A Proofs

The proof of Proposition 1 relies on the following variational form of Kullback–Leibler divergence, which is given in Theorem 5.2.1 of Robert Gray's textbook *Entropy and Information Theory* [10].

Fact 1. *Fix two probability measures* $\mathbf{P}$ *and* $\mathbf{Q}$ *defined on a common measureable space* $(\Omega, \mathcal{F})$*. Suppose that* $\mathbf{P}$ *is absolutely continuous with respect to* $\mathbf{Q}$*. Then*

$$D(\mathbf{P}||\mathbf{Q}) = \sup_X \left\{ \mathbf{E}_{\mathbf{P}}[X] - \log \mathbf{E}_{\mathbf{Q}}[e^X] \right\},$$

where the supremum is taken over all random variables X *such that the expectation of* X *under* $\mathbf{P}$ *is well defined, and* e^X *is integrable under* $\mathbf{Q}$*.*

Proof of Proposition 1.

$$\begin{aligned} I(T;\boldsymbol{\phi}) &= \sum_{i=1}^{n} \mathbf{P}(T=i) D\left(\mathbf{P}(\boldsymbol{\phi} = \cdot | T=i) \,||\, \mathbf{P}(\boldsymbol{\phi} = \cdot)\right) \\ &\geq \sum_{i=1}^{n} \mathbf{P}(T=i) D\left(\mathbf{P}(\phi_i = \cdot | T=i) \,||\, \mathbf{P}(\phi_i = \cdot)\right) \end{aligned}$$

Applying Fact 1 with $\mathbf{P} = \mathbf{P}(\phi_i = \cdot | T = i)$, $\mathbf{Q} = \mathbf{P}(\phi_i = \cdot)$, and $X = \lambda(\phi_i - \mu_i)$, we have

$$D\left(\mathbf{P}(\phi_i = \cdot | T=i) \,||\, \mathbf{P}(\phi_i = \cdot)\right) \geq \sup_\lambda \lambda \Delta_i - \lambda^2 \sigma^2 / 2$$

where $\Delta_i \equiv \mathbf{E}[\phi_i | T = i] - \mu_i$. Taking the derivative with respect to λ, we find that the optimizer is $\lambda = \Delta_i/\sigma^2$. This gives

$$2\sigma^2 I(T;\boldsymbol{\phi}) \geq \sum_{i=1}^{n} \mathbf{P}(T=i)\Delta_i^2 = \mathbf{E}[\Delta_T^2].$$

By the Tower property of conditional expectation and Jensen's inequality

$$\mathbf{E}[\phi_T - \mu_T] = \mathbf{E}[\Delta_T] \leq \sqrt{\mathbf{E}[\Delta_T^2]} \leq \sigma\sqrt{2I(T;\boldsymbol{\phi})}.$$

□

Proof of Proposition 2. Set

$$M_k = \max_{i \leq k} \phi_i = \sigma \max_{i \leq k} \frac{\phi_i}{\sigma}.$$

Basic facts about the maximum of independent standard Gaussian random variables then imply

$$\sigma^{-1}\mathbf{E}[M_k] \geq c\sqrt{2k} \qquad \& \qquad \sigma^{-1}\mathbf{E}[M_k] - \sqrt{2k} \underset{k\to\infty}{\longrightarrow} 0,$$

where c is a numerical constant that does not depend on k or σ. Now we have,

$$\mathbf{E}[\phi_{T_B}] = \mathbf{E}[M_{\lfloor B \rfloor}.]$$

The result then follows by observing that for $B \geq 2$

$$B \leq 2\lfloor B \rfloor$$

and

$$\sqrt{B} - \sqrt{\lfloor B \rfloor} \underset{B\to\infty}{\longrightarrow} 0.$$

□

Proof of Proposition 3. Following the same analysis as in the sub-Gaussian setting, we have

$$D\left(\mathbf{P}(\phi_i = \cdot | T = i) \,||\, \mathbf{P}(\phi_i = \cdot)\right) \geq \sup_{\lambda < 1/b} \lambda \Delta_i - \lambda^2 \sigma^2 / 2$$

The RHS is greater than the value from setting $\lambda = 1/b$. Therefore, we have

$$D\left(\mathbf{P}(\phi_i = \cdot | T = i) \,||\, \mathbf{P}(\phi_i = \cdot)\right) \geq \frac{\Delta_i}{b} - \frac{\sigma^2}{2b^2}.$$

Summing over $P(T = i)$ gives

$$\mathbf{E}[\phi_T - \mu_T] \leq bI(T; \boldsymbol{\phi}) + \frac{\sigma^2}{2b}.$$

When $b > 1$, $\lambda = 1/\sqrt{b} < 1/b$ is also a feasible point. Putting in this value of λ into the calculations above gives the second bound

$$\mathbf{E}[\phi_T - \mu_T] \leq \sqrt{b} I(T; \boldsymbol{\phi}) + \frac{\sigma^2}{2\sqrt{b}}.$$

□

Proof of Corollary 1. Let $Z_i = (\phi_i - \mu_i)^2$. Then, $Z_i \in [-C^2, C^2]$ and hence is C^2–sub-Gaussian. This implies that

$$\mathbf{E}[Z_T - \mu_T] \leq C\sqrt{2I(T; \boldsymbol{Z})} \leq C\sqrt{2I(T; \boldsymbol{\phi})}$$

where the last step follows from the data-processing inequality for mutual-information, and the fact that Z_i is a deterministic function of ϕ_i. □

Proof of Proposition 4.

$$\begin{aligned}
D\left(\mathbf{P}(\phi_T = \cdot) \,||\, \mathbf{P}(\tilde{\phi}_T = \cdot)\right) &\leq D\left(\mathbf{P}(\phi_T = \cdot, T = \cdot) \,||\, \mathbf{P}(\tilde{\phi}_T = \cdot, T = \cdot)\right) \\
&= \sum_{T=1}^{m} \mathbf{P}(T = i) D\left(\mathbf{P}(\phi_T = \cdot | T = i) \,||\, \mathbf{P}(\tilde{\phi}_T = \cdot | T = i)\right) \\
&= \sum_{T=1}^{m} \mathbf{P}(T = i) D\left(\mathbf{P}(\phi_i = \cdot | T = i) \,||\, \mathbf{P}(\phi_i = \cdot)\right) \\
&\leq \sum_{T=1}^{m} \mathbf{P}(T = i) D\left(\mathbf{P}(\boldsymbol{\phi} = \cdot | T = i) \,||\, \mathbf{P}(\boldsymbol{\phi} = \cdot)\right) \\
&= I(T; \boldsymbol{\phi}),
\end{aligned}$$

where both inequalities follow from the data-processing inequality for KL divergence. □

Proof of Proposition 5. Since $\phi_i \sim \text{Uniform}(0, 1)$, $Z_{\epsilon,i} = \mathbf{1}(\phi_i < \epsilon)$ is a Bernoulli random variable with parameter ϵ and $E[Z_i] = \epsilon$. We use the fact that a probability p Bernoulli random variable is sub-Gaussian with parameter [4]

$$\sigma = \sqrt{\frac{1 - 2p}{2\log((1-p)/p)}} \leq \sqrt{\frac{1}{2\log(1/2p)}}.$$

Combining this with Proposition 1, we have the desired result

$$E[Z_T] - E[\mu_T] = \mathbb{P}(p_T < \epsilon) - \epsilon \leq \sqrt{\frac{I(T; \mathbf{Z}_\epsilon)}{\log(1/2\epsilon)}}.$$

The second inequality follows by the data-processing inequality. □

Proof of Proposition 6. Note that

$$\max_{x\in\mathcal{X}} \mu(x) \geq \mathbf{E}[\mu(X^*)]$$

and

$$\mathbf{E}[\max_{x\in\mathcal{X}} f_\theta(x)] = \mathbf{E}[f_\theta(X^*)]$$

Therefore,

$$\mathbf{E}[\max_{x\in\mathcal{X}} f_\theta(x)] - \max_{x\in\mathcal{X}} \mu(x) \leq \mathbf{E}[f_\theta(X^*)] - E[\mu(X^*)] \leq \sigma\sqrt{2I(X^*;\theta)} = \sigma\sqrt{2H(X^*)}$$

□

Proof of Lemma 1. Since, conditional on H_k, T_{k+1} is independent of $\boldsymbol{\phi}$, the data-processing inequality for mutual information implies,

$$I(T_{k+1};\boldsymbol{\phi}) \leq I(H_k;\boldsymbol{\phi}).$$

Now we have,

$$I(H_k;\boldsymbol{\phi}) = \sum_{i=1}^{k} I\left((T_i, Y_{T_i});\boldsymbol{\phi}|H_{i-1}\right).$$

We complete the proof by simplifying the expression for $I\left((T_i, Y_{T_i});\boldsymbol{\phi}|H_{i-1}\right)$. Let $\boldsymbol{\phi}_{(-i)} = (\phi_j : j \neq i)$. Then,

$$\begin{aligned}
I\left((T_i, Y_{T_i});\boldsymbol{\phi}|H_{i-1}\right) &= I\left(T_i;\boldsymbol{\phi}|H_{i-1}\right) + I\left(Y_{T_i};\boldsymbol{\phi}|H_{i-1}, T_i\right) \\
&= I\left(Y_{T_i};\boldsymbol{\phi}|H_{i-1}, T_i\right) \\
&= I(Y_{T_i};\phi_{T_i}|H_{i-1}, T_i) + I(Y_{T_i};\phi_{(-T_i)}|H_{i-1}, T_i, \boldsymbol{\phi}_{T_i}) \\
&= I(Y_{T_i};\phi_{T_i}|H_{i-1}, T_i),
\end{aligned}$$

where the final equality follows because, conditioned on ϕ_{T_i}, Y_{T_i} is independent of $\boldsymbol{\phi}_{(-T_i)}$. □

Proof of Lemma 2.

$$I(X;Y) = -\frac{1}{2}\log\left(1 - \frac{\sigma_1^2}{\sigma_1^2+\sigma_2^2}\right) = -\frac{1}{2}\log\frac{\sigma_2^2}{\sigma_1^2+\sigma_2^2} = \frac{1}{2}\log\left(1+\frac{\sigma_1^2}{\sigma_2^2}\right).$$

□

Proof of Proposition 7. In order to apply our general result, we need to convert the bound on the differences in the expectations of a random variable into a bound on probability. Let

$$Z_i = \mathbf{1}(\phi_i - \mu_i > \tau) \qquad i \in \mathbb{N},$$

so that Z_i is a bernoulli random variable with expectation

$$e_i = \mathbf{E}[Z_i] \leq \exp\{-\frac{n\tau^2}{2}\}$$

It is known [4] that a bernoulli random variable with parameter $p < \frac{1}{2}$ is σ–sub-Gaussian with

$$\sigma = \sqrt{\frac{1-2p}{2\log((1-p)/p)}} \leq \sqrt{\frac{1}{2\log(1/2p)}}$$

and applying this with $p = \mathbf{E}[Z_i]$ shows that Z_i is sub-Gaussian with effective standard deviation less than $\sqrt{\frac{c}{n\tau^2}}$ where c is a universal numerical constant. This shows that

$$\begin{aligned}
\mathbf{E}[Z_{T_{k+1}} - e_{T_{k+1}}] &\leq \frac{\sqrt{c}}{\tau}\sqrt{\frac{2I(T_{k+1};\boldsymbol{\phi})}{n}} \\
&\leq \frac{\sqrt{c}}{\tau}\sqrt{\frac{2\sum_{i=1}^{k} I(Y_{T_i};\phi_{T_i}|Y_{T_1},..,Y_{T_{i-1}})}{n}}
\end{aligned}$$