Bayesian inference and Bayesian model selection

Klaas Enno Stephan

Translational Neuromodeling Unit

Universität Zürich[UZH]

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Lecture as part of "Methods & Models for fMRI data analysis",
University of Zurich & ETH Zurich, 26 November 2019

With slides from and many thanks to:

Kay Brodersen,

Will Penny,

Sudhir Shankar Raman

Bayes' rule

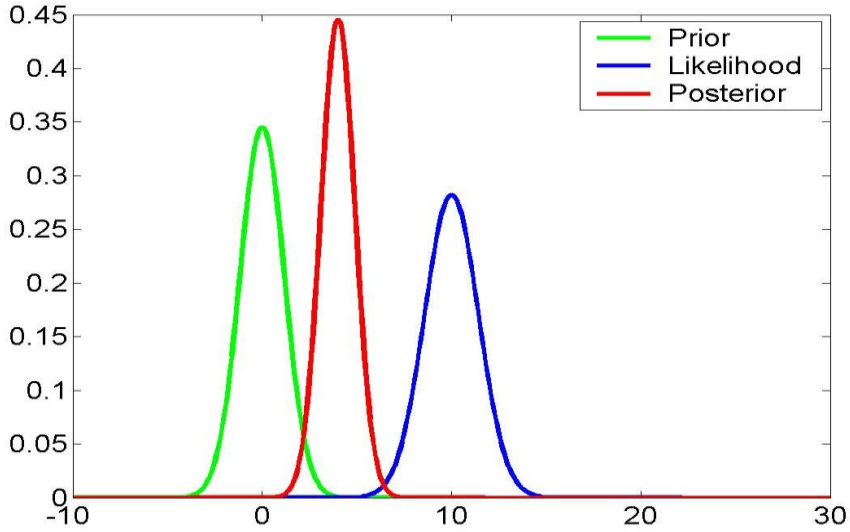

The Reverend Thomas Bayes
(1702-1761)

Likelihood × **prior**: generative model

$$p(\theta \mid y) = \frac{p(y \mid \theta)\, p(\theta)}{p(y)}$$

θ: parameters
y: data

Model evidence: normalisation term and index for model goodness

"... the theorem expresses how a ... degree of belief should rationally change to account for availability of related evidence."

Wikipedia

Bayes' rule

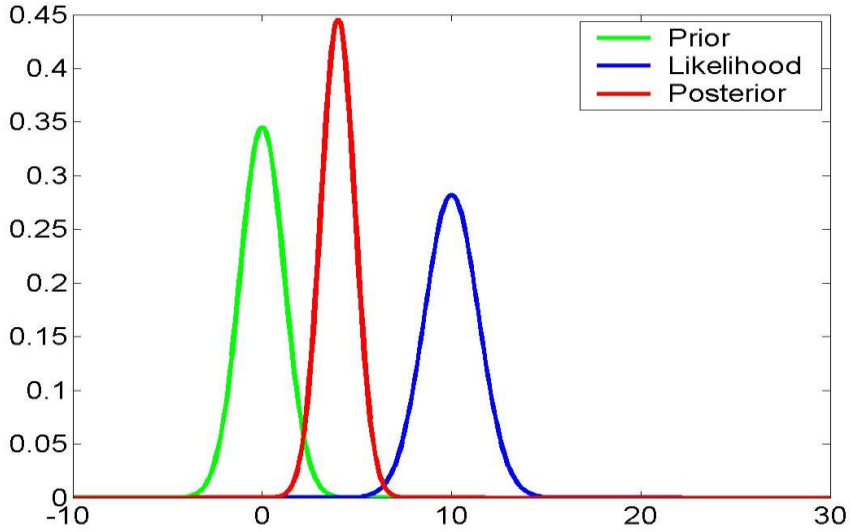

The Reverend Thomas Bayes
(1702-1761)

$$p(\boldsymbol{\theta} \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \boldsymbol{\theta})\, p(\boldsymbol{\theta})}{p(\mathbf{y})}$$

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}$$

Bayesian inference: an animation

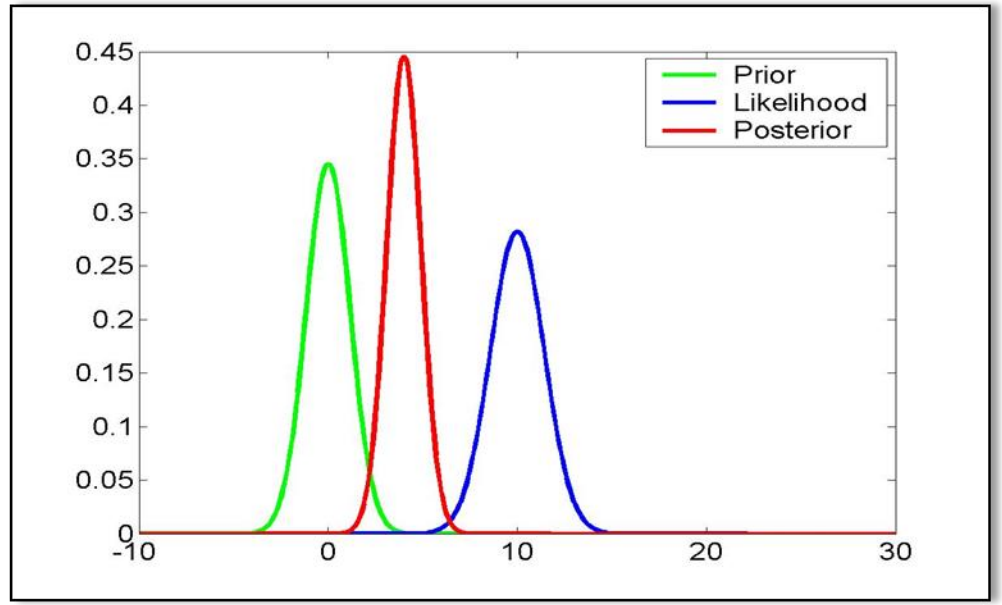

Code courtesy by Guillaume Flandin

But why should I learn about Bayesian inference?

Because Bayesian principles are fundamental for

- **statistical inference** in general
- **system identification**
- **translational neuromodeling** ("computational assays")
  - computational psychiatry
  - computational neurology
  - computational psychosomatics
- contemporary **theories of brain function** (the "Bayesian brain")
  - predictive coding
  - free energy principle
  - active inference

Translational Neuromodeling

❶ Computational assays: Models of disease mechanisms

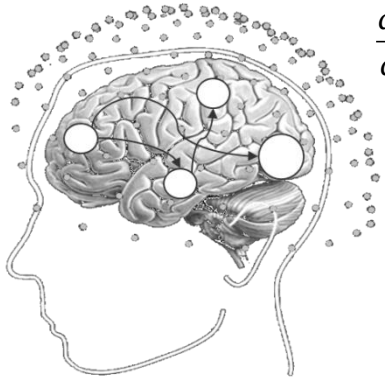

$$\frac{dx}{dt} = f(x,u,\theta) + \omega$$

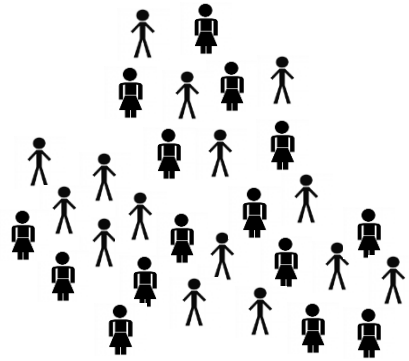

❷ Application to brain activity and behaviour of individual patients

❸ Detecting physiological subgroups (based on inferred mechanisms)

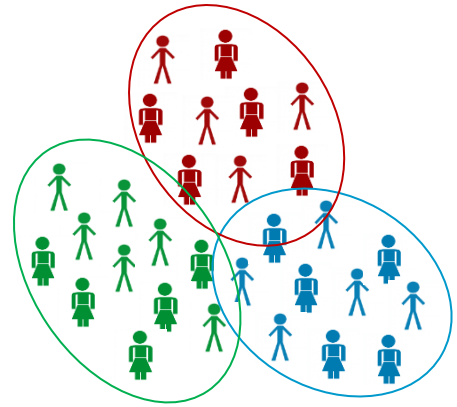

- disease mechanism A
- disease mechanism B
- disease mechanism C

❹ Individual treatment prediction

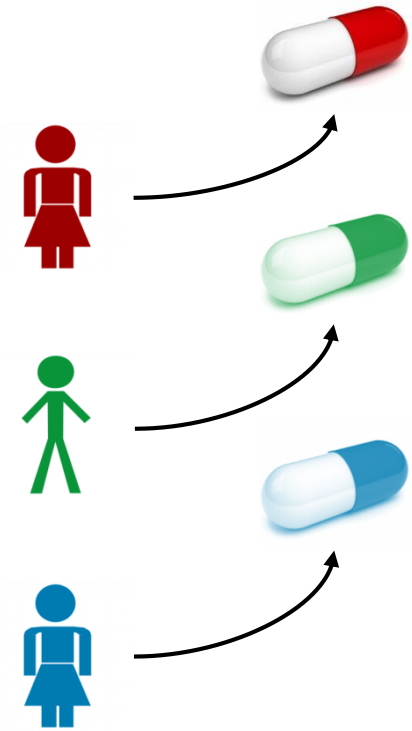

Stephan et al. 2015, *Neuron*

Generative models as "computational assays"

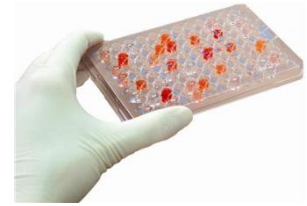

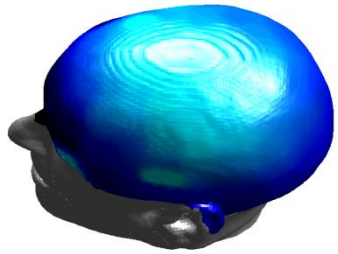

$$p(y \mid \theta, m) \cdot p(\theta \mid m)$$

$$p(\theta \mid y, m)$$

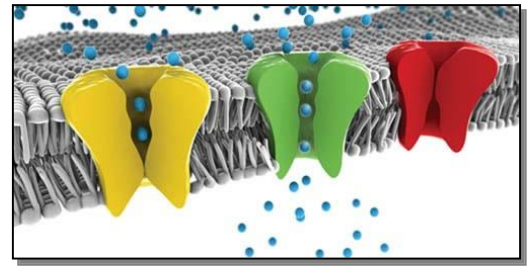

y = data, θ = parameters, m = model

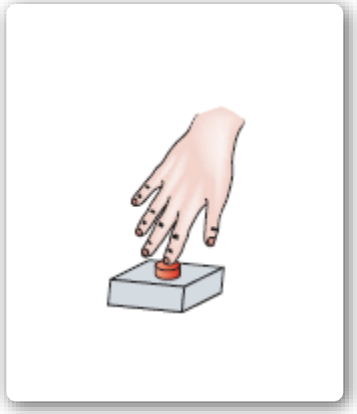

$$p(y \mid \theta, m) \cdot p(\theta \mid m)$$

$$p(\theta \mid y, m)$$

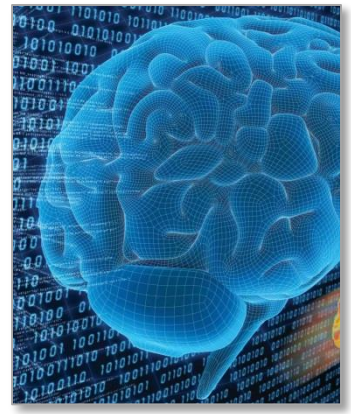

Differential diagnosis by model selection

SYMPTOM
(behaviour
or physiology)

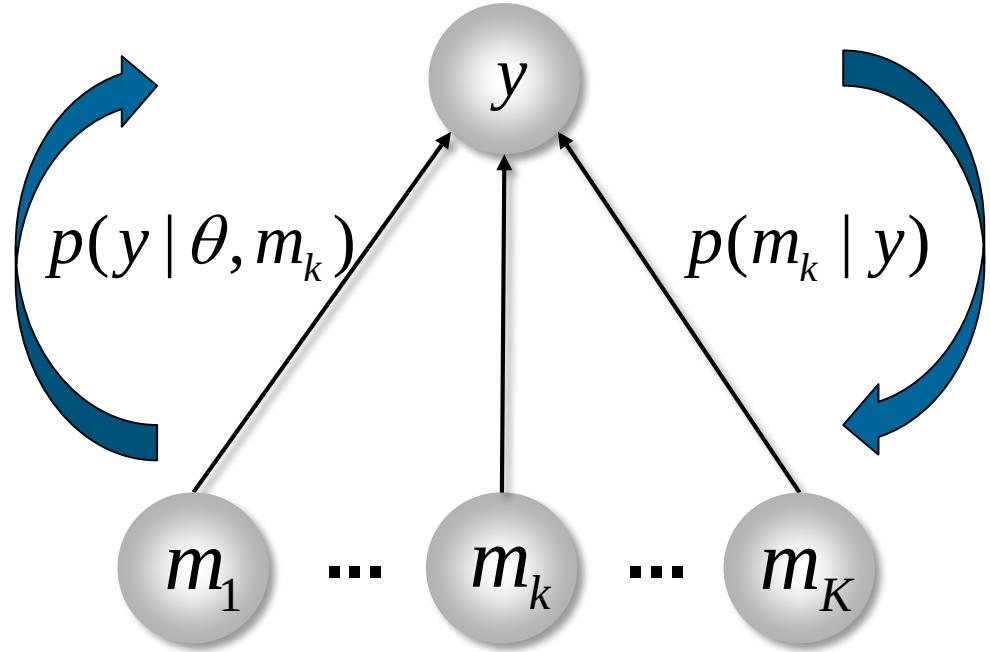

**HYPOTHETICAL
MECHANISM**

$$p(m_k \mid \mathrm{y}) = \frac{p(y \mid m_k)\, p(m_k)}{\sum_k p(y \mid m_k)\, p(m_k)}$$

Stephan et al. 2017, *NeuroImage*

Perception = inversion of a hierarchical generative model

neuronal states

environm. states
others' mental states
bodily states

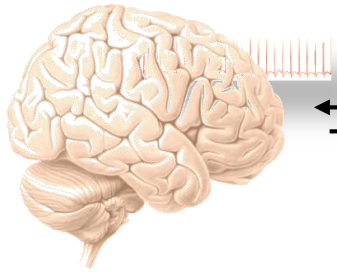

forward model

$$p(y \mid x, m) \cdot p(x \mid m)$$

$$p(x \mid y, m)$$

perception

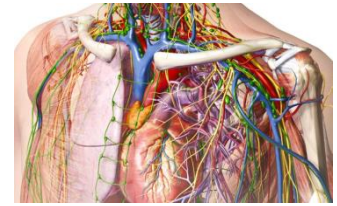

Back to the technicalities...

Bayes' rule

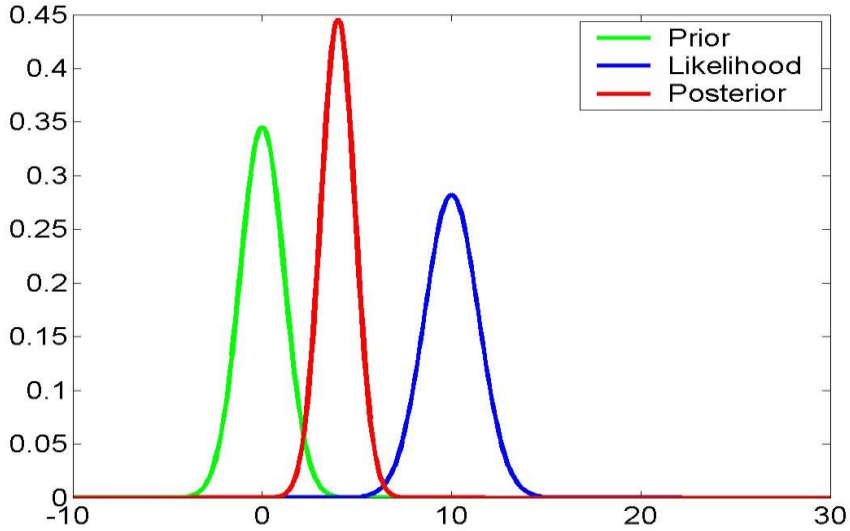

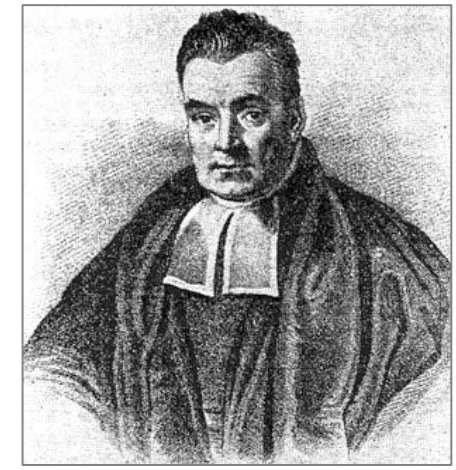

The Reverend Thomas Bayes
(1702-1761)

$$p(\boldsymbol{\theta} \mid \mathbf{y}, m) = \frac{p(\mathbf{y} \mid \boldsymbol{\theta}, m)\, p(\boldsymbol{\theta} \mid m)}{p(\mathbf{y} \mid m)}$$

No change – just making the choice of a particular model m explicit.

Bayes' rule

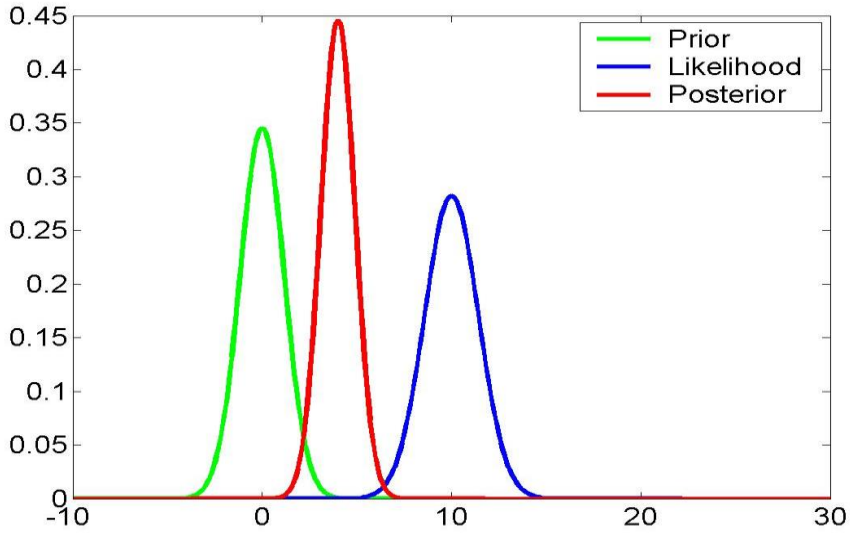

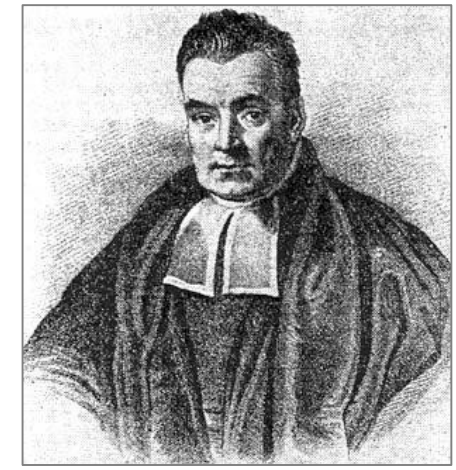

The Reverend Thomas Bayes
(1702-1761)

$$p(\boldsymbol{\theta} \mid \mathbf{y}, m) = \frac{p(\mathbf{y} \mid \boldsymbol{\theta}, m)\, p(\boldsymbol{\theta} \mid m)}{\int p(\mathbf{y} \mid \boldsymbol{\theta}, m)\, p(\boldsymbol{\theta} \mid m)}$$

posterior = likelihood • prior / evidence

Evidence:
probability that data were generated by model m, averaging over all possible parameter values (as weighted by the prior).

The evidence term

continuous θ

$$p(\theta \mid y) = \frac{p(y \mid \theta)\, p(\theta)}{\int p(y \mid \theta)\, p(\theta)}$$

discrete θ

$$p(\theta \mid y) = \frac{p(y \mid \theta)\, p(\theta)}{\sum_{\theta \in \Theta} p(y \mid \theta)\, p(\theta)}$$

Bayesian inference: A clinical example

- *"The probability of breast cancer is 1% for women aged forty who participate in routine screening. If a woman has breast cancer, the probability is 80% that she will get a positive mammogram. If a woman does not have breast cancer, the probability is 9.6% that she will also get a positive mammogram. A woman in this age group has a positive mammogram in a routine screening. What is the probability that she actually has breast cancer?"* (Gigerenzer & Hoffrage 1995)

- From this information, we can deduce:
  - p(C+) = 0.01 $\rightarrow$ p(C-) = 0.99
  - p(M+|C+) = 0.8
  - p(M+|C-) = 0.096

- We can now apply Bayes' rule to compute the posterior probability:

$$p(C+ \mid M+) = \frac{p(M+ \mid C+)\, p(C+)}{p(M+ \mid C+)\, p(C+) + p(M+ \mid C-)\, p(C-)}$$

Bayesian inference: A clinical example

- *"The probability of breast cancer is 1% for women aged forty who participate in routine screening. If a woman has breast cancer, the probability is 80% that she will get a positive mammogram. If a woman does not have breast cancer, the probability is 9.6% that she will also get a positive mammogram. A woman in this age group has a positive mammogram in a routine screening. What is the probability that she actually has breast cancer?"* (Gigerenzer & Hoffrage 1995)

- From this information, we can deduce:
  - p(C+) = 0.01 $\rightarrow$ p(C-) = 0.99
  - p(M+|C+) = 0.8
  - p(M+|C-) = 0.096

- We can now apply Bayes' rule to compute the posterior probability:

$$p(C+ \mid M+) = \frac{0.8 \cdot 0.01}{0.8 \cdot 0.01 + 0.096 \cdot 0.99} = 0.078$$

Generative models

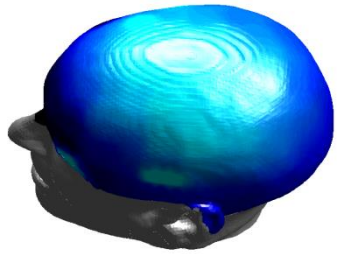

$$p(y \mid \theta, m) \cdot p(\theta \mid m)$$

$$p(\theta \mid y, m)$$

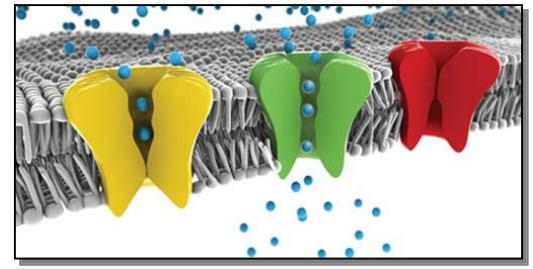

1. specify the joint probability over data (observations) and parameters
2. enforce mechanistic thinking: how could the data have been caused?
3. generate synthetic data (observations) by sampling from the prior – can model explain certain phenomena at all?
4. inference about parameters → p(θ|y)
5. model evidence p(y|m): index of model quality

Bayesian inference in practice

⇨ Formulation of a **generative model**

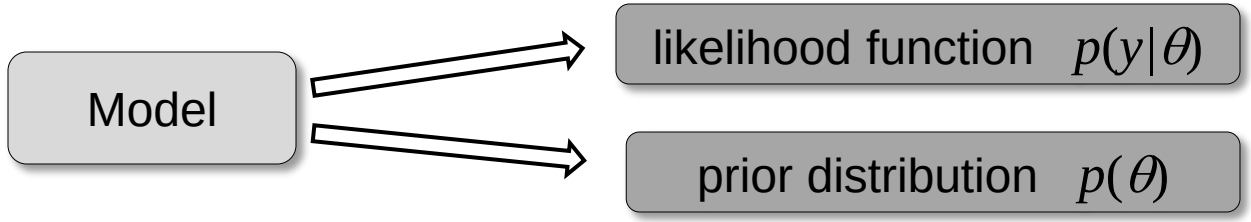

⇨ Observation of **data**

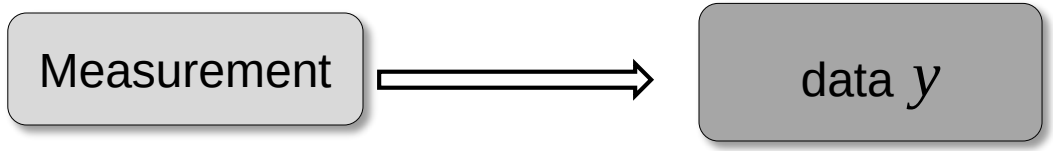

⇨ **Model inversion** – updating one's beliefs

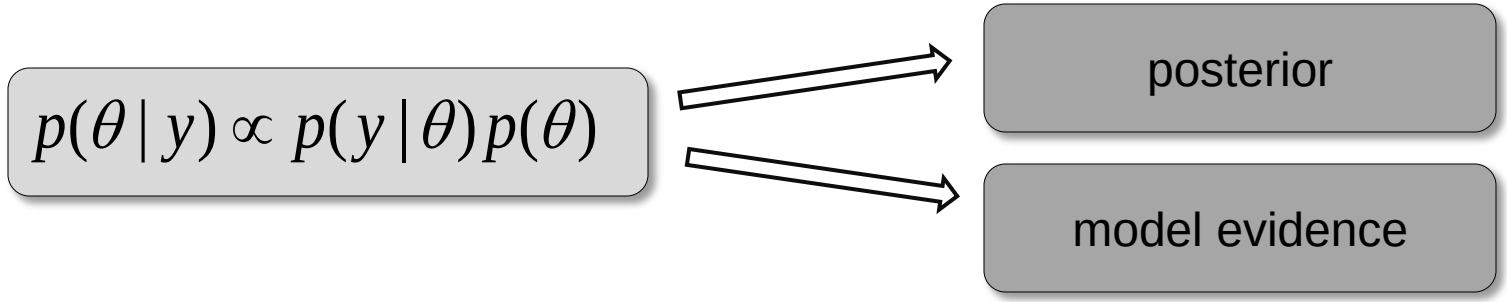

Priors

- Objective priors:
  - "non-informative" priors
  - objective constraints (e.g., non-negativity)
- Subjective priors:
  - subjective but not arbitrary
  - can express beliefs that result from understanding of the problem or system
  - can be result of previous empirical results
- Shrinkage priors:
  - emphasize regularization and sparsity
- Empirical priors:
  - learn parameters of prior distributions from the data ("empirical Bayes")
  - rest on a hierarchical model

Example of a shrinkage prior

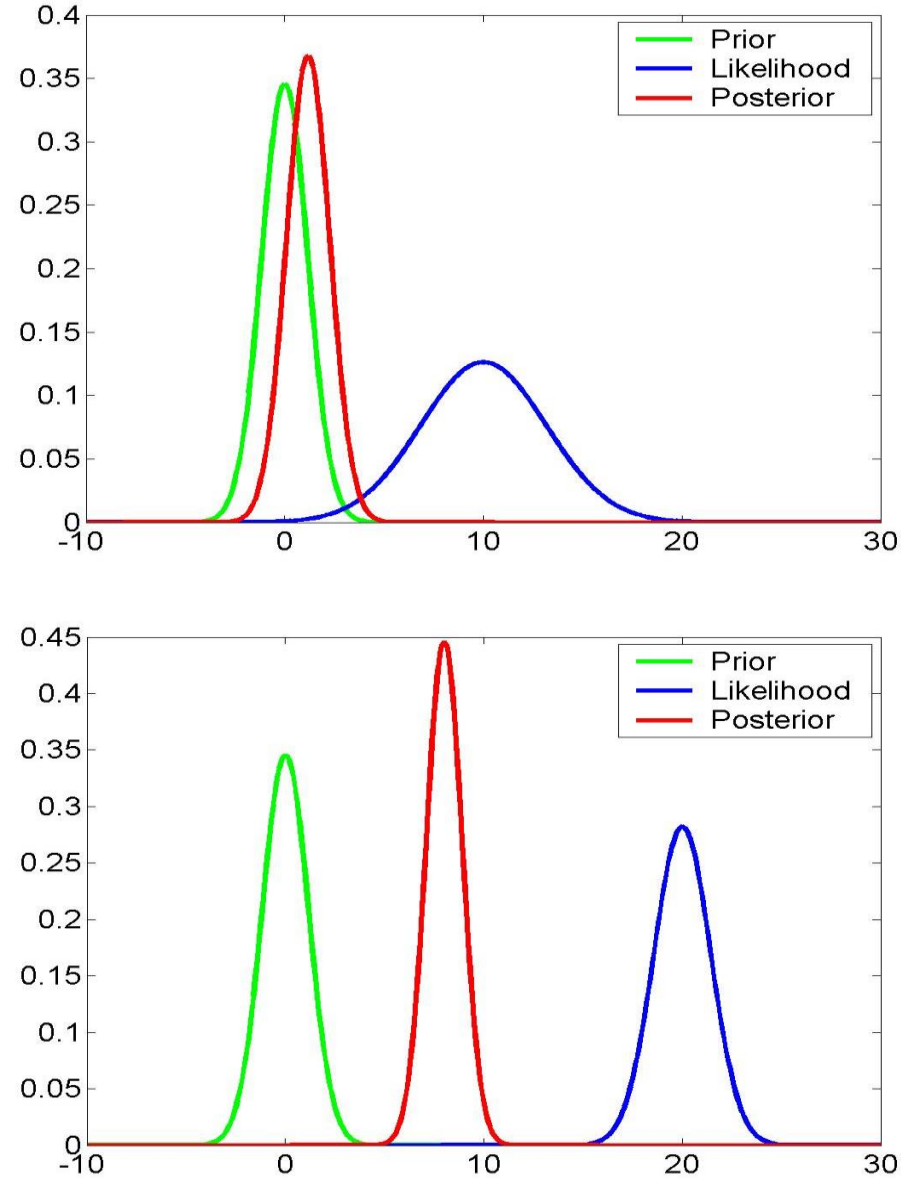

A generative modelling framework for fMRI & EEG: Dynamic causal modeling (DCM)

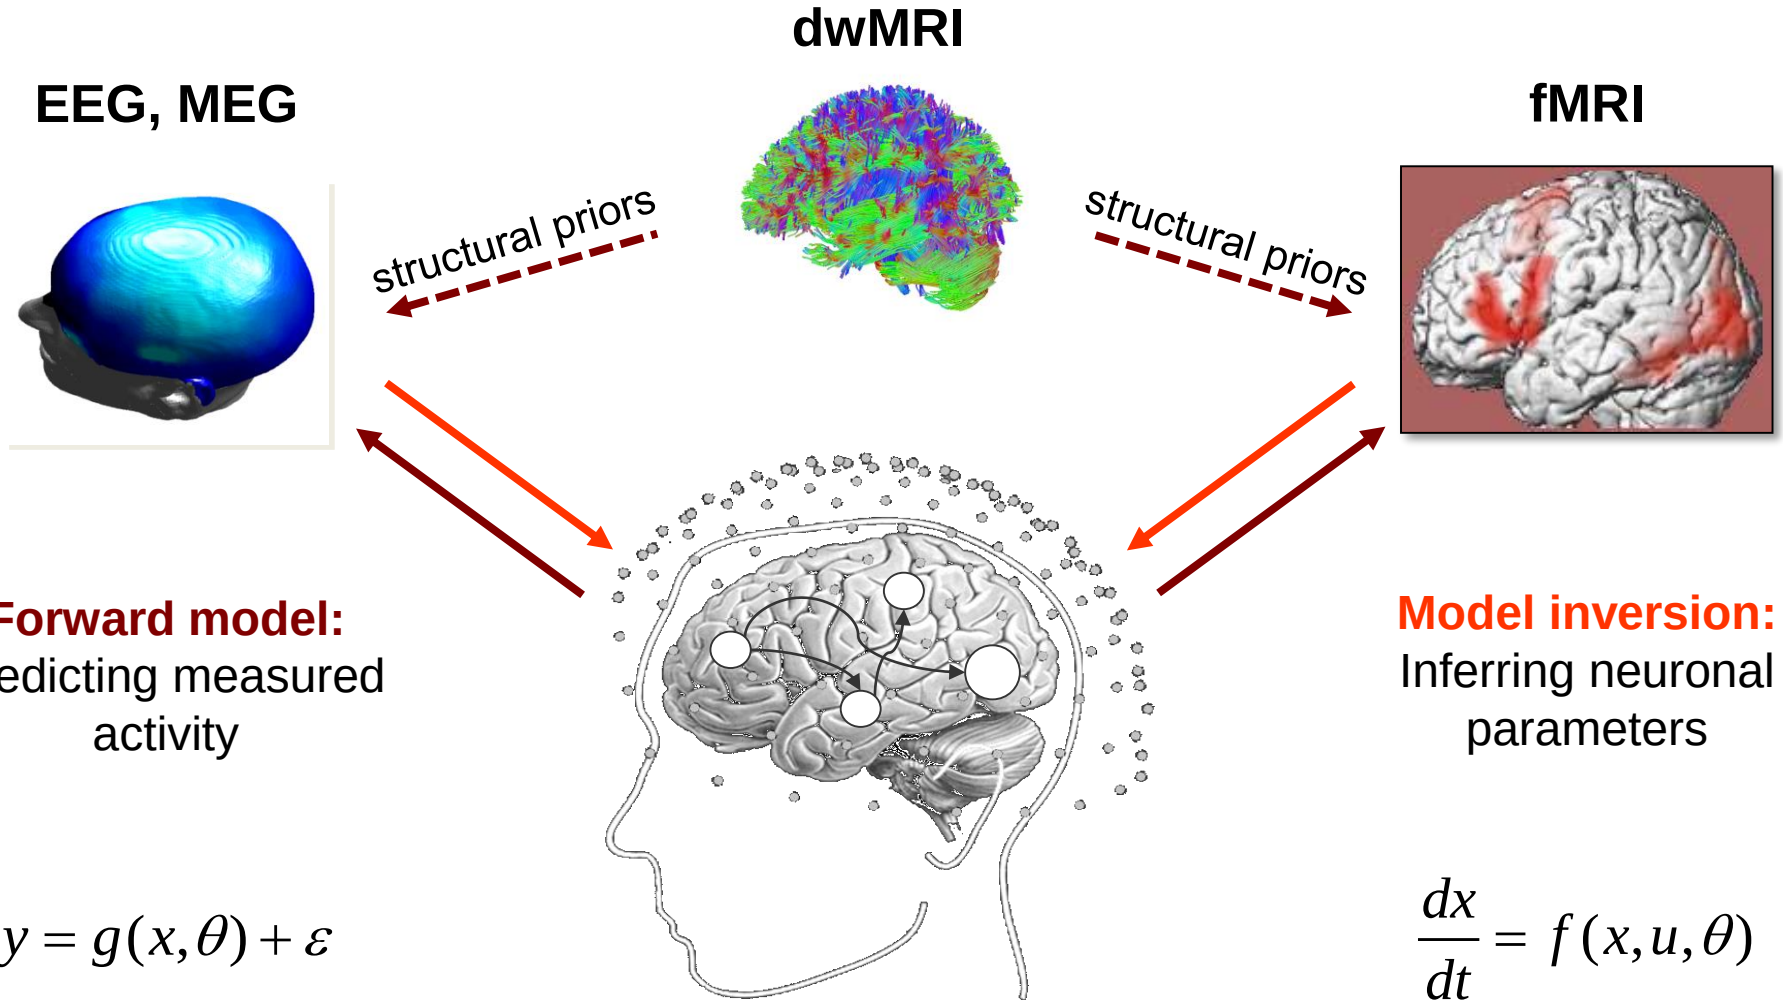

Forward model:
Predicting measured activity

$$y = g(x,\theta) + \varepsilon$$

Model inversion:
Inferring neuronal parameters

$$\frac{dx}{dt} = f(x,u,\theta)$$

Friston et al. 2003, *NeuroImage*

Stephan et al. 2009, *NeuroImage*

DCM for fMRI

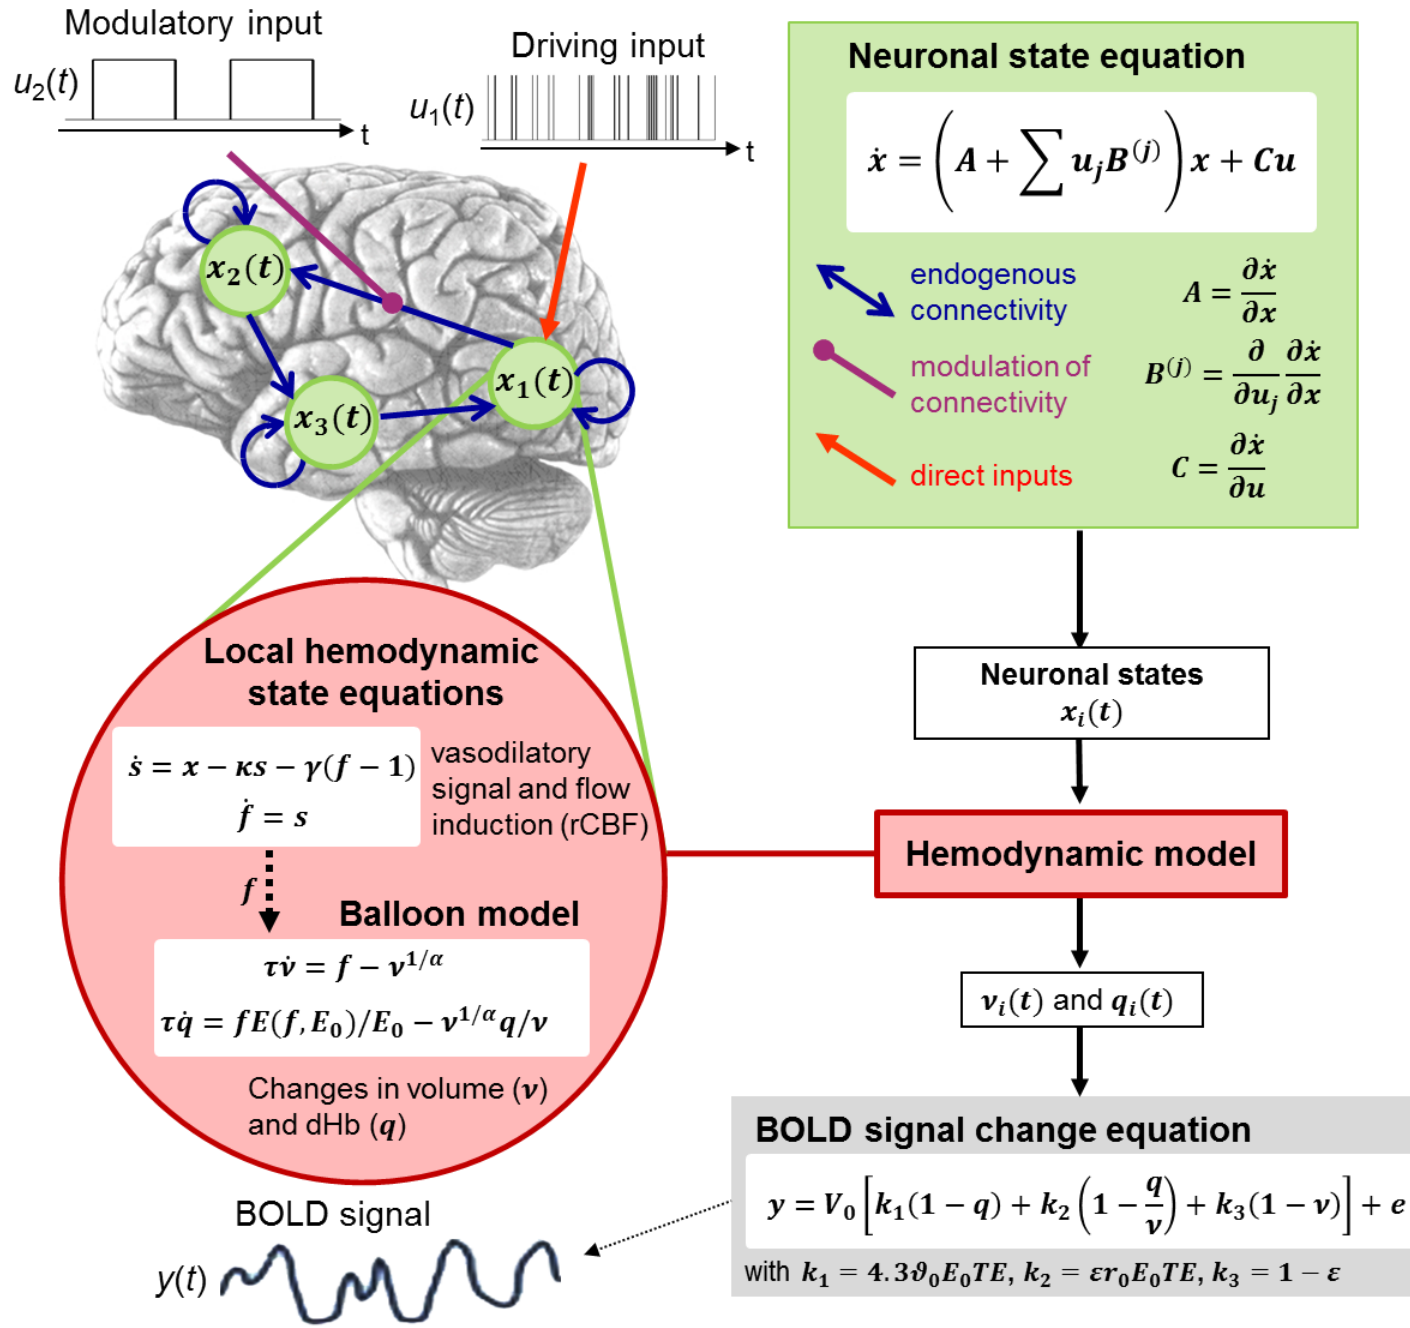

Stephan et al. 2015, *Neuron*

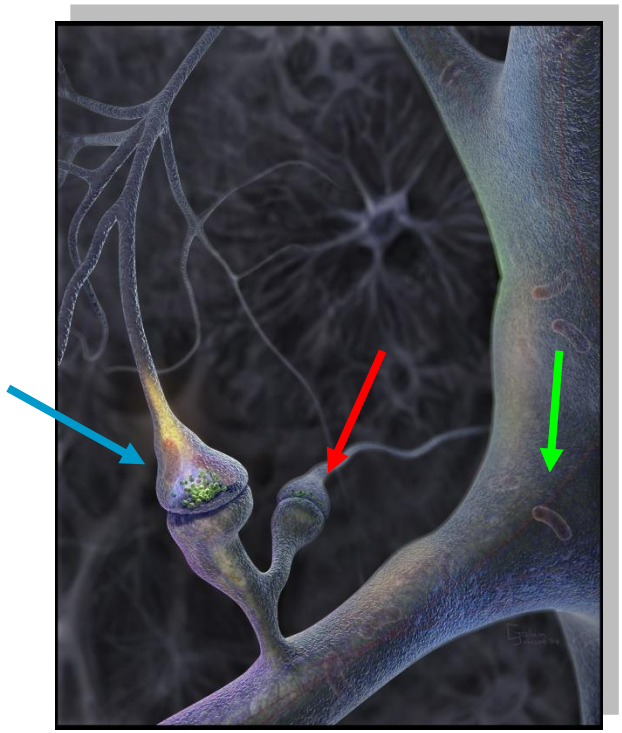

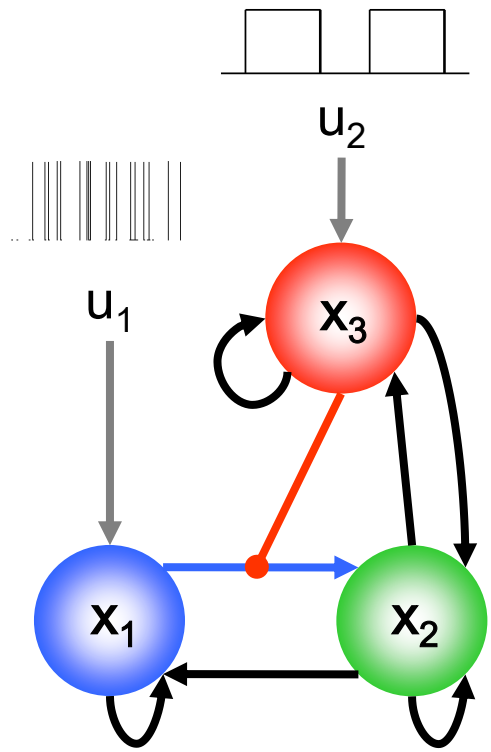

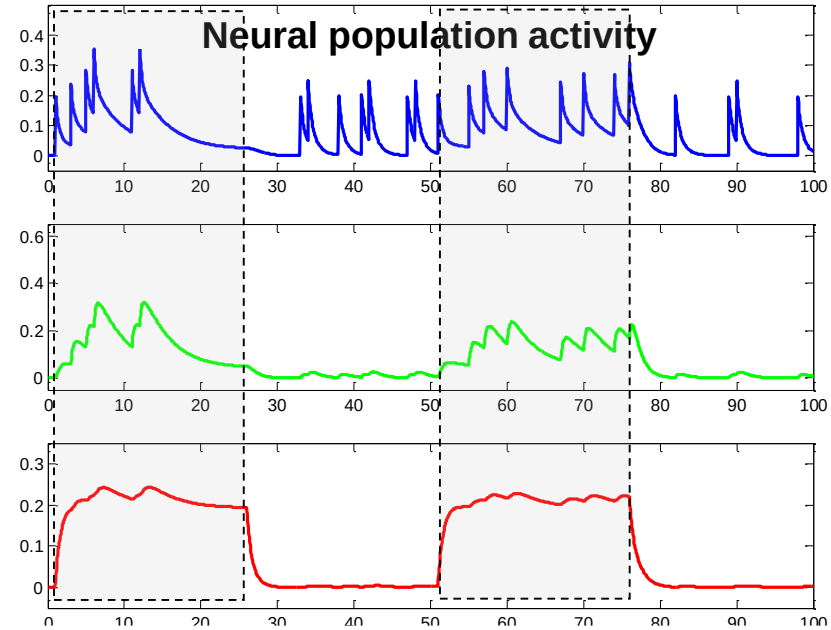

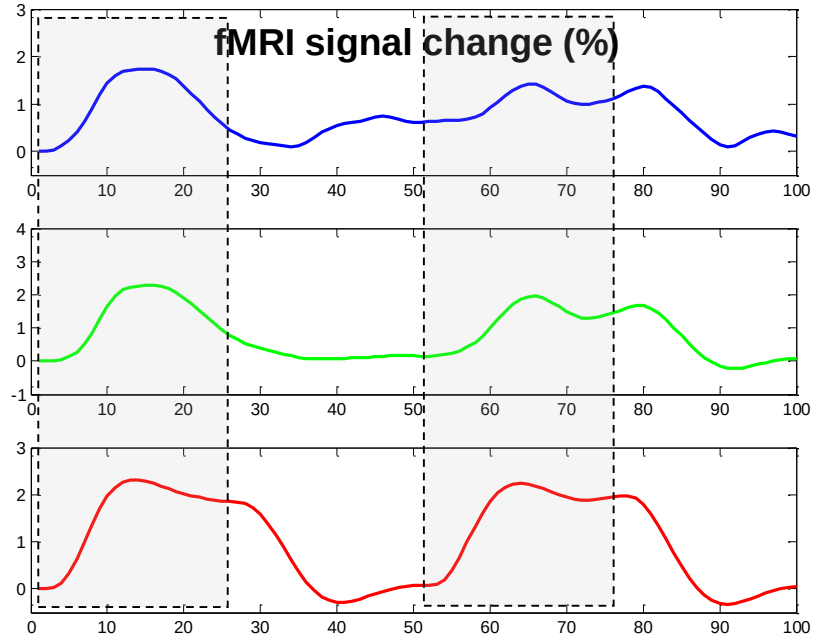

Nonlinear Dynamic Causal Model for fMRI

$$\frac{dx}{dt} = \left(A + \sum_{i=1}^{m} u_i B^{(i)} + \sum_{j=1}^{n} x_j D^{(j)} \right) x + Cu$$

Stephan et al. 2008, *NeuroImage*

Bayesian system identification

$u(t)$

Neural dynamics

$$dx/dt = f(x,u,\theta)$$

Observer function

$$y = g(x,\theta) + \varepsilon$$

$$p(y \mid \theta, m) = N(g(\theta), \Sigma(\theta))$$

$$p(\theta, m) = N(\mu_\theta, \Sigma_\theta)$$

Inference on model structure

$$p(y \mid m) = \int p(y \mid \theta, m)\, p(\theta) d\theta$$

Inference on parameters

$$p(\theta \mid y, m) = \frac{p(y \mid \theta, m)\, p(\theta, m)}{p(y \mid m)}$$

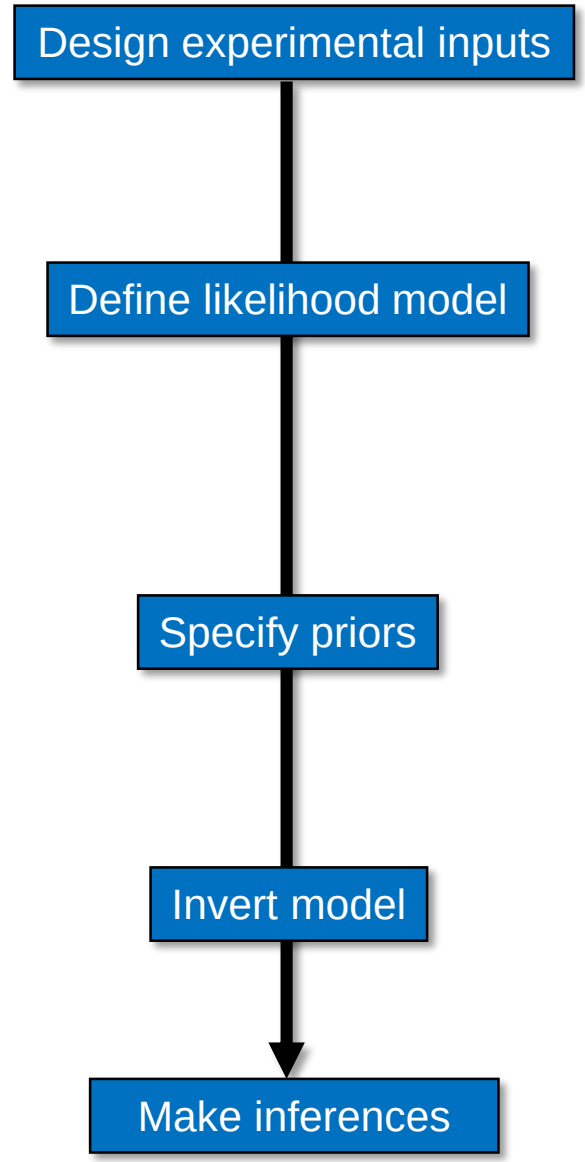

Methods for Bayesian inference

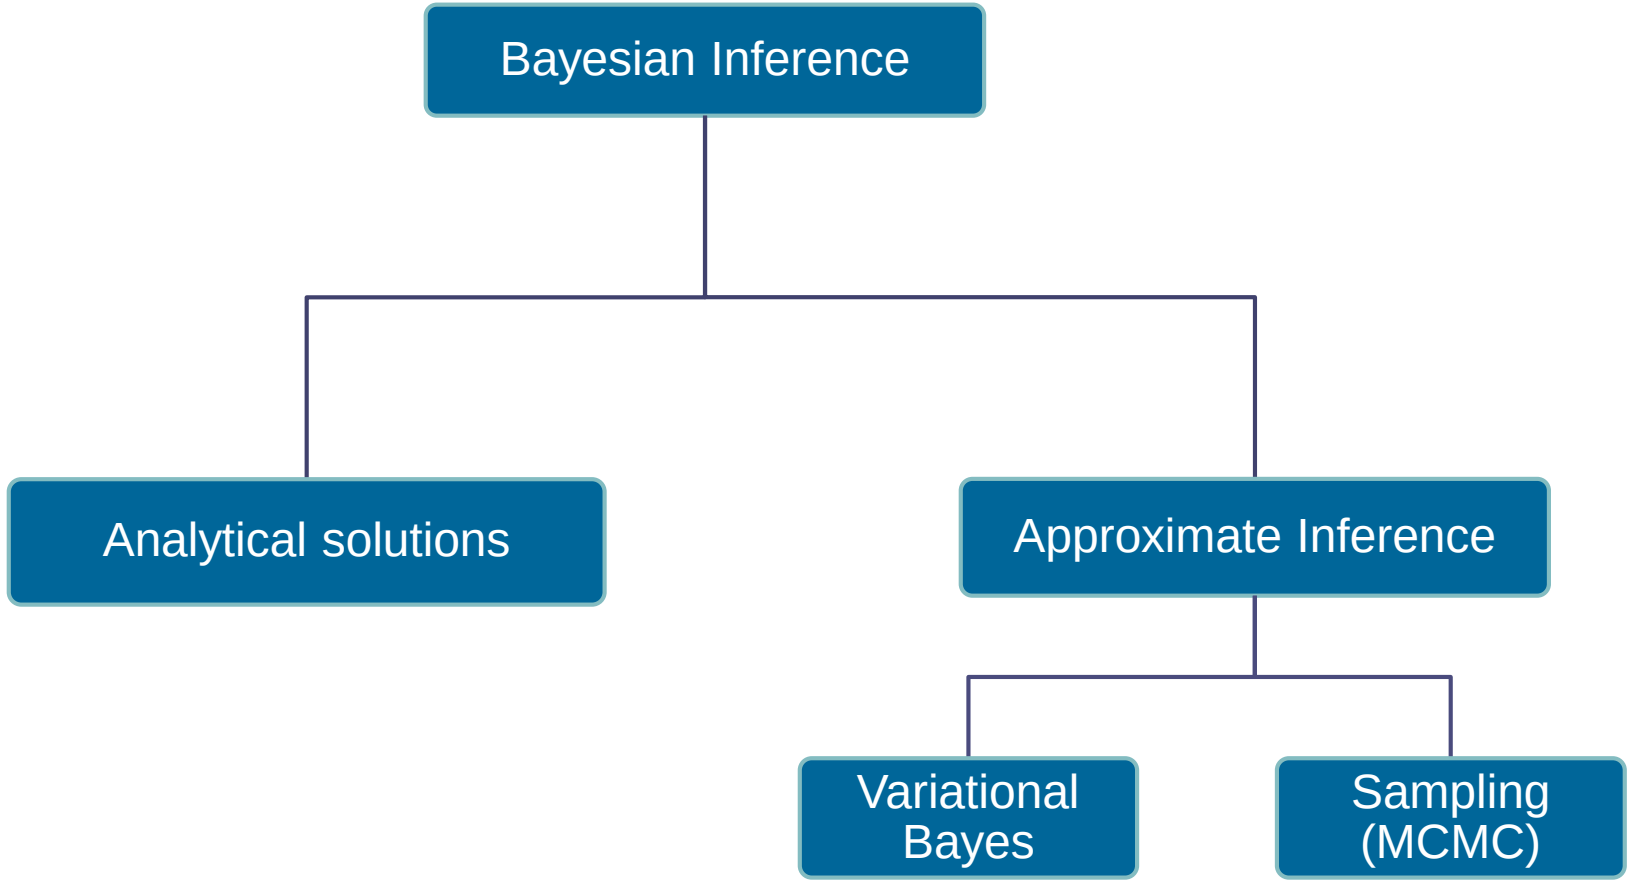

How is the posterior computed = how is a generative model inverted?

- **compute the posterior analytically**
  - requires conjugate priors
- **variational Bayes (VB)**
  - often hard work to derive, but fast to compute
  - uses approximations (approximate posteriors, mean field approx.)
  - problems: local minima, potentially inaccurate approximations
- **sampling methods (MCMC)**
  - theoretically guaranteed to be accurate (for infinite computation time)
  - problems: may require very long run time in practice, only heuristics to decide about convergence in practice

Conjugate priors

- for a given likelihood function, the choice of prior determines the algebraic form of the posterior
- for some probability distributions a prior can be found such that the posterior has the same algebraic form as the prior
- such a prior is called “conjugate” to the likelihood
- examples:
  - Normal $\propto$ Normal × Normal
  - Beta $\propto$ Binomial × Beta
  - Dirichlet $\propto$ Multinomial × Dirichlet

$$p(\boldsymbol{\theta} \mid \mathbf{y}) \propto p(\mathbf{y} \mid \boldsymbol{\theta})\, p(\boldsymbol{\theta})$$

same form

A simple example: univariate Gaussian belief update

Likelihood & prior

$$p(y \mid \theta) = N(\theta, \sigma_e^2)$$

$$p(\theta) = N(\mu_{prior}, \sigma_{prior}^2)$$

Posterior $p(\theta \mid y) = N\left(\mu_{post}, \lambda_{post}^{-1}\right)$

(for a single observation y)

$$\frac{1}{\sigma_{post}^2} = \frac{1}{\sigma_e^2} + \frac{1}{\sigma_{prior}^2}$$

$$\mu_{post} = \sigma_{post}^2 \left(\frac{1}{\sigma_e^2} y + \frac{1}{\sigma_{prior}^2} \mu_{prior} \right)$$

posterior mean = variance-weighted combination of prior mean and data

$$\mathbf{y} = \theta + \boldsymbol{\varepsilon}$$

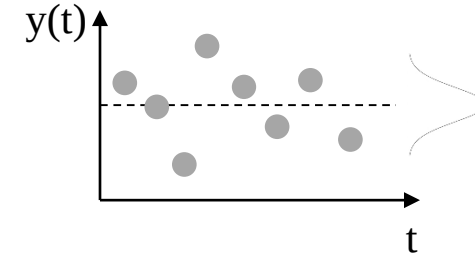

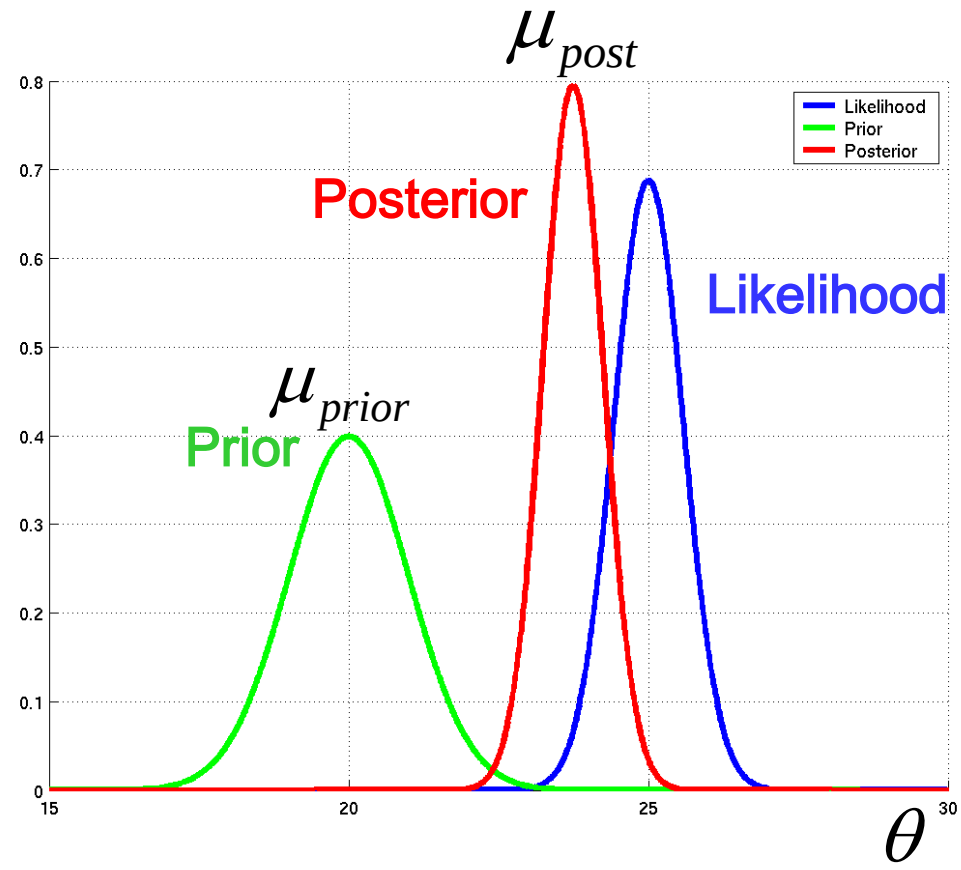

Slides on Gaussian belief updates adapted from Will Penny.

Same thing – but expressed as precision weighting

Likelihood & prior

$$p(y \mid \theta) = N(\theta, \lambda_e^{-1})$$

$$p(\theta) = N(\mu_{prior}, \lambda_{prior}^{-1})$$

Posterior $p(\theta \mid y) = N(\mu_{post}, \lambda_{post}^{-1})$

(for a single observation y)

$$\lambda_{post} = \lambda_e + \lambda_{prior}$$

$$\mu_{post} = \frac{\lambda_e}{\lambda_{post}} y + \frac{\lambda_{prior}}{\lambda_{post}} \mu_{prior}$$

relative precision weighting:
posterior mean = precision-weighted combination of prior mean and data

$$\mathbf{y} = \theta + \boldsymbol{\varepsilon}$$

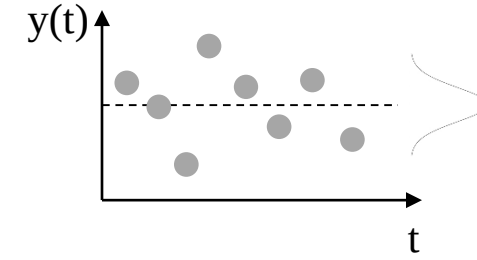

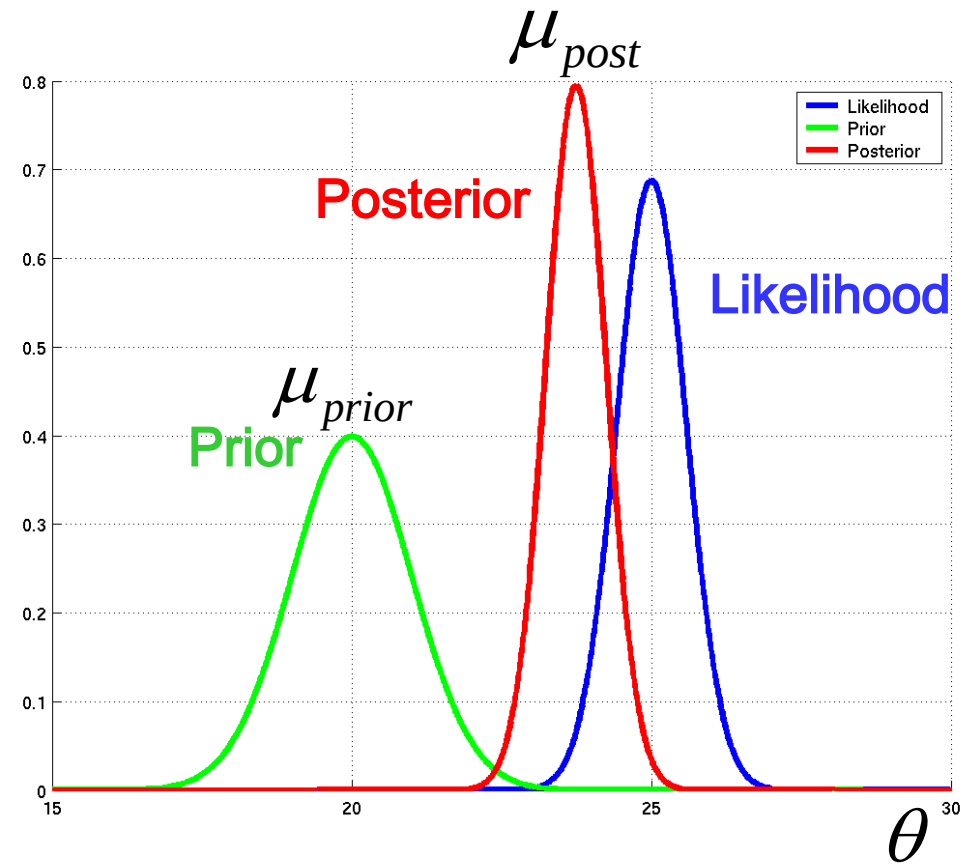

Variational Bayes (VB)

Idea: find an approximate density $q(\theta)$ that is maximally similar to the true posterior $p(\theta|y)$.

This is often done by assuming a particular form for q (fixed form VB) and then optimizing its sufficient statistics.

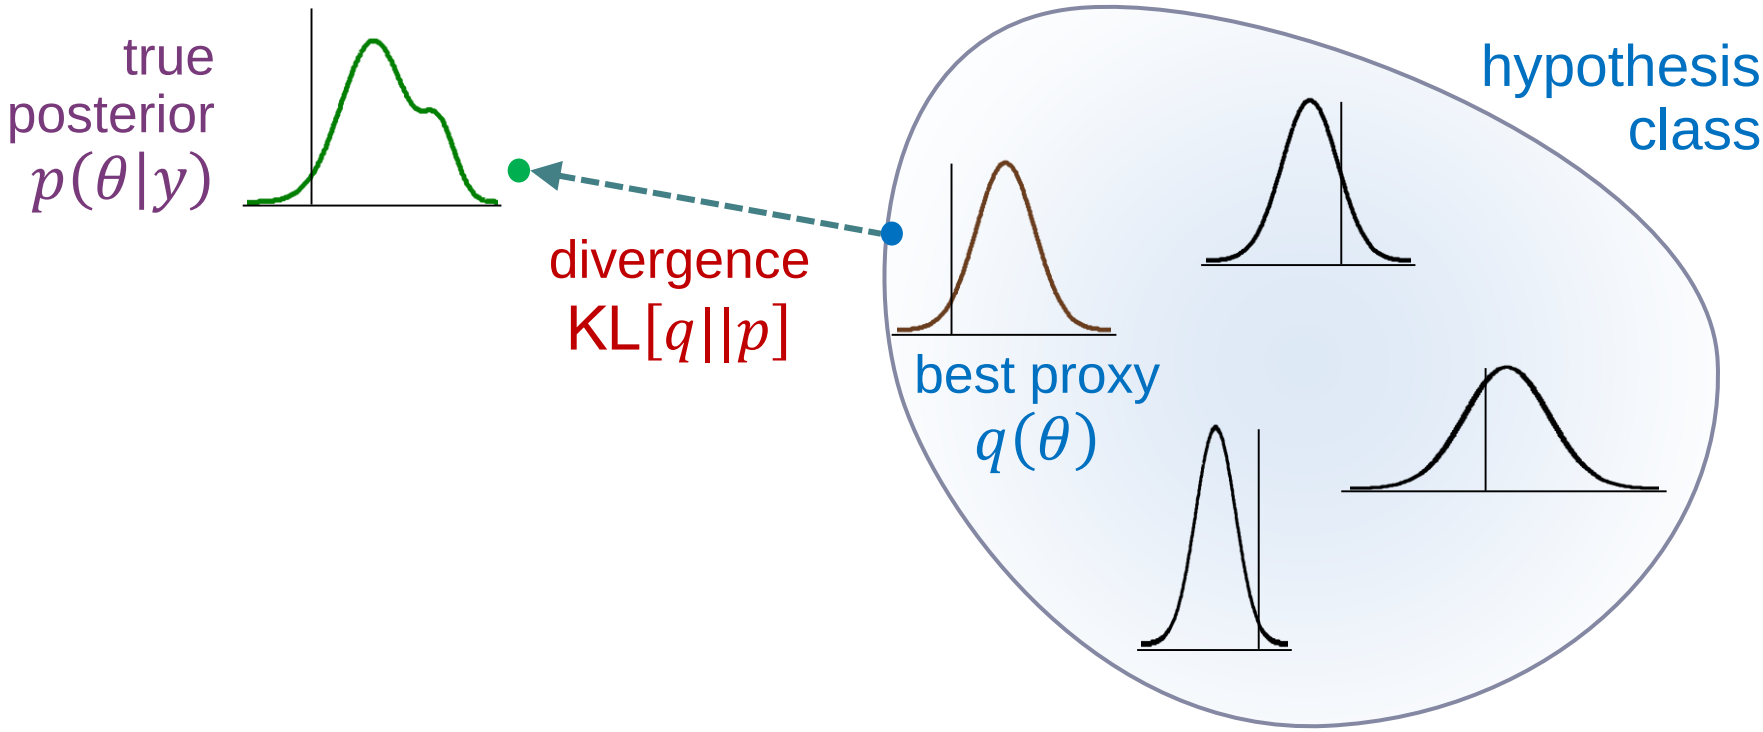

Kullback–Leibler (KL) divergence

- asymmetric measure of the difference between two probability distributions P and Q
- Interpretations of $D_{KL}(P\|Q)$:
  - "Bayesian surprise" when Q=prior, P=posterior: measure of the information gained when one updates one's prior beliefs to the posterior P
  - a measure of the information lost when Q is used to approximate P
- non-negative: ≥ 0 (zero when P=Q)

$$D_{KL}\left(P \,||\, Q\right) = \sum_i P\left(i\right) \ln \frac{P\left(i\right)}{Q\left(i\right)}$$

$$D_{KL}\left(P \,||\, Q\right) = \int p\left(x\right) \ln \frac{p\left(x\right)}{q\left(x\right)} dx$$

Variational calculus

Standard calculus
Newton, Leibniz, and others

- functions $f: x \mapsto f(x)$
- derivatives $\frac{\mathrm{d}f}{\mathrm{d}x}$

Example: maximize the likelihood expression $p(y|\theta)$ w.r.t. θ

Variational calculus
Euler, Lagrange, and others

- functionals $F: f \mapsto F(f)$
- derivatives $\frac{\mathrm{d}F}{\mathrm{d}f}$

Example: maximize the entropy $H[p]$ w.r.t. a probability distribution $p(x)$

Leonhard Euler
(1707 – 1783)

Swiss mathematician, 'Elementa Calculi Variationum'

Variational Bayes

$$\ln p(y) = \underbrace{\mathrm{KL}[q||p]}_{\text{divergence}} + \underbrace{F(q,y)}_{\text{neg. free energy}}$$

divergence ≥ 0 (unknown)

neg. free energy (easy to evaluate for a given q)

$F(q)$ is a functional wrt. the approximate posterior $q(\theta)$.

Maximizing $F(q,y)$ is equivalent to:

- minimizing $\mathrm{KL}[q||p]$
- tightening $F(q,y)$ as a lower bound to the log model evidence

When $F(q,y)$ is maximized, $q(\theta)$ is our best estimate of the posterior.

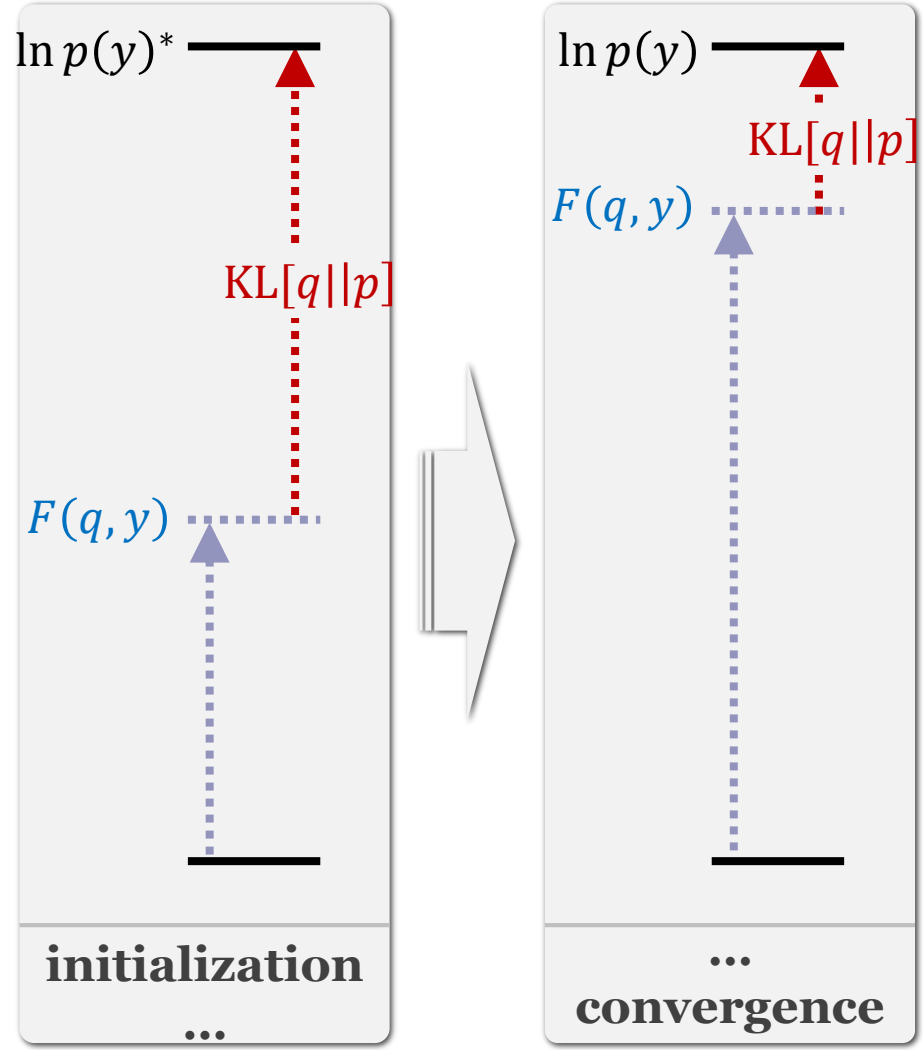

Derivation of the (negative) free energy approximation

- See whiteboard!
- (or Appendix to Stephan et al. 2007, NeuroImage 38: 387-401)

Mean field assumption

Factorize the approximate posterior $q(\theta)$ into independent partitions:

$$q(\theta) = \prod_i q_i(\theta_i)$$

where $q_i(\theta_i)$ is the approximate posterior for the i^{th} subset of parameters.

For example, split parameters and hyperparameters:

$$p(\theta, \lambda \mid y) \approx q(\theta, \lambda) = q(\theta) q(\lambda)$$

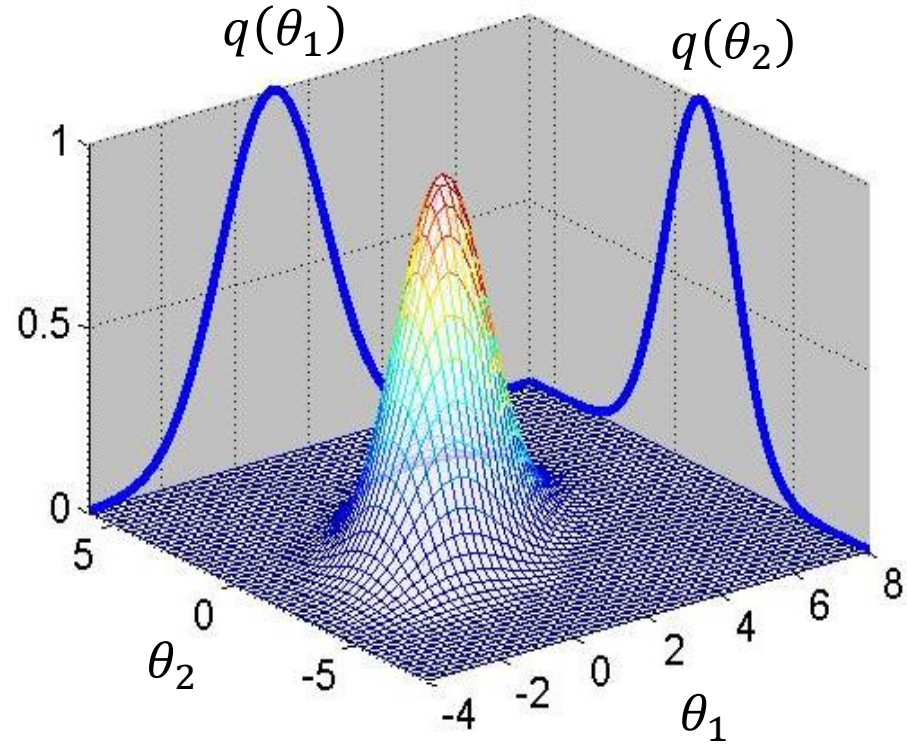

Jean Daunizeau, www.fil.ion.ucl.ac.uk/~jdaunize/presentations/Bayes2.pdf

VB in a nutshell (under mean-field approximation)

❶ Neg. free-energy approx. to model evidence.

$$\ln p(y \mid m) = F + KL\left[q(\theta,\lambda), p(\theta,\lambda \mid y)\right]$$

$$F = \left\langle \ln p(y \mid \theta,\lambda) \right\rangle_q - KL\left[q(\theta,\lambda), p(\theta,\lambda \mid m)\right]$$

❷ Mean field approx.

$$p(\theta,\lambda \mid y) \approx q(\theta,\lambda) = q(\theta)q(\lambda)$$

❸ Maximise neg. free energy wrt. q = minimise divergence, by maximising variational energies

$$q(\theta) \propto \exp(I_\theta) = \exp\left[\left\langle \ln p(y,\theta,\lambda) \right\rangle_{q(\lambda)}\right]$$

$$q(\lambda) \propto \exp(I_\lambda) = \exp\left[\left\langle \ln p(y,\theta,\lambda) \right\rangle_{q(\theta)}\right]$$

❹ Iterative updating of sufficient statistics of approx. posteriors by gradient ascent.

Model comparison and selection

Given competing hypotheses on structure & functional mechanisms of a system, which model is the best?

⇩

Which model represents the best balance between model fit and model complexity?

⇩

For which model m does $p(y|m)$ become maximal?

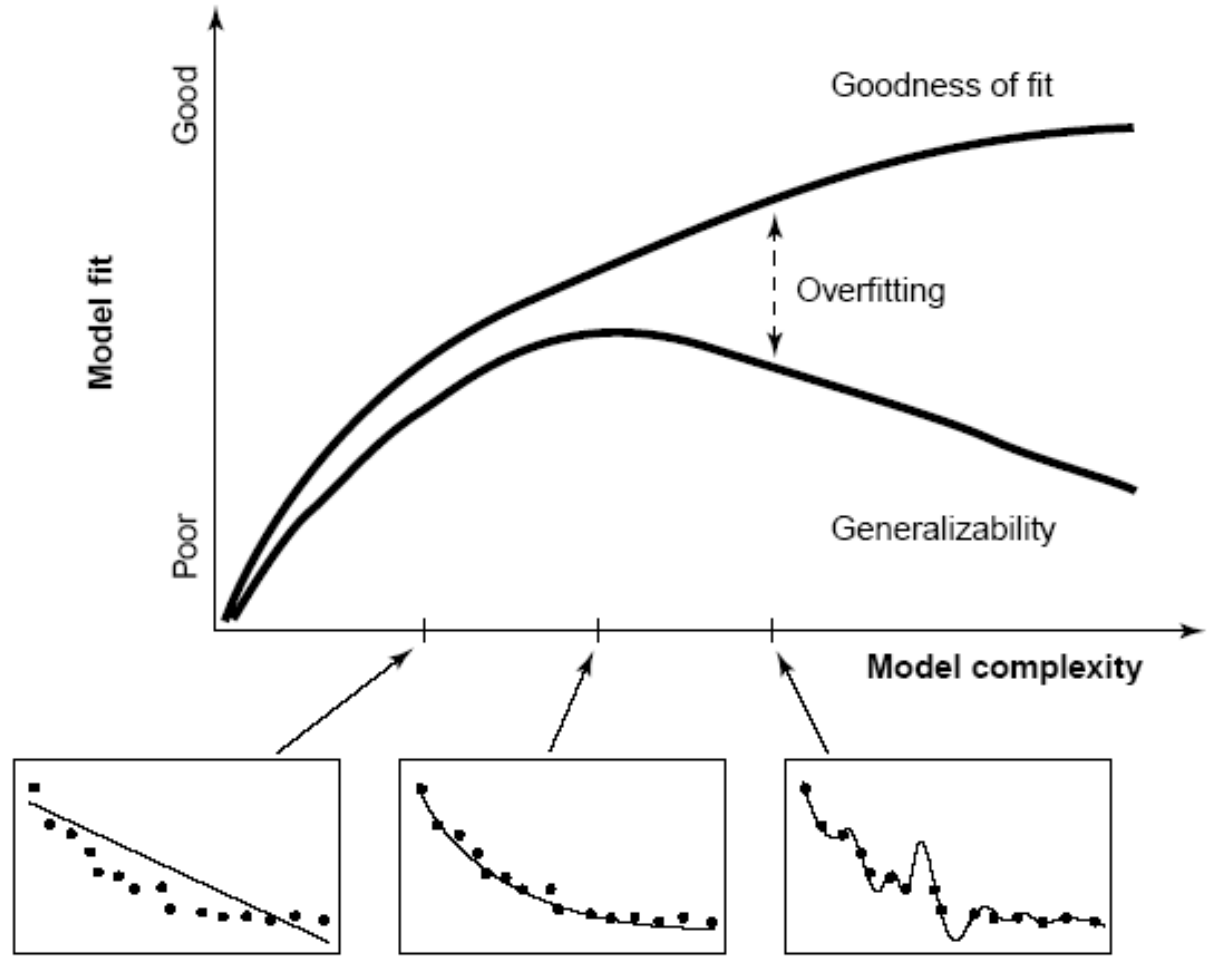

Pitt & Miyung (2002) *TICS*

Bayesian model selection (BMS)

- First step of inference: define model space M

$$|M| \in [1, \infty[$$

- Inference on model structure m:

Posterior model probability

$$p(m \mid y) = \frac{p(y \mid m)\, p(m)}{p(y)}$$

$$= \frac{p(y \mid m)\, p(m)}{\sum_m p(y \mid m)\, p(m)}$$

- For a uniform prior on m, model evidence sufficient for model selection

Model evidence:

$$p(y \mid m) = \int p(y \mid \theta, m)\, p(\theta \mid m)\, d\theta$$

Bayesian model selection (BMS)

Model evidence:

$$p(y \mid m) = \int p(y \mid \theta, m)\, p(\theta \mid m)\, d\theta$$

⟹ probability that data were generated by model m, averaging over all possible parameter values (with probability weights as specified by the prior)

⟹ accounts for both accuracy and complexity of the model

Ghahramani 2004

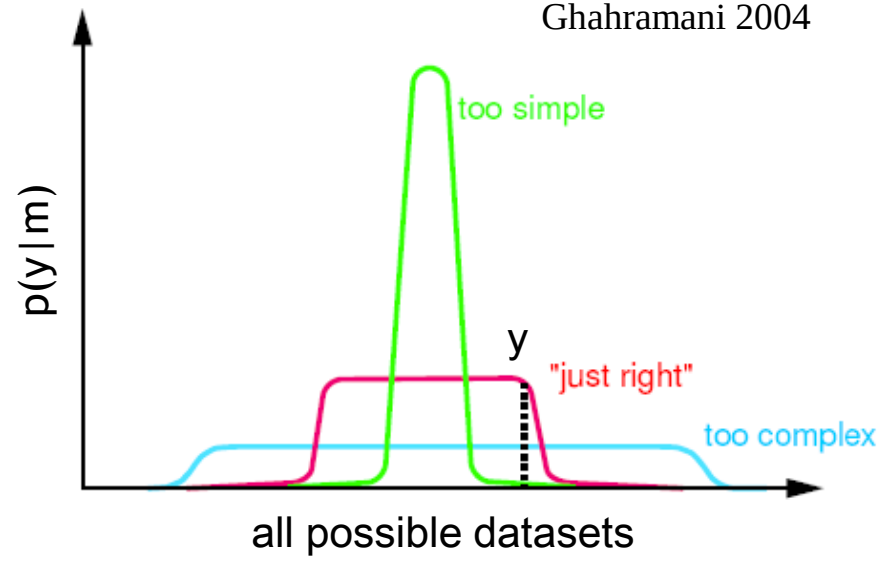

Various approximations:

- negative free energy (F)
- Akaike Information Criterion (AIC)
- Bayesian Information Criterion (BIC)

Bayesian model selection (BMS)

Model evidence:

$$p(y \mid m) = \int p(y \mid \theta, m)\, p(\theta \mid m)\, d\theta$$

⟹ "If I randomly sampled from my prior and plugged the resulting value into the likelihood function, how close would the predicted data be – on average – to my observed data?"

⟹ accounts for both accuracy and complexity of the model

Ghahramani 2004

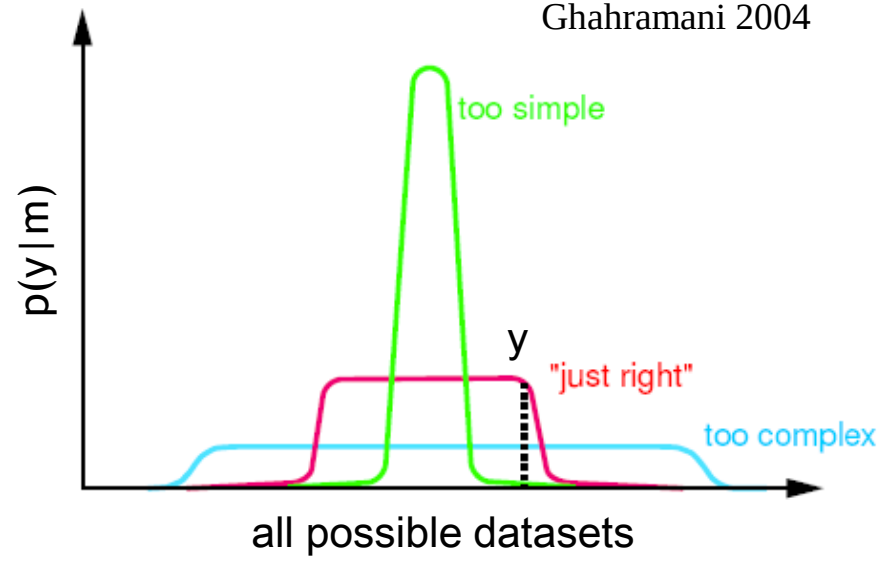

Various approximations:

- negative free energy (F)
- Akaike Information Criterion (AIC)
- Bayesian Information Criterion (BIC)

Approximations to the model evidence

Logarithm is a monotonic function → Maximizing log model evidence = Maximizing model evidence

Log model evidence = balance between fit and complexity

$$\log p(y \mid m) = accuracy(m) - complexity(m)$$
$$= \log p(y \mid \theta, m) - complexity(m)$$

Akaike Information Criterion: $AIC = \log p(y \mid \theta, m) - p$

Bayesian Information Criterion: $BIC = \log p(y \mid \theta, m) - \frac{p}{2} \log N$

No. of parameters

No. of data points

The (negative) free energy approximation F

F is a lower bound on the log model evidence:

$$\log p(y \mid m) = F + KL\left[q(\theta), p(\theta \mid y, m)\right]$$

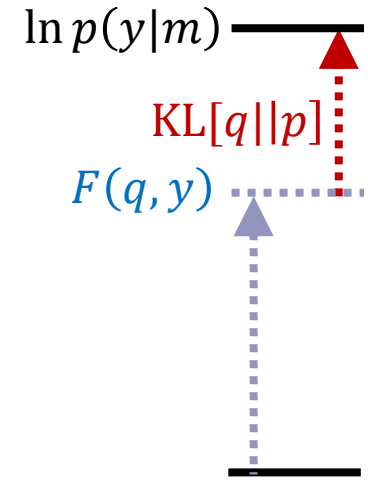

Like AIC/BIC, F is an accuracy/complexity tradeoff:

$$F = \underbrace{\left\langle \log p(y \mid \theta, m) \right\rangle}_{accuracy} - \underbrace{KL\left[q(\theta), p(\theta \mid m)\right]}_{complexity}$$

The (negative) free energy approximation

- Log evidence is thus expected log likelihood (wrt. q) plus 2 KL's:

$$\log p(y \mid m)$$
$$= \left\langle \log p(y \mid \theta, m) \right\rangle - KL\left[q(\theta), p(\theta \mid m) \right] + KL\left[q(\theta), p(\theta \mid y, m) \right]$$

$$F = \log p(y \mid m) - KL\left[q(\theta), p(\theta \mid y, m) \right]$$
$$= \underbrace{\left\langle \log p(y \mid \theta, m) \right\rangle}_{accuracy} - \underbrace{KL\left[q(\theta), p(\theta \mid m) \right]}_{complexity}$$

The complexity term in *F*

- In contrast to AIC & BIC, the complexity term of the negative free energy *F* accounts for parameter interdependencies.
  Under Gaussian assumptions about the posterior (Laplace approximation):

$$KL\left[q(\theta), p(\theta \mid m)\right]$$
$$= \frac{1}{2}\ln\left|C_\theta\right| - \frac{1}{2}\ln\left|C_{\theta|y}\right| + \frac{1}{2}\left(\mu_{\theta|y} - \mu_\theta\right)^T C_\theta^{-1}\left(\mu_{\theta|y} - \mu_\theta\right)$$

- The complexity term of *F* is higher
  - the more independent the prior parameters (↑ effective DFs)
  - the more dependent the posterior parameters
  - the more the posterior mean deviates from the prior mean

Bayes factors

To compare two models, we could just compare their log evidences.

But: the log evidence is just some number – not very intuitive!

A more intuitive interpretation of model comparisons is made possible by Bayes factors:

$$B_{12} = \frac{p(y \mid m_1)}{p(y \mid m_2)}$$

positive value, [0; ∞[

Kass & Raftery classification:

B_{12}	$p(m_1 \mid y)$	Evidence
1 to 3	50-75%	weak
3 to 20	75-95%	positive
20 to 150	95-99%	strong
≥ 150	≥ 99%	Very strong

Kass & Raftery 1995, *J. Am. Stat. Assoc.*

Fixed effects BMS at group level

Group Bayes factor (GBF) for 1...*K* subjects:

$$GBF_{ij} = \prod_k BF_{ij}^{(k)}$$

Average Bayes factor (ABF):

$$ABF_{ij} = \sqrt[K]{\prod_k BF_{ij}^{(k)}}$$

Problems:

- blind with regard to group heterogeneity
- sensitive to outliers

Random effects BMS for heterogeneous groups

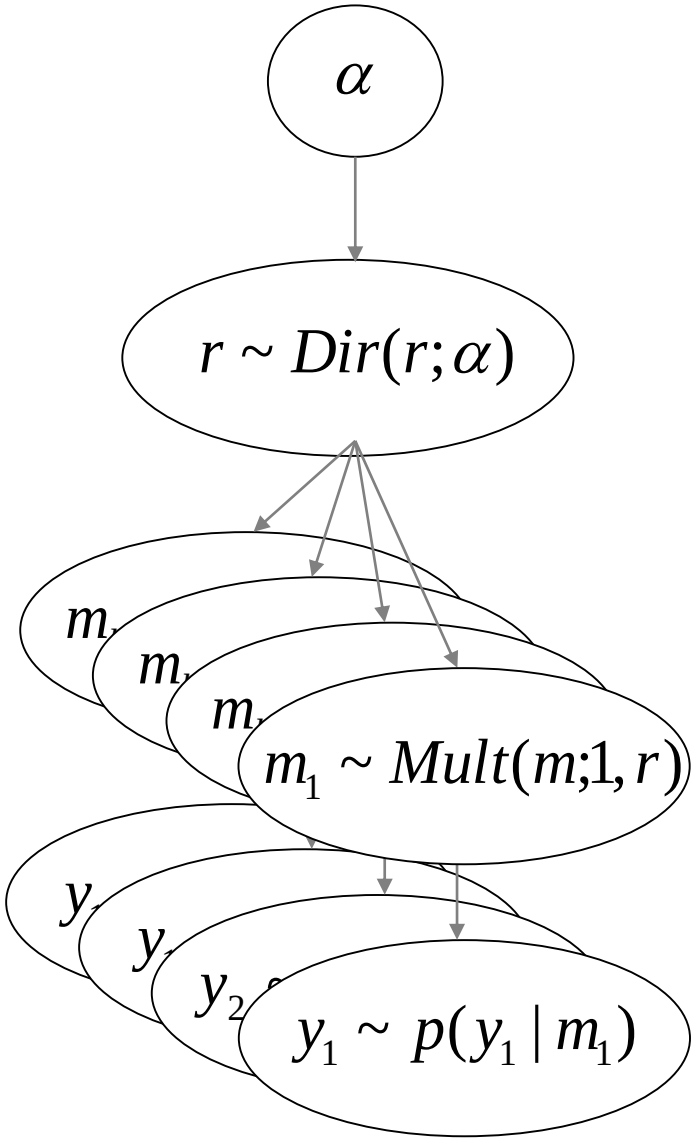

Dirichlet parameters α
= “occurrences” of models in the population

Dirichlet distribution of model probabilities r

Multinomial distribution of model labels m

Model inversion by Variational Bayes or MCMC

Measured data y

Stephan et al. 2009, *NeuroImage*

Random effects BMS for heterogeneous groups

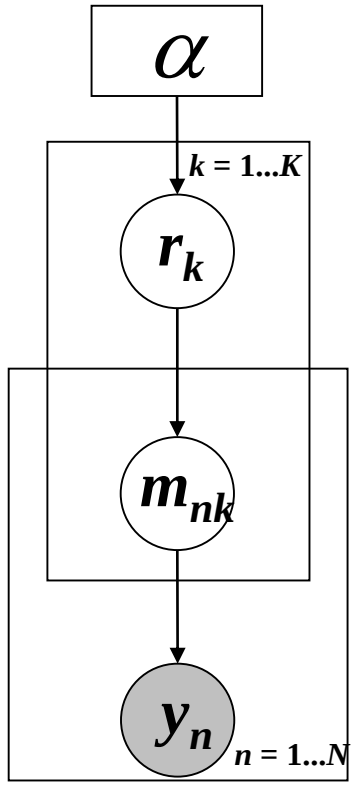

Dirichlet parameters α
= "occurrences" of models in the population

Dirichlet distribution of model probabilities r

Multinomial distribution of model labels m

Model inversion by Variational Bayes or MCMC

Measured data y

Stephan et al. 2009, *NeuroImage*

Four equivalent options for reporting model ranking by random effects BMS

1. **Dirichlet parameter estimates**

$$\alpha$$

2. **expected posterior probability** of obtaining the k-th model for any randomly selected subject

$$\left\langle r_k \right\rangle_q = \alpha_k / (\alpha_1 + \ldots + \alpha_K)$$

3. **exceedance probability** that a particular model *k* is more likely than any other model (of the *K* models tested), given the group data

$$\exists k \in \{1...K\}, \forall j \in \{1...K \mid j \neq k\}:$$

$$\varphi_k = p(r_k > r_j \mid y; \alpha)$$

4. **protected exceedance probability**: see below

Example: Hemispheric interactions during vision

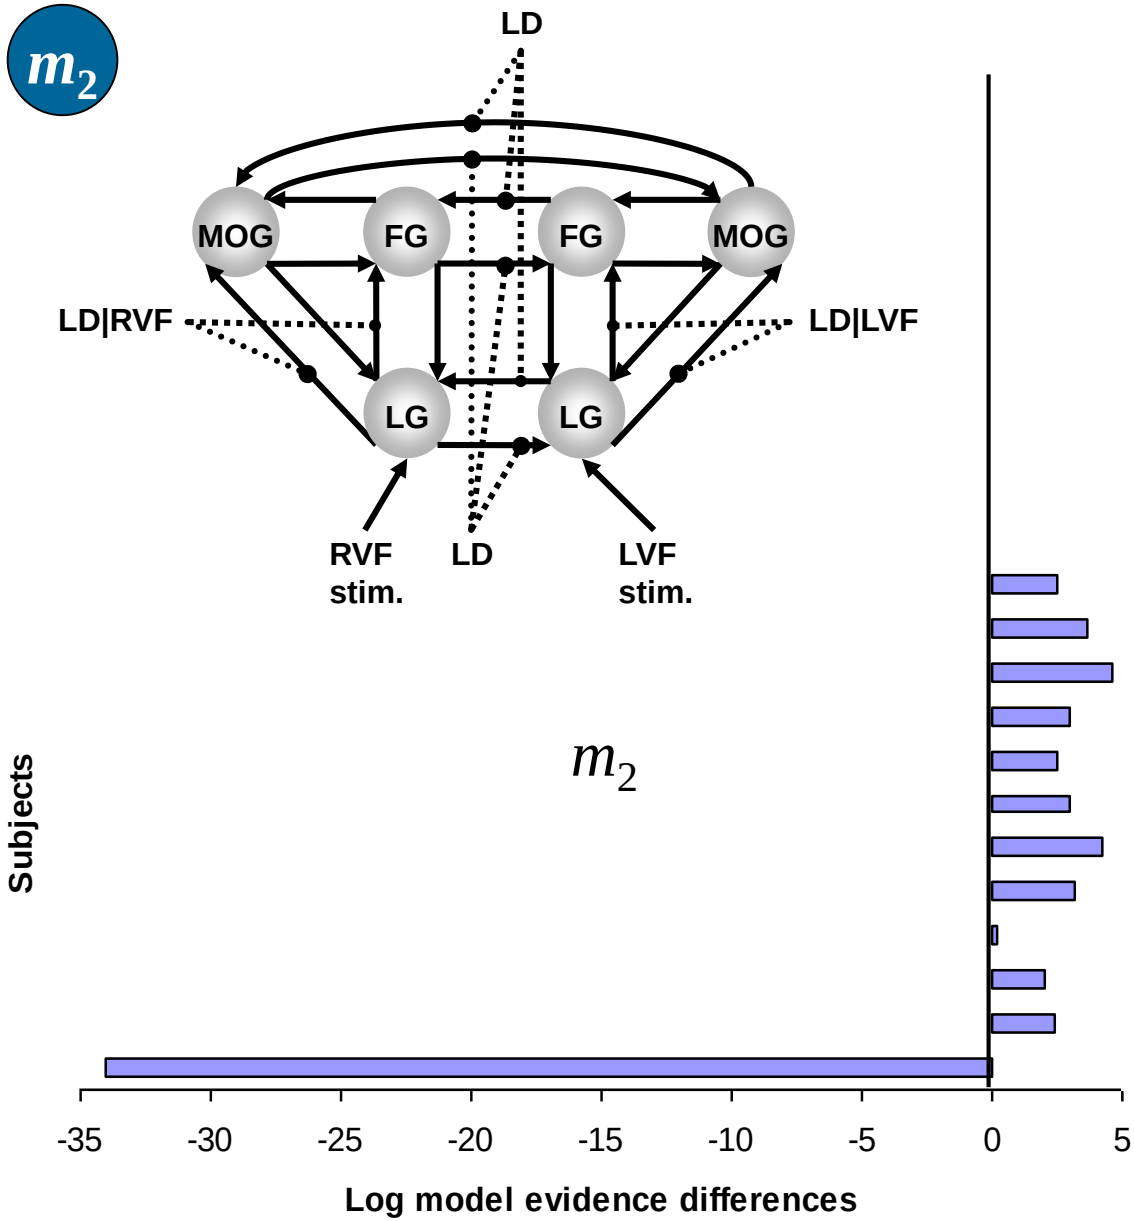

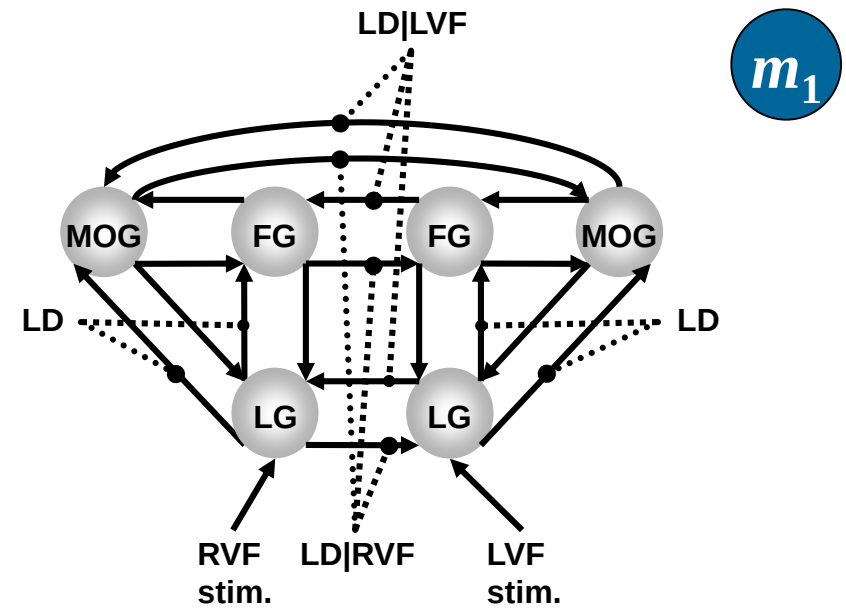

Data: Stephan et al. 2003, *Science*
Models: Stephan et al. 2007, *J. Neurosci.*

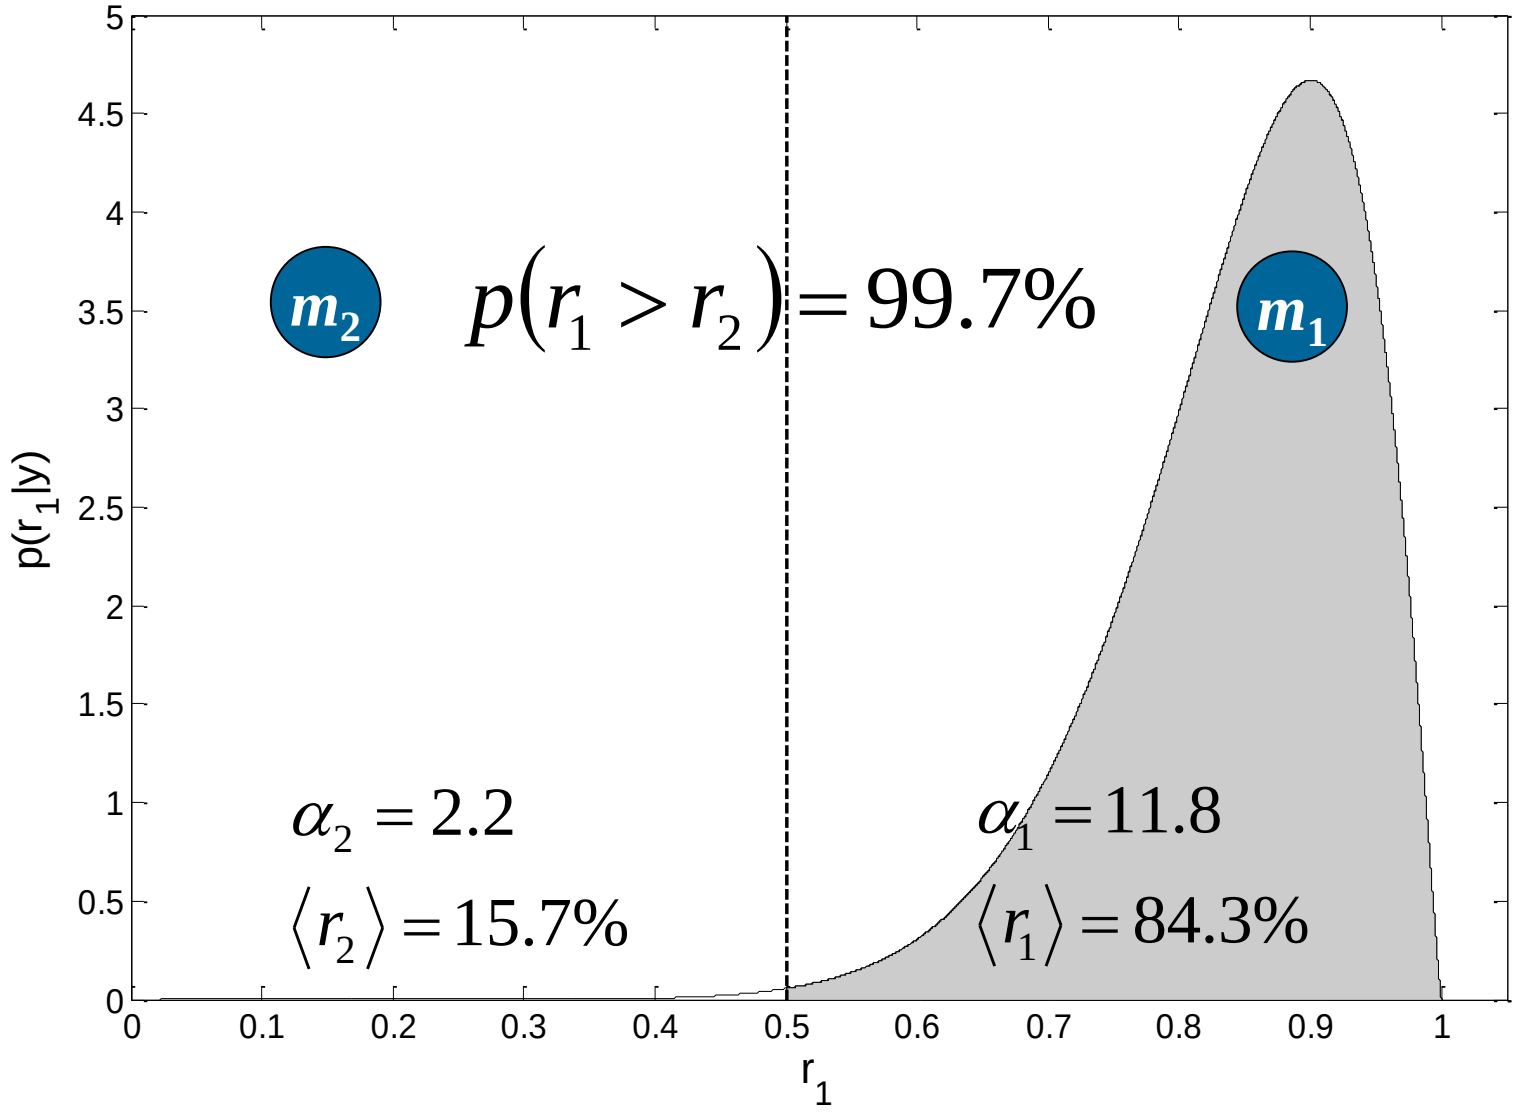

Stephan et al. 2009a, *NeuroImage*

Example: Synaesthesia

- “projectors” experience color externally colocalized with a presented grapheme
- “associators” report an internally evoked association
- across all subjects: no evidence for either model
- but BMS results map precisely onto projectors (bottom-up mechanisms) and associators (top-down)

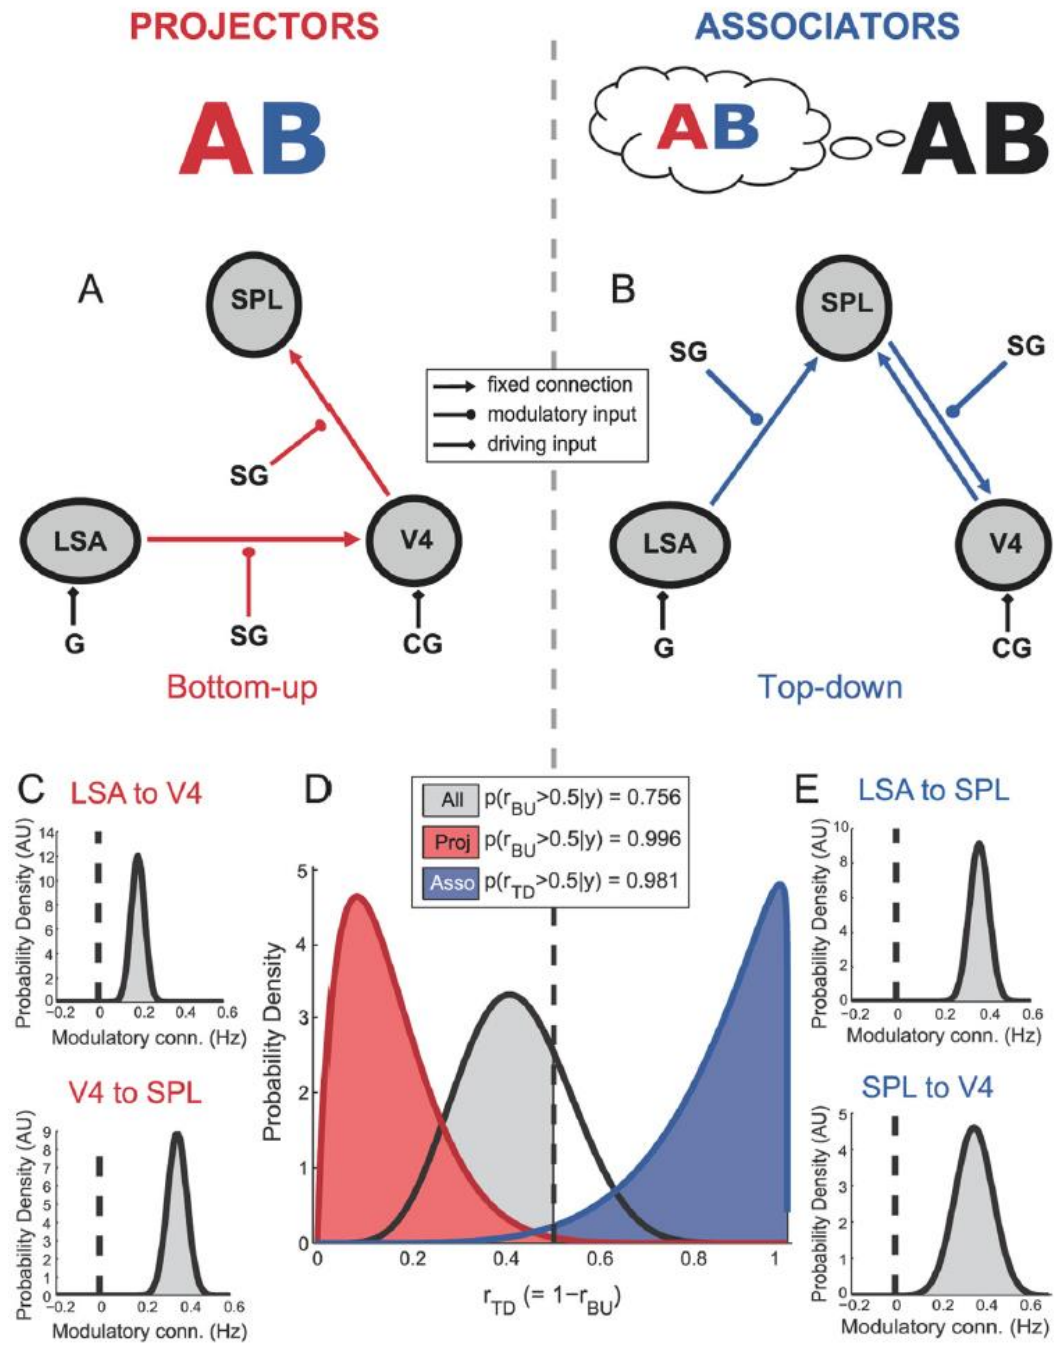

van Leeuwen et al. 2011, *J. Neurosci.*

Overfitting at the level of models

- ↑ #models ⇒ ↑ risk of overfitting
- solutions:
  - regularisation: definition of model space = choosing priors p(m)
  - family-level BMS
  - Bayesian model averaging (BMA)

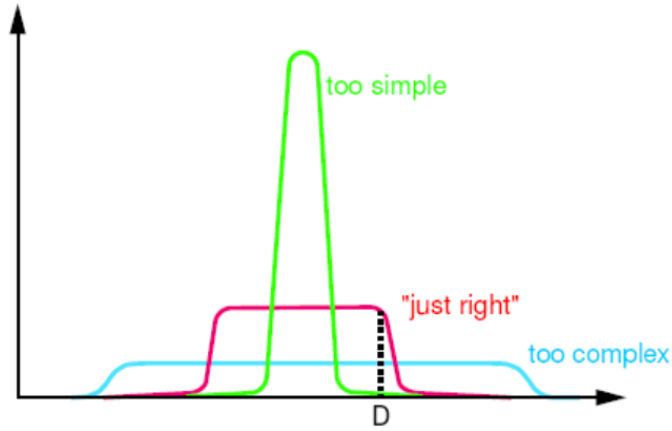

posterior model probability:

$$p(m \mid y)$$

$$= \frac{p(y \mid m)\, p(m)}{\sum_m p(y \mid m)\, p(m)}$$

BMA:

$$p(\theta \mid y)$$

$$= \sum_m p(\theta \mid y, m) p(m \mid y)$$

Model space partitioning: comparing model families

- partitioning model space into K subsets or families:

$$M = \{ f_1, \ldots, f_K \}$$

- pooling information over all models in these subsets allows one to compute the probability of a model family, given the data

$$p(f_k)$$

- effectively removes uncertainty about any aspect of model structure, other than the attribute of interest (which defines the partition)

Stephan et al. 2009, *NeuroImage*
Penny et al. 2010, *PLoS Comput. Biol.*

Family-level inference: random effects – a special case

- When the families are of equal size, one can simply sum the posterior model probabilities within families by exploiting the agglomerative property of the Dirichlet distribution:

$$\left(r_1, r_2, ..., r_K\right) \sim Dir\left(\alpha_1, \alpha_2, ..., \alpha_K\right)$$

$$\Rightarrow r_1^* = \sum_{k \in N_1} r_k, r_2^* = \sum_{k \in N_2} r_k, ..., r_J^* = \sum_{k \in N_J} r_k$$

$$\sim Dir\left(\alpha_1^* = \sum_{k \in N_1} \alpha_k, \alpha_2^* = \sum_{k \in N_2} \alpha_k, ..., \alpha_J^* = \sum_{k \in N_J} \alpha_k\right)$$

Stephan et al. 2009, *NeuroImage*

Model space partitioning: comparing model families

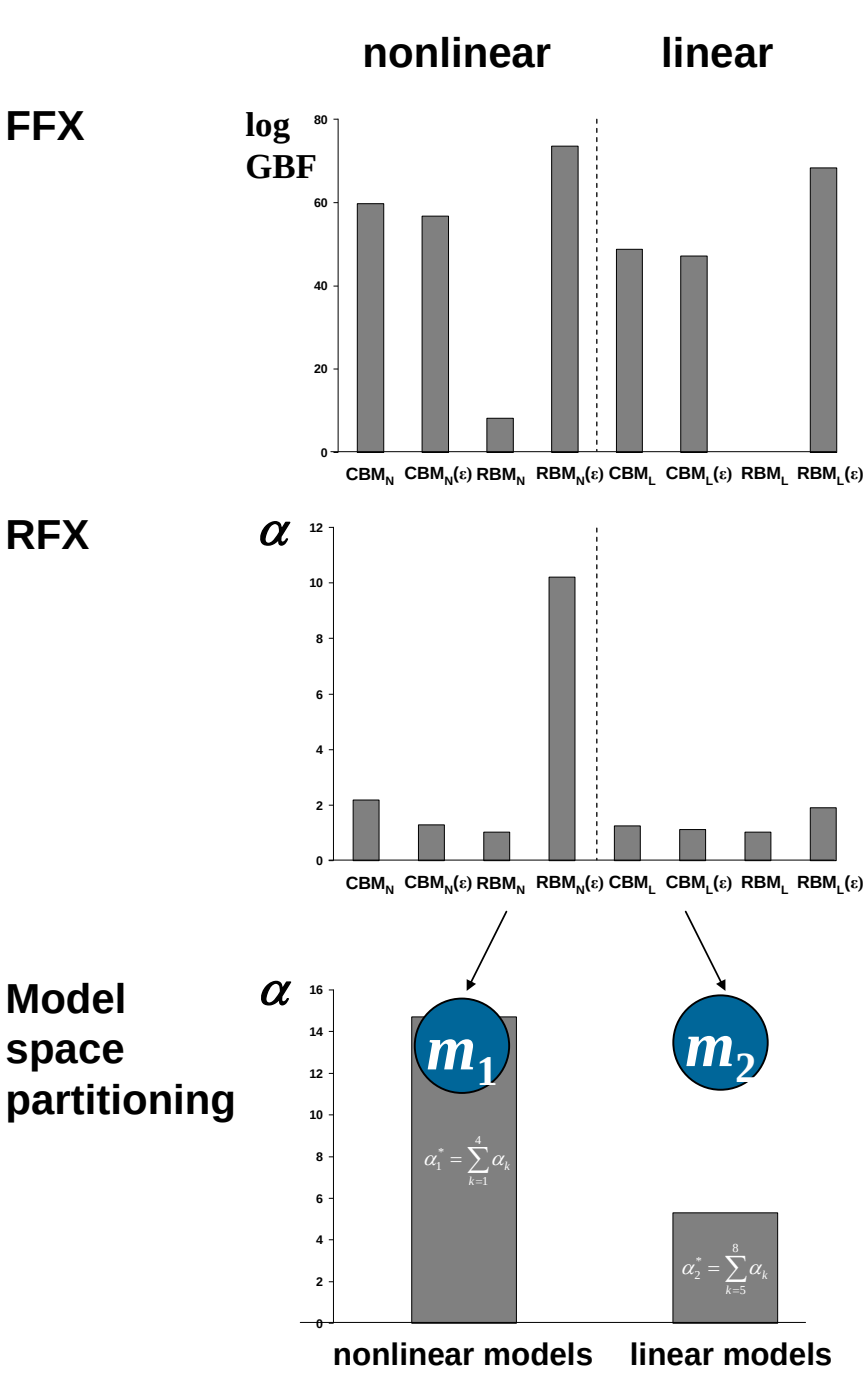

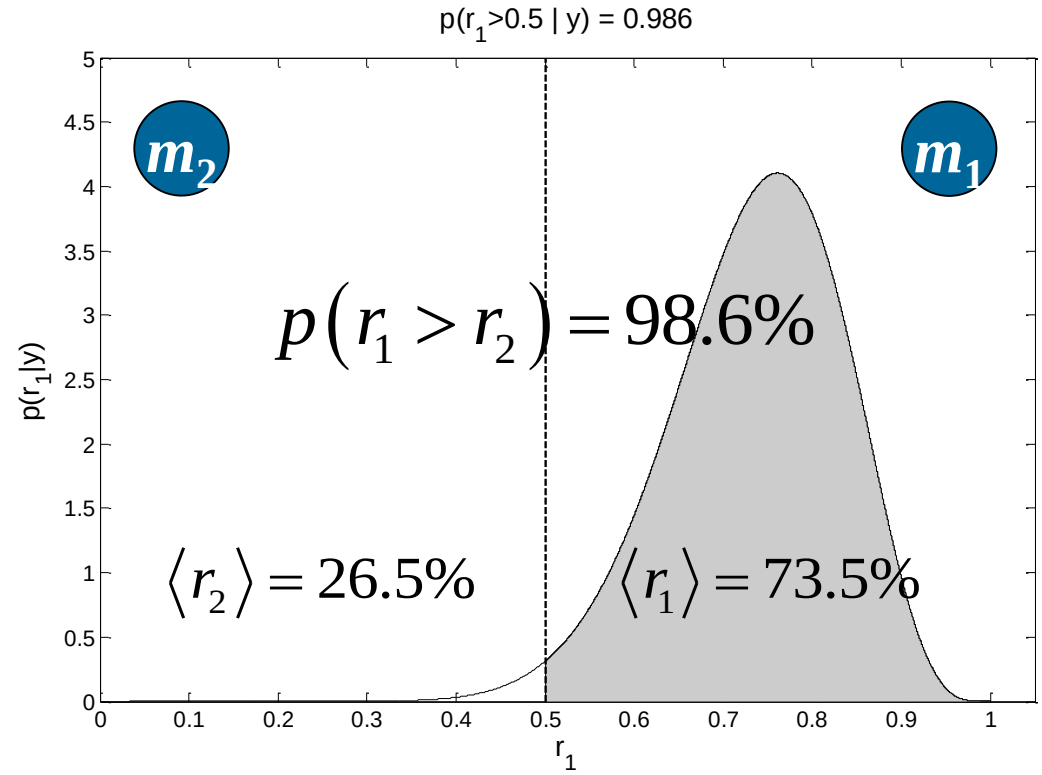

Stephan et al. 2009, *NeuroImage*

Mismatch negativity (MMN)

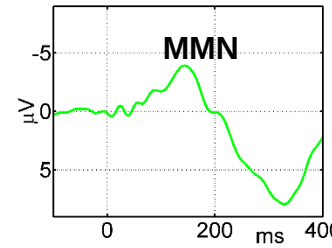

- elicited by surprising stimuli (scales with unpredictability)
- ↓ in schizophrenic patients
- classical interpretations:
  - pre-attentive change detection
  - neuronal adaptation
- current theories:
  - reflection of (hierarchical) Bayesian inference

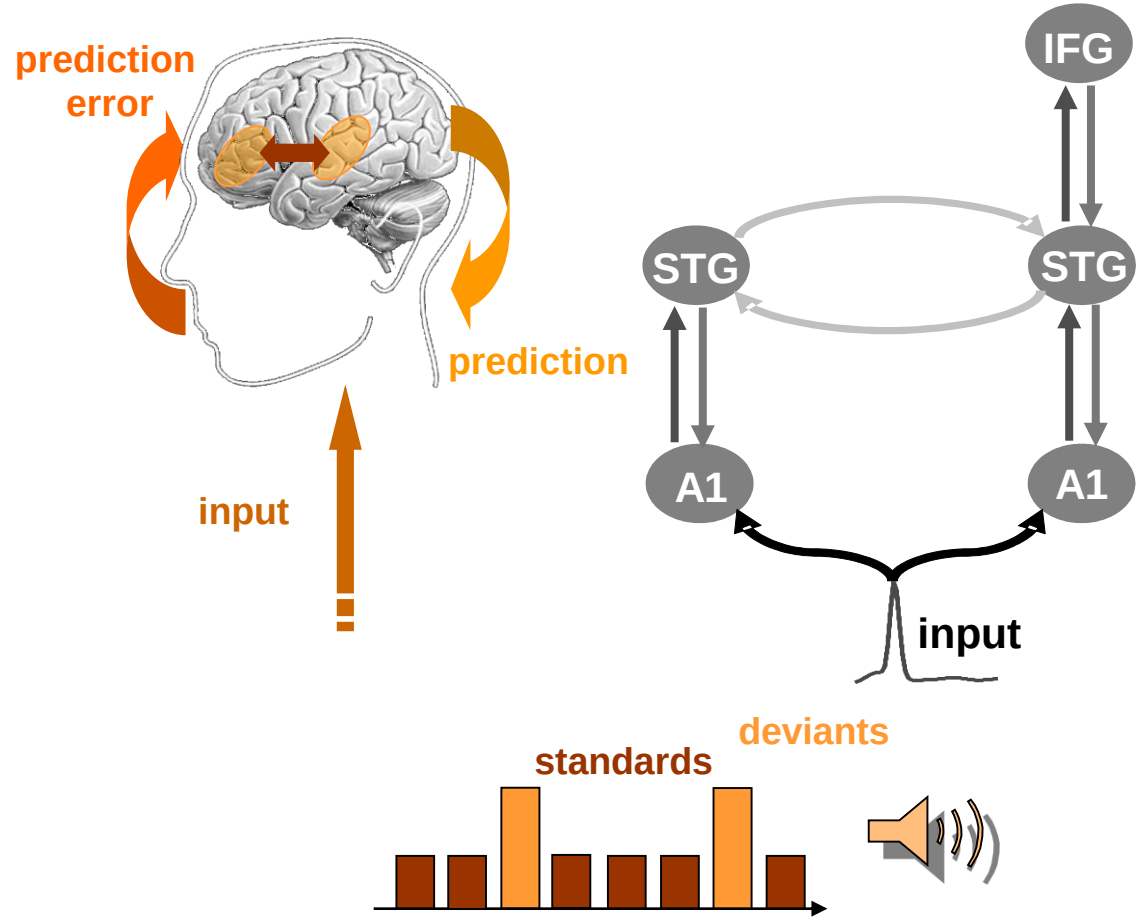

Garrido et al. 2009, *Clin. Neurophysiol.*

Modelling Trial-by-Trial Changes of the Mismatch Negativity (MMN)

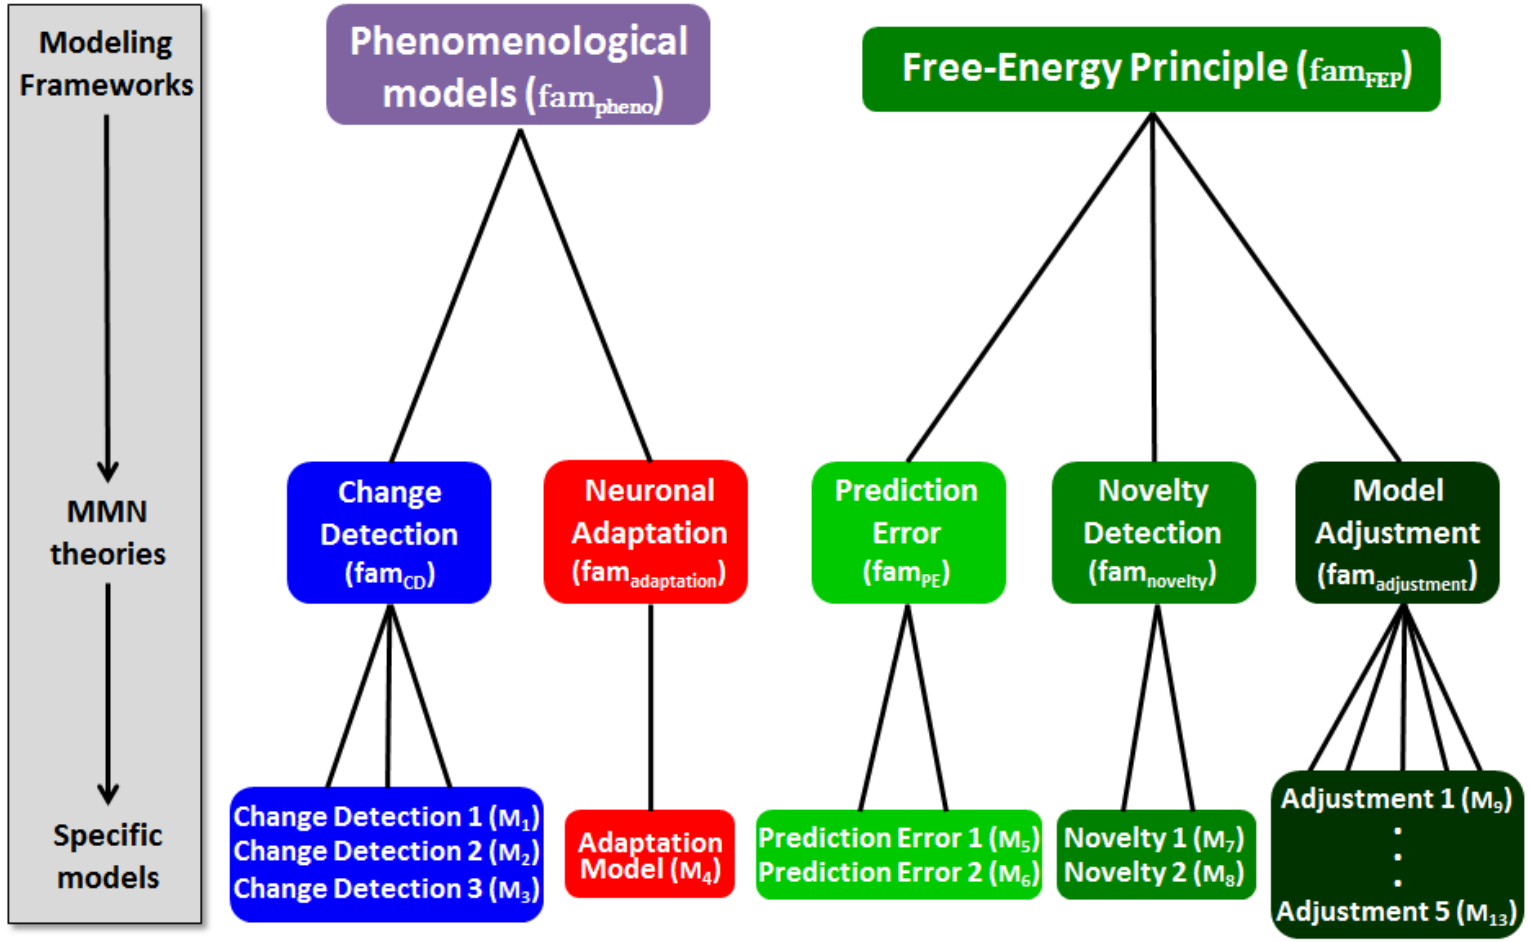

Lieder et al. 2013, *PLoS Comput. Biol.*

MMN model comparison at multiple levels

❷ **Comparing MMN theories**

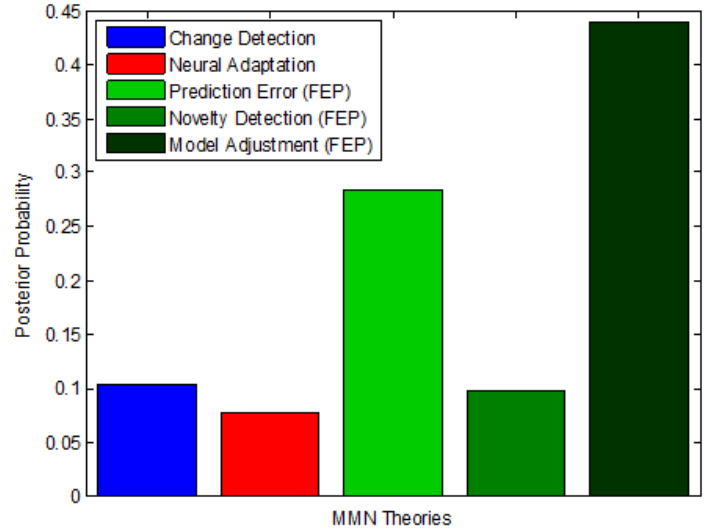

❶ **Comparing individual models**

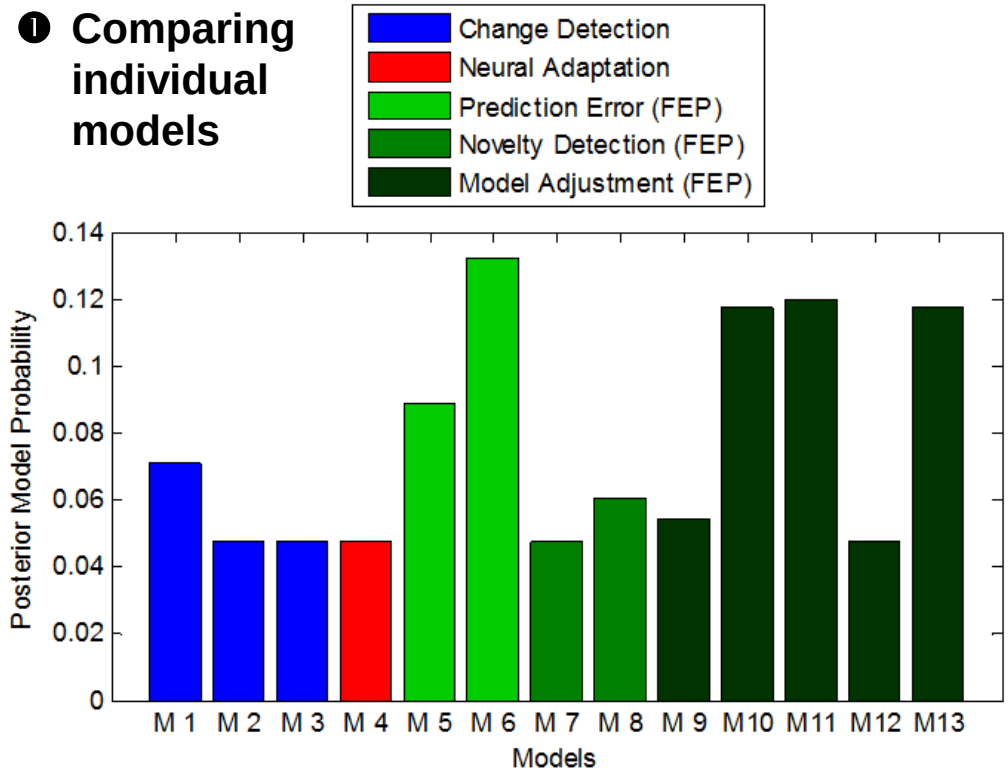

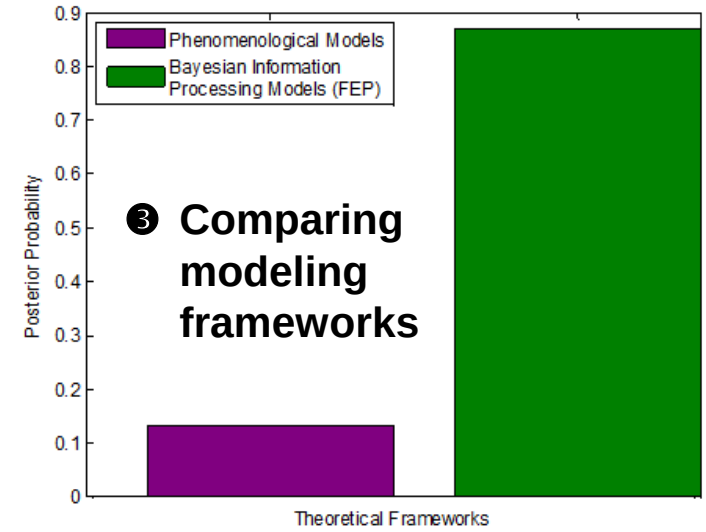

❸ **Comparing modeling frameworks**

Lieder et al. 2013, *PLoS Comput. Biol.*

Bayesian Model Averaging (BMA)

- abandons dependence of parameter inference on a single model and takes into account model uncertainty
- uses the entire model space considered (or an optimal family of models)
- averages parameter estimates, weighted by posterior model probabilities
- represents a particularly useful alternative
  - when none of the models (or model subspaces) considered clearly outperforms all others
  - when comparing groups for which the optimal model differs

single-subject BMA:

$$p(\theta \mid y)$$
$$= \sum_m p(\theta \mid y, m) p(m \mid y)$$

group-level BMA:

$$p(\theta_n \mid y_{1..N})$$
$$= \sum_m p(\theta_n \mid y_n, m) p(m \mid y_{1..N})$$

NB: $p(m|y_{1..N})$ can be obtained by either FFX or RFX BMS

Penny et al. 2010, *PLoS Comput. Biol.*

Prefrontal-parietal connectivity during working memory in schizophrenia

- 17 at-risk mental state (ARMS) individuals
- 21 first-episode patients (13 non-treated)
- 20 controls

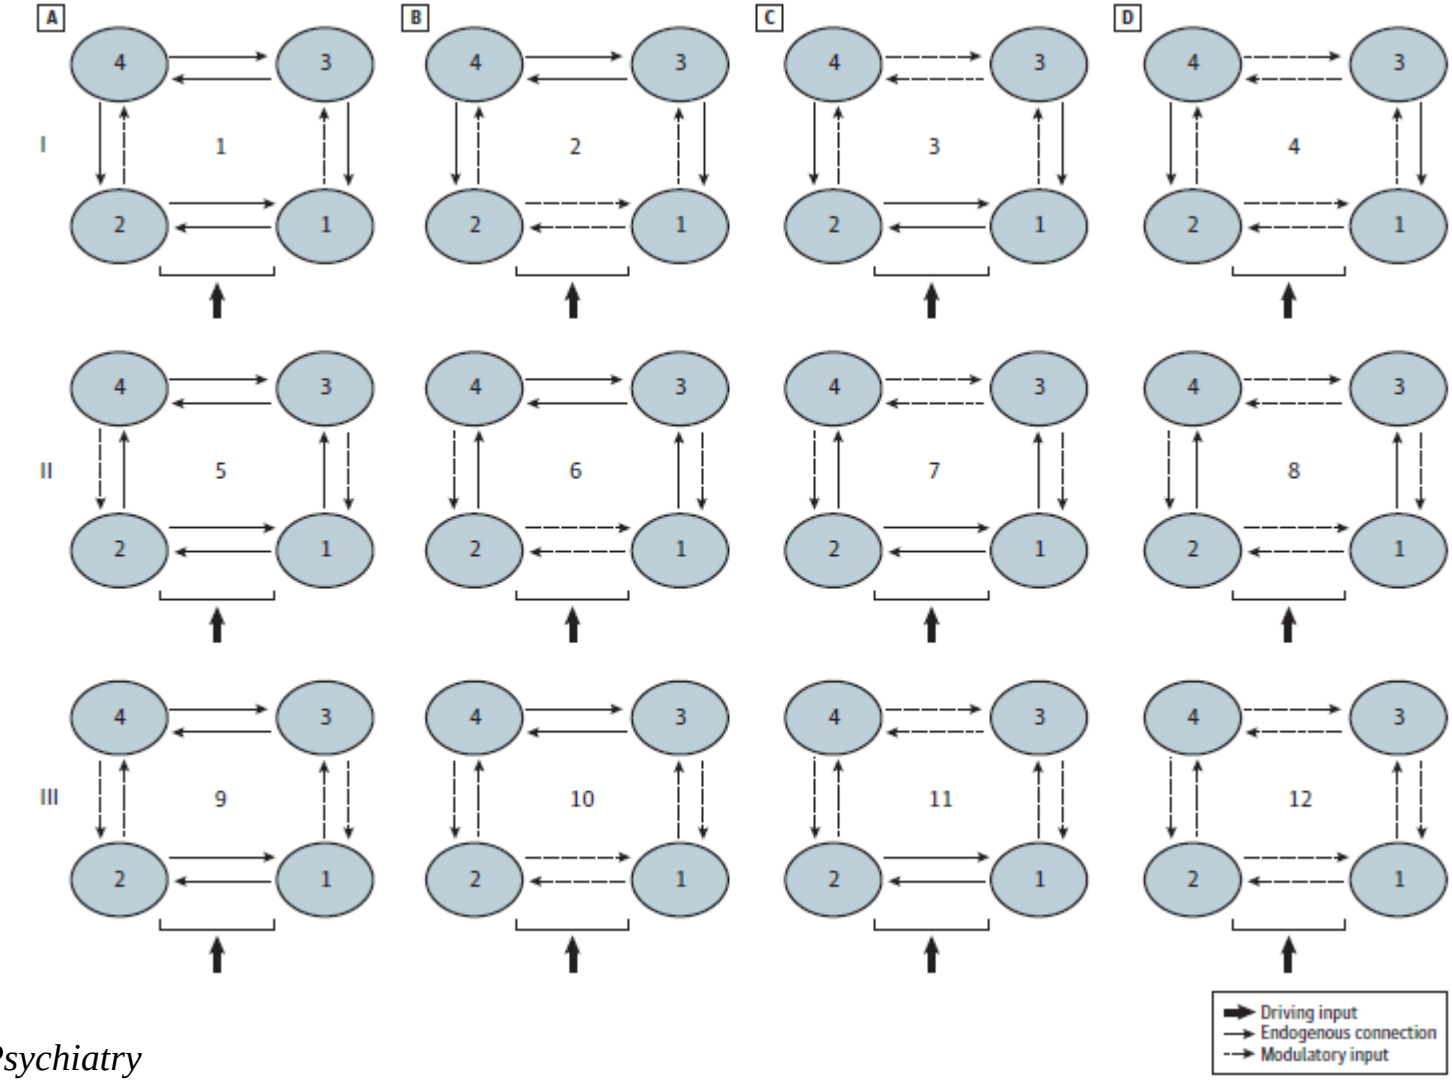

Schmidt et al. 2013, *JAMA Psychiatry*

BMS results for all groups

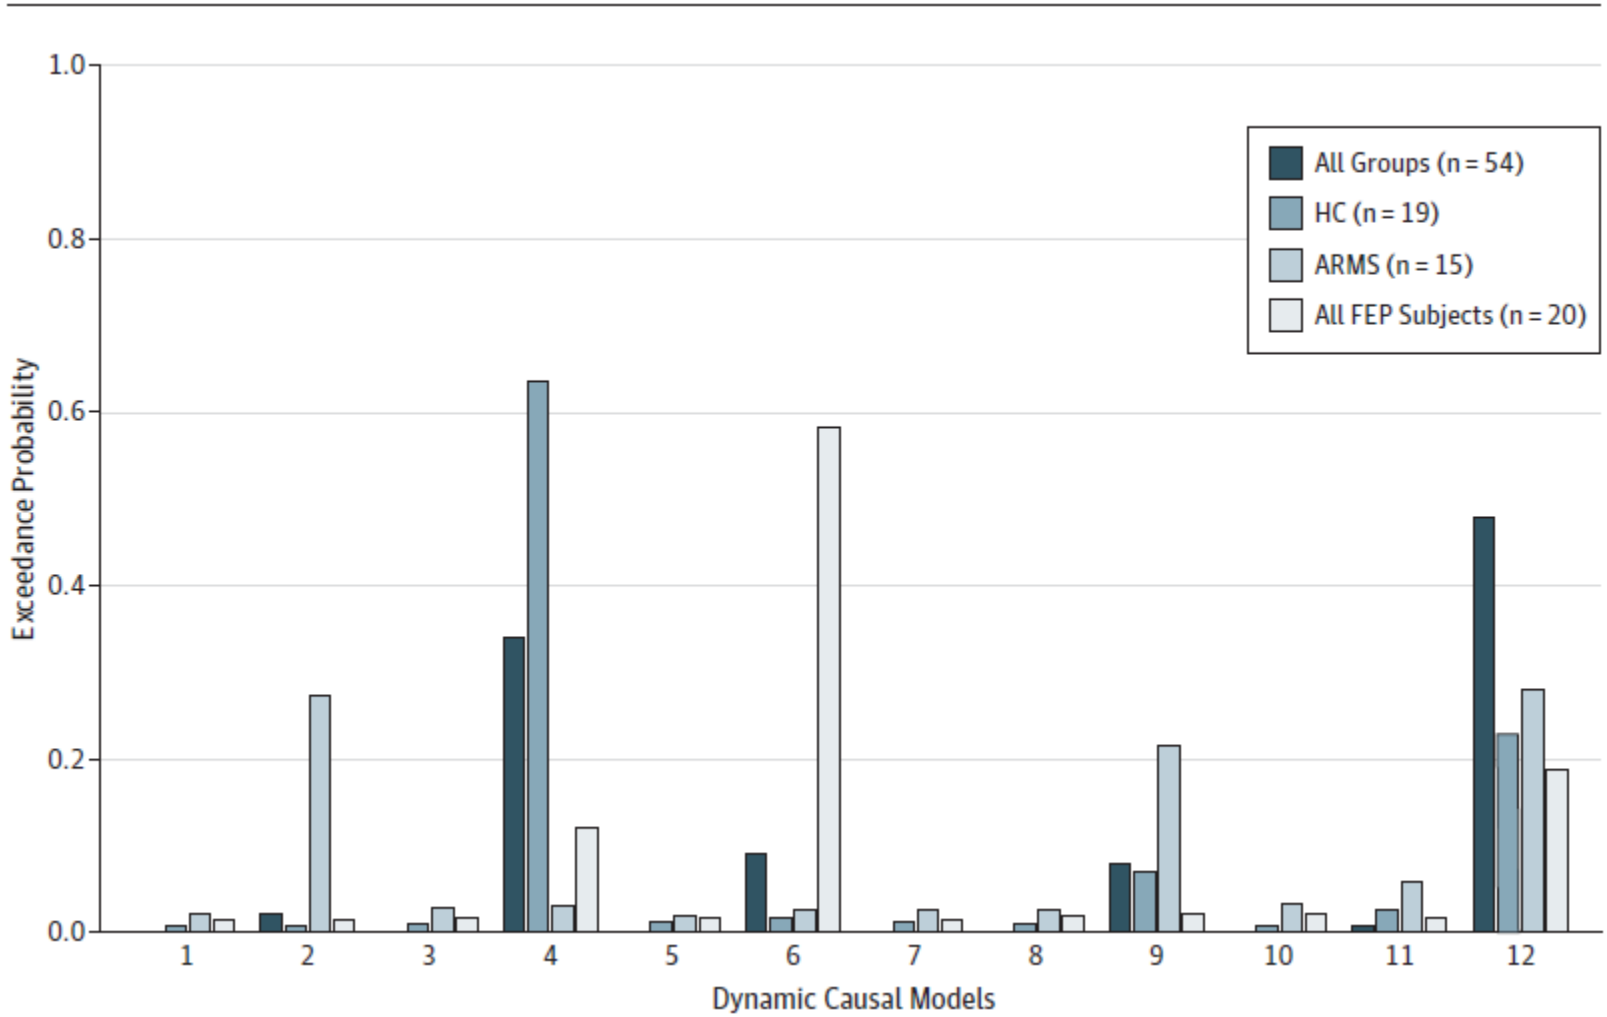

Schmidt et al. 2013, *JAMA Psychiatry*

BMA results: PFC → PPC connectivity

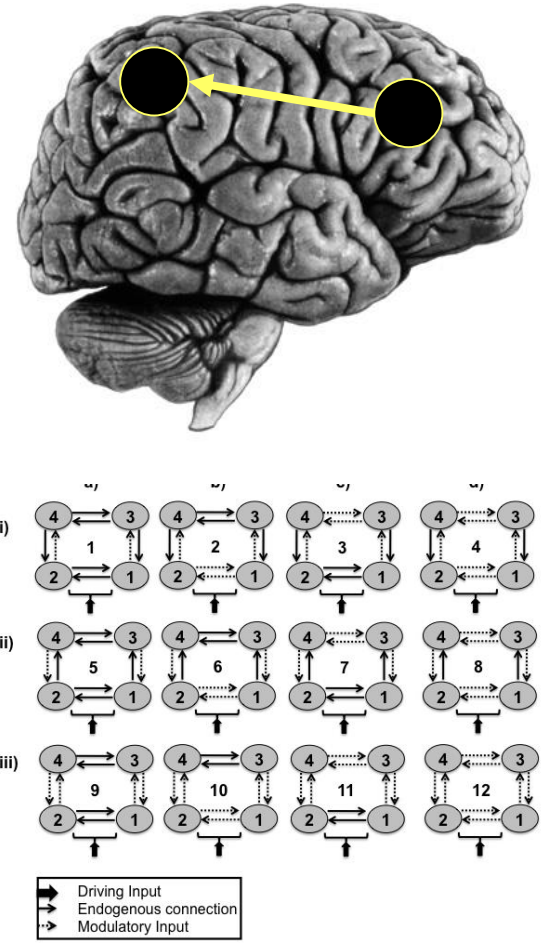

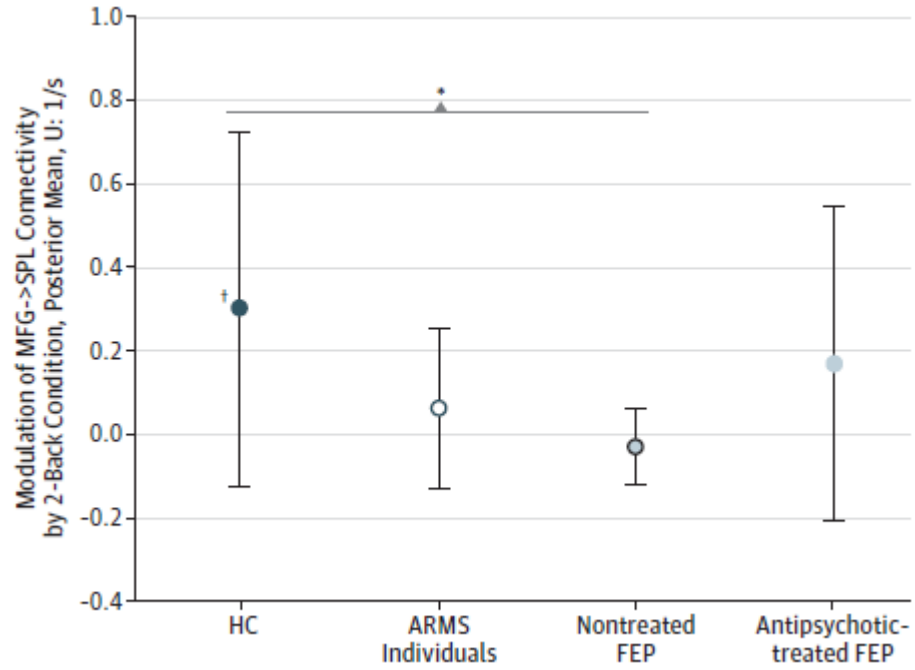

17 ARMS, 21 first-episode (13 non-treated), 20 controls

Schmidt et al. 2013, *JAMA Psychiatry*

Protected exceedance probability: Using BMA to protect against chance findings

- EPs express our confidence that the posterior probabilities of models are different – under the hypothesis H_1 that models differ in probability: $r_k \neq 1/K$
- does not account for possibility "null hypothesis" H_0: $r_k = 1/K$
- **Bayesian omnibus risk (BOR)** of wrongly accepting H_1 over H_0:

$$P_o = \frac{1}{1 + \frac{p(m|H_1)}{p(m|H_0).}}$$

- **protected EP**: Bayesian model averaging over H_0 and H_1:

$$\begin{aligned}\tilde{\varphi}_k &= P(r_k \geq r_{k' \neq k}|y) \\ &= P(r_k \geq r_{k' \neq k}|y, H_1)P(H_1|y) + P(r_k \geq r_{k' \neq k}|y, H_0)P(H_0|y) \\ &= \varphi_k(1 - P_0) + \frac{1}{K}P_0\end{aligned}$$

Rigoux et al. 2014, *NeuroImage*

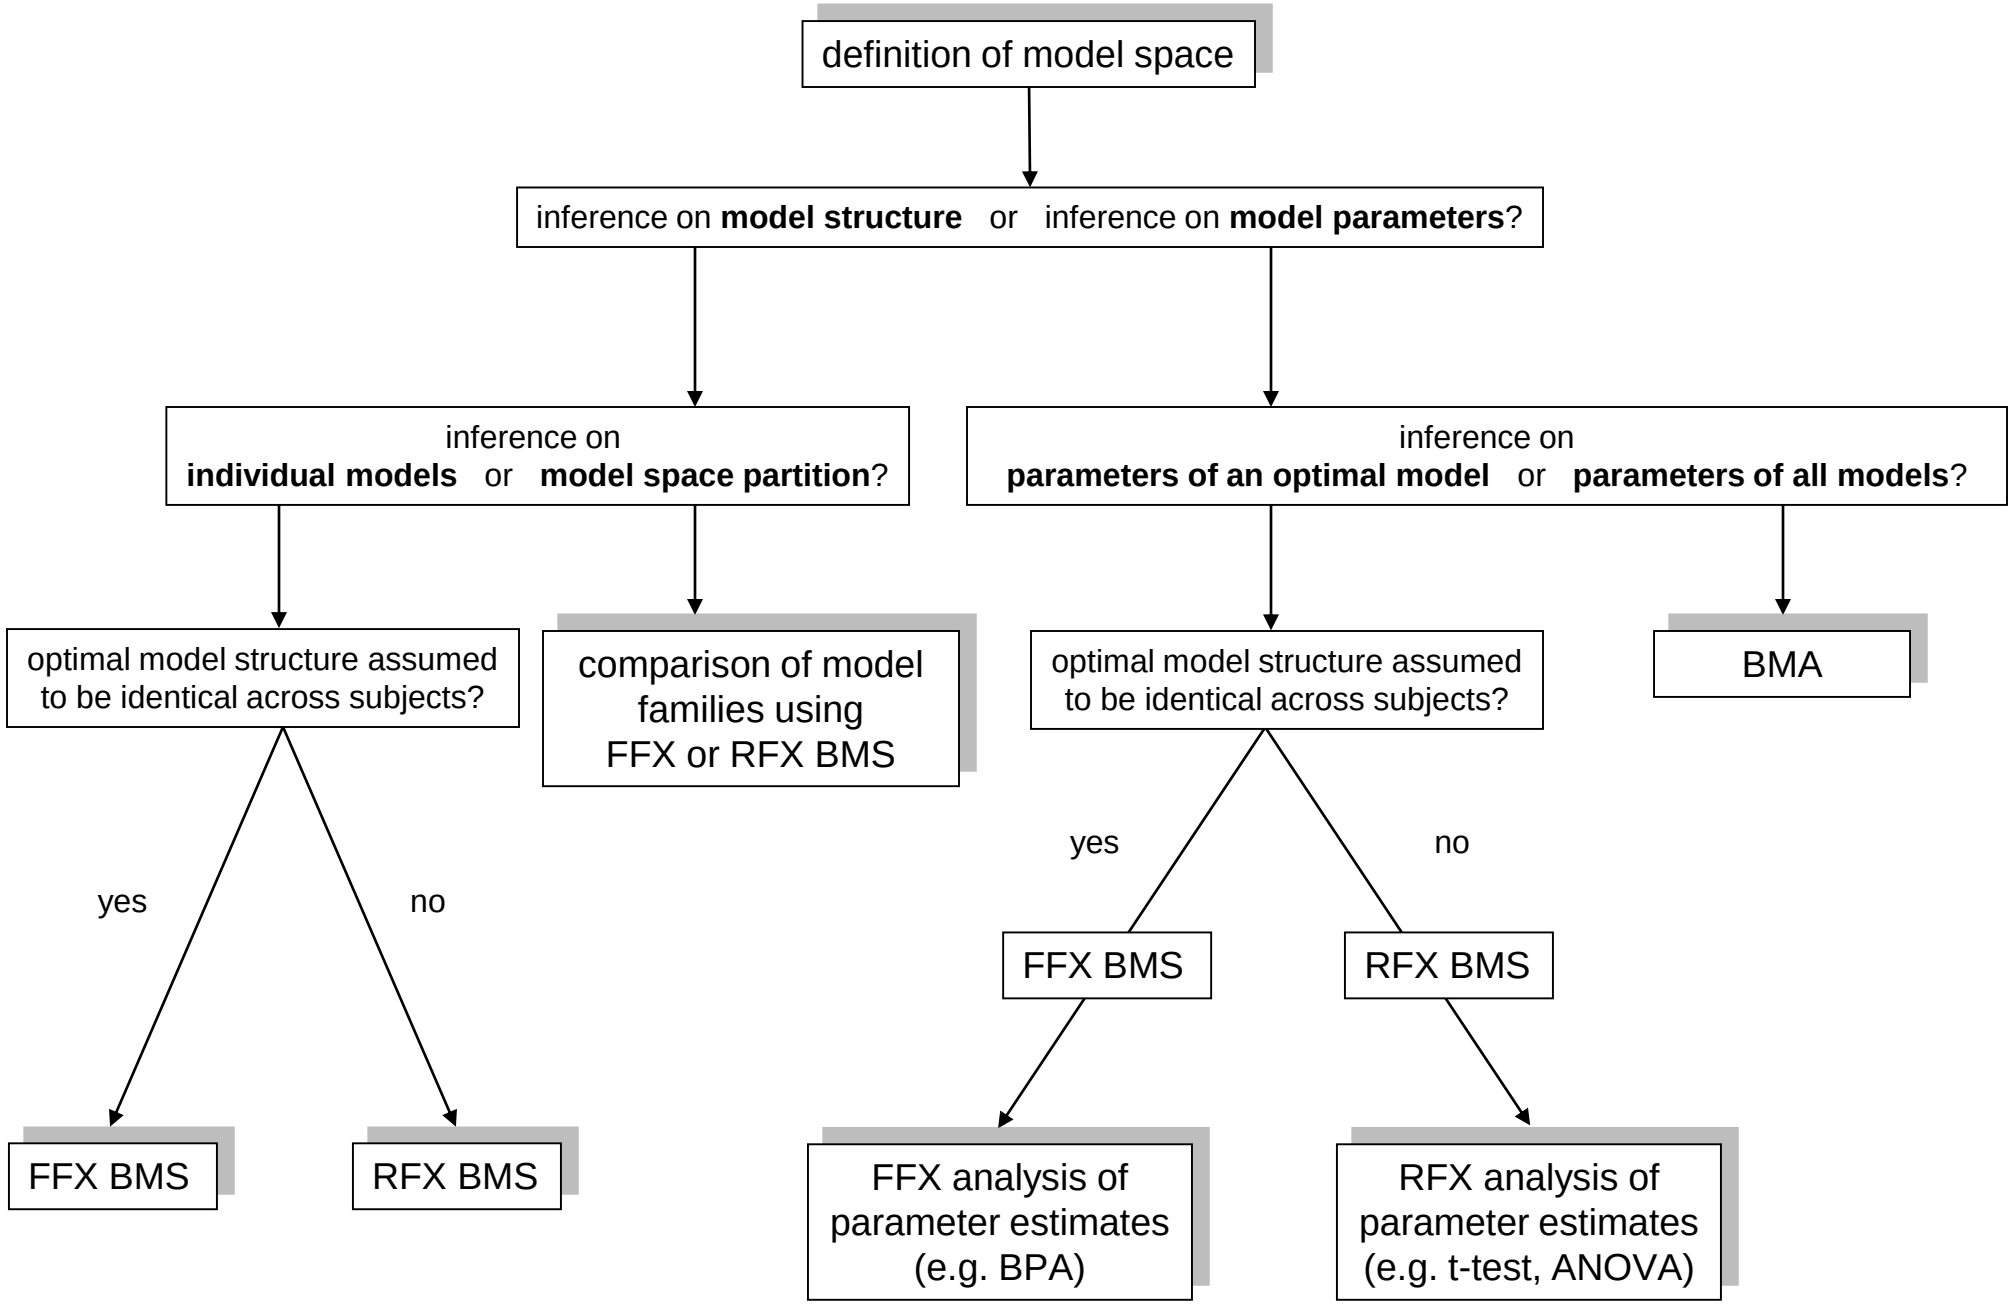

Stephan et al. 2010, *NeuroImage*

Further reading

- Penny WD, Stephan KE, Mechelli A, Friston KJ (2004) Comparing dynamic causal models. NeuroImage 22:1157-1172.
- Penny WD, Stephan KE, Daunizeau J, Joao M, Friston K, Schofield T, Leff AP (2010) Comparing Families of Dynamic Causal Models. PLoS Computational Biology 6: e1000709.
- Penny WD (2012) Comparing dynamic causal models using AIC, BIC and free energy. Neuroimage 59: 319-330.
- Rigoux L, Stephan KE, Friston KJ, Daunizeau J (2014) Bayesian model selection for group studies – revisited. NeuroImage 84: 971-985.
- Stephan KE, Weiskopf N, Drysdale PM, Robinson PA, Friston KJ (2007) Comparing hemodynamic models with DCM. NeuroImage 38:387-401.
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. NeuroImage 46:1004-1017.
- Stephan KE, Penny WD, Moran RJ, den Ouden HEM, Daunizeau J, Friston KJ (2010) Ten simple rules for Dynamic Causal Modelling. NeuroImage 49: 3099-3109.
- Stephan KE, Iglesias S, Heinzle J, Diaconescu AO (2015) Translational Perspectives for Computational Neuroimaging. Neuron 87: 716-732.
- Stephan KE, Schlagenhauf F, Huys QJM, Raman S, Aponte EA, Brodersen KH, Rigoux L, Moran RJ, Daunizeau J, Dolan RJ, Friston KJ, Heinz A (2017) Computational Neuroimaging Strategies for Single Patient Predictions. NeuroImage 145: 180-199.

Thank you