A Note on divergences indexed by α

Many divergences in the literature are indexed with the parameter α (see Table 2). These divergences turn out to be equivalent to the Rényi divergence as we can identify one-to-one correspondences between them.

Divergence	Formulation
I-divergence (Nielsen and Nock (2011))	$D^I_\alpha[p\|\|q] = \int_{\mathcal{S}} p^\alpha q^{(1-\alpha)} ds$
Amari's α divergence (S. Amari (2009))	$D^{AM}_\alpha[p\|\|q] = \frac{4}{1-\alpha^2}(1 - \int_{\mathcal{S}} p^{\frac{1+\alpha}{2}} q^{\frac{1-\alpha}{2}} ds)$
Tsallis' divergence (Nielsen and Nock (2011))	$D^T_\alpha[p\|\|q] = \frac{1}{\alpha-1}(\int_{\mathcal{S}} p^\alpha q^{1-\alpha} ds - 1)$
Rényi divergence	$D_\alpha[p\|\|q] = \frac{1}{(\alpha-1)} \log \int_{\mathcal{S}} p^\alpha q^{(1-\alpha)}\, \mathrm{d}s$

Table 2: Divergence families indexed with α. Amari's α-divergence plays an important role in information geometry as it induces a dually-flat geometry on the space of probability measures, and furthermore, when extended to positive measures, it is the only intersection between f-divergences and Bregman divergences, two important families of divergences (S. Amari (2009); Ay and Gibilisco (2016)).

All of the divergences shown in Table 2 are equivalent, in the sense that there are one-to-one mappings between them.

The Tsallis and Amari's divergences are linear functions of the I-divergence:

$$D^T_\alpha[p||q] = \frac{1}{\alpha - 1}(D^I_\alpha[p||q] - 1) \tag{43}$$

$$D^{AM}_\alpha[p||q] = \frac{4}{1-\alpha^2}(1 - D^I_{\frac{1+\alpha}{2}}[p||q]) \tag{44}$$

As a consequence, the Amari α divergence is a scalar multiple of the Tsallis divergence, under the correspondence $\beta = \frac{1+\alpha}{2}$:

$$D_{\alpha}^{AM}[p||q] = \frac{1}{\beta} D_{\beta}^{T}[p||q] \tag{45}$$

Finally, the Rényi divergence is a monotonic function of the I-divergence:

$$D_{\alpha}[p||q] = \frac{1}{\alpha - 1} \log D_{\alpha}^{I}[p||q] \tag{46}$$

B Derivations

B.1 Negative variational free energy for Gaussian-Gamma distribution

Here, we work through the variational free energy for the system described in Section 4. s, o are the random variables of interest, x the parameter governing the mean and λ_k is the precision parameter:

$$p(s, \lambda_p) = \mathcal{N}(0, (\lambda_p \sigma_p)^{-1}) Gam(\alpha_p, \beta_p) \tag{47}$$

$$p(o|s) = \mathcal{N}(sx, \Sigma_l) \tag{48}$$

$$q(s) = \mathcal{N}(\mu_q, \Sigma_q) \tag{49}$$

where $\Sigma_k = (\lambda_k \sigma_k)^{-1}$,The probability density functions are defined as:

$$p(s,\lambda_p) = \frac{|\lambda_p \sigma_p|^{1/2}}{(2\pi)^{1/2}} \exp\Big[\frac{-\lambda_p}{2} s^T \sigma_p s\Big] \frac{\beta_p^{\alpha_p}}{\Gamma(\alpha_p)} \lambda_p^{\alpha_p - 1} \exp\Big[-\lambda_p \beta_p\Big] \tag{50}$$

$$p(o|s) = \frac{|\Sigma_l|^{-1/2}}{(2\pi)^{n/2}} \exp\Big[-\frac{1}{2}(o - sx)^T \Sigma_l^{-1} (o - sx)\Big] \tag{51}$$

$$q(s) = \frac{|\Sigma_q|^{-1/2}}{(2\pi)^{1/2}} \exp\Big[-\frac{1}{2}(s - \mu_q)^T \Sigma_q^{-1} (s - \mu_q)\Big] \tag{52}$$

We use probability distributions to derive the quantity of interest: $\mathbb{E}_{q(s)}[\log p(s,o) - \log q(s)] = -D_{KL}[q(s)||p(s,o)]$:

$$-D_{KL}[q(s)||p(s,o)] = -\int_{\mathcal{S}} q(s) \log\left(\frac{q(s)}{p(s,o)}\right) \mathrm{d}s = \tag{53}$$

$$= -\int_{\mathcal{S}} q(s) \log\left[\frac{(2\pi)^{\frac{n+1}{2}} |\Sigma_p|^{1/2} |\Sigma_l|^{1/2}}{(2\pi)^{1/2} |\Sigma_q|^{1/2}} \frac{\beta_p^{\alpha_p} \lambda_p^{\alpha_p - 1} \exp(-\lambda_p \beta_p)}{\Gamma(\alpha_p)}\right] \mathrm{d}s + \tag{54}$$

$$+ \int_{\mathcal{S}} q(s) \left[\frac{1}{2}(s - \mu_q)^2 \Sigma_q^{-1} - \frac{1}{2}\Big((o - sx)^T \Sigma_l^{-1} (o - sx) + s^T \Sigma_p^{-1} s\Big)\right] \mathrm{d}s = \tag{55}$$

$$= \frac{1}{2} \log\left[\frac{|\Sigma_q|}{(2\pi)^n |\Sigma_p||\Sigma_l|}\right] + \log\left[\frac{\beta_p^{\alpha_p} \lambda_p^{\alpha_p - 1}}{\Gamma(\alpha_p)}\right] - \lambda_p \beta_p + \tag{56}$$

$$+ \int_{\mathcal{S}} q(s) \left[-\frac{1}{2}[s^2(\Sigma_p^{-1} + x^T \Sigma_l^{-1} x - \Sigma_q^{-1}) - 2s(-\mu_q \Sigma_q^{-1} + x^T \Sigma_l^{-1} o) - \mu_q^2 \Sigma_q^{-1} + o^T \Sigma_l^{-1} o]\right] \mathrm{d}s \tag{57}$$

Consider the last integral:

$$-\frac{1}{2}(\Sigma_p^{-1} + x^T \Sigma_l^{-1} x - \Sigma_q^{-1}) \int_{\mathcal{S}} s^2 q(s)\, \mathrm{d}s + (-\mu_q \Sigma_q^{-1} + x^T \Sigma_l^{-1} o) \int_{\mathcal{S}} s q(s)\, \mathrm{d}s - \frac{1}{2}(-\mu_q^2 \Sigma_q^{-1} + o^T \Sigma_l^{-1} o) \int_{\mathcal{S}} q(s)\, \mathrm{d}s = \tag{58}$$

$$-\frac{1}{2}(\Sigma_p^{-1} + x^T \Sigma_l^{-1} x - \Sigma_q^{-1})(\Sigma_q + \mu_q^2) + (-\mu_q \Sigma_q^{-1} + x^T \Sigma_l^{-1} o)\mu_q - \frac{1}{2}(-\mu_q^2 \Sigma_q^{-1} + o^T \Sigma_l^{-1} o) = \tag{59}$$

$$-\frac{1}{2}(\Sigma_q \Sigma_p^{-1} + \Sigma_q x^T \Sigma_l^{-1} x - 1 + \mu_q^2 \Sigma_p^{-1} + \mu_q^2 x^T \Sigma_l^{-1} x - 2\mu_q x^T \Sigma_l^{-1} o + o^T \Sigma_l^{-1} o) \tag{60}$$

Combining the results we have:

$$-D_{KL}[q(s)||p(s,o)] = \frac{1}{2}\log\left(\frac{|\Sigma_q|}{(2\pi)^n|\Sigma_p||\Sigma_l|}\right) \tag{61}$$

$$-\frac{1}{2}\left(o^T\Sigma_l^{-1}o + \mu_q^2\Sigma_p^{-1} + \mu_q^2 x^T\Sigma_l^{-1}x - 2\mu_q x^T\Sigma_l^{-1}o\right) \tag{62}$$

$$-\frac{1}{2}\left(\Sigma_q x^T\Sigma_l^{-1}x + \Sigma_q\Sigma_p^{-1} - 1\right) \tag{63}$$

$$+\log\left[\frac{\beta_p^{\alpha_p}\lambda_p^{\alpha_{p-1}}}{\Gamma(\alpha_p)}\right] - \lambda_p\beta_p \tag{64}$$

B.2 Rényi bound for Gaussian distribution

The probability density function for the random variables, s, o and x is the parameter governing the means:

$$p(s) = \frac{1}{(2\pi)^{\frac{1}{2}}|\Sigma_p|^{\frac{1}{2}}}\exp\left[-\frac{1}{2}s^T\Sigma_p^{-1}s\right] \tag{65}$$

$$p(o|s) = \frac{1}{(2\pi)^{\frac{n}{2}}|\Sigma_l|^{\frac{1}{2}}}\exp\left[-\frac{1}{2}(o-sx)^T\Sigma_l^{-1}(o-sx)\right] \tag{66}$$

$$q(s) = \frac{1}{(2\pi)^{\frac{1}{2}}|\Sigma_q|^{\frac{1}{2}}}\exp\left[(-\frac{1}{2}(s-\mu_q)^T\Sigma_q^{-1}(s-\mu_q)\right] \tag{67}$$

We now supplement these quantities into the negative Rényi bound, and rewrite using the defined quantities:

$$-D_\alpha[q(s)||p(s,o)] = \frac{1}{1-\alpha}\log\int_{\mathcal{S}} q(s)^\alpha p(s,o)^{1-\alpha}\,\mathrm{d}s \tag{68}$$

$$= \frac{1}{1-\alpha}\log\left(\frac{1}{(2\pi)^{\alpha/2}(2\pi)^{(1-\alpha)\frac{n+1}{2}}|\Sigma_q|^{\alpha/2}|\Sigma_p|^{(1-\alpha)\frac{1}{2}}|\Sigma_l|^{(1-\alpha)\frac{1}{2}}}\right) \tag{69}$$

$$+ \frac{1}{1-\alpha}\log\int_{\mathcal{S}}\exp\Bigg(-\frac{1}{2}\Bigg[\alpha[(s-\mu_q)^T\Sigma_q^{-1}(s-\mu_q)]+ \tag{70}$$

$$(1-\alpha)[o^T\Sigma_l^{-1}o - 2s^Tx^T\Sigma_l^{-1}o + s^T[\Sigma_p^{-1}+x^T\Sigma_l^{-1}x]s]\Bigg]\Bigg)\,\mathrm{d}s \tag{71}$$

$$= \frac{1}{2}\log\left(\frac{\Sigma_q^{\frac{\alpha}{\alpha-1}}}{(2\pi)^{\frac{n(1-\alpha)-1}{1-\alpha}}|\Sigma_p||\Sigma_l|}\right) \tag{72}$$

$$-\frac{1}{\alpha-1}\log\int_{\mathcal{S}}\exp\Bigg(-\frac{1}{2}\Bigg[s^T(\alpha\Sigma_q^{-1}+x^T\Sigma_l^{-1}x(1-\alpha)+\Sigma_p^{-1}(1-\alpha))s \tag{73}$$

$$-2s((1-\alpha)x^T\Sigma_l^{-1}o+\alpha\mu_q\Sigma_q^{-1})+\alpha\mu_q^2\Sigma_q^{-1}+(1-\alpha)o^T\Sigma_l^{-1}o\Bigg]\Bigg)\,\mathrm{d}s \tag{74}$$

First, let us focus on the term inside the integral. To avoid clutter we replace: $\Sigma_\alpha^{-1} := \alpha\Sigma_q^{-1} + x^T\Sigma_l^{-1}x(1-\alpha) + \Sigma_p^{-1}(1-\alpha)$ and assume it is invertible. We define $\mu_\alpha := \Sigma_\alpha(\alpha\mu_q\Sigma_q^{-1} + (1-\alpha)x^T\Sigma_l^{-1}o)$. Then Eq. 73 and 74 can be rewritten as:

$$-\frac{1}{2}\frac{-(1-\alpha)o^T\Sigma_l^{-1}o - \alpha\mu_q^2\Sigma_q^{-1}+\mu_\alpha^2\Sigma_\alpha^{-1}}{\alpha-1} \tag{75}$$

$$-\frac{1}{\alpha-1}\log\int_{\mathcal{S}}\frac{(2\pi)^{1/2}|\Sigma_\alpha|^{1/2}}{(2\pi)^{1/2}|\Sigma_\alpha|^{1/2}}\exp\left(-\frac{1}{2}(s-\mu_\alpha)^T\Sigma_\alpha^{-1}(s-\mu_\alpha)\right)\mathrm{d}s \tag{76}$$

$$= -\frac{1}{2(1-\alpha)}\left[(1-\alpha)o^T\Sigma_l^{-1}o+\alpha\mu_q^2\Sigma_q^{-1}-\mu_\alpha^2\Sigma_\alpha^{-1}\right]+\frac{1}{2}\log((2\pi)^{\frac{1}{1-\alpha}}\Sigma_\alpha^{\frac{1}{1-\alpha}}) \tag{77}$$

Putting it all together:

$$D_\alpha[q(s)||p(s,o)] = \frac{1}{2}\log\left(\frac{\Sigma_q^{\frac{\alpha}{\alpha-1}}|\Sigma_\alpha^{-1}|^{\frac{1}{\alpha-1}}}{(2\pi)^n|\Sigma_p||\Sigma_l|}\right) \tag{78}$$

$$-\frac{1}{2}\left[o^T\Sigma_l^{-1}o - \frac{\alpha}{(\alpha-1)}\mu_q^2\Sigma_q^{-1} + \frac{1}{(\alpha-1)}\mu_\alpha^2\Sigma_\alpha^{-1}\right] \tag{79}$$

With this formulation, we turn to the first term:

$$\frac{1}{2}\log\left[\frac{|\Sigma_q|^{\frac{\alpha}{\alpha-1}}|\Sigma_\alpha^{-1}|^{\frac{1}{\alpha-1}}}{(2\pi)^n|\Sigma_p||\Sigma_l|}\right] \tag{80}$$

$$= \frac{1}{2}\log\left[\frac{|\Sigma_q|}{(2\pi)^n|\Sigma_p||\Sigma_l|}\right] - \frac{1}{2}\log(\Sigma_q\Sigma_\alpha^{-1})^{\frac{1}{1-\alpha}} \tag{81}$$

$$= \frac{1}{2}\log\left[\frac{|\Sigma_q|}{(2\pi)^n|\Sigma_p||\Sigma_l|}\right] \tag{82}$$

$$-\frac{1}{2(1-\alpha)}\log\left(1 + (1-\alpha)(\Sigma_q x^T\Sigma_l^{-1}x + \Sigma_q\Sigma_p^{-1} - 1)\right) \tag{83}$$

Now, let us consider the second term:

$$-\frac{1}{2}\left[o^T\Sigma_l^{-1}o-\frac{\alpha}{(\alpha-1)}\mu_q^2\Sigma_q^{-1}+\frac{1}{(\alpha-1)}\mu_\alpha^2\Sigma_\alpha^{-1}\right] \tag{84}$$

$$=-\frac{1}{2}\left[o^T\Sigma_l^{-1}o-\frac{\alpha}{(\alpha-1)}\mu_q^2\Sigma_q^{-1}\right. \tag{85}$$

$$\left.+\frac{1}{(\alpha-1)}\frac{\alpha^2\mu_q^2(\Sigma_q^{-1})^2+(1-\alpha)^2(x^T\Sigma_l^{-1}o)^2+2\alpha(1-\alpha)\mu_q\Sigma_q^{-1}x^T\Sigma_l^{-1}o}{\Sigma_\alpha^{-1}}\right] \tag{86}$$

$$=-\frac{1}{2}\left[o^T\Sigma_l^{-1}o+\frac{-\alpha^2\mu_q^2(\Sigma_q^{-1})^2-(1-\alpha)\alpha\mu_q^2\Sigma_q^{-1}\Sigma_p^{-1}}{(\alpha-1)(\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x))}\right. \tag{87}$$

$$-\frac{(1-\alpha)\alpha\mu_q^2\Sigma_q^{-1}x^T\Sigma_l^{-1}x+\alpha^2\mu_q^2(\Sigma_q^{-1})^2+(1-\alpha)^2(x^T\Sigma_l^{-1}o)^2}{(\alpha-1)(\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x))} \tag{88}$$

$$\left.+\frac{2\alpha(1-\alpha)\mu_q\Sigma_q^{-1}x^T\Sigma_l^{-1}o}{(\alpha-1)(\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x))}\right] \tag{89}$$

$$=-\frac{1}{2}\left[o^T\Sigma_l^{-1}o+\frac{\alpha\mu_q^2\Sigma_q^{-1}\Sigma_p^{-1}+\alpha\mu_q^2\Sigma_q^{-1}x^T\Sigma_l^{-1}x}{\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x)}\right. \tag{90}$$

$$\left.+\frac{-(1-\alpha)(x^T\Sigma_l^{-1}o)^2-2\alpha\mu_q\Sigma_q^{-1}x^T\Sigma_l^{-1}o}{\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x)}\right] \tag{91}$$

$$=-\frac{1}{2}\left[\frac{\alpha\Sigma_q^{-1}o^T\Sigma_l^{-1}o+(1-\alpha)\Sigma_p^{-1}o^T\Sigma_l^{-1}o+(1-\alpha)o^T\Sigma_l^{-1}ox^T\Sigma_l^{-1}x}{\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x)}\right. \tag{92}$$

$$\left.+\frac{\alpha\mu_q^2\Sigma_q^{-1}\Sigma_p^{-1}+\alpha\mu_q^2\Sigma_q^{-1}x^T\Sigma_l^{-1}x-(1-\alpha)(x^T\Sigma_l^{-1}o)^2-2\alpha\mu_q\Sigma_q^{-1}x^T\Sigma_l^{-1}o}{\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x)}\right] \tag{93}$$

$$=-\frac{1}{2}\left[\frac{\alpha\Sigma_q^{-1}o^T\Sigma_l^{-1}o+(1-\alpha)\Sigma_p^{-1}o^T\Sigma_l^{-1}o+\alpha\mu_q^2\Sigma_q^{-1}\Sigma_p^{-1}}{\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x)}\right. \tag{94}$$

$$\left.+\frac{\alpha\mu_q^2\Sigma_q^{-1}x^T\Sigma_l^{-1}x-2\alpha\mu_q\Sigma_q^{-1}x^T\Sigma_l^{-1}o}{\alpha\Sigma_q^{-1}+(1-\alpha)(\Sigma_p^{-1}+x^T\Sigma_l^{-1}x)}\right] \tag{95}$$

$$=-\frac{\alpha}{2\Sigma_q\Sigma_\alpha^{-1}}\left[o^T\Sigma_l^{-1}o+\Sigma_q\frac{(1-\alpha)}{\alpha}\Sigma_p^{-1}o^T\Sigma_l^{-1}o+\mu_q^2\Sigma_p^{-1}\right. \tag{96}$$

$$\left.+\mu_q^2x^T\Sigma_l^{-1}x-2\mu_qx^T\Sigma_l^{-1}o\right] \tag{97}$$

From this, the simplified formulation for the Rényi bound is:

$$D_\alpha[q(s)||p(s,o)] = \frac{1}{2}\log\left(\frac{|\Sigma_q|}{(2\pi)^n|\Sigma_p||\Sigma_l|}\right) \tag{98}$$

$$-\frac{\alpha}{2(\Sigma_q\Sigma_\alpha^{-1})}\left(o^T\Sigma_l^{-1}o + \mu_q^2\Sigma_p^{-1} + \mu_q^2 x^T\Sigma_l^{-1}x - 2\mu_q x^T\Sigma_l^{-1}o\right) \tag{99}$$

$$-\frac{1}{2(1-\alpha)}\log\left(1 + (1-\alpha)(\Sigma_q x^T\Sigma_l^{-1}x + \Sigma_q\Sigma_p^{-1} - 1)\right) \tag{100}$$

$$-\frac{1}{2\Sigma_\alpha^{-1}}\left((1-\alpha)\Sigma_p^{-1}o^T\Sigma_l^{-1}o\right) \tag{101}$$

C MAB experiment details

We implemented the MAB simulations as described in Algorithm 1.

Algorithm 1 MAB optimisation using Rényi variational inference

---

Input: Variational density $q(s)$ for each arm. Empty observation buffer D_i for each arm i
Output : Optimal arm selection
Initialise μ_q^i, Σ_q^i for each arm i
repeat:
  for each arm i **do**:
    Sample one $s_i \sim q(\cdot|\mu_q^i, \Sigma_q^i)$
  end for
  Compute $i^* = \arg\max_i \frac{s_i}{\Sigma_q^i}$
  Pull arm i^*, receive reward R_i^* and store it in D_{i^*}
  for each arm i **do**:
    Update variational parameters by stochastic gradient descent (ADAM):
    $\nabla_{\mu_q^i} D_\alpha(q(s)||p(s,o))$
    $\nabla_{\Sigma_q^i} D_\alpha(q(s)||p(s,o))$
  end for
until convergence

For learning, the experiments were parameterised as:

- 4000 iterations for each simulation.
- Rényi bound was optimised using ADAM (Kingma and Ba (2014)) with a learning rate of $2e-2$ and 10 gradient steps for each update.

- 300 Monte-Carlo samples were used to update the variational posterior, $q(s)$ at each iteration

For each simulation, the prior specification is shown in Table 3, and generative process in Table 4.

	μ	Σ	Weights
$q(s)$:	25	$1e-8$	$\cdot$
Arm 1	$13, 20$	$1.5, 1.5$	0.5
Arm 2	$16, 14$	$1.5, 1.5$	0.5
Arm 3	$10, 17$	$1.5, 1.5$	0.5

Table 3: MAB multi-modal priors.

	μ	Σ	Weights
Arm 1	$10, 22$	$1, 1$	$0.97, 0.03$
Arm 2	16	3	$\cdot$
Arm 3	$10, 10$	$1, 1$	$0.97, 0.03$

Table 4: MAB multi-modal generative process.