 + \nu_2 \operatorname{Erf}(\nu_2)}{2}. \\ \sup_{\pi} \mathbb{E} \left[\sum_{t=1}^n X_t \middle| I_1 = 2 \right] &= \nu_2 + \int_{-\infty}^{\infty} \max \{ 0, \nu_2 + \delta \} \frac{1}{\sqrt{\pi/3}} \exp(-3\delta^2) d\delta \\ &= \nu_2 + \frac{\exp(-3\nu_2^2)}{\sqrt{2\pi}} + \frac{\nu_2 + \nu_2 \operatorname{Erf}(\sqrt{3}\nu_2)}{2}. \end{aligned}$$

Solving leads to the conclusion that $I_2 = 2$ is optimal only if $\nu_2 \gtrsim 0.116462$ and hence the Gittins strategy is not Bayesian optimal (it does not minimise Eq. (2)). Despite this, the following result shows that the regret analysis for the Gittins index strategy can also be applied to the intractable fully Bayesian algorithm when the number of arms is $d = 2$. The idea is to show that the Bayesian algorithm will never choose an arm for which a UCB-like upper confidence bound is smaller than the largest Gittins index. Then the analysis in Section 3 may be repeated to show that that Theorems 5 and 12 also hold for the Bayesian algorithm when $d = 2$.

Theorem 13 *Assume $d = 2$ and the horizon is n . Let Q be the multivariate Gaussian prior measure with mean $\nu \in \mathbb{R}^d$ and covariance matrix $\Sigma = \text{diag}(\sigma^2)$. If the Bayesian optimal action is $I_1 = 1$, then*

$$\nu_1 + \sqrt{2c\sigma_1^2 \log n} \geq \gamma(\nu_2, \sigma_2^2, n),$$

where $c > 0$ is a universal constant.

The proof of Theorem 13 is surprisingly tricky and may be found in Appendix H. Empirically the behaviour of the Bayesian algorithm and the Gittins index strategy is almost indistinguishable, at least for two arms and small horizons (see Appendix C).

Appendix B. Computing the Gittins Index

I briefly describe a method of computing the Gittins index in the Gaussian case. A variety of authors have proposed sophisticated methods for computing the Gittins index, both in the discounted and finite horizon settings (Niño-Mora, 2011; Chakravorty and Mahajan, 2013, and references therein). The noise model used here is Gaussian and so continuous, which seems to make prior work inapplicable. Fortunately the Gaussian model is rather special, mostly due to the shift invariance (Lemma 3), which can be exploited to compute and store the index quite efficiently.

Let $\nu \in \mathbb{R}$ and $\sigma^2 > 0$ be the prior mean and variance respectively and m be the number of rounds remaining. For $t \in \{1, \dots, m\}$ define independent random variables

$$\eta_t \sim \mathcal{N}\left(0, \frac{\sigma^2}{1 + (t-1)\sigma^2} \cdot \frac{\sigma^2}{1 + t\sigma^2}\right).$$

Let $\nu_1 = \nu$ and $\nu_t = \nu_{t-1} + \eta_{t-1}$ for $t > 1$. Then

$$\gamma(\nu, \sigma^2, m) = \max \left\{ \gamma : \sup_{1 \leq \tau \leq m} \mathbb{E} \left[\sum_{t=1}^{\tau} (\nu_t - \gamma) \right] = 0 \right\} \quad (20)$$

where the stopping time is with respect to the filtration generated by η_1, \dots, η_m . Eq. (20) is the Bayesian view of Eq. (5) with the integral over μ in that equation incorporated into the posterior means. As in the proof of Theorem 1, let ν be such that $\gamma(\nu, \sigma^2, m) = 0$. Then

$$0 = \gamma(\nu, \sigma^2, m) = \sup_{1 \leq \tau \leq m} \mathbb{E} \left[\sum_{t=1}^{\tau} \nu_t \right].$$

The optimal stopping problem above can be solved by finding the root of the following Bellman equation for $t = m$.

$$\mathcal{V}_t(x, \sigma^2) = \begin{cases} x + \mathbb{E}_{\eta \sim \mathcal{N}(0, \sigma^4/(1+\sigma^2))} \left[\max \left\{ 0, \mathcal{V}_{t-1} \left(x + \eta, \frac{\sigma^2}{1+\sigma^2} \right) \right\} \right] & \text{if } t \geq 1 \\ 0 & \text{otherwise.} \end{cases} \quad (21)$$

Then by Lemma 3 the Gittins index satisfies

$$\gamma(0, \sigma^2, m) = \gamma(\nu, \sigma^2, m) - \nu = -\nu = -\max \{x : \mathcal{V}_m(x, \sigma^2) = 0\}. \quad (22)$$

Now Eq. (21) is a Bellman equation and conveniently there is a whole field devoted to solving such equations (Bertsekas and Tsitsiklis, 1995, and references therein). An efficient algorithm that computes the Gittins index to arbitrary precision by solving Eq. (21) using backwards induction and approximating $\max\{0, \mathcal{V}_t(x, \sigma^2)\}$ using quadratic splines may be found at <https://github.com/tor/libbandit>. The motivation for choosing the quadratic is because the expectation in Eq. (21) can be computed explicitly in terms of the error function and because $\mathcal{V}_t(x, \sigma^2)$ is convex in x , so the quadratic leads to a relatively good fit with only a few splines. The convexity of \mathcal{V} also means that Eq. (22) can be computed efficiently from a sufficiently good approximation of \mathcal{V}_m . Finally, in order to implement Algorithm 1 up to a horizon of n we may need $\gamma(\nu, 1/T, m) = \nu + \gamma(0, 1/T, m)$ for all T and m satisfying $T + m \leq n$. This can be computed by solving n copies of Eq. (21) with $\sigma^2 = 1$ and $m \in \{1, \dots, n-1\}$. The total running time is $O(n^2N)$ where N is the maximum number of splines required for sufficient accuracy. The results can be stored in a lookup table of size $O(n^2)$, which makes the actual simulation of Algorithm 1 extremely fast. Computation time for $n = 10^4$ was approximately 17 hours using 8 cores of a Core-i7 machine.

Appendix C. Experiments

I compare the Gittins strategy given in Algorithm 1 with UCB, OCUCB (Lattimore, 2015a) and Thompson sampling (TS) with a flat Gaussian prior (Thompson, 1933; Agrawal and Goyal, 2012). Due to the computational difficulties in calculating the Gittins index we are limited to modest horizons. For longer horizons the Gittins index may be approximation by the lower bound in Theorem 1. I also compare the Gittins index strategy to the Bayesian optimal strategy that can be computed with reasonable accuracy for the two-armed case and a horizon of $n = 2000$. Error bars are omitted from all plots because they are too small to see. All code is available from <https://github.com/tor/libbandit>.

Algorithm	Index
UCB	$\hat{\mu}_i(t-1) + \sqrt{\frac{2}{T_i(t-1)} \log t}$
OCUCB	$\hat{\mu}_i(t-1) + \sqrt{\frac{3}{T_i(t-1)} \log \left(\frac{2n}{t}\right)}$
TS	$\sim \mathcal{N}(\hat{\mu}_i(t-1), (T_i(t-1) + 1)^{-1})$

Table 1: Comparison Algorithms

Worst-case regret I start with the worst case regret for horizons $n \in \{10^3, 10^4\}$. In all experiments the first arm is optimal and has mean $\mu_1 = 0$. The remaining arms have $\mu_i = -\Delta$. In Fig. 1 I plot the expected regret of various algorithms with $d \in \{2, 5, 10\}$ and varying Δ . The results demonstrate that the Gittins strategy significantly outperforms both UCB and Thompson sampling, and is a modest improvement on OCUCB in most cases.

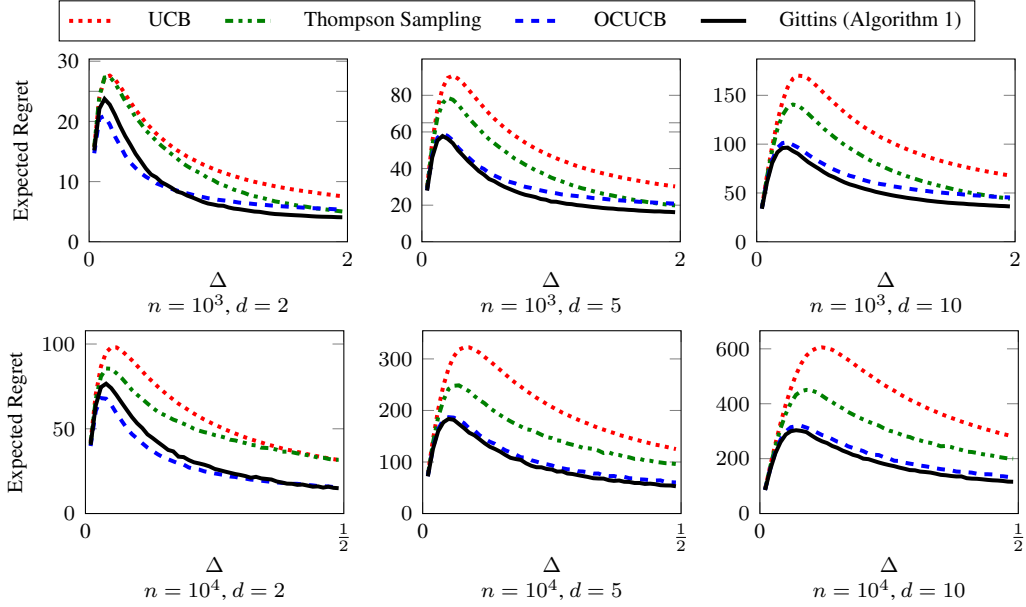


Figure 1: Worst case regret comparison

Approximating the Gittins index Current methods for computing the Gittins index are only practical for horizons up to $O(10^4)$. For longer horizons it seems worthwhile to find a closed form approximation. Taking inspiration from Theorem 1, define $\tilde{\gamma}(\nu, \sigma^2, m) \approx \gamma(\nu, \sigma^2, m)$ by

$$\tilde{\gamma}(\nu, \sigma^2, m) = \nu + \sqrt{2\sigma^2 \log \beta(\sigma^2, m)} \quad \text{where}$$

$$\beta(\sigma^2, m) = \max \left\{ 1, \frac{1}{4} \min \left\{ \frac{m}{\log^{\frac{3}{2}}(m)}, \frac{m\sigma^2}{\log^{\frac{1}{2}}(m\sigma^2)} \right\} \right\}.$$

This is exactly the lower bound in Theorem 1, but with the leading constant chosen (empirically) to be 1/4. I compare the indices in two key regimes. First fix $\nu = 0$ and $\sigma^2 = 1$ and vary the horizon. In the second regime the horizon is fixed to $m = 10^3$ and $\sigma^2 = 1/T$ is varied. The results (Fig. 2) suggest a relatively good fit.




Figure 2: Approximation of the index

The approximated index strategy is reasonably competitive with the Gittins strategy (Fig. 3). For longer horizons the Gittins index cannot be computed in reasonable time, but the approximation can be compared to other efficient algorithms such as OCUCB, Thompson sampling and UCB (Fig. 3).

The approximated version Gittins index is performing well compared to Thompson sampling and UCB, and is only marginally worse than OCUCB. As an added bonus, by cloning the proofs of Theorems 5 and 12 it can be shown that the approximate index enjoys the same regret guarantees as the Gittins strategy/UCB.

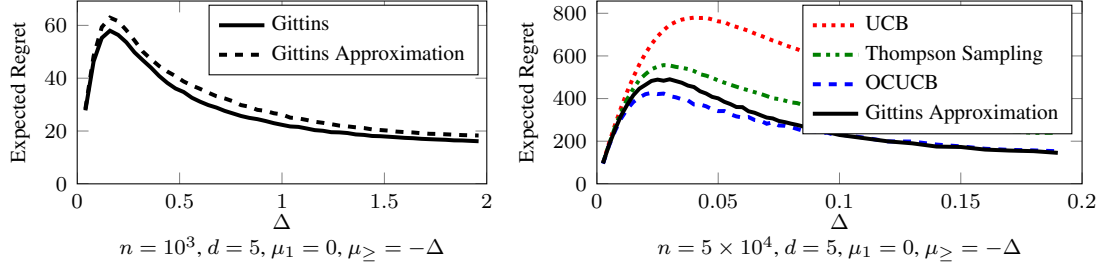


Figure 3: Regret for approximate Gittins index strategy

Comparing Gittins and Bayes Unfortunately the Bayesian optimal strategy is not typically an index strategy. The two natural exceptions occur when the rewards are discounted geometrically or when there are only two arms and the return of the second arm is known. Despite this, the Bayesian optimal strategy can be computed for small horizons and a few arms using a similar method as the Gittins index. Below I compare the Gittins index against the Bayesian optimal and OCUCB for $n = 2000$ and $d = 2$ and $\mu_1 = 0$ and $\mu_2 = -\Delta$. The results show the similar behaviour of the Gittins index strategy and the Bayesian optimal strategy, while OCUCB is making a different trade-off (better minimax regret at the cost of worse long-term regret).




Figure 4: Comparing Gittins and Bayesian optimal

Appendix D. Proof of Lemma 2

For parts (1) and (2) we have $\theta \in (-\infty, 0)$. Define $t_\theta \in \{1, 2, \dots\}$ by

$$t_\theta = \min \left\{ t \geq \frac{1}{\nu^2} : \sqrt{\frac{4}{t} \log(4t\nu^2)} \leq -\theta/2 \right\}.$$

Then

$$\begin{aligned}
\mathbb{E}[\tau|\mu = \theta] &= \sum_{t=1}^m \mathbb{P}\{\tau \geq t|\mu = \theta\} \leq t_\theta + \sum_{t=t_\theta+1}^m \mathbb{P}\{\tau \geq t|\mu = \theta\} \\
&\leq t_\theta + \sum_{t=t_\theta}^{m-1} \mathbb{P}\left\{\hat{\mu}_t + \sqrt{\frac{4}{t} \log(4t\nu^2)} > 0 \middle| \mu = \theta\right\} \\
&\leq t_\theta + \sum_{t=t_\theta}^{m-1} \mathbb{P}\{\hat{\mu}_t \geq \theta/2|\mu = \theta\} \leq t_\theta + \sum_{t=t_\theta}^{\infty} \exp\left(-\frac{t\theta^2}{8}\right) \leq t_\theta + \frac{8}{\theta^2}.
\end{aligned}$$

The results follow by naively bounding t_θ . For part 3, assume $\mu \geq 0$ and define $S_t = t\hat{\mu}_t$ and $t_k = \lceil 1/\nu^2 \rceil 2^k$. By using the peeling device we obtain

$$\begin{aligned}
\mathbb{P}\{\tau < m|\mu = \theta\} &= \mathbb{P}\left\{\exists t \geq \frac{1}{\nu^2} : \hat{\mu}_t + \sqrt{\frac{4}{t} \log(4t\nu^2)} \leq 0 \middle| \mu = \theta\right\} \\
&\leq \sum_{k=0}^{\infty} \mathbb{P}\left\{\exists t \leq t_{k+1} : S_t + \sqrt{2 \cdot t_{k+1} \log(4t_k\nu^2)} \leq 0 \middle| \mu = \theta\right\} \\
&\leq \sum_{k=0}^{\infty} \frac{1}{4t_k\nu^2} \leq \sum_{k=0}^{\infty} \frac{1}{4 \cdot 2^k} = \frac{1}{2}.
\end{aligned}$$

Therefore $\mathbb{E}[\tau|\mu = \theta] \geq m/2$.

Appendix E. Concentration of Gaussian Random Variables

These lemmas are not new, but are collected here with proofs for the sake of completeness.

Lemma 14 *Let $\sigma \geq 0$ and $\nu < 0$ and $X \sim \mathcal{N}(\nu, \sigma^2)$, then for $x \geq \nu$*

1. $\mathbb{P}\{X \geq x\} \leq \exp\left(-\frac{(\nu - x)^2}{2\sigma^2}\right)$.
2. $\mathbb{E}[X\mathbb{1}\{X \leq 0\}] \geq \nu - \sigma/\sqrt{2\pi}$.

Proof The first is well known. The second follows since

$$\begin{aligned}
\mathbb{E}[X\mathbb{1}\{X \leq 0\}] &= \int_{-\infty}^0 \frac{x}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\nu - x)^2}{2\sigma^2}\right) dx \\
&= \frac{\nu}{2} \operatorname{Erfc}\left(\frac{\nu}{\sqrt{2\sigma^2}}\right) - \sqrt{\frac{\sigma^2}{2\pi}} \exp\left(-\frac{\nu^2}{2\sigma^2}\right) \\
&\geq \nu - \sigma/\sqrt{2\pi}
\end{aligned}$$

as required. ■

Lemma 15 *Let $\delta \in (0, 1)$ and $\Delta \geq 0$ and X_1, \dots, X_n be a sequence of i.i.d. random variables with $X_t \sim \mathcal{N}(0, 1)$ and $S_t = \sum_{s=1}^t X_s$. Then*

$$\mathbb{P} \left\{ \exists t \leq n : S_t \geq n\Delta + \sqrt{2n \log \left(\frac{1}{\delta} \right)} \right\} \leq \frac{\delta}{\sqrt{\pi \log(1/\delta)}} \exp \left(-\frac{n\Delta^2}{2} \right).$$

Proof Let $\varepsilon > 0$. By the reflection principle we have $\mathbb{P} \{ \exists t \leq n : S_t \geq \varepsilon \} \leq 2\mathbb{P} \{ S_n \geq \varepsilon \}$ and

$$\begin{aligned} 2\mathbb{P} \{ S_n \geq \varepsilon \} &= 2 \int_{\varepsilon}^{\infty} \frac{1}{\sqrt{2\pi n}} \exp \left(-\frac{x^2}{2n} \right) dx \\ &\leq \frac{2}{\varepsilon} \int_{\varepsilon}^{\infty} \frac{x}{\sqrt{2\pi n}} \exp \left(-\frac{x^2}{2n} \right) dx \\ &= \frac{1}{\varepsilon} \sqrt{\frac{2n}{\pi}} \exp \left(-\frac{\varepsilon^2}{2n} \right). \end{aligned}$$

There result is completed by substituting $\varepsilon = n\Delta + \sqrt{2n \log(1/\delta)}$ and naive simplification. \blacksquare

Appendix F. Proof of Lemma 6

For the first part we can apply Lemma 15 and the union bound.

$$\begin{aligned} &\mathbb{P} \left\{ \exists t \leq n : S_t \geq t\Delta + \sqrt{2t \log \left(\frac{cn}{2 \log_+^{\frac{3}{2}}(n/2)} \right)} \right\} \\ &\leq \sum_{t=1}^n \mathbb{P} \left\{ S_t \geq t\Delta + \sqrt{2t \log \left(\frac{cn}{2 \log_+^{\frac{3}{2}}(n/2)} \right)} \right\} \\ &\leq \frac{2 \log_+^{\frac{3}{2}}(n/2)}{cn \sqrt{\pi \log \left(\frac{cn}{2 \log_+^{\frac{3}{2}}(n/2)} \right)}} \sum_{t=1}^{\infty} \exp \left(-\frac{t\Delta^2}{2} \right) \\ &\leq \frac{\log(n)}{n\Delta^2} \cdot \frac{4 \log_+^{\frac{3}{2}}(n/2)}{c \log(n) \sqrt{\pi \log \left(\frac{cn}{2 \log_+^{\frac{3}{2}}(n/2)} \right)}}. \end{aligned}$$

The result follows by noting that

$$\lim_{n \rightarrow \infty} \frac{4 \log_+^{\frac{3}{2}}(n/2)}{c \log(n) \sqrt{\pi \log \left(\frac{cn}{2 \log_+^{\frac{3}{2}}(n/2)} \right)}} = \frac{4}{c\sqrt{\pi}}$$

and then choosing n_0 sufficiently large. Moving to the second part. Let $\eta \in (0, 1]$ be a constant to be chosen later and $t_k = (1 + \eta)^k$ and

$$\alpha_k = \frac{2 \log_+^{\frac{1}{2}} \left(\frac{n}{2t_k} \right)}{c} \quad K = \max \left\{ k : t_k \leq 2en \quad \text{and} \quad \sqrt{\frac{n}{t_k}} \geq \alpha_k \geq 1 \right\}.$$

An easy computation shows that there exists a universal constant $c' > 0$ such that $t_K \geq c'n$. Then

$$\begin{aligned} & \mathbb{P} \left\{ \exists t \leq n : S_t \geq t\Delta + \max \left\{ 0, \sqrt{2t \log \left(\frac{cn}{2t \log_+^{\frac{1}{2}}(n/(2t))} \right)} \right\} \right\} \\ & \leq \sum_{k=0}^K \mathbb{P} \left\{ \exists t \leq t_k : S_t \geq \frac{t_k \Delta}{(1 + \eta)} + \sqrt{\frac{2t_k}{1 + \eta} \log \left(\frac{n}{t_k \alpha_k} \right)} \right\} + \mathbb{P} \{ \exists t_K \leq t \leq n : S_t \geq t\Delta \}. \end{aligned} \quad (23)$$

The second term in Eq. (23) is easily bounded by a peeling argument.

$$\begin{aligned} \mathbb{P} \{ \exists t_K \leq t \leq n : S_t \geq t\Delta \} & \leq \sum_{i=1}^{\infty} \mathbb{P} \{ \exists t \leq (i + 1)t_K : S_t \geq it_K \Delta \} \\ & \leq \sum_{i=1}^{\infty} \exp \left(-\frac{i^2 t_K^2 \Delta^2}{2(i + 1)t_K} \right) \\ & \leq \frac{c''}{n\Delta^2}, \end{aligned} \quad (24)$$

where $c'' > 0$ is some sufficiently large universal constant and we have used the fact that $t_K \geq c'n$ and naive algebra. The first term in Eq. (23) is slightly more involved. Using peeling and Lemmas

15 and 21 we have universal constants c' and c'' such that

$$\begin{aligned}
 & \sum_{k=0}^K \mathbb{P} \left\{ \exists t \leq t_k : S_t \geq \frac{t_k \Delta}{(1+\eta)} + \sqrt{\frac{2t_k}{1+\eta} \log \left(\frac{n}{t_k \alpha_k} \right)} \right\} \\
 & \stackrel{(a)}{\leq} \sum_{k=0}^K \sqrt{\frac{1+\eta}{\pi \log \left(\frac{n}{t_k \alpha_k} \right)}} \left(\frac{t_k \alpha_k}{n} \right)^{1/(1+\eta)} \exp \left(-\frac{t_k \Delta^2}{2(1+\eta)^2} \right) \\
 & \stackrel{(b)}{\leq} \sum_{k=0}^K \alpha_k \sqrt{\frac{2}{\pi \log \left(\frac{n}{t_k \alpha_k} \right)}} \left(\frac{t_k}{n} \right)^{1/(1+\eta)} \exp \left(-\frac{t_k \Delta^2}{2(1+\eta)^2} \right) \\
 & \stackrel{(c)}{\leq} \sum_{k=0}^K \frac{2 \log_+^{\frac{1}{2}} \left(\frac{n}{t_k} \right)}{c} \sqrt{\frac{4}{\pi \log \left(\frac{n}{t_k} \right)}} \left(\frac{t_k}{n} \right)^{1/(1+\eta)} \exp \left(-\frac{t_k \Delta^2}{2(1+\eta)^2} \right) \\
 & \stackrel{(d)}{=} \frac{4}{c\sqrt{\pi}} \sum_{k=0}^K \left(\frac{t_k}{n} \right)^{1/(1+\eta)} \exp \left(-\frac{t_k \Delta^2}{2(1+\eta)^2} \right) \\
 & \stackrel{(e)}{\leq} \frac{3 \cdot 4}{c\eta} \left(\frac{2(1+\eta)}{n\Delta^2} \right)^{1/(1+\eta)} \\
 & \stackrel{(f)}{\leq} \frac{4 \cdot 3 \cdot 4}{c\eta} \left(\frac{1}{n\Delta^2} \right)^{1/(1+\eta)}.
 \end{aligned}$$

where (a) follows from Lemma 15, (b) since $\eta \in (0, 1]$ and $\alpha_k \geq 1$, (c) by substituting the value of α_k and because $\sqrt{n/t_k} \geq \alpha_k$, (d) and (f) are trivial while (e) follows from Lemma 21. Then choose

$$\eta = \frac{1}{\log_+(n\Delta^2)} \in (0, 1],$$

which satisfies

$$\frac{1}{\eta} \left(\frac{1}{n\Delta^2} \right)^{1/(1+\eta)} \leq \frac{e}{n\Delta^2} \log_+(n\Delta^2).$$

The result by combining the above reasoning with Eq. (24) and substituting into Eq. (23).

Appendix G. Proof of Theorem 12

As before I assume for convenience that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_d$. Let n be fixed and abbreviate $\gamma_i(t) = \gamma(\nu_i(t-1), \sigma_i^2(t-1), n-t+1)$. Like the proof of Theorem 5 there are two main difficulties. First, showing for suboptimal arms i that $\gamma_i(t) < \mu_i + \Delta_i/2$ occurs with sufficiently high probability once $T_i(t)$ is sufficiently large. Second, showing that $\gamma_1(t) \geq \mu_1 - \Delta_1/2$ occurs with sufficiently high probability. Let F be the event that there exists an arm i and round $t \in \{1, \dots, n\}$ such that

$$|\hat{\mu}_i(t) - \mu_i| \geq \sqrt{\frac{2}{T_i(t)} \log(2dn^2)}.$$

As in the proof of Theorem 5 we have $\mathbb{P}\{F\} \leq 1/n$. The regret due to F occurring is negligible, so from now on assume that F does not hold. By Eq. (18) we have

$$\nu_i(t) = \eta_i(t)\nu_i + (1 - \eta_i(t))\hat{\mu}_i(t),$$

where

$$\eta_i(t) = \frac{1}{1 + \sigma_i^2 T_i(t)} = O\left(\frac{1}{T_i(t)}\right).$$

Therefore for sufficiently large n and $u_i = \lceil \frac{32}{\Delta^2} \log(dn^2) \rceil$ and $T_i(t-1) = u_i$

$$\begin{aligned} \gamma_i(t) &\leq \eta_i(t-1)\nu_i + (1 - \eta_i(t-1))\hat{\mu}_i(t-1) + \sqrt{2\sigma_i^2(t-1)\log(2n)} \\ &\leq \eta_i(t-1)(\nu_i - \mu_i) + \mu_i + \sqrt{\frac{2}{T_i(t-1)}\log(2n)} + \sqrt{\frac{2}{T_i(t-1)}\log(2dn^2)} \\ &\leq \eta_i(t-1)(\nu_i - \mu_i) + \mu_i + \Delta_i/2 \leq \mu_i + 3\Delta_i/4. \end{aligned}$$

For the first arm we have for $t \leq n/2$ that

$$\gamma_1(t) \geq \eta_1(t-1)\nu_1 + (1 - \eta_1(t-1))\hat{\mu}_1(t-1) + \sqrt{2\sigma_1^2(t-1)\log\beta_t}.$$

where by Theorem 1

$$\beta_t \geq c \min \left\{ \frac{n/2}{\log_+^{\frac{3}{2}}(n/2)}, \frac{n\sigma_1^2(t-1)/2}{\log_+^{\frac{1}{2}}(n\sigma_1^2(t-1)/2)} \right\}.$$

For any $\Delta > 0$ define F_Δ to be the event that there exists a $t \leq n/2$ for which

$$\hat{\mu}_1(t-1) + \sqrt{\frac{2}{T_1(t-1)}\log\beta_t} \leq \mu_1 - \Delta.$$

By Lemma 6 there is a universal constant $c' > 0$ such that

$$\mathbb{P}\{F_\Delta\} \leq c'(\log(n) + \log_+(n\Delta^2))/(n\Delta^2)$$

and when F_Δ does not occur we have for $t \leq n/2$ that

$$\begin{aligned} \gamma_1(t) &\geq \eta_1(t-1)\nu_1 + (1 - \eta_1(t-1)) \left(\mu_1 - \Delta - \sqrt{\frac{2}{T_1(t-1)}\log\beta_t} \right) \\ &\quad + \sqrt{2\sigma_1^2(t-1)\log\beta_t} \\ &\geq \mu_1 - \Delta - O\left(\frac{1}{T_1(t-1)}\right) - (1 - \eta_1(t-1))\sqrt{\frac{2}{T_1(t-1)}\log\beta_t} \\ &\quad + \sqrt{2\sigma_1^2(t-1)\log\beta_t}. \end{aligned}$$

Comparing the ratio of the last two terms gives

$$\frac{(1 - \eta_1(t-1))\sqrt{\frac{2}{T_1(t-1)} \log \beta_t}}{\sqrt{2\sigma_1^2(t-1)} \log \beta_t} = \sqrt{\frac{\sigma_1^2 T_1(t-1)}{1 + \sigma_1^2 T_1(t-1)}} < 1.$$

Since $\log \beta_t = \Theta(\log n)$ for $t \leq n/2$, if F_Δ does not hold and n is sufficiently large, then

$$\gamma_1(t) \geq \mu_1 - 2\Delta \quad \text{for all } t \leq n/2.$$

The proof is completed by duplicating the argument given in the proof of Theorem 5 and taking the limit as n tends to infinity.

Appendix H. Proof of Theorem 13

Lemma 16 *Let $X \sim \mathcal{N}(0, \sigma^2)$ and A be an arbitrary event. Then*

$$\mathbb{E}\mathbb{1}\{A\} X \leq \mathbb{P}\{A\} \sqrt{2\sigma^2 \log \left(\frac{1}{\mathbb{P}\{A\}} \right)}.$$

Lemma 17 *Let B_t be the standard Brownian motion. Let $\varepsilon > 0$ and define stopping time $\tau = \min\{t : B_t = -\varepsilon\}$. Then*

$$\mathbb{P}\{\tau > t\} = \Phi\left(\frac{\varepsilon}{\sqrt{t}}\right) - \Phi\left(\frac{-\varepsilon}{\sqrt{t}}\right).$$

The proofs of Lemmas 16 and 17 are omitted, but the former follows by letting $\delta = \mathbb{P}\{A\}$ and noting that the worst-case occurs when $A = \mathbb{1}\{X \geq \alpha\}$ where α is such that

$$\int_{\alpha}^{\infty} \frac{\exp(-x^2/(2\sigma^2))}{\sqrt{2\pi\sigma^2}} dx = \delta.$$

The second lemma is well known (Lerche, 1986, for example). For what follows we need an alternative view of the Gittins index. Let $\nu \in \mathbb{R}$ and $1 \geq \sigma^2 \geq 0$. Furthermore define

$$\sigma_t^2 = \frac{\sigma^2}{1 + (t-1)\sigma^2} \cdot \frac{\sigma^2}{1 + t\sigma^2},$$

which satisfies $\sum_{t=1}^{\infty} \sigma_t^2 = \sigma^2$. I abbreviate $\sigma_{\leq t}^2 = \sum_{s=1}^t \sigma_s^2$. Let $\nu_1 = \nu$ and $\nu_{t+1} = \nu_t + \eta_t$ where $\eta_t \sim \mathcal{N}(0, \sigma_t^2)$. Then the Gittins index is given by

$$\gamma(\nu, \sigma^2, n) = \sup_{1 \leq \tau \leq n} \frac{\mathbb{E}[\sum_{t=1}^{\tau} \nu_t]}{\mathbb{E}[\tau]},$$

where $\mathbb{1}\{\tau = t\}$ is measurable with respect to the σ -algebra generated by $\nu_1, \nu_2, \dots, \nu_{t+1}$. This is equivalent to the definition in Eq. (5), but written in terms of the evolution of the posterior mean rather than the observed reward sequences (these can be computed from each other in a deterministic way).

Lemma 18 *There exists a universal constant $1 \geq c > 0$ such that $\gamma(\nu, \sigma^2, n) \geq c\sqrt{2\sigma_{\leq n}^2}$.*

The proof is left as an exercise (hint: choose a stopping time $\tau = \lceil n/2 \rceil$ unless $\nu_{\lceil n/2 \rceil}$ is large enough when $\tau = n$). Recall that $\tau = \tau(\nu, \sigma^2, n)$ is the stopping rule defining the Gittins index for mean ν , variance σ^2 and horizon n . The following lemma shows that $\tau = n$ with reasonable probability and is crucial for the proof of Theorem 13.

Lemma 19 *Let $\nu \in \mathbb{R}$ and $\sigma^2 > 0$ and let $\tau = \tau(\nu, \sigma^2, n)$. Then there exists a universal constant $c' > 0$ such that*

$$\mathbb{P}\{\tau = n\} \geq \frac{c'}{n^2 \log(n)}.$$

Proof Assume without loss of generality (by translation and Lemma 3) that $\gamma(\nu, \sigma^2, n) = 0$ and let $\mathcal{E} = -\nu \geq c\sqrt{2\sigma_{\leq n}^2}$ by Lemma 18. By definition of the Gittins index we have

$$\mathbb{E}\left[\sum_{t=1}^{\tau} \nu_t\right] = 0.$$

Let $\delta_t = \mathbb{P}\{\tau \geq t \text{ and } \nu_t \geq \mathcal{E}/(2n)\}$ and $\alpha_t = \mathbb{E}[\nu_t | \tau \geq t \text{ and } \nu_t \geq \mathcal{E}/(2n)]$. By Lemma 16 we have $\alpha_t \leq \sqrt{2\sigma_{\leq t}^2 \log(1/\delta_t)} - \mathcal{E} \leq \sqrt{2\sigma_{\leq t}^2 \log(1/\delta_t)}$. Therefore

$$0 = \mathbb{E}\left[\sum_{t=1}^{\tau} \nu_t\right] \leq -\mathcal{E}/2 + \sum_{t=2}^n \delta_t \alpha_t.$$

Therefore there exists a t such that $\delta_t \alpha_t \geq \mathcal{E}/(2n)$ and so

$$\delta_t \sqrt{2\sigma_{\leq t}^2 \log \frac{1}{\delta_t}} \geq \delta_t \alpha_t \geq \frac{\mathcal{E}}{2n} \geq \frac{c\sqrt{2\sigma_{\leq n}^2}}{2n}.$$

Straightforward analysis shows that

$$\delta_t \geq \frac{c}{2n \log\left(\frac{2n}{c}\right)}.$$

Now we apply Lemma 17. First note that if $\tau \geq t$ and $\nu_s \geq 0$ for all $s \geq t$, then $\tau = n$. Therefore

$$\begin{aligned} \mathbb{P}\{\tau = n\} &\geq \mathbb{P}\{\tau \geq t \text{ and } \nu_s \geq 0 \text{ for all } s \geq t\} \\ &\geq \frac{c}{2n \log\left(\frac{2n}{c}\right)} \left(\Phi\left(\frac{\mathcal{E}}{2n\sigma_{\leq n}}\right) - \Phi\left(-\frac{\mathcal{E}}{2n\sigma_{\leq n}}\right) \right) \\ &\geq \frac{c}{2n \log\left(\frac{2n}{c}\right)} \left(\Phi\left(\frac{c}{2n}\right) - \Phi\left(-\frac{c}{2n}\right) \right) \\ &\geq \frac{c^2}{8n^2 \log\left(\frac{2n}{c}\right)} \end{aligned}$$

as required. ■

Remark 20 *It is rather clear that the above proof is quite weak. Empirically it appears that $\mathbb{P}\{\tau = n\} = \Omega(1/n)$, but the result above is sufficient for our requirements.*

Proof of Theorem 13 Let π^* be the Bayesian optimal policy, which by assumption chooses $I_1 = 1$. Define π be the policy that chooses arm 2 until t such that $\gamma_2(t) < \gamma_2(1)$ and thereafter follows π^* as if no samples had been observed and substituting the observed values for the second arm as it is chosen. Let $\tau = \tau(\nu_2, \sigma_2^2, n)$, then π chooses the second time until at least $t = \tau$. For any policy π let $T_i^\pi(n)$ be the number of plays of arm i by policy π , which is a random variable. Let $\nu_{i,s}$ be the posterior mean of the i th arm after $s - 1$ plays. Then the value of policy π is

$$V^\pi = \mathbb{E} \left[\sum_{i=1}^d \sum_{s=1}^{T_i^\pi(n)} \nu_{i,s} \right],$$

which is maximised by the Bayesian optimal policy π^* . Therefore if we abbreviate $T_i^{\pi^*}(n)$ by $T_i(n)$ and $\gamma_2(1)$ by γ , then

$$\begin{aligned} 0 &\geq V^\pi - V^{\pi^*} \\ &\geq \mathbb{E} \left[\mathbf{1}\{\tau > T_2(n)\} \sum_{s=T_1(n)-\tau+T_2(n)}^{T_1(n)} (\gamma - \nu_{1,s}) \right] \\ &\geq \sum_{s=1}^n \mathbb{E} \left[\mathbf{1}\{\tau - T_2(n) \geq s\} (\gamma - \max_{s \leq n} \nu_{1,n}) \right] \\ &\geq \sum_{s=1}^n \mathbb{P}\{\tau - T_2(n) \geq s\} \left(\gamma - \nu_1 - \sqrt{2\sigma_1^2 \log \left(\frac{1}{\mathbb{P}\{\tau - T_2(n) \geq s\}} \right)} \right) \\ &= \sum_{s=1}^n \delta_s \left(\gamma - \nu_1 - \sqrt{2\sigma_1^2 \log \left(\frac{1}{\delta_s} \right)} \right), \end{aligned} \tag{25}$$

where $\delta_s = \mathbb{P}\{\tau - T_2(n) \geq s\}$. Since the optimal policy chooses action 1 in the first round by assumption, if $\tau = n$, then $\tau - T_2(n) \geq 1$ is guaranteed. Therefore by Lemma 19 we have

$$\delta_1 \geq \frac{c'}{n^2 \log(n)}.$$

By rearranging Eq. (25) we have for any $\delta \leq \delta_1$ that

$$\begin{aligned} \gamma &\leq \nu_1 + \frac{\sum_{s=1}^n \delta_s \sqrt{2\sigma_1^2 \log \left(\frac{1}{\delta_s} \right)}}{\sum_{s=1}^n \delta_s} \\ &\leq \nu_1 + \sqrt{2\sigma_1^2 \log \left(\frac{1}{\delta} \right)} + \frac{\sum_{s:\delta_s < \delta} \delta_s \sqrt{-2\sigma_1^2 \log \delta_s}}{\delta_1} \\ &\leq \nu_1 + \sqrt{2\sigma_1^2 \log \left(\frac{1}{\delta} \right)} + \frac{\sum_{s:\delta_s < \delta} \sqrt{2\sigma_1^2 \delta_s}}{\delta_1} \\ &\leq \nu_1 + \sqrt{2\sigma_1^2 \log \left(\frac{1}{\delta} \right)} + \frac{n\sqrt{2\sigma_1^2 \delta}}{\delta_1}. \end{aligned}$$

The result is completed by choosing $\delta = \delta_1^2 n^{-2}$. ■

Appendix I. Technical Lemmas

Lemma 21 *Let $\eta \in (0, 1)$ and $\Delta > 0$. Then*

$$\sum_{k=0}^{\infty} (1+\eta)^{\frac{k}{1+\eta}} \exp\left(- (1+\eta)^k \Delta^2\right) \leq \frac{3}{\eta} \left(\frac{1}{\Delta^2}\right)^{1/(1+\eta)}.$$

Proof Let $L = \min \{L : (1+\eta)^L \geq \frac{1}{\Delta^2}\}$. Then

$$\begin{aligned} & \sum_{k=0}^{\infty} (1+\eta)^{\frac{k}{1+\eta}} \exp\left(- (1+\eta)^k \Delta^2\right) \\ &= \sum_{k=0}^{L-1} (1+\eta)^{\frac{k}{1+\eta}} \exp\left(- (1+\eta)^k \Delta^2\right) + \sum_{k=L}^{\infty} (1+\eta)^{\frac{k}{1+\eta}} \exp\left(- (1+\eta)^k \Delta^2\right) \\ &\leq \sum_{k=0}^{L-1} (1+\eta)^{\frac{k}{1+\eta}} + \sum_{k=L}^{\infty} (1+\eta)^{\frac{k}{1+\eta}} \exp\left(- (1+\eta)^k \Delta^2\right) \\ &\leq \frac{1}{\eta} \left(\frac{1}{\Delta^2}\right)^{1/(1+\eta)} + \sum_{k=L}^{\infty} (1+\eta)^{\frac{k}{1+\eta}} \exp\left(- (1+\eta)^k \Delta^2\right) \\ &\leq \frac{1}{\eta} \left(\frac{1}{\Delta^2}\right)^{1/(1+\eta)} + \left(\frac{1+\eta}{\Delta^2}\right)^{1/(1+\eta)} \sum_{k=0}^{\infty} (1+\eta)^{\frac{k}{1+\eta}} \exp\left(- (1+\eta)^k\right) \\ &\leq \frac{1}{\eta} \left(\frac{1}{\Delta^2}\right)^{1/(1+\eta)} + \left(\frac{1+\eta}{\Delta^2}\right)^{1/(1+\eta)} \sum_{k=0}^{\infty} (1+\eta)^k \exp\left(- (1+\eta)^k\right) \\ &\leq \frac{1}{\eta} \left(\frac{1}{\Delta^2}\right)^{1/(1+\eta)} + \left(\frac{1+\eta}{\Delta^2}\right)^{1/(1+\eta)} \left(\frac{1}{e} + \int_0^{\infty} (1+\eta)^k \exp\left(- (1+\eta)^k\right) dk\right) \\ &\leq \frac{1}{\eta} \left(\frac{1}{\Delta^2}\right)^{1/(1+\eta)} + \left(\frac{1+\eta}{\Delta^2}\right)^{1/(1+\eta)} \frac{1}{e} \left(1 + \frac{1}{\log(1+\eta)}\right) \\ &\leq \frac{2}{\eta} \left(\frac{1+\eta}{\Delta^2}\right)^{1/(1+\eta)} \\ &\leq \frac{3}{\eta} \left(\frac{1}{\Delta^2}\right)^{1/(1+\eta)} \end{aligned}$$

as required. ■

Appendix J. Completing Proof of Theorem 1

First we need an easy lemma.

Lemma 22 $f(\beta) = \frac{1}{\beta} \sqrt{\frac{1}{2\pi}} - \sqrt{\frac{\log \beta}{2}} \operatorname{Erfc}(\sqrt{\log \beta})$ satisfies:

1. $\lim_{\beta \rightarrow \infty} f(\beta) \beta \log(\beta) = \frac{1}{\sqrt{8\pi}}$.
2. $f(\beta) \geq \frac{1}{10\beta \log(\beta)}$ for all $\beta \geq 3$.
3. $f(\beta) \leq \frac{1}{\sqrt{8\pi}} \cdot \frac{1}{\beta \log \beta}$.

Completing Theorem 1 To begin we choose $c = (40c''')^4$. Recall that we need to show that

$$m\sigma f(\beta) \geq c''' \max \left\{ \sigma, \sqrt{\sigma^2 \log(\beta)}, \frac{1}{\sqrt{\sigma^2 \log(\beta)}}, \sqrt{mW(m\sigma^2 \log(\beta))} \beta^{-\frac{7}{8}} \right\},$$

where

$$\beta_1 = \frac{m}{c \log_+^{\frac{3}{2}}(m)} \quad \beta_2 = \frac{m\sigma^2}{c \log_+^{\frac{1}{2}}(m\sigma^2)} \quad \beta = \min \{\beta_1, \beta_2\}$$

and $\beta \geq 3$ is assumed. This is equivalent to showing:

- (1) $m\sigma f(\beta) \geq c''' \sigma$ and
- (2) $m\sigma f(\beta) \geq c''' \sqrt{\sigma^2 \log(\beta)}$ and
- (3) $m\sigma f(\beta) \geq c''' / \sqrt{\sigma^2 \log(\beta)}$ and
- (4) $m\sigma f(\beta) \geq c''' \sqrt{mW(m\sigma^2 \log(\beta))} \beta^{-\frac{7}{8}}$.

By Theorem 22 we have

$$\begin{aligned} m\sigma f(\beta) &\geq \frac{m\sigma}{10\beta \log_+(\beta)} \geq \frac{m\sigma}{10 \cdot \frac{m}{c \log_+^{\frac{3}{2}}(m)} \log_+ \left(\frac{m}{c \log_+^{\frac{3}{2}}(m)} \right)} \\ &\geq c''' \sigma \sqrt{\log_+(m)} \geq c''' \sqrt{\sigma^2 \log(\beta)} \geq c''' \sigma. \end{aligned}$$

Therefore (1) and (2) hold. For (3) we have

$$\begin{aligned} m\sigma f(\beta) &\geq \frac{m\sigma}{10\beta \log_+(\beta)} \geq \frac{m\sigma}{10 \cdot \frac{m\sigma^2}{c \log_+^{\frac{1}{2}}(m\sigma^2)} \log_+^{\frac{1}{2}} \left(\frac{m\sigma^2}{c \log_+^{\frac{1}{2}}(m\sigma^2)} \right) \sqrt{\log_+(\beta)}} \\ &\geq \frac{c'''}{\sqrt{\sigma^2 \log(\beta)}}. \end{aligned}$$

Therefore (3) holds. Finally

$$\begin{aligned} m\sigma f(\beta) \beta^{\frac{7}{8}} &\stackrel{(a)}{\geq} \frac{m\sigma}{10\beta^{\frac{1}{8}} \log_+(\beta)} \stackrel{(b)}{\geq} \frac{m\sigma}{40\beta^{\frac{1}{4}}} \stackrel{(c)}{\geq} \frac{m\sigma}{40 \left(\frac{m\sigma^2}{(40c''')^4} \right)^{\frac{1}{4}}} \stackrel{(d)}{=} c''' \sqrt{m} (m\sigma^2)^{\frac{1}{4}} \\ &\stackrel{(e)}{\geq} c''' \sqrt{mW(m\sigma^2 \log(m\sigma^2))} \stackrel{(f)}{\geq} c''' \sqrt{mW(m\sigma^2 \log(\beta))}, \end{aligned}$$

where (a) follows from Lemma 22. (b) since $\beta^{\frac{1}{8}} \log(\beta) \leq 4\beta^{\frac{1}{4}}$ for $\beta \geq 3$. (c) by substituting and naively upper bounding $\beta \leq \beta_2$. (d) is trivial. (e) since $x^{\frac{1}{4}} \geq \sqrt{W(x \log(x))}$ for all $x \geq 0$. (f) is true because $\beta \leq \beta_2 \leq m\sigma^2$.

Appendix K. A Brief History of Multi-Armed Bandits

This list is necessarily not exhaustive. It contains in (almost) chronological order the papers that are most relevant for the present work, with a special focus on earlier papers that are maybe less well known. Please contact me if there is a serious omission in this list.

- [Thompson \(1933\)](#) proposes a sampling approach for finite-horizon undiscounted Bernoulli bandits now known as Thompson sampling. No theoretical results are given.
- [Robbins \(1952\)](#) sets out the general finite-horizon bandit problem and gives some Hannan consistent algorithms and general discussion. Robbins seems to have been unaware of Thompson's work.
- [Bradt et al. \(1956\)](#) present the optimal solution for the Bayesian one-armed bandit in the finite-horizon undiscounted setting. They also prove many of the counter-intuitive properties of the Bayesian optimal solution for multi-armed bandits. The index used in the present article could just as well be called the Bradt–Johnson–Karlin index.
- [Gittins \(1979\)](#) presents the optimal solution for the Bayesian multi-armed bandit in the infinite-horizon geometrically discounted setting. Gittins seems to have been unaware of the work by [Bradt et al. \(1956\)](#) who defined a similar index. Gittins' result does not apply in the finite-horizon setting.
- [Whittle \(1980\)](#); [Weber \(1992\)](#); [Tsitsiklis \(1994\)](#) all give alternative proofs and/or characterisations of the Gittins index for the infinite-horizon discounted case.
- [Bather \(1983\)](#) gives an asymptotic approximation of the infinite-horizon Gittins index (see also the work by [Yao \(2006\)](#)).
- The textbook by [Berry and Fristedt \(1985\)](#) summarises many of the earlier results on Bayesian multi-armed bandits. They also prove that geometric discounting is essentially necessary for the Gittins index theorem to hold and give counter-examples in the Bernoulli case.
- [Lai and Robbins \(1985\)](#) develop asymptotically optimal frequentist algorithms for the multi-armed bandits and prove the first lower bounds.
- [Agrawal \(1995\)](#) and [Katehakis and Robbins \(1995\)](#) independently develop the UCB algorithm.
- [Burnetas and Katehakis \(1997\)](#) give asymptotic approximations for the finite-horizon Gittins index (more specific results in a similar vein are given by [Chernoff and Ray \(1965\)](#)).
- [Auer et al. \(2002\)](#) proves finite-time guarantees for UCB.

- [Niño-Mora \(2011\)](#) presents some methods for computing the finite-horizon Gittins index in the discrete case, and also suggest that index algorithm is a good approximation for the intractable Bayesian solution (no reference given).
- [Kaufmann et al. \(2012b\)](#) gave promising empirical results for the finite-horizon Gittins index strategy for the Bernoulli case, but did not study its theoretical properties. They also incorrectly claim that the Gittins strategy is Bayesian optimal. Their article played a large role in motivating the present work.

Appendix L. Table of Notation

d	number of arms
n	horizon
t	current round
μ_i	expected return of arm i
μ^*	maximum mean reward, $\mu^* = \max_i \mu_i$
$\hat{\mu}_i(t)$	empirical estimate of return of arm i after round t
$T_i(t)$	number of times arm i has been chosen after round t
I_t	arm chosen in round t
X_t	reward in round t
ν_i	prior mean of arm i
σ_i^2	prior variance for mean of arm i
Δ_i	gap between the expected returns of the best arm and the i th arm
Δ_{\max}	maximum gap, $\max_i \Delta_i$
Δ_{\min}	minimum non-zero gap, $\min \{\Delta_i : \Delta_i > 0\}$
$\log_+(x)$	$\max \{1, \log(x)\}$
$W(x)$	product logarithm $x = W(x)\exp(W(x))$
$\text{diag}(\sigma_1^2, \dots, \sigma_d^2)$	diagonal matrix with entries $\sigma_1^2, \dots, \sigma_d^2$
$\gamma(\nu, \sigma^2, m)$	Gittins index for mean/variance (ν, σ^2) and horizon m
$\tau(\nu, \sigma^2, m)$	stopping time that determines the Gittins index
$\text{Erfc}(x)$	complementary error function

$$\text{Erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^{\infty} \exp(-y^2) dy$$

$\text{Erf}(x)$ error function $\text{Erf}(x) = 1 - \text{Erfc}(x)$

$\mathcal{N}(\mu, \sigma^2)$ Gaussian with mean μ and variance σ^2

$\Phi(x)$ standard Gaussian cdf

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp\left(-\frac{t^2}{2}\right) dt$$