

On Information Inequalities and the Complexity of Sequential Decision Problems

Orsay Machine Learning Master Seminar

Aurélien Garivier

November 5th, 2019

UMPA
ENS DE LYON



Table of contents

Introduction: Information for Deviation Lower Bounds

Intermediate: Sequential Test

Active Learning: Finding the Distribution with Largest Mean

Introduction: Information for Deviation Lower Bounds

Chernoff Bound for Bernoulli variables

Let $\mu \in (0, 1)$. Let $X_1, X_2, \dots, X_n \stackrel{iid}{\sim} \mathcal{B}(\mu)$, and $\bar{X}_n = (X_1 + \dots + X_n)/n$.

Theorem

For all $\mu \leq x \leq 1$,

$$\mathbb{P}_\mu (\bar{X}_n \geq x) \leq e^{-n \text{kl}(x, \mu)}$$

where $\text{kl}(x, y) = x \log \frac{x}{y} + (1 - x) \log \frac{1 - x}{1 - y}$ is the binary relative entropy. Similarly, for all $0 \leq x \leq \mu$,

$$\mathbb{P}_\mu (\bar{X}_n \leq x) \leq e^{-n \text{kl}(x, \mu)} .$$


Corollary

For every $\delta > 0$,

$$\mathbb{P}_\mu \left(n \text{kl}(\bar{X}_n, \mu) \geq \log \frac{1}{\delta} \right) \leq 2\delta .$$

A Divergence on the Set of Possible Means

$$\mathbb{P}_\mu \left(n \text{kl}(\bar{X}_n, \mu) \geq \log \frac{1}{\delta} \right) \leq 2\delta$$



Proof: Fenchel-Legendre transform of log-Laplace

For every $\lambda > 0$,

$$\begin{aligned}\mathbb{P}_\mu(\bar{X}_n \geq x) &= \mathbb{P}_\mu\left(e^{\lambda(X_1 + \dots + X_n)} \geq e^{\lambda nx}\right) \\ &\leq \frac{\mathbb{E}_\mu\left[e^{\lambda(X_1 + \dots + X_n)}\right]}{e^{\lambda nx}} \quad \text{by Markov's inequality} \\ &= e^{-n(\lambda x - \log \mathbb{E}_\mu[\exp \lambda X_1])}.\end{aligned}$$

$$\begin{aligned}\text{Thus, } -\frac{1}{n} \log \mathbb{P}_\mu(\bar{X}_n \geq x) &\geq \sup_{\lambda > 0} \left\{ \lambda x - \log \mathbb{E}_\mu[\exp \lambda X_1] \right\} \\ &= \sup_{\lambda > 0} \left\{ \lambda x - \log(1 - \mu + \mu e^\lambda) \right\} \\ &= \text{kl}(x, \mu).\end{aligned}$$

kl = binary Kullback-Leibler divergence: $\text{kl}(x, \mu) = \text{KL}(\mathcal{B}(x), \mathcal{B}(\mu))$

$$\text{where } \text{KL}(P, Q) = \begin{cases} \int \log \frac{dP}{dQ} dP & \text{if } P \ll Q, \\ +\infty & \text{otherwise.} \end{cases}$$

Properties: $0 \leq \text{KL}(P, Q) \leq +\infty$, and $\text{KL}(P, Q) = 0$ iff $P = Q$.

Lower Bound: Change of Measure

For all $\epsilon > 0$ and all $\alpha > 0$,

$$\begin{aligned}\mathbb{P}_\mu(\bar{X}_n \geq x) &= \mathbb{E}_\mu[\mathbb{1}\{\bar{X}_n \geq x\}] \\ &= \mathbb{E}_{x+\epsilon} \left[\mathbb{1}\{\bar{X}_n \geq x\} \times \frac{d\mathbb{P}_\mu}{d\mathbb{P}_{x+\epsilon}}(X_1, \dots, X_n) \right] \\ &= \mathbb{E}_{x+\epsilon} \left[\mathbb{1}\{\bar{X}_n \geq x\} \times e^{-\sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu}(X_i)} \right] \\ &\geq \mathbb{E}_{x+\epsilon} \left[\mathbb{1}\{\bar{X}_n \geq x\} \mathbb{1}\left\{ \frac{1}{n} \sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu}(X_i) \leq \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu}(X_1) \right] + \alpha \right\} \right. \\ &\quad \left. \times e^{-\sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu}(X_i)} \right] \\ &\geq e^{-n \left\{ \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu}(X_1) \right] + \alpha \right\}} \left[1 - \mathbb{P}_{x+\epsilon}(\bar{X}_n < x) \right. \\ &\quad \left. - \mathbb{P}_{x+\epsilon} \left(\frac{1}{n} \sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu}(X_i) > \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu}(X_1) \right] + \alpha \right) \right] \\ &= e^{-n \{ \text{kl}(x+\epsilon, \mu) + \alpha \}} (1 - o_n(1)).\end{aligned}$$

Lower Bound: Change of Measure

For all $\epsilon > 0$ and all $\alpha > 0$,

$$\begin{aligned}
 \mathbb{P}_\mu (\bar{X}_n \geq x) &= \mathbb{E}_\mu [\mathbb{1}\{\bar{X}_n \geq x\}] \\
 &\geq \mathbb{E}_{x+\epsilon} \left[\mathbb{1}\{\bar{X}_n \geq x\} \mathbb{1}\left\{ \frac{1}{n} \sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu} (X_i) \leq \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu} (X_1) \right] + \alpha \right\} \right. \\
 &\quad \left. \times e^{-\sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu} (X_i)} \right] \\
 &\geq e^{-n \left\{ \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu} (X_1) \right] + \alpha \right\}} \left[1 - \mathbb{P}_{x+\epsilon} (\bar{X}_n < x) \right. \\
 &\quad \left. - \mathbb{P}_{x+\epsilon} \left(\frac{1}{n} \sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu} (X_i) > \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu} (X_1) \right] + \alpha \right) \right] \\
 &= e^{-n \{ \text{kl}(x+\epsilon, \mu) + \alpha \}} (1 - o_n(1)) .
 \end{aligned}$$

Asymptotic Optimality (Large Deviation Lower Bound)

$$\liminf_n \frac{1}{n} \log \mathbb{P}_\mu (\bar{X}_n \geq x) \geq -\text{kl}(x, \mu) .$$

Lower Bound: Change of Measure

For all $\epsilon > 0$ and all $\alpha > 0$,

$$\mathbb{P}_\mu (\bar{X}_n \geq x) = \mathbb{E}_\mu [\mathbb{1}\{\bar{X}_n \geq x\}]$$

$$\geq \mathbb{E}_{x+\epsilon} \left[\mathbb{1}\{\bar{X}_n \geq x\} \mathbb{1}\left\{ \frac{1}{n} \sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu} (X_i) \leq \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu} (X_1) \right] + \alpha \right\} \right. \\ \left. \times e^{-\sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu} (X_i)} \right]$$

$$\geq e^{-n \left\{ \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu} (X_1) \right] + \alpha \right\}} \left[1 - \mathbb{P}_{x+\epsilon} (\bar{X}_n < x) \right]$$

$$- \mathbb{P}_{x+\epsilon} \left(\frac{1}{n} \sum_{i=1}^n \log \frac{dP_{x+\epsilon}}{dP_\mu} (X_i) > \mathbb{E}_{x+\epsilon} \left[\log \frac{dP_{x+\epsilon}}{dP_\mu} (X_1) \right] + \alpha \right)$$

$$= e^{-n \{ \text{kl}(x+\epsilon, \mu) + \alpha \}} (1 - o_n(1)) .$$

Asymptotic Optimality (Large Deviation Principle)

$$\frac{1}{n} \log \mathbb{P}_\mu (\bar{X}_n \geq x) \xrightarrow{n \rightarrow \infty} -\text{kl}(x, \mu) .$$

Lower Bound: the Entropic Way

Let $\Omega = \{0, 1\}^n$, $X_i(\omega) = \omega_i$

Probability laws on Ω : $\mathbb{P}_p = \mathcal{B}(p)^{\otimes n}$.

For all $\epsilon > 0$,

$$n \text{kl}(x + \epsilon, \mu) = \text{KL}(\mathbb{P}_{x+\epsilon}, \mathbb{P}_\mu) \quad \text{KL}(P \otimes P', Q \otimes Q') = \text{KL}(P, Q) + \text{KL}(P', Q')$$


$$\geq \text{KL}\left(\mathbb{P}_{x+\epsilon}^{\mathbf{1}\{\bar{X}_n \geq x\}}, \mathbb{P}_\mu^{\mathbf{1}\{\bar{X}_n \geq x\}}\right) \quad \begin{array}{l} \text{KL}(P, Q) \geq \text{KL}(P^X, Q^X) \\ \text{contraction of entropy} \\ = \text{data-processing inequality} \end{array}$$

$$= \text{kl}\left(\mathbb{P}_{x+\epsilon}(\bar{X}_n \geq x), \mathbb{P}_\mu(\bar{X}_n \geq x)\right)$$

$$\geq \mathbb{P}_{x+\epsilon}(\bar{X}_n \geq x) \log \frac{1}{\mathbb{P}_\mu(\bar{X}_n \geq x)} - \log(2) \quad \text{kl}(p, q) \geq p \log \frac{1}{q} - \log 2$$

A non-asymptotic lower bound

$$\forall \epsilon > 0, \quad \mathbb{P}_\mu(\bar{X}_n \geq x) \geq e^{-\frac{n \text{kl}(x+\epsilon, \mu) + \log(2)}{1 - e^{-2n\epsilon^2}}}.$$



Intermediate: Sequential Test

Sequential Test : $\mu > 1/2$ versus $\mu < 1/2$

X_1, X_2, \dots independent random variables with distribution $\mathcal{B}(\mu)$.

- Goal: say if $\mu > 1/2$ or $\mu < 1/2$.
- Sequential: **stopping time** τ , decision rule $\hat{a}_\tau \in \{+, -\}$ is \mathcal{F}_τ -measurable, where $\mathcal{F}_t = \sigma(X_1, \dots, X_t)$.
- **δ -correctness** constraint:

$$\forall \mu < 1/2, \mathbb{P}_\mu(\hat{a}_\tau = +) \leq \delta \quad \text{and} \quad \forall \mu > 1/2, \mathbb{P}_\mu(\hat{a}_\tau = -) \leq \delta .$$

- a good procedure **minimizes** $\mathbb{E}_\mu[\tau]$ for all $\mu \neq 1/2$.

Entropic Lower Bound on the Sample Complexity

If $\mu > 1/2$:

$$\begin{aligned}\mathbb{E}_\mu[\tau] \operatorname{kl}\left(\mu, \frac{1}{2} - \epsilon\right) &= \operatorname{KL}\left(\mathbb{P}_\mu^{(X_1, \dots, X_\tau)}, \mathbb{P}_{\frac{1}{2} - \epsilon}^{(X_1, \dots, X_\tau)}\right) && \text{tensorization works with stopping times} \\ &\geq \operatorname{KL}\left(\mathbb{P}_\mu^{\mathbb{1}\{\hat{a}_\tau = +\}}, \mathbb{P}_{\frac{1}{2} - \epsilon}^{\mathbb{1}\{\hat{a}_\tau = +\}}\right) && \text{contraction just like before} \\ &= \operatorname{kl}\left(\mathbb{P}_\mu(\hat{a}_\tau = +), \mathbb{P}_{\frac{1}{2} - \epsilon}(\hat{a}_\tau = +)\right) \\ &\geq \operatorname{kl}(1 - \delta, \delta) && \text{by } \delta\text{-correctness.}\end{aligned}$$

Theorem

For every δ -correct stopping time τ and every $\mu \in [0, 1]$,

$$\mathbb{E}_\mu[\tau] \geq \frac{1}{\operatorname{kl}\left(\mu, \frac{1}{2}\right)} \operatorname{kl}(1 - \delta, \delta).$$

Remark: $\operatorname{kl}(\delta, 1 - \delta) \underset{\delta \rightarrow 0}{\sim} \log \frac{1}{\delta}$ and $\operatorname{kl}(\delta, 1 - \delta) \geq \log \frac{1}{2.4\delta}$.

Tight: possible to choose τ s.t. $\mathbb{E}_\mu[\tau] \leq \frac{\log \frac{1}{\delta}}{\operatorname{kl}\left(\mu, \frac{1}{2}\right)} \left(1 + o_\delta\left(\log \frac{1}{\delta}\right)\right)$.

Active Learning: Finding the Distribution with Largest Mean

Best-Arm Identification with Fixed Confidence

K options = $(\mathcal{B}(\mu_a))_{1 \leq a \leq K}$. Parameter $\mu = (\mu_1, \dots, \mu_K)$.



μ_1



μ_2



μ_3



μ_4



μ_5

At round t , you may:

- choose an option $A_t = \phi_t(A_1, X_1, \dots, A_{t-1}, X_{t-1}) \in \{1, \dots, K\}$
- observe a new independent sample $X_t \sim \mathcal{B}(\mu_{A_t})$


so as to identify the best arm $a^*(\mu) = \operatorname{argmax}_a \mu_a$ as fast as possible:
stopping time τ , decision \hat{a}_τ .

Goal: minimize $\mathbb{E}_\mu[\tau]$

under δ -correctness constraint: $\forall \mu, \mathbb{P}_\mu(\hat{a}_\tau \neq a^*(\mu)) \leq \delta$.

Best-Arm Identification with Fixed Confidence

K options = $(\mathcal{B}(\mu_a))_{1 \leq a \leq K}$. Parameter $\mu = (\mu_1, \dots, \mu_K)$.



At round t , you may:

- choose an option $A_t = \phi_t(A_1, X_1, \dots, A_{t-1}, X_{t-1}) \in \{1, \dots, K\}$
- observe a new independent sample $X_t \sim \mathcal{B}(\mu_{A_t})$

so as to identify the best arm $a^*(\mu) = \operatorname{argmax}_a \mu_a$ as fast as possible:
stopping time τ , decision \hat{a}_τ .

Goal: minimize $\mathbb{E}_\mu[\tau]$

under δ -correctness constraint: $\forall \mu, \mathbb{P}_\mu(\hat{a}_\tau \neq a^*(\mu)) \leq \delta$.

Racing Strategy see [Kaufmann & Kalyanakrishnan '13]

$\mathcal{R} := \{1, \dots, K\}$ set of **remaining arms**.

$r := 0$ current round

while $|\mathcal{R}| > 1$

- $r := r + 1$
- draw each $a \in \mathcal{R}$, compute $\hat{\mu}_{a,r}$, the empirical mean of the r samples observed so far
- compute the **empirical best** and **empirical worst** arms:

$$b_r = \operatorname{argmax}_{a \in \mathcal{R}} \hat{\mu}_{a,r} \quad w_r = \operatorname{argmin}_{a \in \mathcal{R}} \hat{\mu}_{a,r}$$


- Elimination step: if

$$\ell_{b_r}(r) > u_{w_r}(r),$$

then eliminate w_r : $\mathcal{R} := \mathcal{R} \setminus \{w_r\}$

end

Output: \hat{a} the single element in \mathcal{R} .



Entropic Lower Bound

Let $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)$ and $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_K)$ be two parameters with different maxima: $a^*(\boldsymbol{\mu}) \neq a^*(\boldsymbol{\lambda})$. Then

$$\begin{aligned} \sum_{a=1}^K \mathbb{E}_{\boldsymbol{\mu}} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) &= \text{KL} \left(\mathbb{P}_{\boldsymbol{\mu}}^{(X_1, \dots, X_\tau)}, \mathbb{P}_{\boldsymbol{\lambda}}^{(X_1, \dots, X_\tau)} \right) \\ &\geq \text{KL} \left(\mathbb{P}_{\boldsymbol{\mu}}^{\mathbb{1}\{\hat{a}_\tau = a^*(\boldsymbol{\mu})\}}, \mathbb{P}_{\boldsymbol{\lambda}}^{\mathbb{1}\{\hat{a}_\tau = a^*(\boldsymbol{\mu})\}} \right) \\ &\geq \text{kl} \left(\mathbb{P}_{\boldsymbol{\mu}}(\hat{a}_\tau = a^*(\boldsymbol{\mu})), \mathbb{P}_{\boldsymbol{\lambda}}(\hat{a}_\tau = a^*(\boldsymbol{\mu})) \right) \\ &\geq \text{kl}(1 - \delta, \delta) . \end{aligned}$$

Entropic Lower Bound

[Kaufmann, Cappé, G.'15],[G., Ménard, Stoltz '16]

For every δ -correct procedure, if $a^*(\boldsymbol{\mu}) \neq a^*(\boldsymbol{\lambda})$ then

$$\sum_{a=1}^K \mathbb{E}_{\boldsymbol{\mu}} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(1 - \delta, \delta) .$$

Using the Entropic Lower Bound

Let $\mu = (\mu_1, \dots, \mu_K)$ and $\lambda = (\lambda_1, \dots, \lambda_K)$ be two bandit models.

Entropic Lower Bound

If $a^*(\mu) \neq a^*(\lambda)$, any δ -correct algorithm satisfies

$$\sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta).$$

Using it for each arm separately, one obtains:

Proposition [Kaufmann, Cappé, G.'15]

For any δ -correct algorithm,


$$\mathbb{E}_{\mu}[\tau] \geq \left(\frac{1}{\text{kl}(\mu_1, \mu_2)} + \sum_{a=2}^K \frac{1}{\text{kl}(\mu_a, \mu_1)} \right) \text{kl}(\delta, 1 - \delta).$$

Combining the Entropic Lower Bounds

Entropic Lower Bound

If $a^*(\mu) \neq a^*(\lambda)$, any δ -correct algorithm satisfies

$$\sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta).$$



Let $\text{Alt}(\mu) = \{\lambda : a^*(\lambda) \neq a^*(\mu)\}$.

$$\inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$

$$\mathbb{E}_{\mu} [\tau] \times \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \frac{\mathbb{E}_{\mu} [N_a(\tau)]}{\mathbb{E}_{\mu} [\tau]} \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$


$$\mathbb{E}_{\mu} [\tau] \times \left(\sup_{w \in \Sigma_K} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K w_a \text{kl}(\mu_a, \lambda_a) \right) \geq \text{kl}(\delta, 1 - \delta)$$

Combining the Entropic Lower Bounds

Entropic Lower Bound

If $a^*(\mu) \neq a^*(\lambda)$, any δ -correct algorithm satisfies

$$\sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta).$$



Let $\text{Alt}(\mu) = \{\lambda : a^*(\lambda) \neq a^*(\mu)\}$.

$$\inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$

$$\mathbb{E}_{\mu} [\tau] \times \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \frac{\mathbb{E}_{\mu} [N_a(\tau)]}{\mathbb{E}_{\mu} [\tau]} \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$


$$\mathbb{E}_{\mu} [\tau] \times \left(\sup_{w \in \Sigma_K} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K w_a \text{kl}(\mu_a, \lambda_a) \right) \geq \text{kl}(\delta, 1 - \delta)$$

Combining the Entropic Lower Bounds

Entropic Lower Bound

If $a^*(\mu) \neq a^*(\lambda)$, any δ -correct algorithm satisfies

$$\sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta).$$



Let $\text{Alt}(\mu) = \{\lambda : a^*(\lambda) \neq a^*(\mu)\}$.

$$\inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$

$$\mathbb{E}_{\mu} [\tau] \times \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \frac{\mathbb{E}_{\mu} [N_a(\tau)]}{\mathbb{E}_{\mu} [\tau]} \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$


$$\mathbb{E}_{\mu} [\tau] \times \left(\sup_{w \in \Sigma_K} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K w_a \text{kl}(\mu_a, \lambda_a) \right) \geq \text{kl}(\delta, 1 - \delta)$$

Combining the Entropic Lower Bounds

Entropic Lower Bound

If $a^*(\mu) \neq a^*(\lambda)$, any δ -correct algorithm satisfies

$$\sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta).$$



Let $\text{Alt}(\mu) = \{\lambda : a^*(\lambda) \neq a^*(\mu)\}$.

$$\inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \mathbb{E}_{\mu} [N_a(\tau)] \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$

$$\mathbb{E}_{\mu} [\tau] \times \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K \frac{\mathbb{E}_{\mu} [N_a(\tau)]}{\mathbb{E}_{\mu} [\tau]} \text{kl}(\mu_a, \lambda_a) \geq \text{kl}(\delta, 1 - \delta)$$

$$\mathbb{E}_{\mu} [\tau] \times \left(\sup_{w \in \Sigma_K} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K w_a \text{kl}(\mu_a, \lambda_a) \right) \geq \text{kl}(\delta, 1 - \delta)$$

Lower Bound: the Complexity of BAI

Theorem [G. and Kaufmann, 2016]

For any δ -correct algorithm,

$$\mathbb{E}_{\mu}[\tau] \geq T^*(\mu) \text{kl}(\delta, 1 - \delta),$$

where

$$\begin{aligned} T^*(\mu)^{-1} &= \sup_{w \in \Sigma_K} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K w_a \text{kl}(\mu_a, \lambda_a) \\ &= \max_{w \in \Sigma_K} \min_{b \neq a^*} \inf_{\mu_b \leq \lambda \leq \mu_{a^*}} w_{a^*} \text{kl}(\mu_{a^*}, \lambda) + w_b \text{kl}(\mu_b, \lambda). \end{aligned}$$

- A kind of **game** : you choose the proportions of draws $(w_a)_a$, the opponent chooses the alternative.
- the **optimal proportions of arm draws** are

$$w^*(\mu) = \operatorname{argmax}_{w \in \Sigma_K} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=1}^K w_a \text{kl}(\mu_a, \lambda_a).$$

What if I want only an ϵ -optimal distribution?

Now, at round t , you may:

- choose an option $A_t = \phi_t(A_1, X_1, \dots, A_{t-1}, X_{t-1}) \in \{1, \dots, K\}$
- observe a new independent sample $X_t \sim \nu_{A_t}$

so as to identify **any ϵ -optimal** $a \in \mathcal{A}_\epsilon = \{a : \mu_a \geq \mu^* - \epsilon\}$ where $\mu^* = \max_a \mu_a$ as fast as possible: stopping time $\tau_{\delta, \epsilon}$.

\implies **minimize $\mathbb{E}[\tau_{\delta, \epsilon}]$** under the **PAC constraint** $\mathbb{P}_\mu(\mu_{\hat{a}_\tau} < \mu^* - \epsilon) \leq \delta$.

- PAC constraint: Probably Approximately Correct;
- more natural objective, especially in the context of discretized optimization;
- permits to avoid infinite loops in case of draws.

Apparently more complicated

Theorem [G. and Kaufmann'19]

For any ϵ, δ -PAC algorithm with converging proportions of draws,

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_{\mu}[\tau]}{\log \frac{1}{\delta}} \geq T_{\epsilon}^*(\mu),$$

where

$$T_{\epsilon}^*(\mu)^{-1} = \sup_{w \in \Sigma_K} \max_{a \in \mathcal{A}_{\epsilon}} \min_{b \neq a} \inf_{(\lambda_a, \lambda_b): \lambda_a \leq \lambda_b - \epsilon} w_a \text{kl}(\mu_a, \lambda_a) + w_b \text{kl}(\mu_b, \lambda_b).$$

- Asymptotic result (only).
- Assumption on the algorithm (convergence of the proportions of draws) – can be dropped.
- We do not manage to use the information-theoretic technique! We have to go back to the change of measure... even for the sequential test with multiple correct answers!

