Supplementary Methods

Hierarchical Bayesian inference for concurrent model fitting and comparison for group studies

Payam Piray[1,*], Amir Dezfouli[2], Tom Heskes[3], Michael J. Frank[4] and Nathaniel D. Daw[1]

1. Princeton Neuroscience Institute, Princeton University
2. University of New South Wales Sydney
3. Institute for Computing and Information Sciences, Radboud University
4. Department of Cognitive, Linguistics, and Psychological Sciences, Brown University

* Corresponding author: Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08540, USA. Email: ppiray@princeton.edu.

Simulation analyses

Here, we explain three models that were used frequently in our simulation analyses. For simulation analyses presented in Figures 1-2, data have been generated by two different models, Kalman filter (KF) and a reinforcement learning (RL) model. These models were also used in other simulations. Both models estimate the value of each action a_t, $Q_t(a_t)$, and take an action among two choices on every trial according to a softmax rule with the parameter $\beta > 0$. Both models learn using a prediction error signal, δ_t, the difference between the seen and expected reward:

$$\delta_t = r_t - Q_t(a_t)$$

where r_t is the reward. The RL model then updates its action value according to the prediction error:

$$Q_{t+1}(a_t) = Q_t(a_t) + \alpha\delta_t$$

where $0 < \alpha < 1$ is the learning rate. The update rule of the KL model is slightly different as it contains a dynamic learning rate (or Kalman gain), K_t, depending on the current estimate of variance, $V_t(a_t)$:

$$Q_{t+1}(a_t) = Q_t(a_t) + K_t\delta_t$$

where

$$K_t = \frac{V_t(a_t)}{V_t(a_t) + \omega}$$

where $\omega > 0$ is the observation noise. The variance also gets updated on every trial:

$$V_{t+1}(a_t) = (1 - K_t)V_t(a_t)$$

Note that all methods assume that parameters are normally distributed. Therefore, we used the sigmoid function, $\alpha = \sigma(\bar{\alpha}) = (1 + e^{-\bar{\alpha}})^{-1}$, to transform the normally-distributed learning rate, $\bar{\alpha}$, to the unit range. We used the exponential function to transform normally-distributed parameters corresponding to ω and β. For simulation analyses presented in Figures 1-2, these values were used as the group mean: RL, we used $\alpha = 0.1$, $\beta = 2$; Kalman filter $\omega = 2$, $\beta = 2$. For generating synthetic datasets for simulations, the parameters of the group of subjects assigned to each model was drawn from a normal distribution (with the standard deviation of 0.5). Next, for generating each individual dataset, we first generated two Gaussian random-walk (in the range of [-1 1]) sequences of 100 trials (corresponding to two actions). The learning rule of the corresponding model with the generated individual parameters were then used to generate a sequence of reward time-series for each action (both action values initialized at zero). Those reward time-series were then binarized and used to generate choice data given the corresponding model according to the softmax rule with the individual softmax parameter. The number of trials for all simulation was T=100, unless it is specified otherwise. The same procedure was employed for generating individual synthetic datasets in all simulations.

The dual-α RL model is similar to the RL model, with the only difference that two different learning rates, α^+ and α^- (both in the unit range), have been used for updating action values depending on whether the prediction error is positive or negative. Again, the sigmoid function was used to transform the normally distributed parameters into the unit range. For simulation analyses presented in Figures 3-6, these values were used as the group mean: RL, $\alpha = 0.1$, $\beta = 1$; dual-α RL: $\alpha^+ = 0.8$, $\alpha^- = 0.4$, $\beta = 3$.

For simulation analysis presented in Figure 7, these parameters used for the RL: $\alpha = 0.3$, $\beta = 3$. In scenario 1, the outliers were generated using the RL model with a very low decision noise parameter, i.e. $\alpha = 0.3$, $\beta = 0.1$. In scenario 2, the outliers were generated using RL with a low learning rate and a low decision noise parameter $\alpha = 0.1$, $\beta = 0.1$.

These parameters were used for all scenarios presented in Figure 8: the same parameters for the RL and the dual-α RL that used for generating data in Figures 3-6, the same parameters for the Kalman filter that used for generating data in Figure 1. These parameters used for the actor-critic RL: $\alpha^c = 0.1$, $\alpha^a = 0.6$, $\beta = 3$. For simulation analysis of the two-step task (Figure 9), we used parameters of the hybrid model reported in [1] (or the corresponding subset for the model-based or the model-free accounts) to generate data (201 trials). 30, 10 and 10 artificial subjects were generated using the hybrid, model-based and pure

model-free accounts. For generating the individual datasets, the outcome associated with each of the four second-stage options were generated according to a Gaussian random walk (the same sequences used in [2] were used here).

For the analysis presented in Figures 10A, 11A and 12A, data were generated using the same learning rate and decision parameters used for generating data in Figure 1. For the analysis presented in Figure 11B, data have been generated using the same parameters used for generating data in Figure 3 and the bias $b = 0$. In all simulations presented in Figures 10-11, the bias parameter, b, was generated using random draws form a normal distribution with the standard deviation 1. For the analysis presented in Figure 12, random numbers from the skewed distribution was drawn from the Pearson system of distributions with mean 0, variance 1, skewness –0.5 and kurtosis 3 (i.e. the kurtosis of the normal distribution) using the MATLAB function pearsrnd (Statistics and Machine Learning Toolbox).

Implementation of NHI and HPE

The NHI method is based on a Laplace approximation of the joint distribution of data and parameters, separately for each model and subject. Specifically, NHI relies on a maximum-a-posteriori (MAP) estimate of parameters. Formally, if μ_0 and V_0 are prior mean and variance, respectively, and $\ell(h) = p(x_n|M_k, h)N(h|\mu_0, V_0)$, in which x_n is the nth dataset, M_k denotes the kth model and h is the corresponding parameter vector (with the size D_k), then NHI quantifies model evidence of model k for subject n, L_{NHI}^{kn}, as:

$$L_{NHI}^{kn} = \log f_{kn} + \frac{1}{2} D_k \log 2\pi - \frac{1}{2} \log |A_{kn}|,$$

where,

$$\theta_{kn} = \mathrm{argmax_h}\, \ell(\mathrm{h})$$

$$A_{kn} = -\nabla\nabla \log\, \ell(\mathrm{h})\,|\theta_{kn}$$

$$f_{kn} = \ell(\theta_{kn}).$$

We then used values of L_{NHI}^{kn} across subjects to do the random-effects model selection and obtain model frequency given data and exceedance probability [3,4].

Similar to HBI and NHI, we used the Laplace approximation for quantifying model evidence by the HPE. Suppose that μ_k and V_k are group mean and variance, respectively in the last iteration of the algorithm (see [5,6] for a full explanation of the algorithm). We defined $\ell(h) = p(x_n|M_k, h)N(h|\mu_k, V_k)$

and obtained values of θ_{kn}, A_{kn} and f_{kn} using the same equations as above. Similar to [5,6], we used the Bayesian Information Criterion (BIC) to penalize models due to group parameters. Thus, the model evidence for the kth model, L^k_{HPE}, is given by the sum of (Laplace-approximated) log-evidence across all subjects plus the BIC-penalty of $2D_k$ (the mean and the diagonal variance) group parameters:

$$L^k_{HPE} = -D_k \log N + \sum_{n=1}^{N} \log f_{kn} + \frac{1}{2} D_k \log 2\pi - \frac{1}{2} \log |A_{kn}|,$$

Thus, the NHI relies on prior parameters μ_0 and V_0, which, for all models and parameters, were assumed μ_0=0 and $V_0 = 6.25$. This means that the prior mean for parameters constrained in the unit range, such as learning rate, was in the middle of their theoretical range. The value of prior variance is chosen to ensure that parameters can vary in a wide range with no substantial limitation due to the prior. In particular, a prior variance of 6.25 ensures that the difference between the maximum and minimum penalty is less than one chance-level choice in 90% of the theoretical range of a parameter constrained in the unit range. For such a parameter, the minimum penalty occurs at the mean, $\log N(0|0, V_0)$, and the penalty for $x = \text{logit}\,(0.95)$ is $\log N(x|0, V_0)$, where the logit function is the inverse of sigmoid function used for transforming unit-range parameters. Therefore, the relative penalty is given by

$$C(V_0) = \log N(x|0, V_0) - \log N(0|0, V_0)$$

Assuming that $C(V_0) = \log 0.5$, we obtain $V_0 = 6.25$. This value for prior variance also indicates that exponentially-transformed variables (such as decision noise) can vary between 0.05 and 19.00 with no substantial effect of prior (with the same criterion defined above).

These MAP estimates were also used to initialize both HPE and HBI methods. Non-linear derivative-based optimization method (Quasi-Newton algorithm as implemented in the fminunc routine in MATLAB, ©Mathwork, version R2014b) was used.

References

1. Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ. Model-based influences on humans' choices and striatal prediction errors. Neuron. 2011;69: 1204–1215. doi:10.1016/j.neuron.2011.02.027
2. Piray P, Toni I, Cools R. Human Choice Strategy Varies with Anatomical Projections from Ventromedial Prefrontal Cortex to Medial Striatum. J Neurosci. 2016;36: 2857–2867. doi:10.1523/JNEUROSCI.2033-15.2016
3. Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ. Bayesian model selection for group studies. NeuroImage. 2009;46: 1004–1017. doi:10.1016/j.neuroimage.2009.03.025
4. Rigoux L, Stephan KE, Friston KJ, Daunizeau J. Bayesian model selection for group studies - revisited. NeuroImage. 2014;84: 971–985. doi:10.1016/j.neuroimage.2013.08.065

5. Huys QJM, Eshel N, O'Nions E, Sheridan L, Dayan P, Roiser JP. Bonsai trees in your head: how the pavlovian system sculpts goal-directed choices by pruning decision trees. PLoS Comput Biol. 2012;8: e1002410. doi:10.1371/journal.pcbi.1002410
6. Huys QJM, Cools R, Gölzer M, Friedel E, Heinz A, Dolan RJ, et al. Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding. PLoS Comput Biol. 2011;7: e1002028. doi:10.1371/journal.pcbi.1002028

S1 Fig. A control analysis assessing the effects of prior variance on NHI performance.} In scenario 1, similar to the analysis presented in the main text (Fig 3), 10 and 30 subjects generated with the RL and dual-α RL models, respectively. Conversely, in scenario 2, the RL model is more likely (30 subjects) than the dual-α RL model (10 subjects). In A and C, protected exceedance probability (PXP) as a function of prior variance is plotted in scenario 1 and 2, respectively. In B and D, estimation error for the learning rate parameter of RL is plotted in scenarios 1 and 2, respectively. The simulations show in general that no single prior is flexible enough to capture the different scenarios. In particular, while narrowing the prior reduces the complexity penalty (and thus somewhat improves model selection in scenario 1, when the more complex model should be favored), it also worsens parameter estimation in both scenarios. This is because the learning rates for the two models are, generatively, different, and a narrow prior cannot support both at once. Here, the true value of the RL learning rate was 0.1, which was quite away from the prior mean (i.e. 0.5), making it difficult for a narrower variance to capture it. Finally, this poor parameter estimation for the RL model has negative consequences also for model selection in scenario 2 (where the RL model should be favored, but the evidence for it is hampered by poor fit to the learning rate with smaller prior variance). The parameters used in this simulation are the same as those used in the original simulation analyses (Figs 3-4). Median across 100 simulations is plotted. Error-bars indicate the first and third quantiles. The prior variance in all simulation analyses of the main text is 6.25.

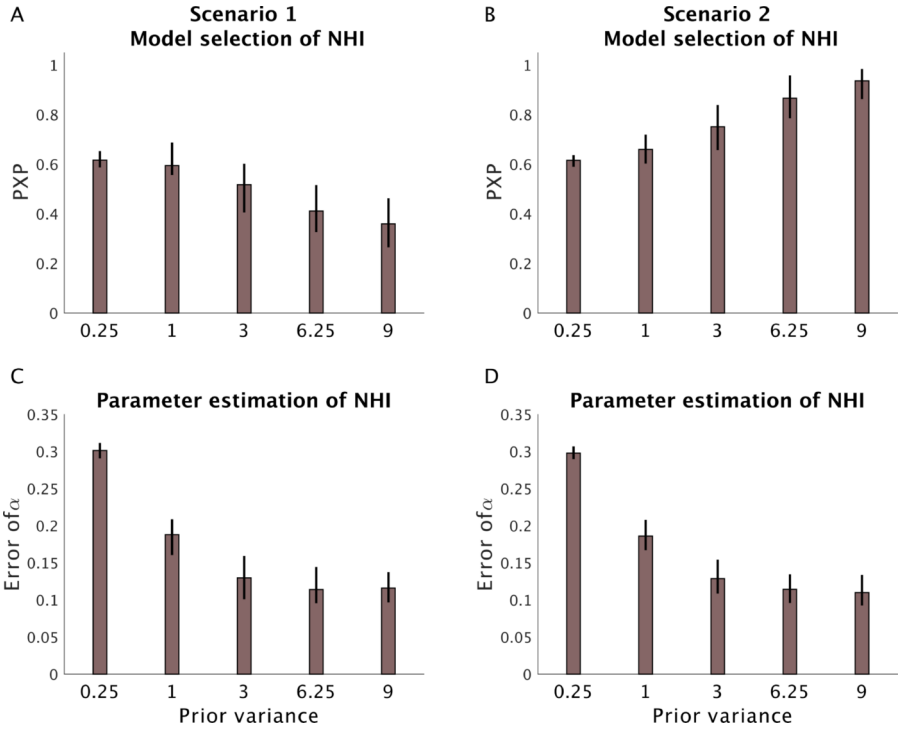

S2 Fig. A control simulation analysis extending that from Fig 3, with different settings of learning rates for simulating data. The same parameters as in Fig 3 were used for simulations here, with the only difference that the learning rate parameter for the RL model was different here. In particular, the true learning rate of the RL was in the middle of those for the dual-α RL (for RL: α=0.6; for dual-α RL: α^{+} =0.8, α^{-}=0.4). The difference between parameter estimation performance of the HPE and HBI is not as pronounced as in Fig 3, which is expected theoretically. A) Model selection by HBI using protected exceedance probabilities (PXP); B) Model frequencies estimated by the HBI. C) Model attribution at the individual level by the HBI. Responsibility estimates are plotted for true attributions (TA) and for false attributions (FA). Inset: percentage of correct assignment of the model by the HBI at the individual level. D) Model selection performance of NHI, HPE, and HBI; E, F) Error in estimating individual parameters of the RL (E) and the dual-α RL model (F). In all plots, error-bars are standard errors of the mean obtained across 20 simulations.

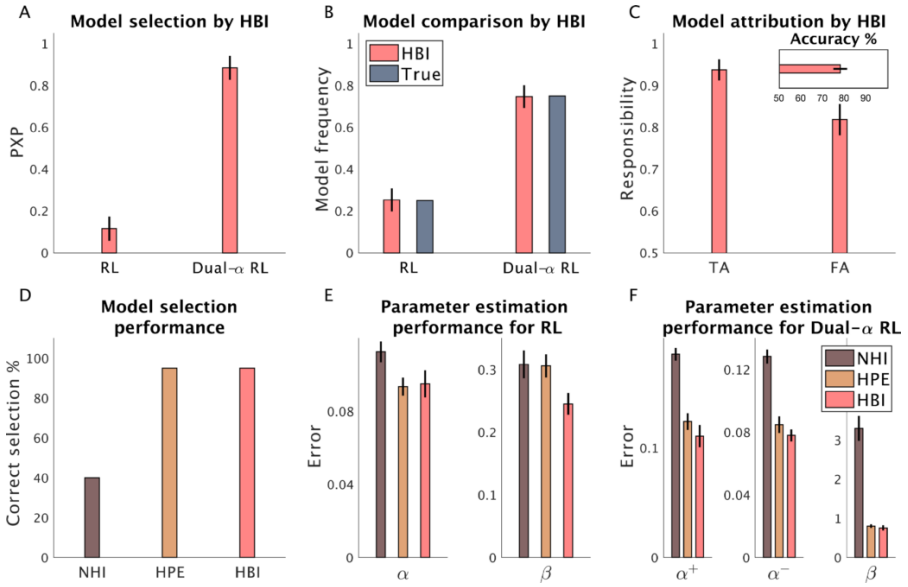

Supplementary Appendix

Hierarchical Bayesian inference for model fitting and comparison

Payam Piray, Amir Dezfouli,
Tom Heskes, Michael J. Frank, Nathaniel D. Daw

1 Formal derivations of the HBI algorithm

In this appendix, we provide the proof of the results given in the main text. The proof is given in three parts by obtaining 1) the functional form of $q(\mathbf{H}, \mathbf{Z})$; 2) the posterior $q(\boldsymbol{\mu}, \boldsymbol{\tau}, \mathbf{m})$ and corresponding update equations; and 3) the update equations for the posterior over latent variables, $q(\mathbf{H}, \mathbf{Z})$.

Recall that the HBI generative model is given by

$$p(\mathbf{Z}|\mathbf{m}) = \prod_n \prod_k m_k^{z_{kn}}. \tag{1}$$

$$p(\mathbf{X}|\mathbf{H}, \mathbf{Z}) = \prod_k \prod_n p(\mathbf{x}_n|\mathbf{h}_{kn}, M_k)^{z_{kn}}, \tag{2}$$

$$p(\mathbf{H}|\mathbf{Z}, \boldsymbol{\mu}, \mathbf{T}) = \prod_k \prod_n \mathcal{N}(\mathbf{h}_{kn}|\boldsymbol{\mu}_k, \mathbf{T}_k^{-1})^{z_{kn}}, \tag{3}$$

$$p(\mathbf{m}) = \mathrm{Dir}(\mathbf{m}|\boldsymbol{\alpha_0}) = C(\alpha_0) \prod_{k=1}^{K} m_k^{\alpha_0 - 1}, \tag{4}$$

where $\mathbf{T}_k$ is a diagonal matrix. By defining $\mathbf{T}_k = \mathrm{diag}(\boldsymbol{\tau}_k)$, in which diag(.) is an operator outputting a diagonal matrix with elements given by $\boldsymbol{\tau}_k$, we have:

$$p(\boldsymbol{\mu}, \boldsymbol{\tau}) = \prod_{k=1}^{K} \mathcal{N}(\boldsymbol{\mu}_k|\mathbf{a}_0, \mathrm{diag}(b\boldsymbol{\tau}_k)^{-1}) \mathcal{G}(\boldsymbol{\tau}_k|v, \mathbf{s}), \tag{5}$$

where we have defined:

$$\mathcal{G}(\boldsymbol{\tau}_k|v,\mathbf{s}) = \prod_{i=1}^{D_k} \mathcal{G}(\tau_{ki}|v,s).$$

Thus, the full probabilistic model is given by,

$$p(\mathbf{X},\mathbf{H},\mathbf{Z},\boldsymbol{\mu},\boldsymbol{\tau},\mathbf{m}) = p(\mathbf{X}|\mathbf{H},\mathbf{Z})p(\mathbf{H}|\mathbf{Z},\boldsymbol{\mu},\boldsymbol{\tau})p(\mathbf{Z}|\mathbf{m})p(\boldsymbol{\mu}|\boldsymbol{\tau})p(\boldsymbol{\tau})p(\mathbf{m}). \quad (6)$$

1.1 The functional form of the posterior over H and Z

Let us first consider the derivation of the functional form for the factor $q(\mathbf{H},\mathbf{Z})$. According to standard results in variational inference [15,16], the log of this factor is given by:

$$\log q(\mathbf{H},\mathbf{Z}) = \mathbb{E}_{\boldsymbol{\mu},\boldsymbol{\tau},\mathbf{m}}[\log p(\mathbf{X},\mathbf{H},\mathbf{Z},\boldsymbol{\mu},\boldsymbol{\tau},\mathbf{m})] + \text{constant}^{\backslash\mathbf{H},\backslash\mathbf{Z}},$$

where the constant term denotes all the terms independent of the corresponding variables. Note that the expectation is taken with respect to the current estimates of $q(\boldsymbol{\mu},\boldsymbol{\tau},\mathbf{m})$. By using equation (6) and absorbing all the terms which are independent of $\mathbf{H}$ and $\mathbf{Z}$ into the additive constant, we have:

$$\log q(\mathbf{H},\mathbf{Z}) = \mathbb{E}_{\boldsymbol{\mu},\boldsymbol{\tau}}[\log p(\mathbf{X},\mathbf{H}|\mathbf{Z},\boldsymbol{\mu},\boldsymbol{\tau})] + \mathbb{E}_{\mathbf{m}}[\log p(\mathbf{Z}|\mathbf{m})] + \text{constant}^{\backslash\mathbf{H},\backslash\mathbf{Z}}.$$

Substituting the two conditional distribution on the right-hand side using equations (1-3), we have:

$$\log q(\mathbf{H},\mathbf{Z}) = \sum_k \sum_n z_{kn}(\log I_{kn} + \mathbb{E}_{\mathbf{m}}[\log m_k]) + \text{constant}^{\backslash\mathbf{H},\backslash\mathbf{Z}},$$

where,

$$\log I_{kn} = \mathbb{E}_{\boldsymbol{\mu},\boldsymbol{\tau}}[\log p(\mathbf{x}_n|\mathbf{h}_{kn},M_k)\mathcal{N}(\mathbf{h}_{kn}|\boldsymbol{\mu}_k,\mathbf{T}_k^{-1})].$$

Note that we have defined $\mathbf{T}_k = \text{diag}(\boldsymbol{\tau}_k)$. We assume that there is a quadratic approximation of I_{kn} with respect to $\mathbf{h}_{kn}$,

$$I_{kn} \propto \exp(-\frac{1}{2}(\mathbf{h}_{kn}-\boldsymbol{\theta}_{kn})^{\mathsf{T}}\mathbf{A}_{kn}(\mathbf{h}_{kn}-\boldsymbol{\theta}_{kn})),$$

which gives,

$$\log q(\mathbf{H}, \mathbf{Z}) = \sum_k \sum_n z_{kn}(-\frac{1}{2}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn})^\mathsf{T} \mathbf{A}_{kn}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn}) + \mathbb{E}_{\mathbf{m}}[\log m_k]) + \text{constant}^{\backslash \mathbf{H}, \backslash \mathbf{Z}}.$$

Since $\log q(\mathbf{H}|\mathbf{Z}) = \log q(\mathbf{H}, \mathbf{Z}) - \log q(\mathbf{Z})$, we can read off terms involving $\mathbf{H}$ in $\log q(\mathbf{H}, \mathbf{Z})$ to obtain $\log q(\mathbf{H}|\mathbf{Z})$:

$$\log q(\mathbf{H}|\mathbf{Z}) = \sum_k \sum_n z_{kn}(-\frac{1}{2}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn})^\mathsf{T} \mathbf{A}_{kn}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn})) + \text{constant}^{\backslash \mathbf{H}}.$$

Requiring that this distribution should be normalized, we obtain:

$$q(\mathbf{H}|\mathbf{Z}) = \prod_k \prod_n \mathcal{N}(\mathbf{h}_{kn}|\boldsymbol{\theta}_{kn}, \mathbf{A}_{kn}^{-1})^{z_{kn}}. \tag{7}$$

Subtracting $\log q(\mathbf{H}|\mathbf{Z})$ from $\log q(\mathbf{H}, \mathbf{Z})$ cancels out the quadratic component and yields $\log q(\mathbf{Z})$, which is a linear function with respect to z_{kn}. Therefore, we have:

$$q(\mathbf{Z}) = \prod_k \prod_n r_{kn}^{z_{kn}}. \tag{8}$$

The functional form of $q(\mathbf{H}, \mathbf{Z})$ is then given by,

$$q(\mathbf{H}, \mathbf{Z}) = \prod_k \prod_n r_{kn}^{z_{kn}} \mathcal{N}(\mathbf{h}_{kn}|\boldsymbol{\theta}_{kn}, \mathbf{A}_{kn}^{-1})^{z_{kn}}. \tag{9}$$

Here our goal was to obtain the functional form of the posterior over latent variables. We will obtain values of r_{kn}, $\boldsymbol{\theta}_{kn}$ and $\mathbf{A}_{kn}$ in section 1.3.

1.2 The posterior over $\boldsymbol{\mu}$, $\boldsymbol{\tau}$ and $\mathbf{m}$

We continue with obtaining the functional form and update equations for the other variational factor $q(\boldsymbol{\mu}, \boldsymbol{\tau}, \mathbf{m})$. The posterior of $\mathbf{m}$ is independent from the posterior over $\boldsymbol{\mu}$ and $\boldsymbol{\tau}$ because the log-posterior decomposes into the terms that

only depend on $\mathbf{m}$ and terms that only depend on $\boldsymbol{\mu}$ and $\boldsymbol{\tau}$:

$$\begin{aligned}\log q(\boldsymbol{\mu}, \boldsymbol{\tau}, \mathbf{m}) &= \mathbb{E}_{\mathbf{H},\mathbf{Z}}[\log p(\mathbf{X}, \mathbf{H}, \mathbf{Z}, \boldsymbol{\mu}, \boldsymbol{\tau}, \mathbf{m})] + \text{constant}^{\backslash \boldsymbol{\mu}, \boldsymbol{\tau}, \mathbf{m}} \\ &= \mathbb{E}_{\mathbf{H},\mathbf{Z}}[\log p(\mathbf{H}|\mathbf{Z}, \boldsymbol{\mu}, \boldsymbol{\tau}) + \log p(\boldsymbol{\mu}, \boldsymbol{\tau})] + \mathbb{E}_{\mathbf{Z}}[\log p(\mathbf{Z}|\mathbf{m}) + p(\mathbf{m})] \\ &+ \text{constant}^{\backslash \boldsymbol{\mu}, \boldsymbol{\tau}, \mathbf{m}},\end{aligned}$$

where we have used equation (6). This implies that the variational posterior $q(\boldsymbol{\mu}, \mathbf{T}, \mathbf{m})$ factorizes to give $q(\boldsymbol{\mu}, \mathbf{T})q(\mathbf{m})$. Thus, the posterior over $\boldsymbol{\mu}$ and $\boldsymbol{\tau}$ is given by:

$$\log q(\boldsymbol{\mu}, \boldsymbol{\tau}) = \mathbb{E}_{\mathbf{H},\mathbf{Z}}[\log p(\mathbf{H}|\mathbf{Z}, \boldsymbol{\mu}, \boldsymbol{\tau})] + \log p(\boldsymbol{\mu}, \boldsymbol{\tau}) + \text{constant}^{\backslash \boldsymbol{\mu}, \boldsymbol{\tau}},$$

in which we absorbed any terms independent of $\boldsymbol{\mu}$, $\boldsymbol{\tau}$ into the additive constant. Substituting for the distributions on the right-hand side, we have:

$$\begin{aligned}\log q(\boldsymbol{\mu}) = &\sum_k \sum_n \mathbb{E}_{\mathbf{H},\mathbf{Z},\boldsymbol{\tau}}[z_{kn} \log \mathcal{N}(\mathbf{h}_{kn}|\boldsymbol{\mu}_k, \mathbf{T}_k^{-1})] + \\ &\sum_k \log \mathcal{N}(\boldsymbol{\mu}_k|\mathbf{a}_0, (b\mathbf{T}_k)^{-1}) + \sum_k \log \mathcal{G}(\boldsymbol{\tau}_k|v, \mathbf{s}) + \text{constant}^{\backslash \boldsymbol{\mu}, \boldsymbol{\tau}}.\end{aligned}$$

Note that we have defined $\mathbf{T}_k = \text{diag}(\boldsymbol{\tau}_k)$. Using equation (9) and by absorbing terms independent of $\boldsymbol{\mu}_k$ and $\boldsymbol{\tau}_k$ into the additive constant, we have:

$$\begin{aligned}\log q(\boldsymbol{\mu}_k, \boldsymbol{\tau}_k) = &\sum_n \frac{1}{2} r_{kn} \log |\mathbf{T}_k| - \sum_n \frac{1}{2} (\boldsymbol{\mu}_k - \boldsymbol{\theta}_{kn})^{\mathsf{T}} r_{kn} \mathbf{T}_k (\boldsymbol{\mu}_k - \boldsymbol{\theta}_{kn}) + \\ &- \sum_n \frac{1}{2} r_{kn} \text{Tr}(\mathbf{A}_{\text{kn}}^{-1} \mathbf{T}_{\text{k}}) + \frac{1}{2} \log |\mathbf{T}_k| + \\ &- \frac{1}{2} (\boldsymbol{\mu}_k - \mathbf{a}_0)^{\mathsf{T}} b \mathbf{T}_k (\boldsymbol{\mu}_k - \mathbf{a}_0) + \\ &+ \sum_{i=1}^{D_k} \log G(v, s) + (v-1) \log \tau_{ki} - s\tau_{ki} + \text{constant}^{\backslash \boldsymbol{\mu}_k, \boldsymbol{\tau}_k}.\end{aligned}$$

As the right-hand side is quadratic with respect to $\boldsymbol{\mu}_k$, the posterior over $\boldsymbol{\mu}_k$ also takes the form of a Gaussian with a variance depending on $\boldsymbol{\tau}_k$:

$$q(\boldsymbol{\mu}_k|\boldsymbol{\tau}_k) = \mathcal{N}(\boldsymbol{\mu}_k|\mathbf{a}_k, (\beta_k \mathbf{T}_k)^{-1}),$$

where

$$\mathbf{a}_k = \frac{1}{\overline{N}_k + b} (\sum_n r_{kn} \boldsymbol{\theta}_{kn} + b\mathbf{a}_0)$$

$$\beta_k = b + \overline{N}_k,$$

and N_k is given by:

$$\overline{N}_k = \sum_n r_{kn}.$$

By subtracting $\log q(\boldsymbol{\mu}_k|\boldsymbol{\tau}_k)$ from $\log q(\boldsymbol{\mu}_k, \boldsymbol{\tau}_k)$, we obtain the posterior over $\boldsymbol{\tau}_k$:

$$q(\boldsymbol{\tau}_k) = \mathcal{G}(\boldsymbol{\tau}_k|\nu_k, \boldsymbol{\sigma}_k),$$

where

$$\begin{aligned}\boldsymbol{\sigma}_k =& \frac{1}{2}\sum_n \mathrm{diag}(r_{kn}[(\boldsymbol{\theta}_{kn} - \overline{\boldsymbol{\theta}}_k)(\boldsymbol{\theta}_{kn} - \overline{\boldsymbol{\theta}}_k)^{\mathsf{T}} + \mathbf{A}_{kn}^{-1}]) \\ &+ \frac{1}{2}\frac{b\overline{N}_k}{b + \overline{N}_k}\mathrm{diag}((\overline{\boldsymbol{\theta}}_k - \mathbf{a}_0)(\overline{\boldsymbol{\theta}}_k - \mathbf{a}_0)^{\mathsf{T}}) + \mathbf{s}\end{aligned}$$

$$\nu_k = v + \frac{1}{2}\overline{N}_k,$$

and $\overline{\boldsymbol{\theta}}_k$ is given by:

$$\overline{\boldsymbol{\theta}}_k = \frac{1}{\overline{N}_k}\sum_n r_{kn}\boldsymbol{\theta}_{kn}.$$

Finally, we consider the factor $q(\mathbf{m})$:

$$\log q(\mathbf{m}) = \mathbb{E}_{\mathbf{Z}}[\log p(\mathbf{Z}|\mathbf{m})] + \log p(\mathbf{m}) + \text{constant}^{\backslash \mathbf{m}}.$$

Substituting for the two distributions on the right-hand side, we have

$$\log q(\mathbf{m}) = \sum_k \sum_n r_{kn} \log m_k + \log C(\alpha_0) + \sum_k (\alpha_0 - 1)\log m_k + \text{constant}^{\backslash \mathbf{m}}.$$

Therefore $q(\mathbf{m})$ takes the form of Dirichlet distribution:

$$q(\mathbf{m}) = \mathrm{Dir}(\mathbf{m}|\boldsymbol{\alpha}),$$

where $\boldsymbol{\alpha}$ has components α_k given by,

$$\alpha_k = \alpha_0 + \overline{N}_k.$$

1.3 The posterior over H and Z

We have already seen in section 1.1 that $q(\mathbf{H}, \mathbf{Z})$ could be written,

$$\log q(\mathbf{H}, \mathbf{Z}) = \sum_k \sum_n z_{kn} (\log I_{kn} + \mathbb{E}_{\mathbf{m}}[\log m_k]) + \text{constant}^{\backslash \mathbf{H}, \backslash \mathbf{Z}},$$

where,

$$\log I_{kn} = \mathbb{E}_{\boldsymbol{\mu}, \boldsymbol{\tau}}[\log p(\mathbf{x}_n | \mathbf{h}_{kn}, M_k) \mathcal{N}(\mathbf{h}_{kn} | \boldsymbol{\mu}_k, \mathbf{T}_k^{-1})].$$

Since we have already obtained $q(\boldsymbol{\mu}, \boldsymbol{\tau})$, we can now compute I_{kn}:

$$\begin{aligned} \log I_{kn} = &+ \log p(\mathbf{x}_n | \mathbf{h}_{kn}, M_k) - \frac{1}{2} D_k \log 2\pi + \frac{1}{2} \mathbb{E}[\log |\mathbf{T}_k|] \\ &- \frac{1}{2} \mathbb{E}_{\mathbf{T}, \boldsymbol{\mu}}[(\mathbf{h}_{kn} - \boldsymbol{\mu}_k)^\mathsf{T} \mathbf{T}_k (\mathbf{h}_{kn} - \boldsymbol{\mu}_k)]. \end{aligned}$$

Therefore, $\log I_{kn}$ is given by:

$$\begin{aligned} \log I_{kn} = &\log p(\mathbf{x}_n | \mathbf{h}_{kn}, M_k) - \frac{1}{2} D_k \log 2\pi + \frac{1}{2} \mathbb{E}[\log |\mathbf{T}_k|] \\ &- \frac{1}{2} (\mathbf{h}_{kn} - \mathbf{a}_k)^\mathsf{T} \mathbb{E}[\mathbf{T}_k] (\mathbf{h}_{kn} - \mathbf{a}_k) - \frac{1}{2} \mathbb{E}[\mathrm{Tr}(\mathbf{T}_\mathrm{k} (\beta_\mathrm{k} \mathbf{T}_\mathrm{k})^{-1})], \end{aligned}$$

which can be written in the form,

$$\log I_{kn} = \log p(\mathbf{x}_n | \mathbf{h}_{kn}, M_k) \mathcal{N}(\mathbf{h}_{kn} | \mathbf{a}_k, \mathbb{E}[\mathbf{T}_k]^{-1}) + \lambda_k,$$

where λ_k is independent of $\mathbf{h}_{kn}$ and is given by

$$\lambda_k = \frac{1}{2} \mathbb{E}[\log |\mathbf{T}_k|] - \frac{1}{2} \log |\mathbb{E}[\mathbf{T}_k]| - \frac{1}{2} \frac{D_k}{\beta_k},$$

Substituting the moments of $\mathbf{T}_k = \mathrm{diag}(\boldsymbol{\tau}_k)$ with their values under $q(\boldsymbol{\tau}_k)$,

$$\log \mathbb{E}[\boldsymbol{\tau}_k] = D_k \log \nu_k - \sum_i \log \sigma_{ki},$$

$$\mathbb{E}[\log \boldsymbol{\tau}_k] = D_k \psi(\nu_k) - \sum_i \log \sigma_{ki},$$

gives

$$\lambda_k = \frac{D_k}{2} \big(\psi(\nu_k) - \log \nu_k - \frac{1}{\beta_k} \big)$$

Now, we make a quadratic approximation of $p(\mathbf{x}_n|\mathbf{h}_{kn}, M_k)\mathcal{N}(\mathbf{h}_{kn}|\mathbf{a}_k, \mathbb{E}[\mathbf{T}_k]^{-1})$ with respect to $\mathbf{h}_{kn}$ (for example using Laplace approximation or any other method):

$$p(\mathbf{x}_n|\mathbf{h}_{kn}, M_k)\mathcal{N}(\mathbf{h}_{kn}|\mathbf{a}_k, \mathbb{E}[\mathbf{T}_k]^{-1}) \simeq \\ f_{kn} \exp(-\frac{1}{2}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn})^\mathsf{T} \mathbf{A}_{kn}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn})).$$

Substituting this approximation into $\log I_{kn}$, we obtain

$$\log I_{kn} = \log f_{kn} - \frac{1}{2}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn})^\mathsf{T} \mathbf{A}_{kn}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn}) + \lambda_k.$$

Therefore, we have:

$$\log q(\mathbf{H}, \mathbf{Z}) = \sum_k \sum_n z_{kn} \Big(+ \log f_{kn} - \frac{1}{2}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn})^\mathsf{T} \mathbf{A}_{kn}(\mathbf{h}_{kn} - \boldsymbol{\theta}_{kn}) \\ + \lambda_k + \mathbb{E}[\log m_k] \Big) + \text{constant}^{\backslash \mathbf{H}, \backslash \mathbf{Z}}.$$

Subtracting this equation from $\log q(\mathbf{H}|Z)$ given by the log of equation (7), we have:

$$\log q(\mathbf{Z}) = \sum_k \sum_n z_{kn} \log \rho_{kn} + \text{constant}^{\backslash \mathbf{H}, \backslash \mathbf{Z}},$$

where

$$\log \rho_{kn} = \log f_{kn} + \frac{1}{2} D_k \log 2\pi - \frac{1}{2} \log |\mathbf{A}_{kn}| + \lambda_k + \mathbb{E}[\log m_k].$$

Requiring that $q(\mathbf{Z})$ be normalized, we obtain equation (8), where

$$r_{kn} = \frac{\rho_{kn}}{\sum_{j=1}^K \rho_{jn}},$$

which completes the proof.