Fitting behavioural data with RL models

Quentin Huys & Michael Frank

UCL & Brown University

Janelia Farm, March 6-9th 2011

Fitting models: matching and noise

- probabilistic policy, e.g. softmax

$$p(a|s) = \frac{e^{\beta Q(s,a)}}{\sum_{a'} e^{\beta Q(s,a')}}$$

- total likelihood

$$\mathcal{L}(\theta) = p(\{a_t\}_{t=1}^{T} | \{s_t\}_{t=1}^{T}, \{r_t\}_{t=1}^{T}, \theta) = \prod_{t=1}^{T} p(a_t | s_t, r_{1\dots t-1}, \theta)$$

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \, \mathcal{L}(\theta)$$

Transforming variables

$$\begin{aligned}\beta &= e^{\beta'} \\ &\Rightarrow \quad \beta' = \log(\beta) \\ \epsilon &= \frac{1}{1+e^{-\epsilon'}} \\ &\Rightarrow \quad \epsilon' = \log\left(\frac{\epsilon}{1-\epsilon}\right)\end{aligned}$$

$$\frac{d \log \mathcal{L}(\theta')}{d\theta'}$$

ML can be noisy

$$\mathcal{L}(\beta = 10) \approx \mathcal{L}(\beta = 100)$$

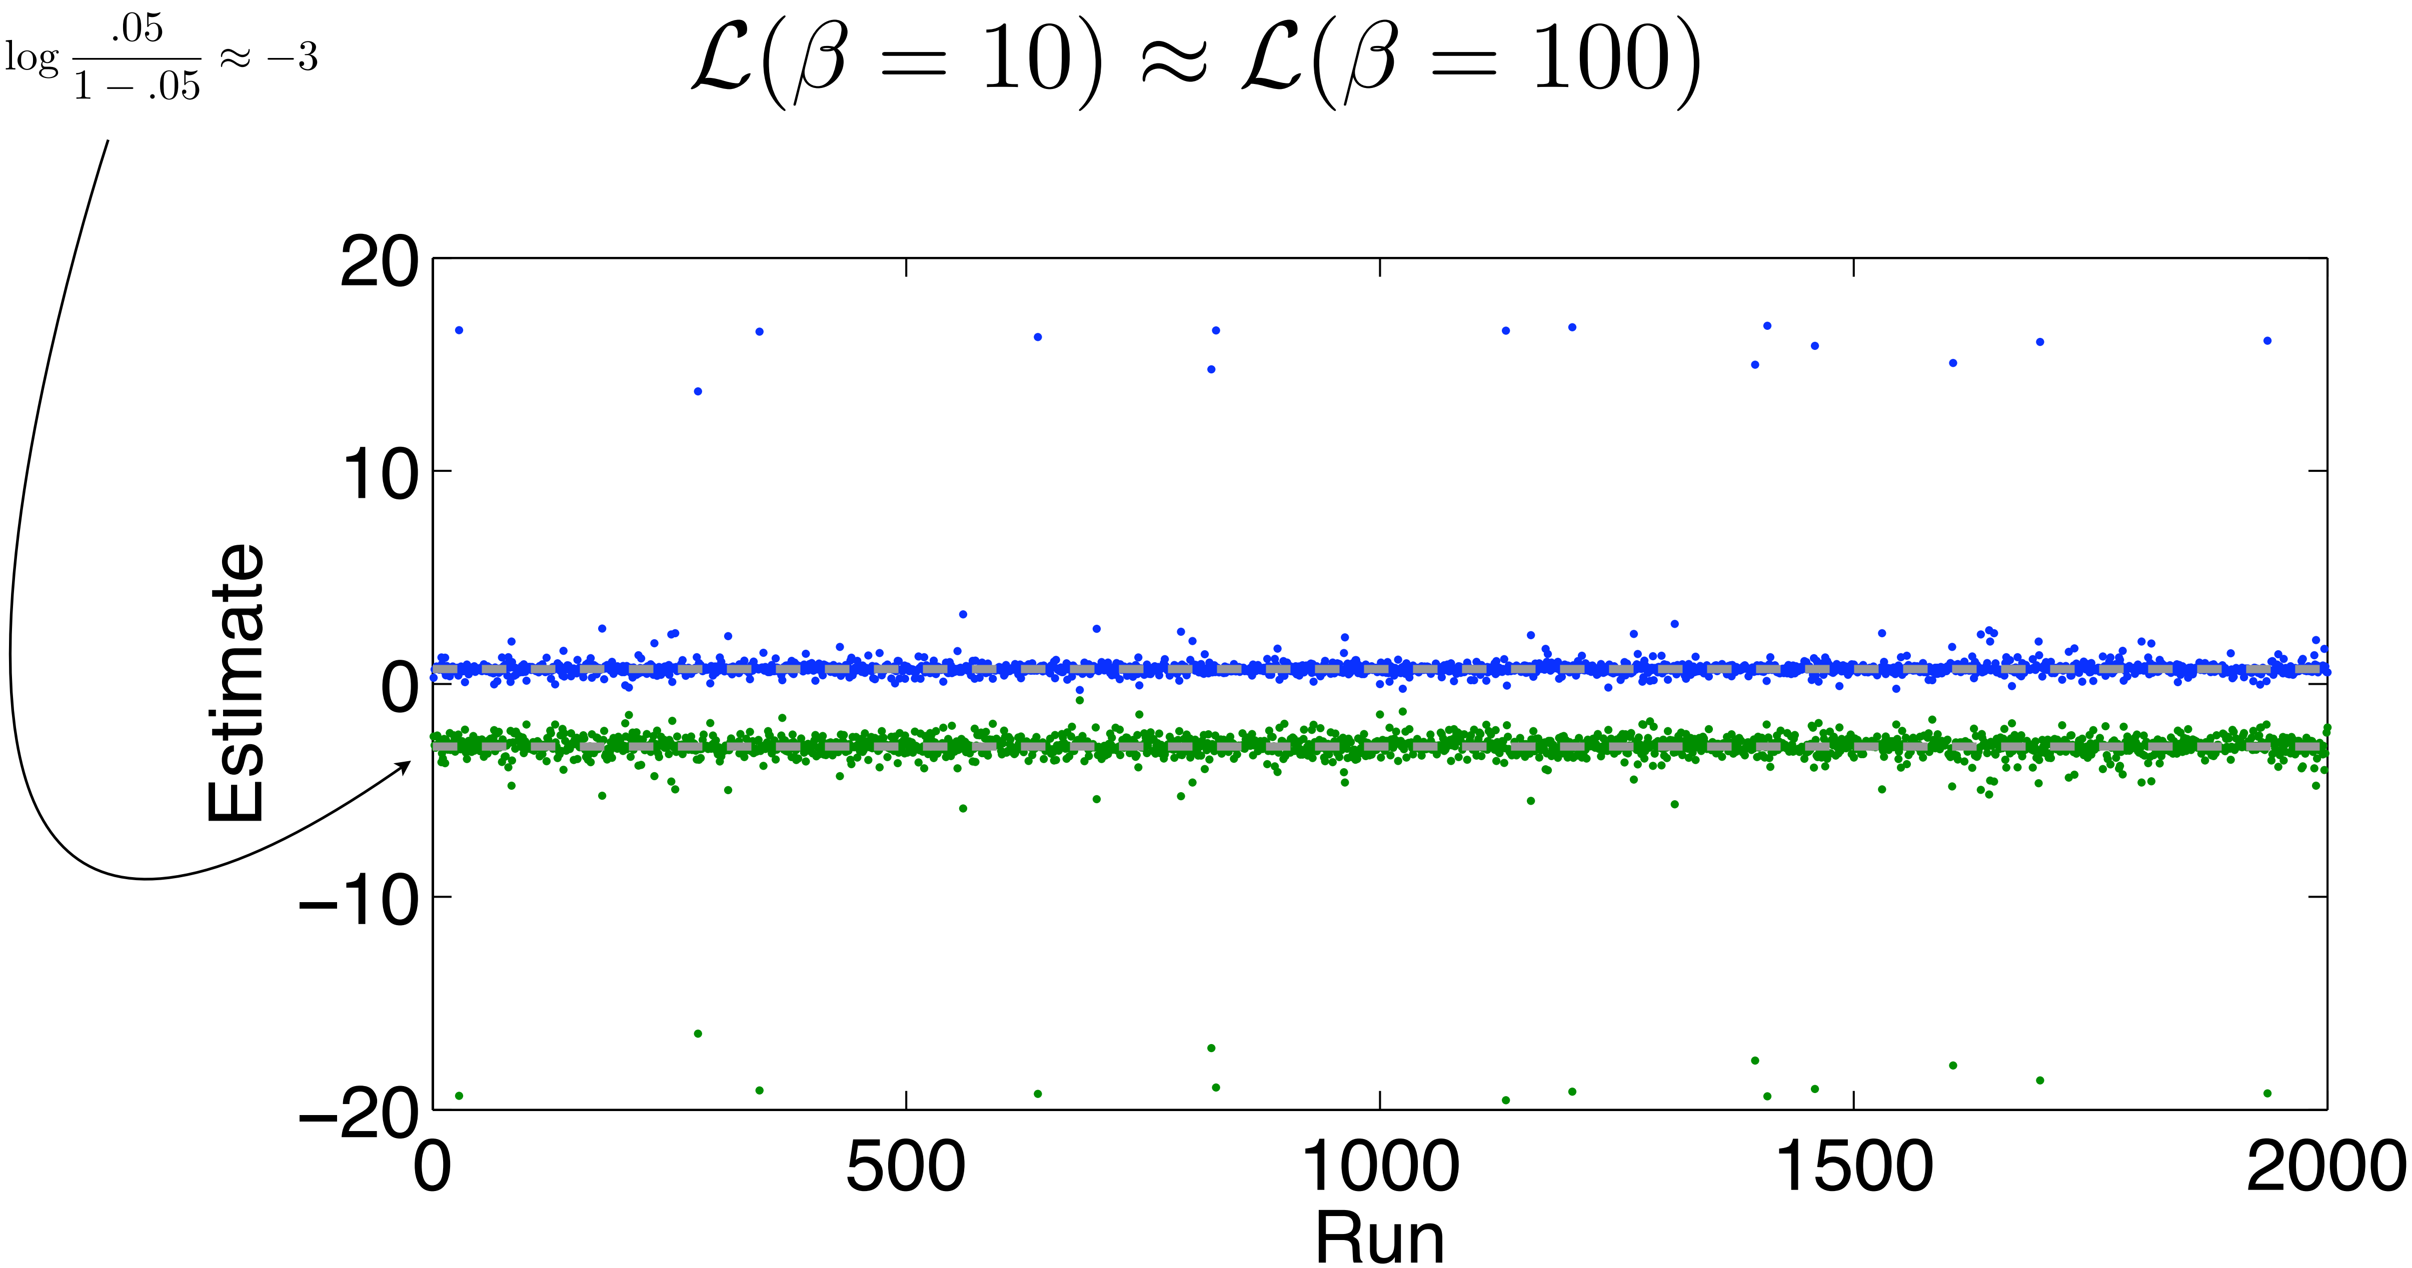

200 trials, 1 stimulus, 10 actions, learning rate = .05, beta=2

Constraining ML

Not so smooth

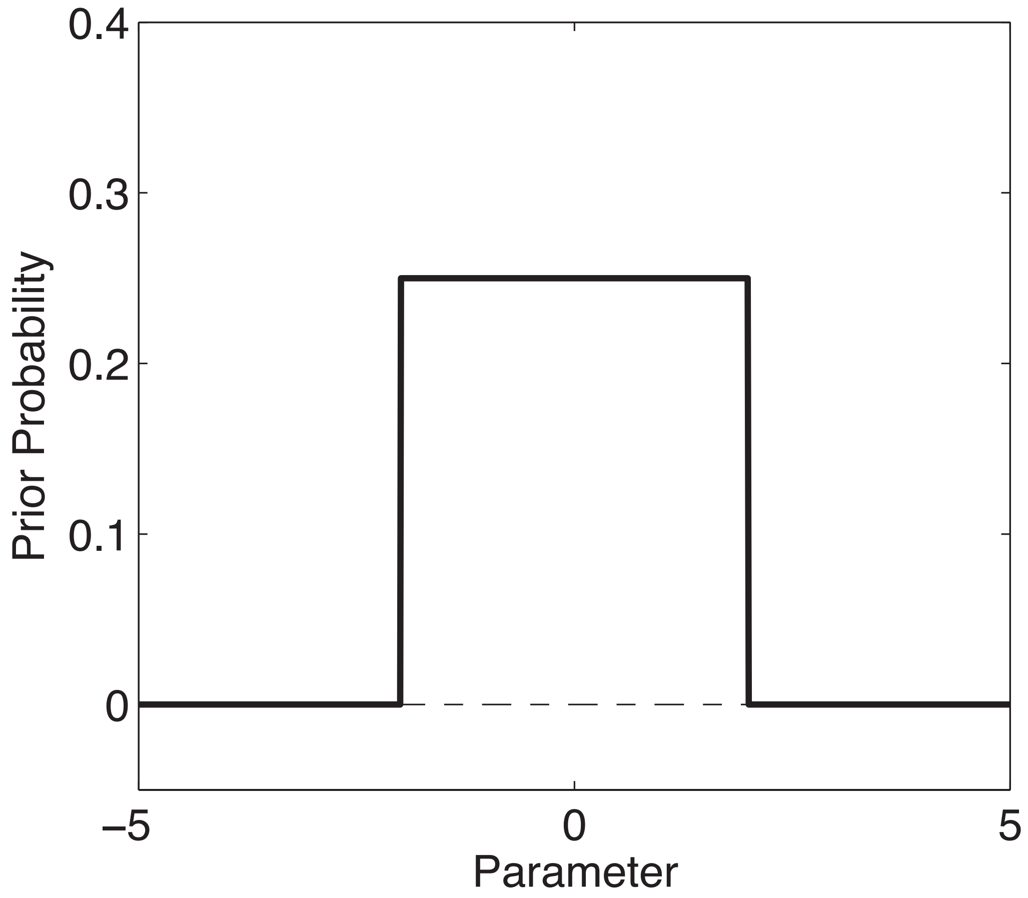

Smooth

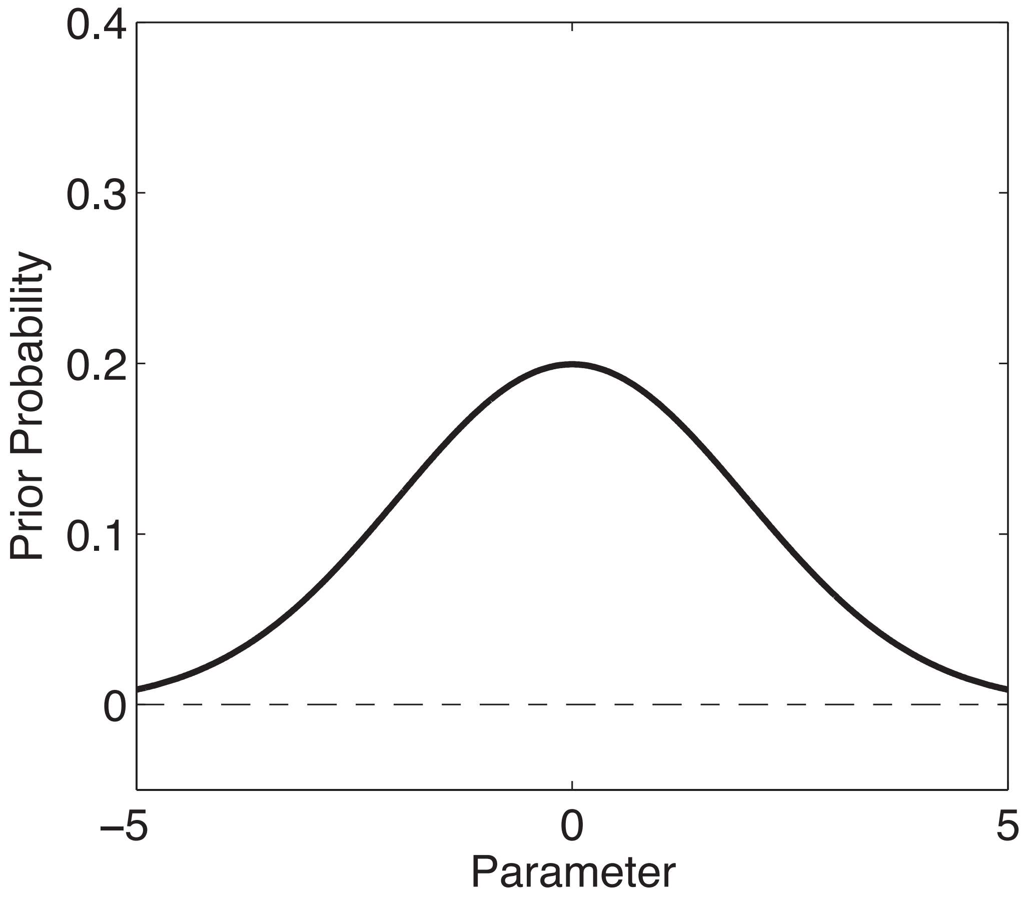

Constraining ML

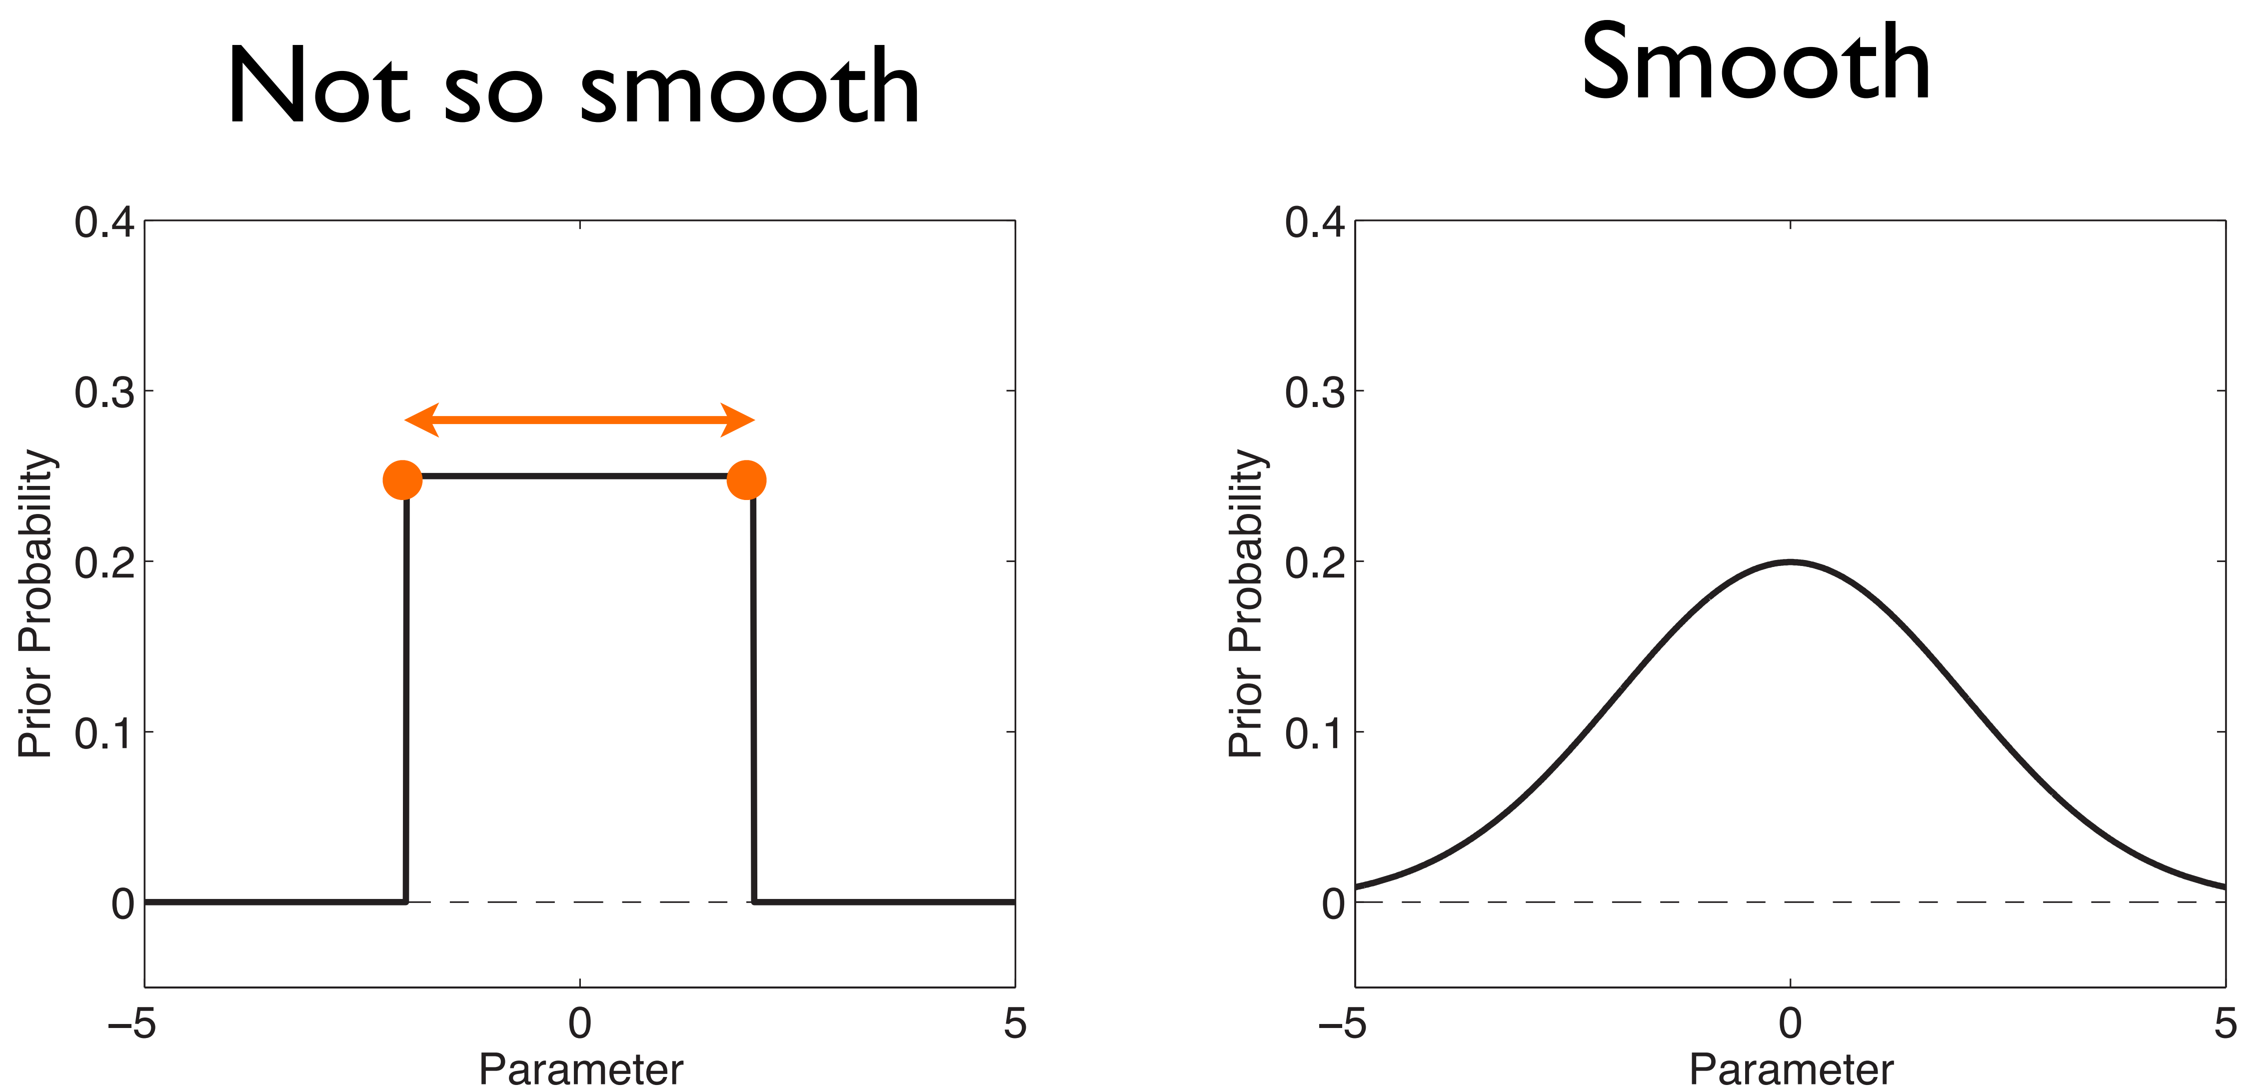

Maximum a posteriori estimate

$$\mathcal{P}(\theta) = p(\theta|a_{1\ldots T}) = \frac{p(a_{1\ldots T}|\theta)p(\theta)}{\int d\theta p(\theta|a_{1\ldots T})p(\theta)}$$

$$\log \mathcal{P}(\theta) = \sum_{t=1}^{T} \log p(a_t|\theta) + \log p(\theta) + const.$$

$$\frac{\log \mathcal{P}(\theta)}{d\alpha} = \frac{\log \mathcal{L}(\theta)}{d\alpha} + \frac{d\,p(\theta)}{d\theta}$$

Maximum a posteriori estimate

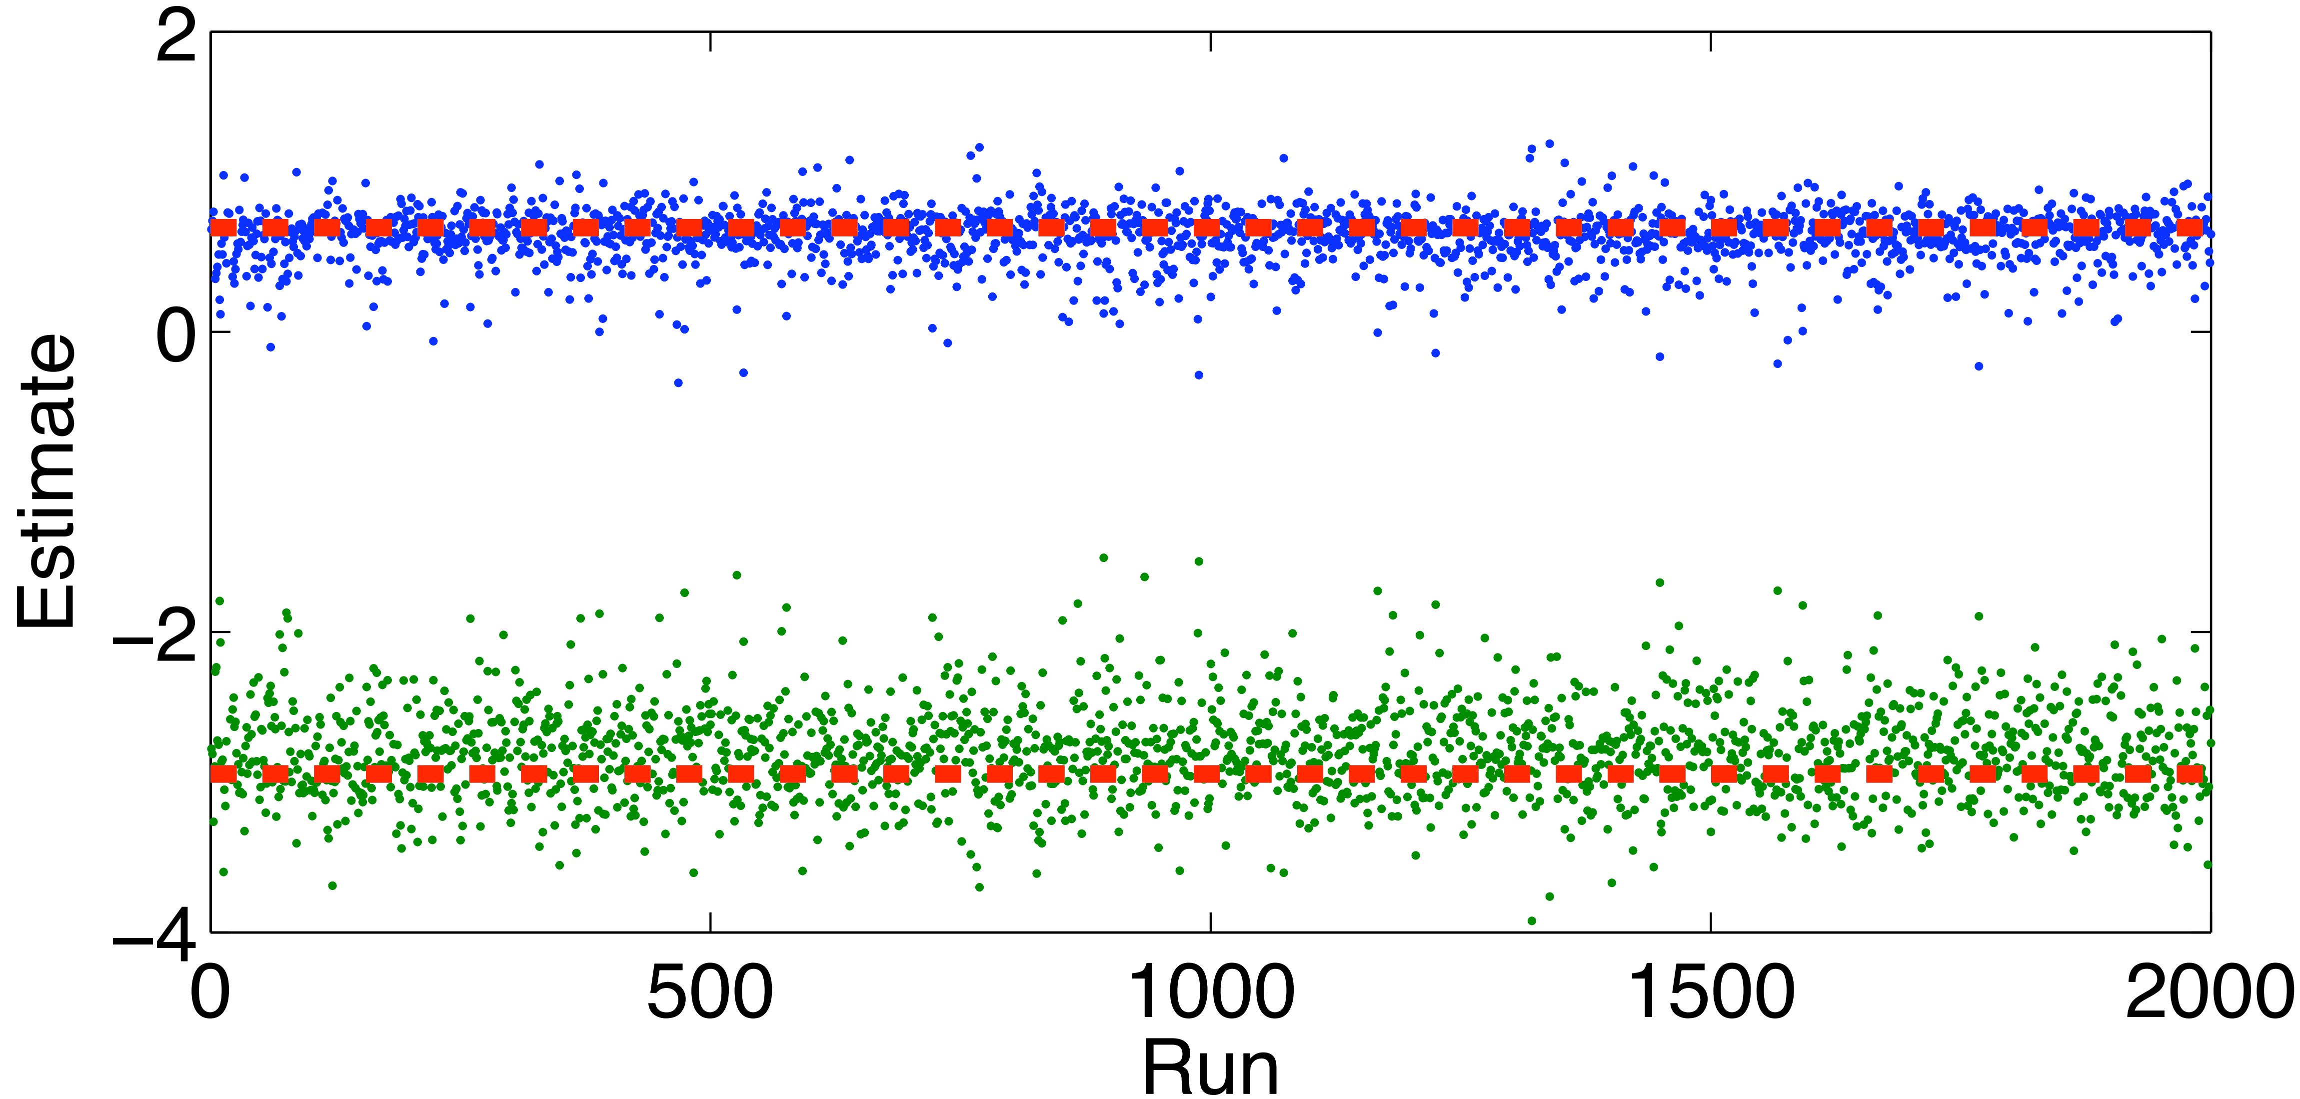

200 trials, 1 stimulus, 10 actions, learning rate = .05, beta=2

m_{beta}=0, m_{eps}=-3, n=1

But

What prior parameters should I use?

Estimating the hyperparameters

- What should the hyperparameters be?

$$\log \mathcal{P}(\theta) = \mathcal{L}(\theta) + \log \underbrace{p(\theta)}_{=p(\theta|\zeta)} + const.$$

- Empirical Bayes: set them to ML estimate

$$\hat{\zeta} \quad = \quad \underset{\zeta}{\operatorname{argmax}}\, p(\mathcal{A}|\zeta)$$

- where we use all the actions by all the *k* subjects

$$\mathcal{A} = \{a^k_{1\ldots T}\}_{k=1}^{K}$$

ML estimate of top-level parameters

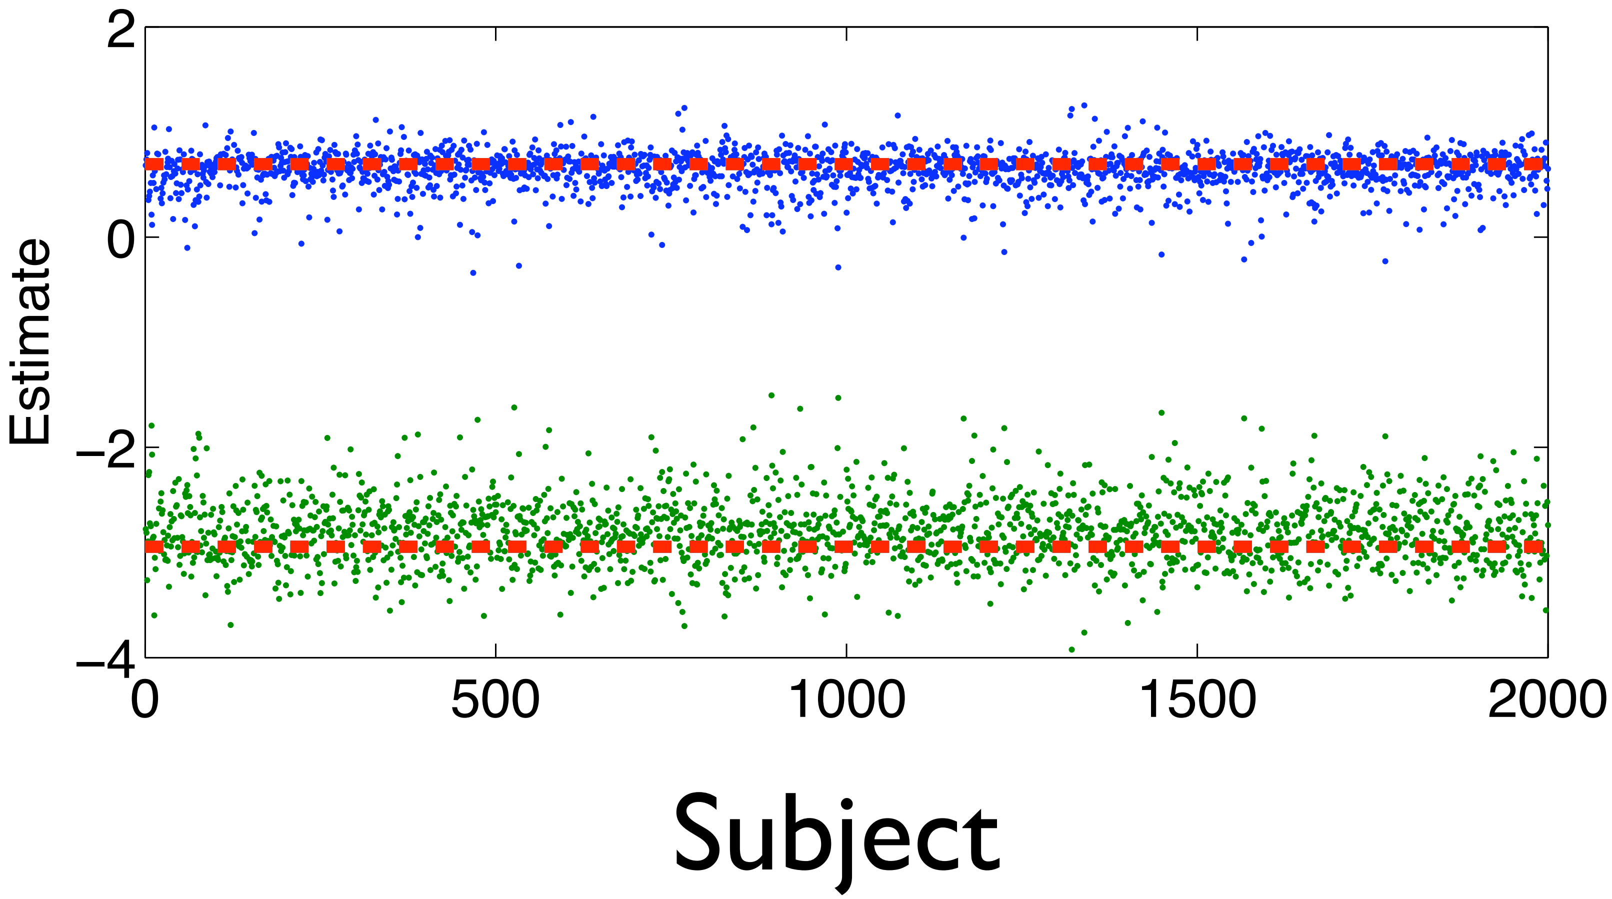

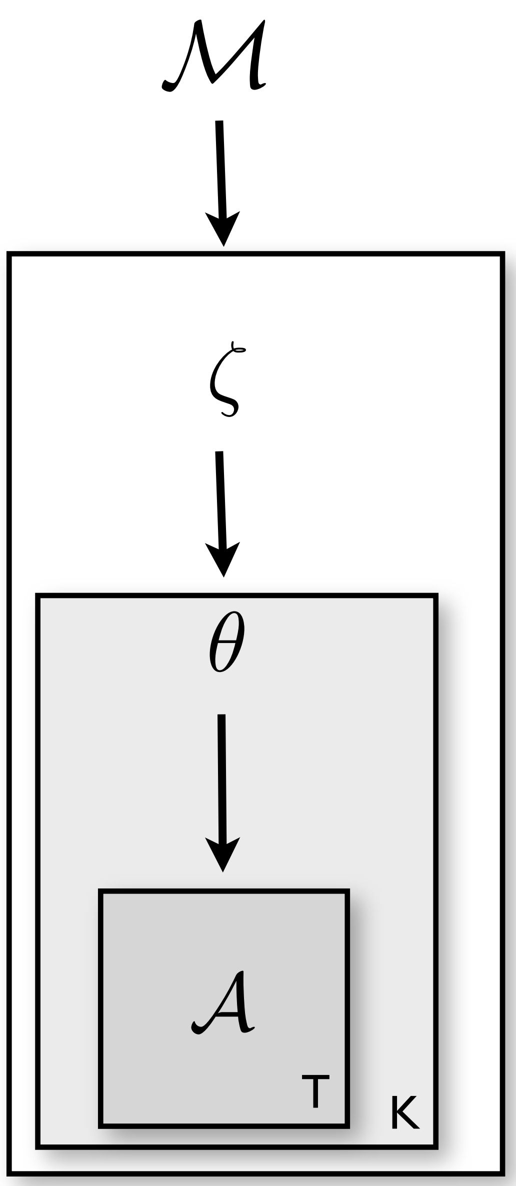

$$\hat{\zeta} = \underset{\zeta}{\operatorname{argmax}}\, p(\mathcal{A}|\zeta)$$

Estimating the hyperparameters

- Can't just do gradient ascent

$$\frac{d}{d\zeta}p(\mathcal{A}|\zeta)$$

- Contains integral over individual parameters:

$$p(\mathcal{A}|\zeta) = \int d\theta p(\mathcal{A}|\theta)\, p(\theta|\zeta)$$

- So we need to:

$$\begin{aligned} \hat{\zeta} &= \operatorname*{argmax}_{\zeta} p(\mathcal{A}|\zeta) \\ &= \operatorname*{argmax}_{\zeta} \int d\,\theta p(\mathcal{A}|\theta)\, p(\theta|\zeta) \end{aligned}$$

Expectation Maximisation

- Iterate between
  - Estimating MAP parameters given prior parameters
  - Estimating prior parameters from MAP parameters

- There are other approaches
  - MCMC
  - Analytical conjugate priors
  - ...

EM with Laplace approximation

- First infer each subject's parameter and the certainty around them

$$
\begin{aligned}
\text{E step:} \quad q_k(\theta) &= \mathcal{N}(\mathbf{m}_k, \mathbf{S}_k) \\
\mathbf{m}_k &\leftarrow \underset{\theta}{\operatorname{argmax}}\, p(\mathbf{a}_k|\theta)p(\theta|\zeta^{(i)}) \\
\mathbf{S}_k^{-1} &\leftarrow \left.\frac{\partial^2 p(\mathbf{a}^k|\theta)p(\theta|\zeta^{(i)})}{\partial\theta^2}\right|_{\theta=\mathbf{m}_k}
\end{aligned}
$$

matlab: `[m,L,,,S]=fminunc(…)`

Just what we had before: MAP inference given some prior parameters

EM with Laplace approximation

- Next update the prior

Prior mean = mean of MAP estimates

$$\text{M step:} \qquad \zeta_\mu^{(i+1)} = \frac{1}{K}\sum_k \mathbf{m}_k$$

$$\zeta_{\nu^2}^{(i+1)} = \frac{1}{N}\sum_i \left[(\mathbf{m}_k)^2 + \mathbf{S}_k\right] - (\zeta_\mu^{(i+1)})^2$$

Prior variance depends on S and variance of MAP estimates

- And now iterate until convergence

Overview

- Empirical prior
  - Infer with approximate EM
- Model comparison
  - Group-level comparison
  - AIC / BIC / Laplacian
  - Error bars on group means
- Parameters
  - Comparisons

Model fit: likelihood

- How well does the model do?
  - choice probabilities:

$$
\begin{aligned}
\mathbb{E}p(correct) &= e^{\mathcal{L}(\hat{\theta})/K/T} \\
&= e^{\log p(\mathcal{A}|\theta)/K/T} \\
&= \left(\prod_{k,t=1}^{K,T} p(a_{k,t}|\theta_k)\right)^{\frac{1}{KT}}
\end{aligned}
$$

"Predictive probabilities"

  - typically around 0.65-0.75 for 2-way choice
  - for 10-armed bandit example
  - pseudo R squared
  - better than chance?

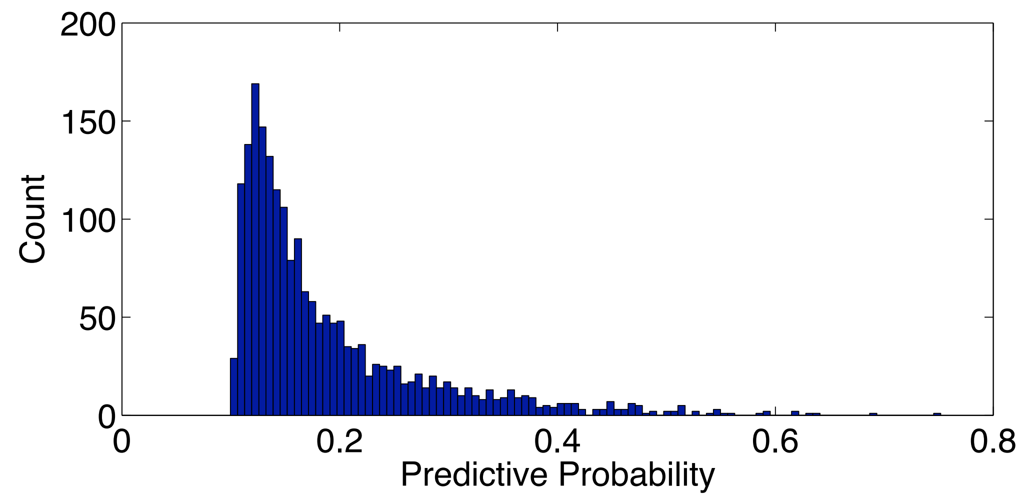

$$
\begin{aligned}
\mathbb{E}[N_k(correct)] &= \mathbb{E}[p_k(correct)]T \\
p_{bin}(\mathbb{E}[N_k(correct)]|N_k d, p_0 = 0.5) &< 1 - \alpha
\end{aligned}
$$

Wednesday, 11 May 2011

Generative test

- Model: probability(actions)
  - simply draw from this distribution, and see what happens

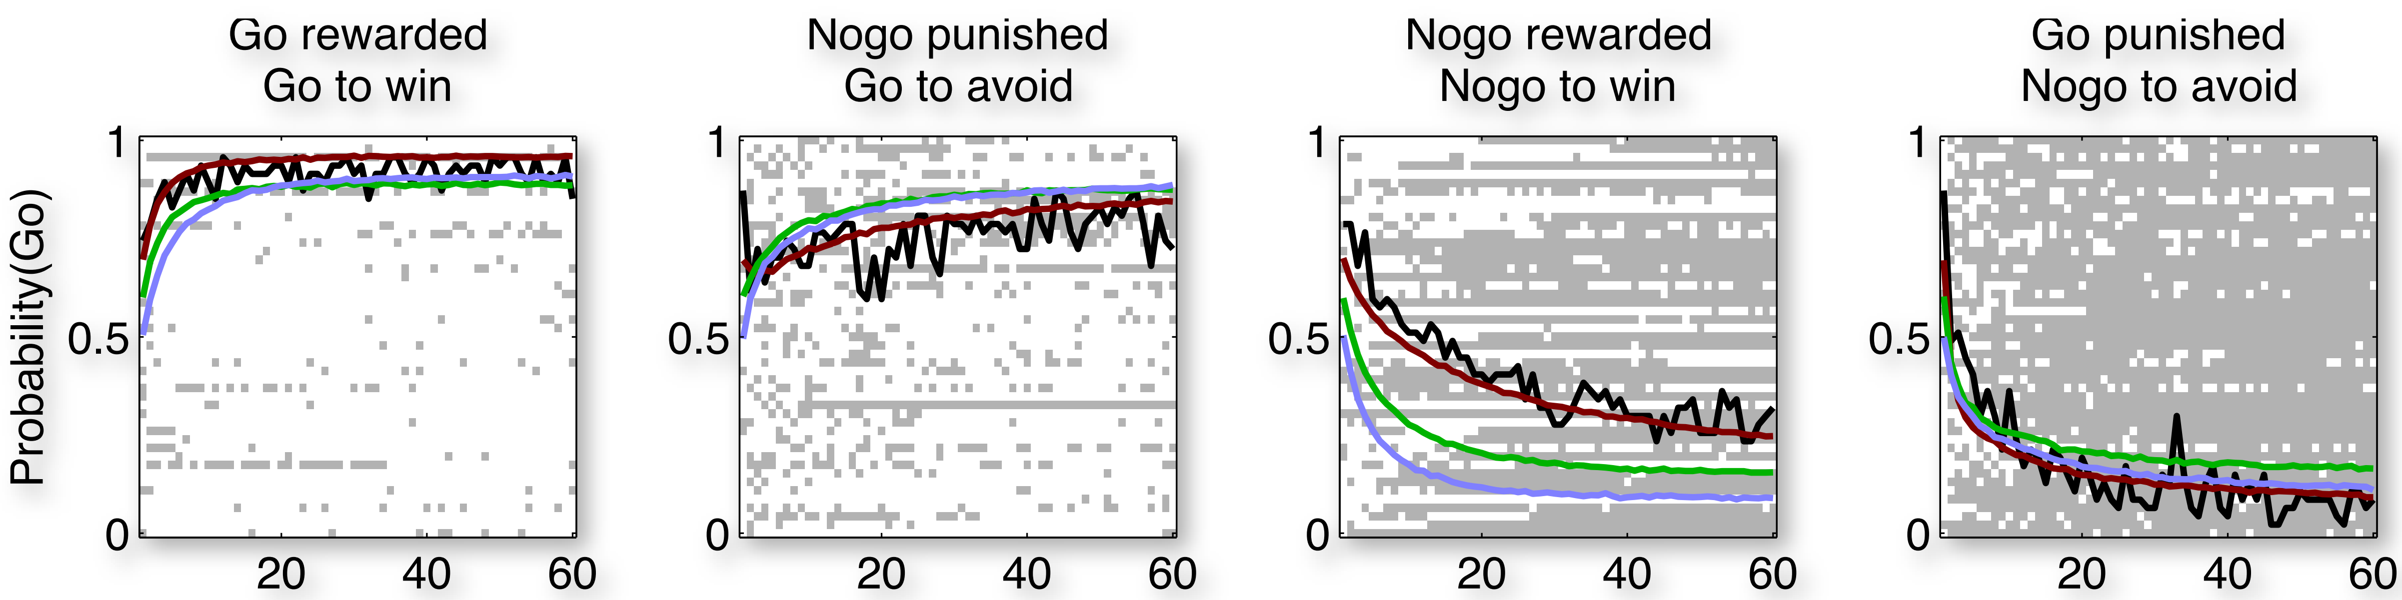

- Another sanity test: can my model fit this data at all?
- BUT: it might still be overfitting!

Overfitting

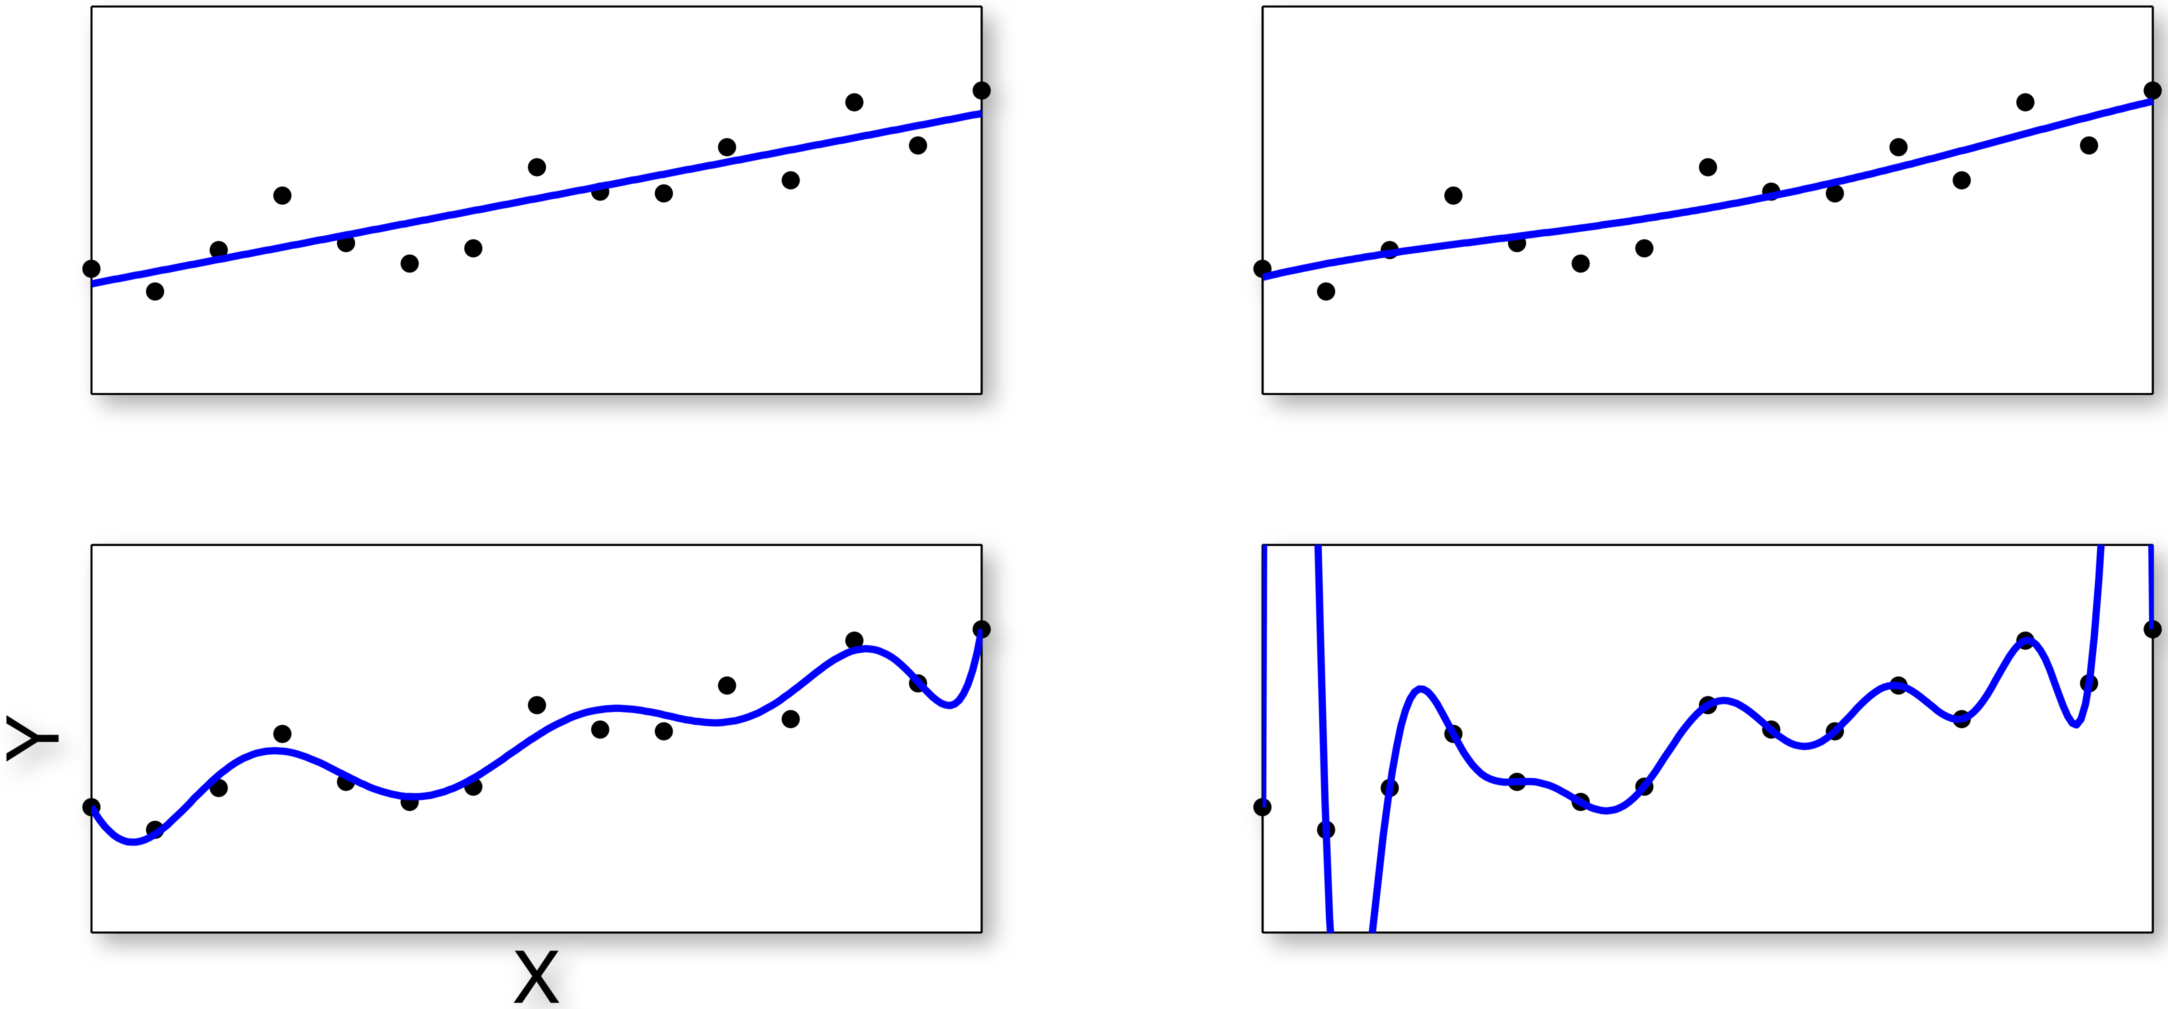

Model comparison

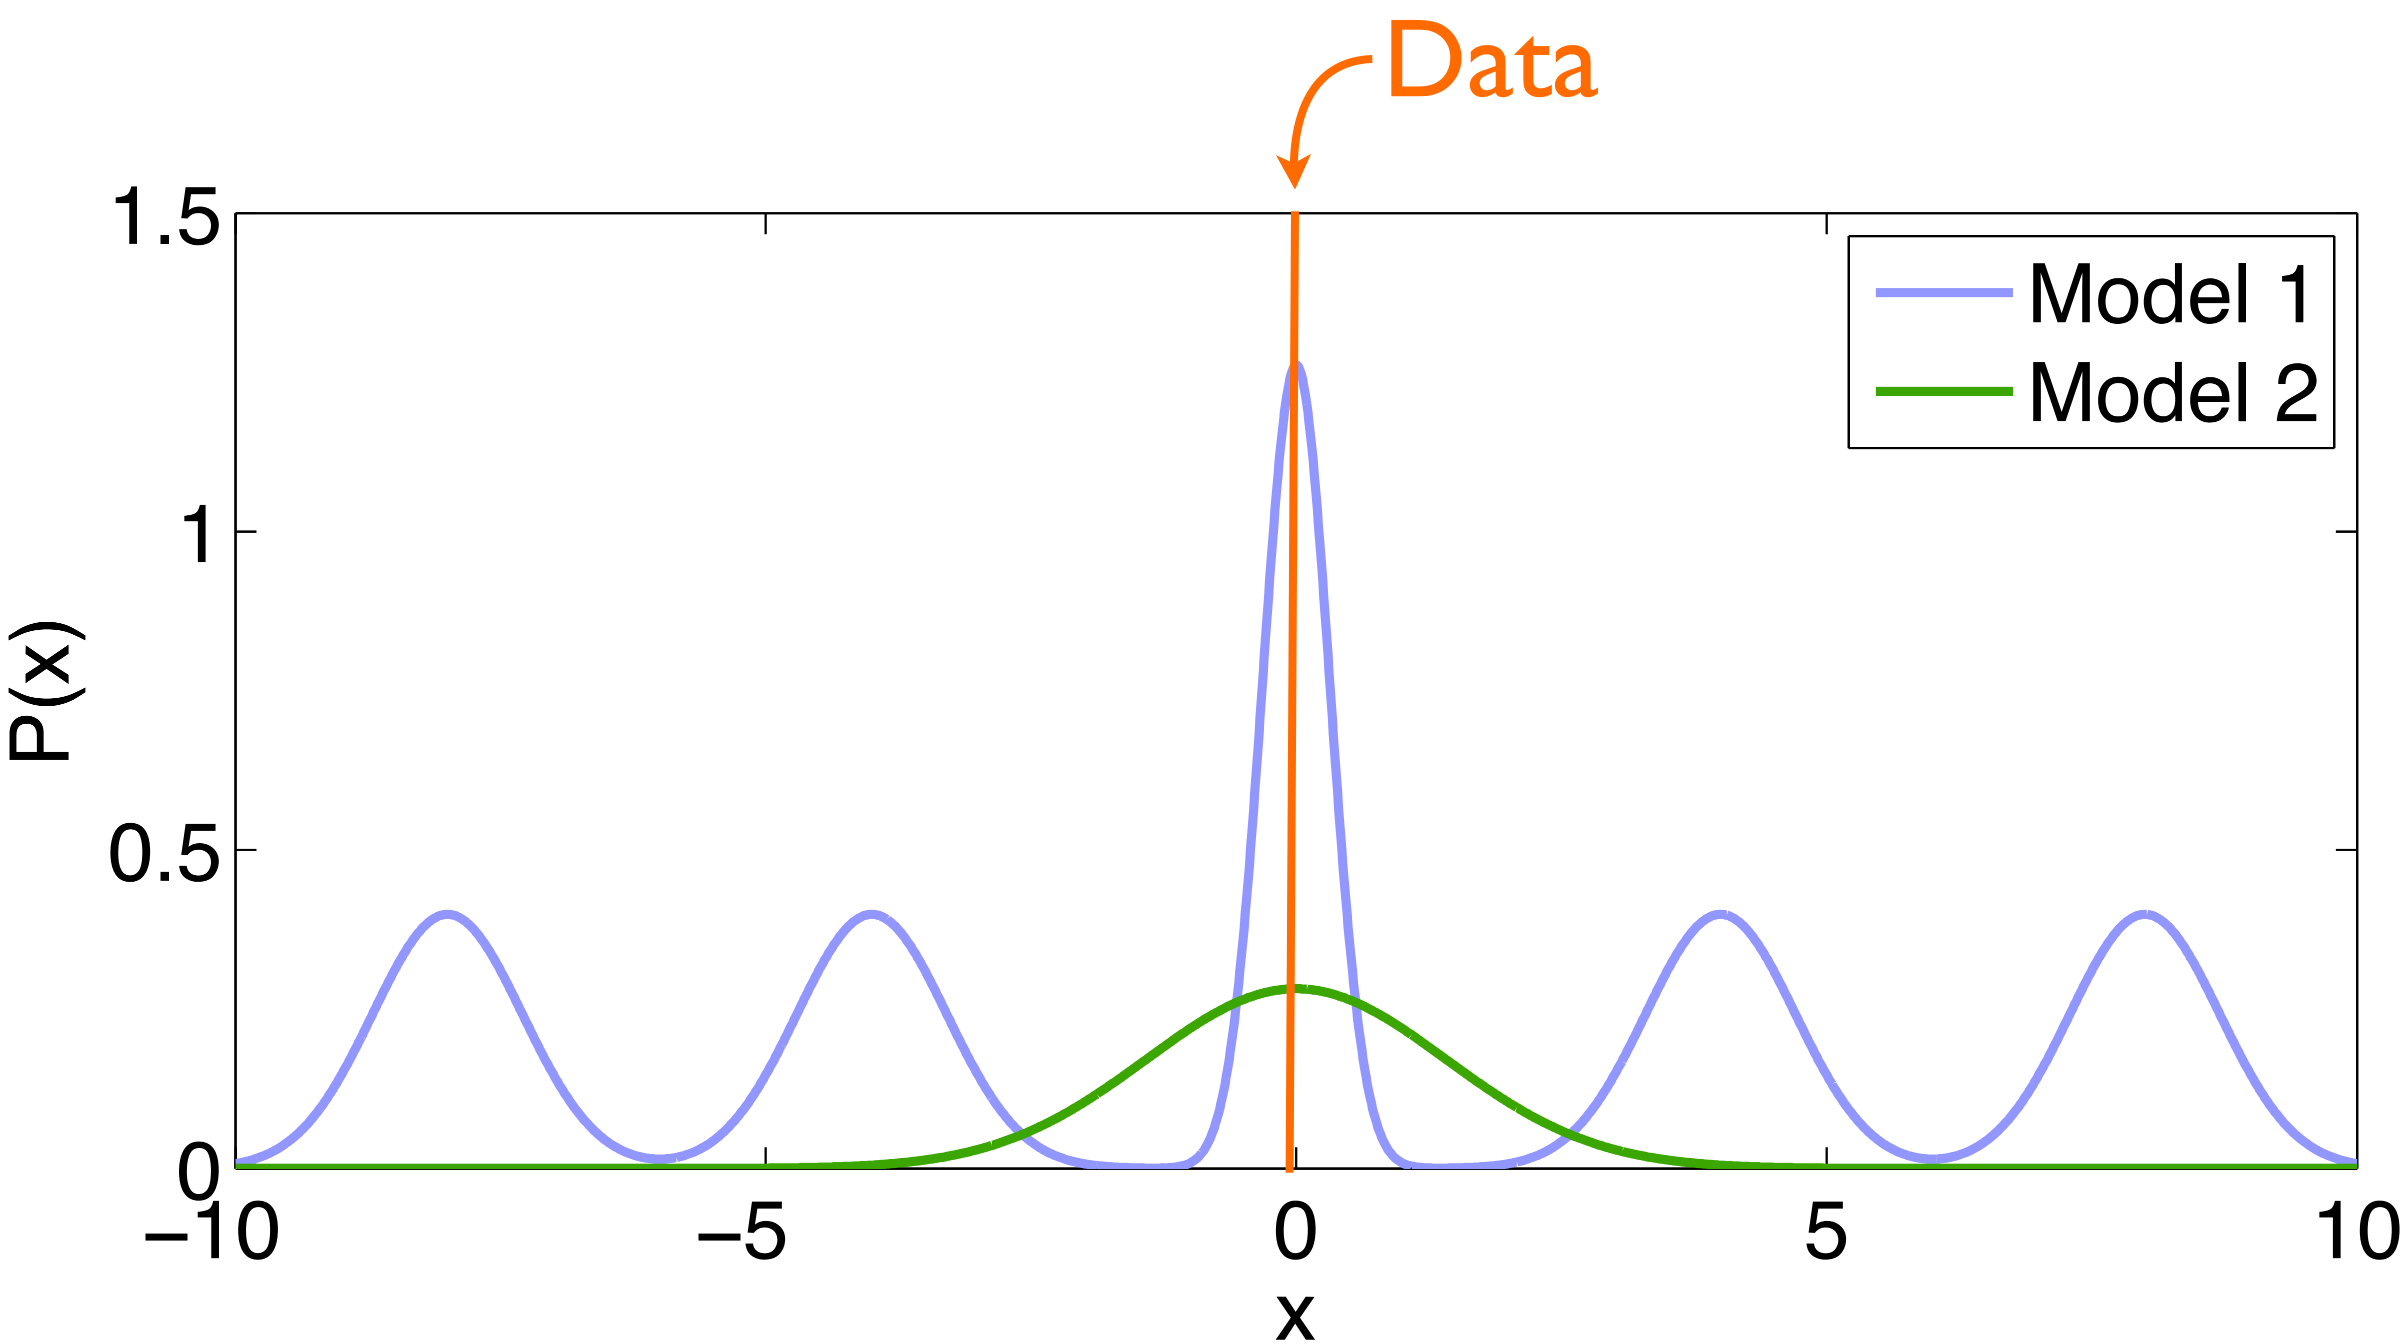

Model comparison

- So far: individual likelihood: $p(\mathbf{a}_k|\theta^k)$

- But can we allow different model for each subject?
  - No: use *all* the data
  - Yes? Forget the group level from now on

Model comparison

- So far: individual likelihood: $p(\mathbf{a}_k|\theta^k)$

- But can we allow different model for each subject?
  - No: use *all* the data: $\mathcal{A} = \{\{\mathbf{a}_{k,t}\}_{t=1}^T\}_{k=1}^K$

- To choose between models at the group level:

$$p(\mathcal{M}|\mathcal{A}) = \frac{p(\mathcal{A}|\mathcal{M})p(\mathcal{M})}{p(\mathcal{A})}$$

- If we have a prior over Models, we should use it:

$$p(\mathcal{M})$$

- Otherwise stick with model likelihood: $p(\mathcal{A}|\mathcal{M})$

Wednesday, 11 May 2011

Evaluating the model likelihood

- **Contains two integrals:**
  - subject parameters
  - prior parameters

$$p(\mathcal{A}|\mathcal{M}) \quad = \quad \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

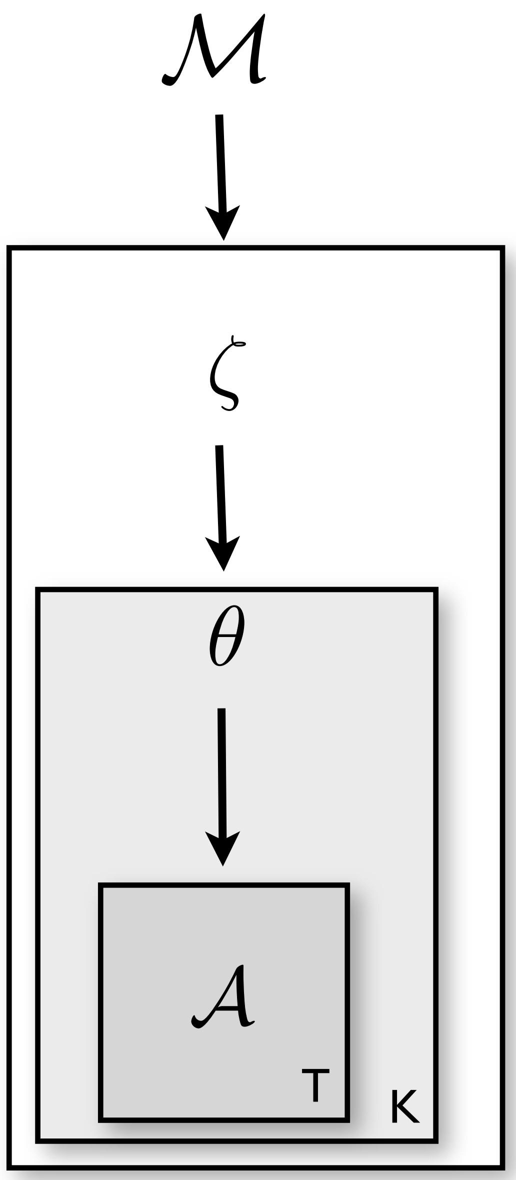

Why integrals?

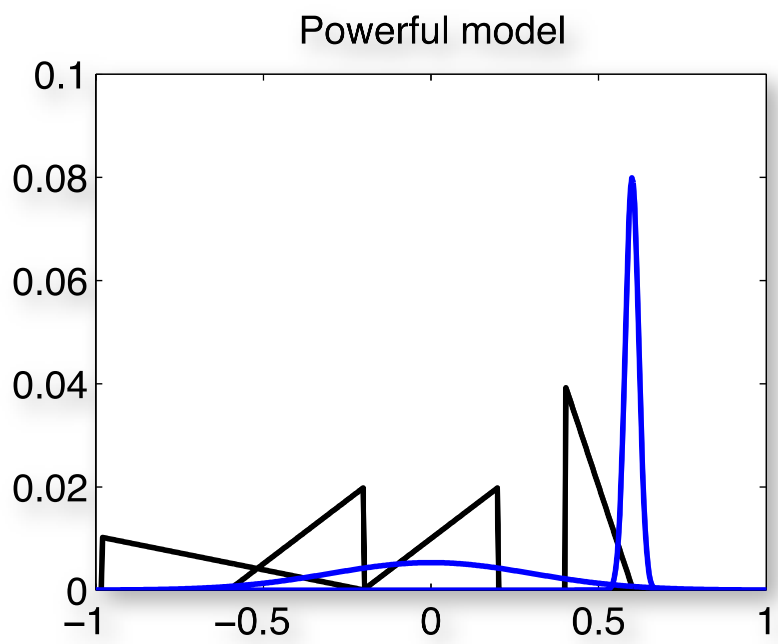

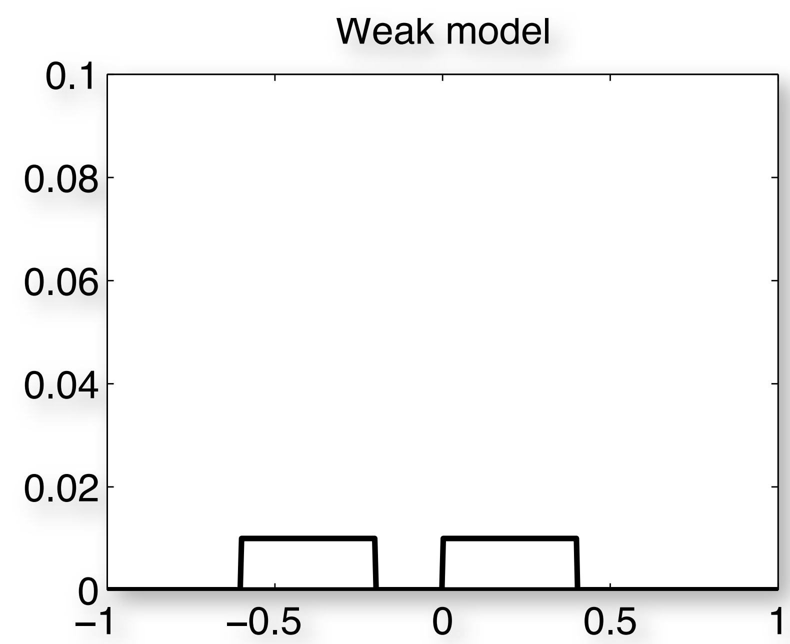

Why integrals?

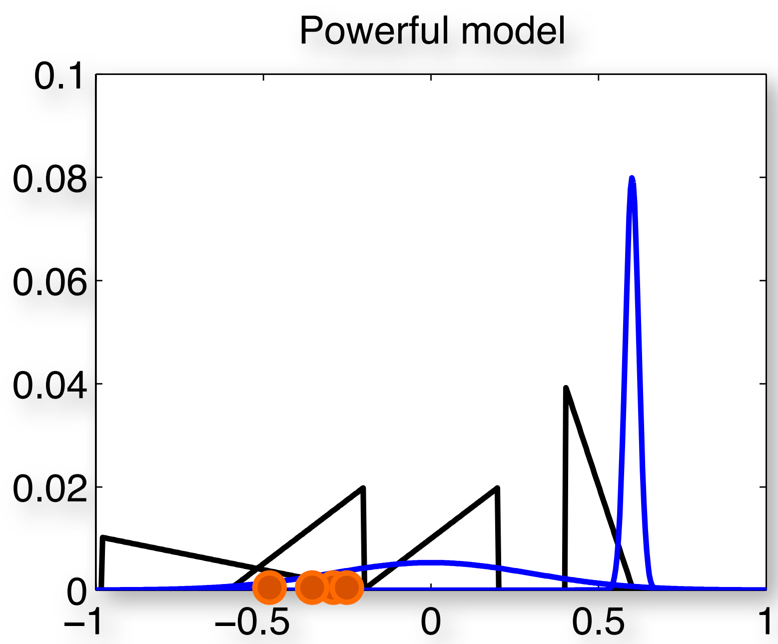

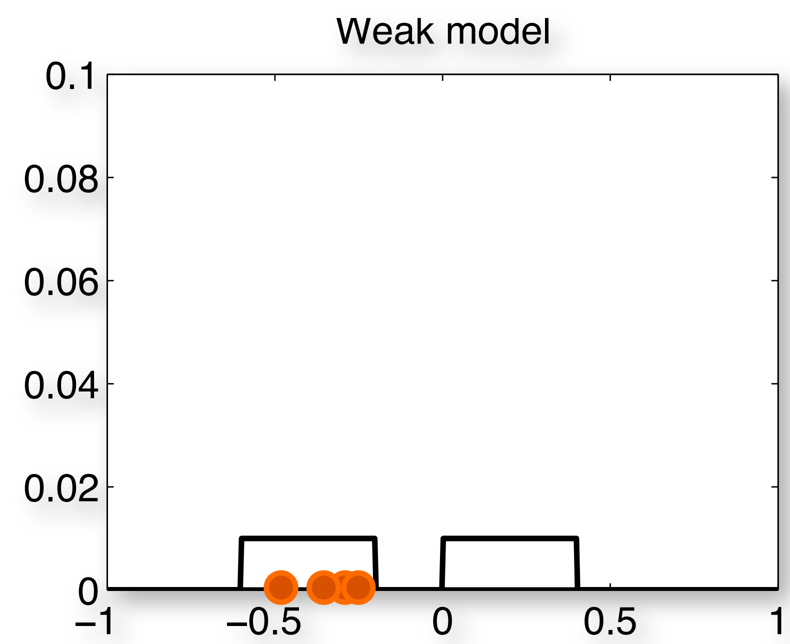

Why integrals?

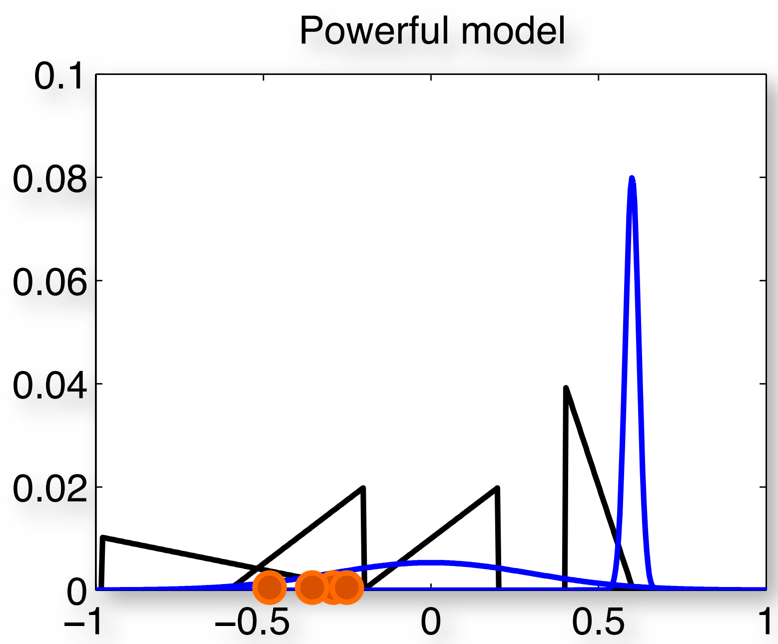

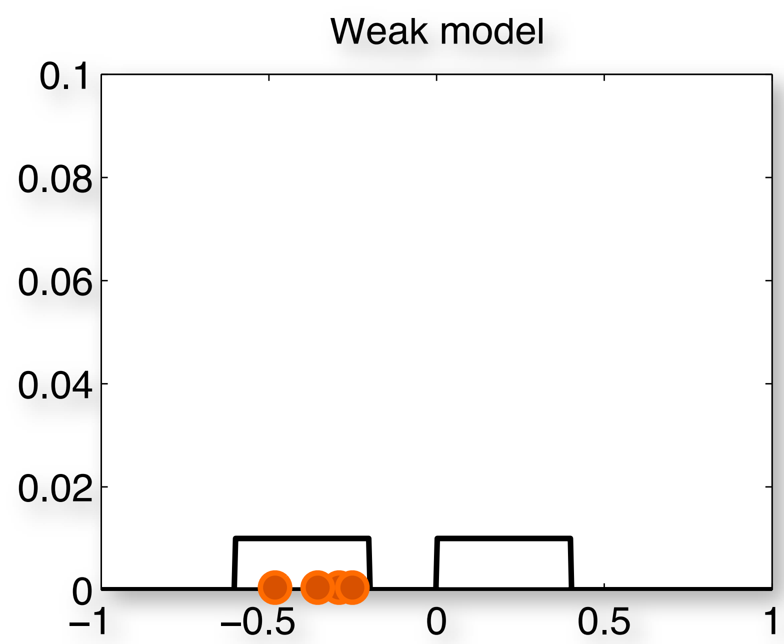

$$\frac{1}{N}(\mathbf{p}(\mathbf{X}|\boldsymbol{\theta}_1) + p(X|\theta_2) + \cdots)$$

These two factors fight it out

Model complexity vs model fit

Evaluating p(A|M)

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

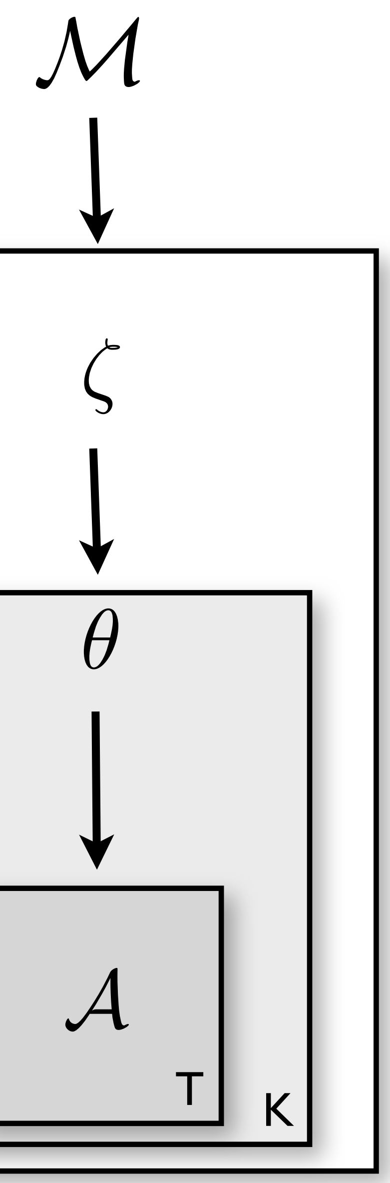

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

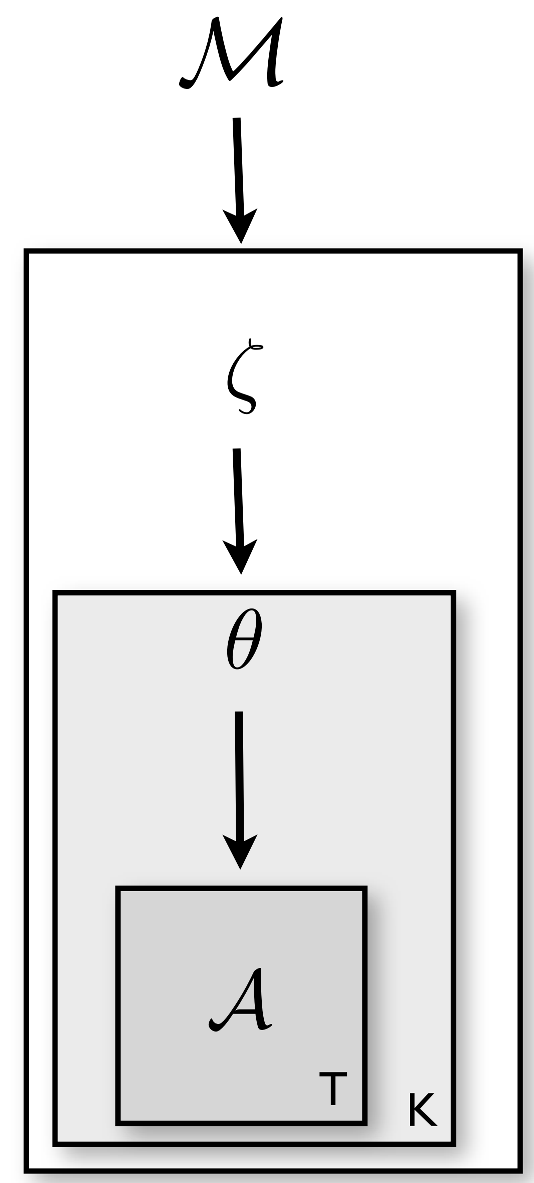

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

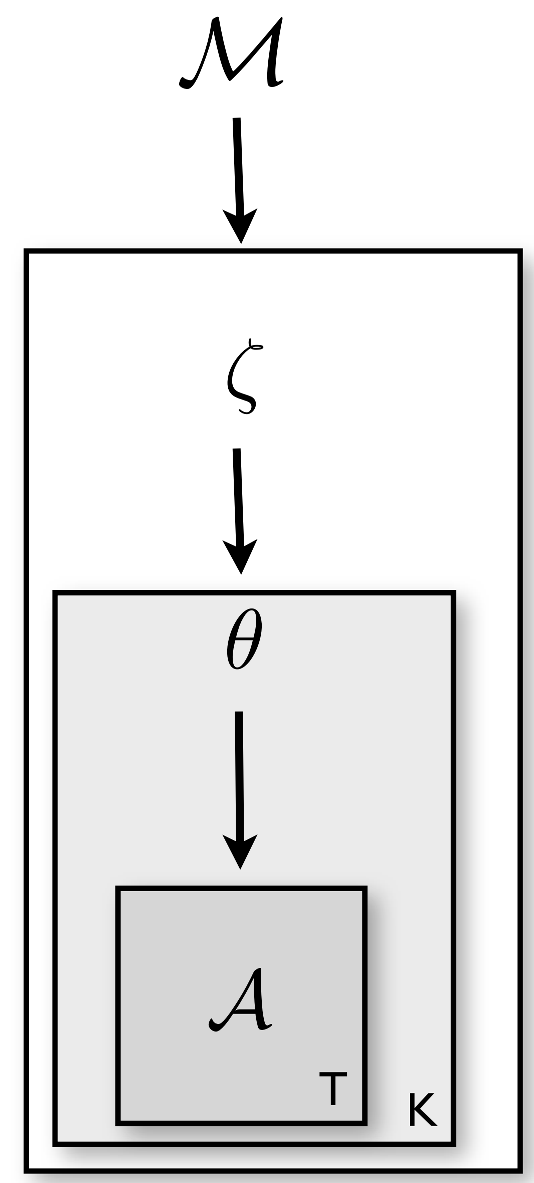

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$$p(\mathcal{A}|\mathcal{M}) = \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\, p(\zeta|\mathcal{M})$$

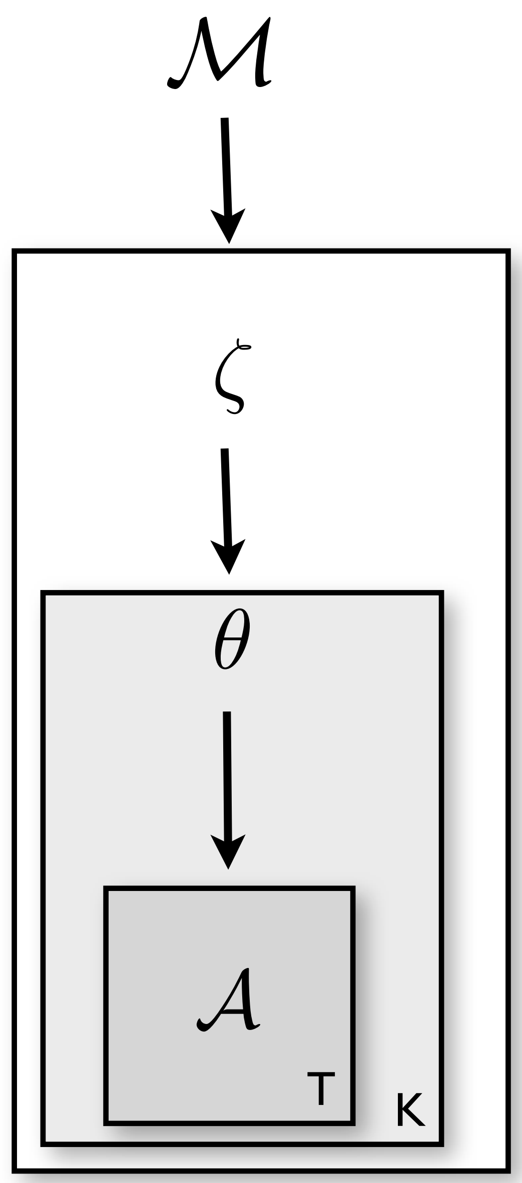

Laplacian approximation

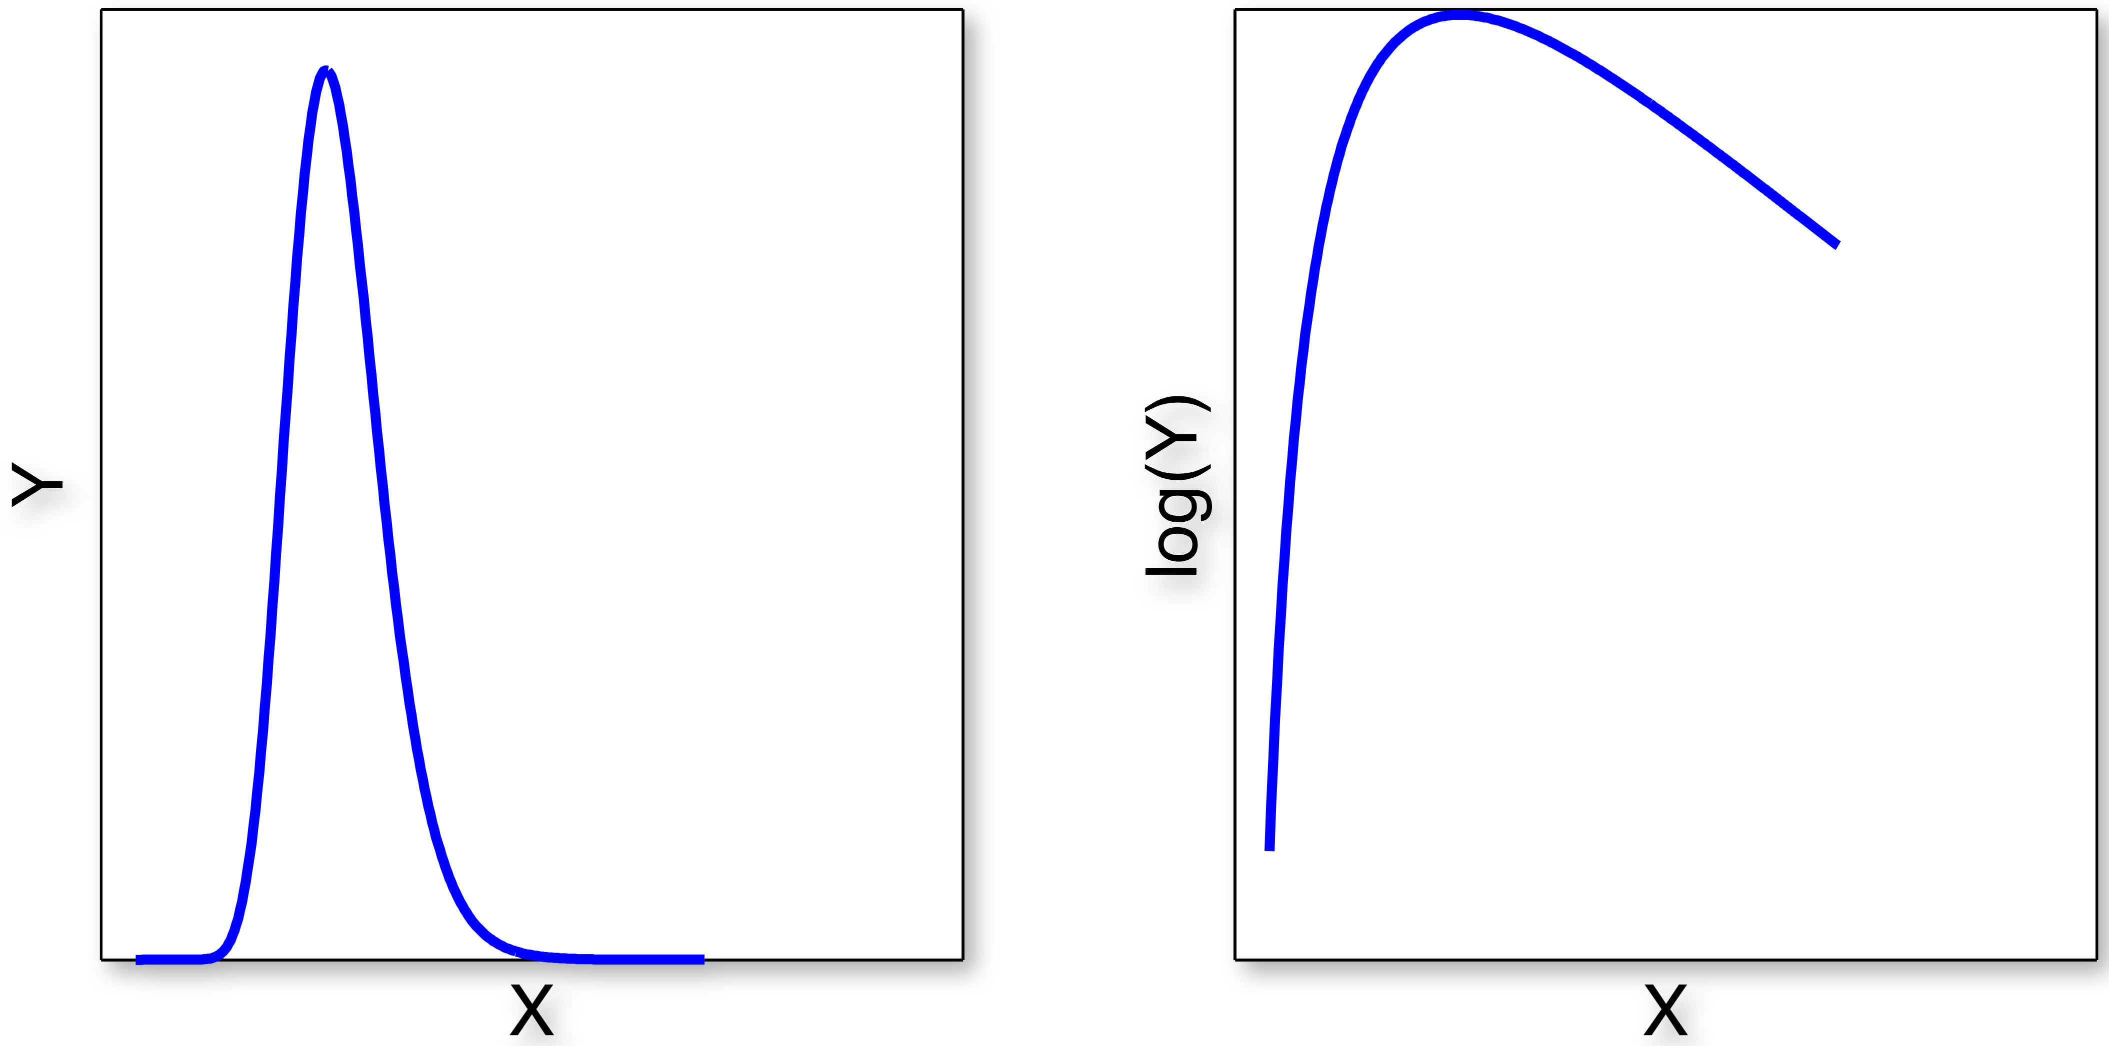

Laplacian approximation

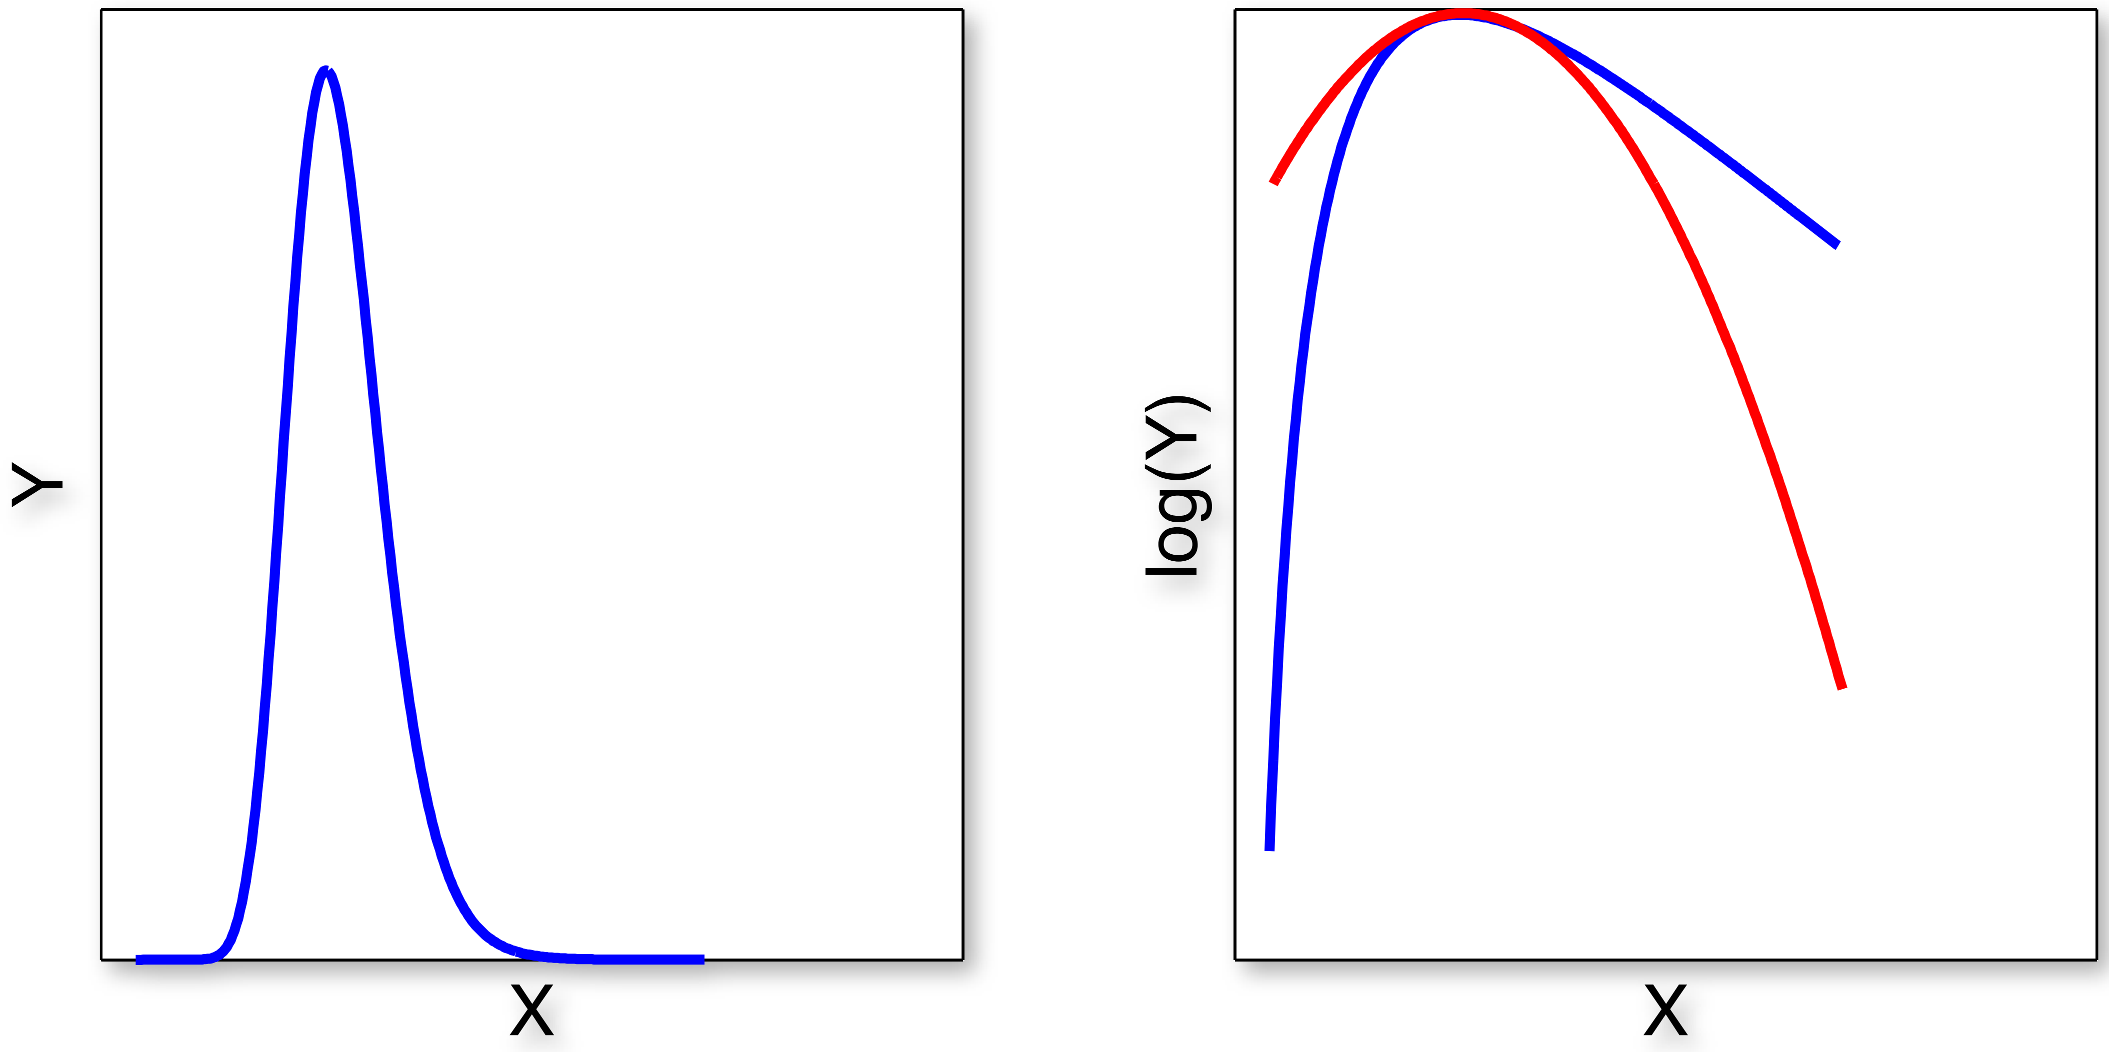

Laplacian approximation

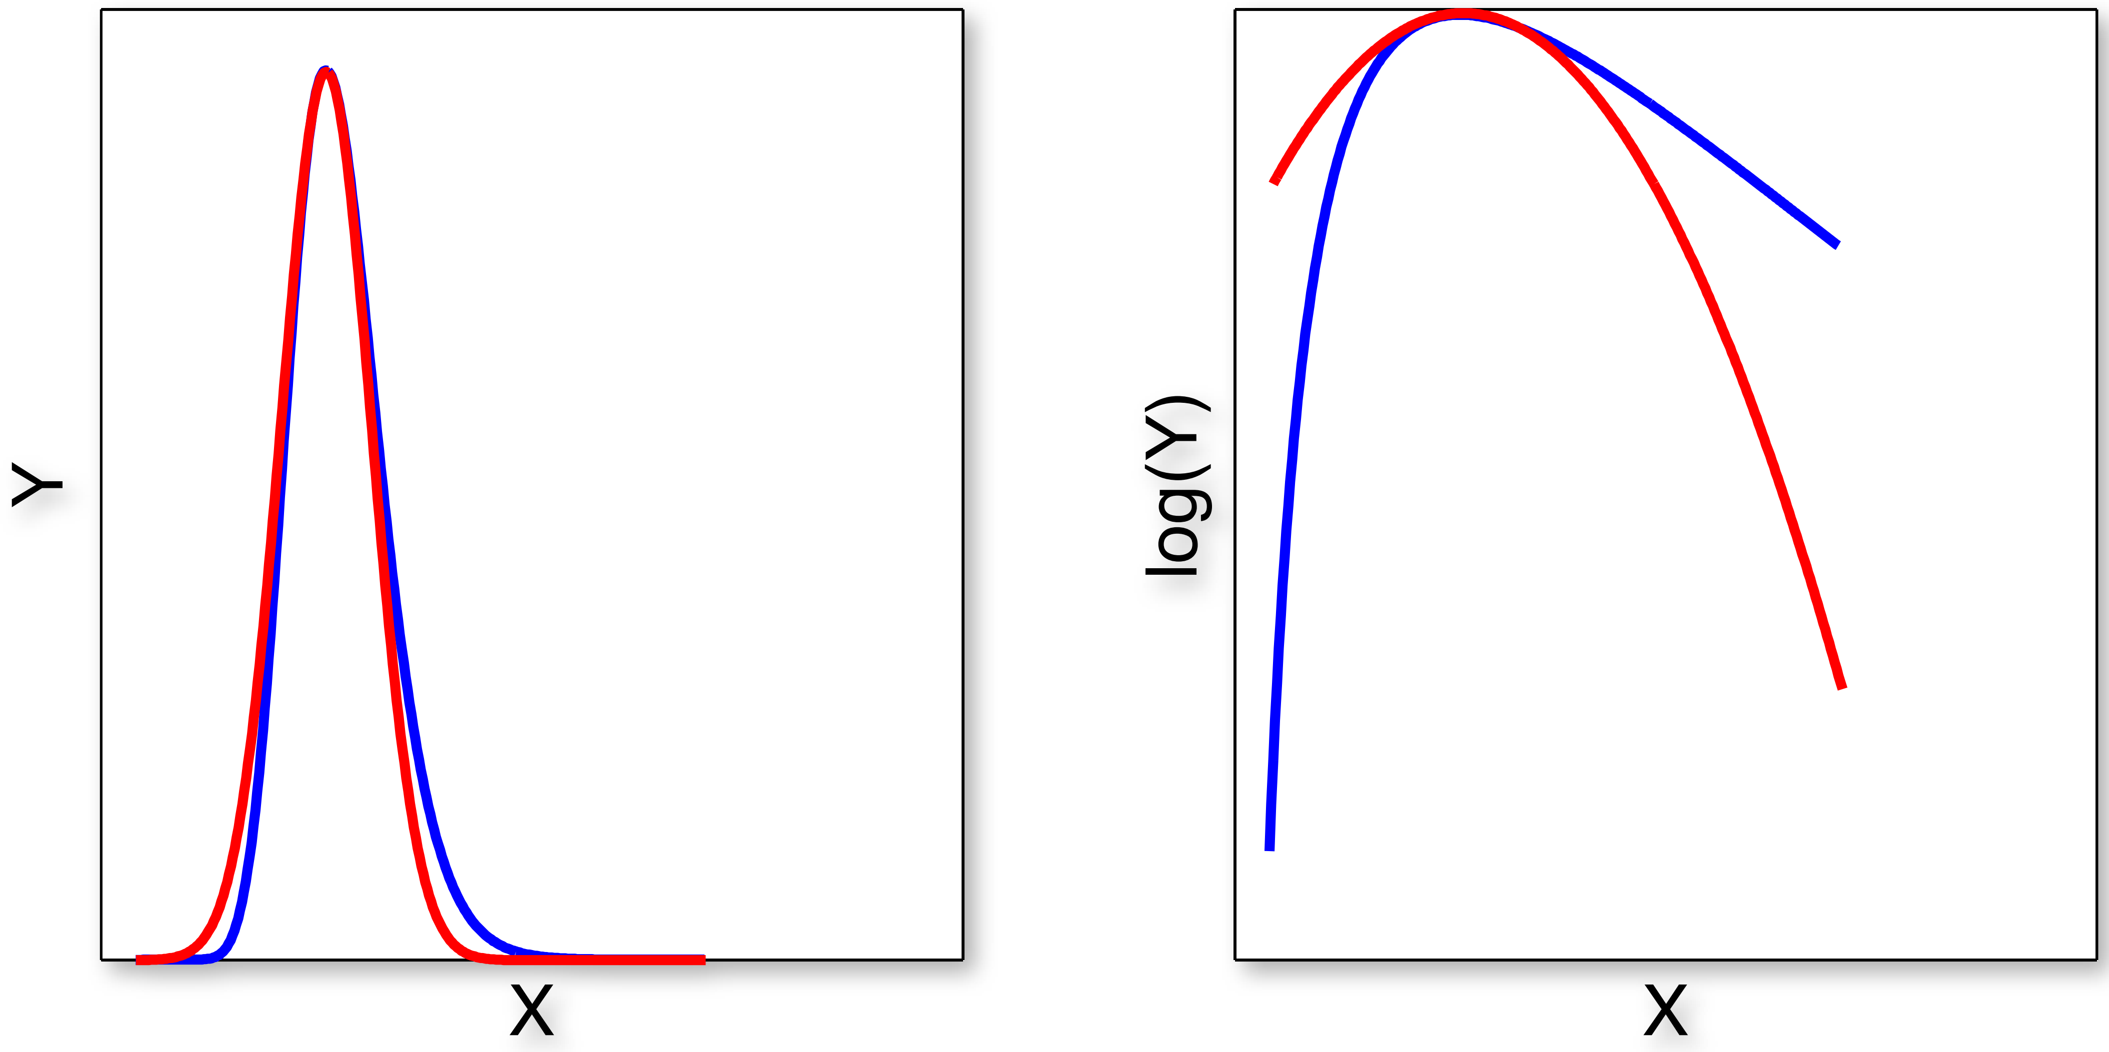

Laplacian approximation

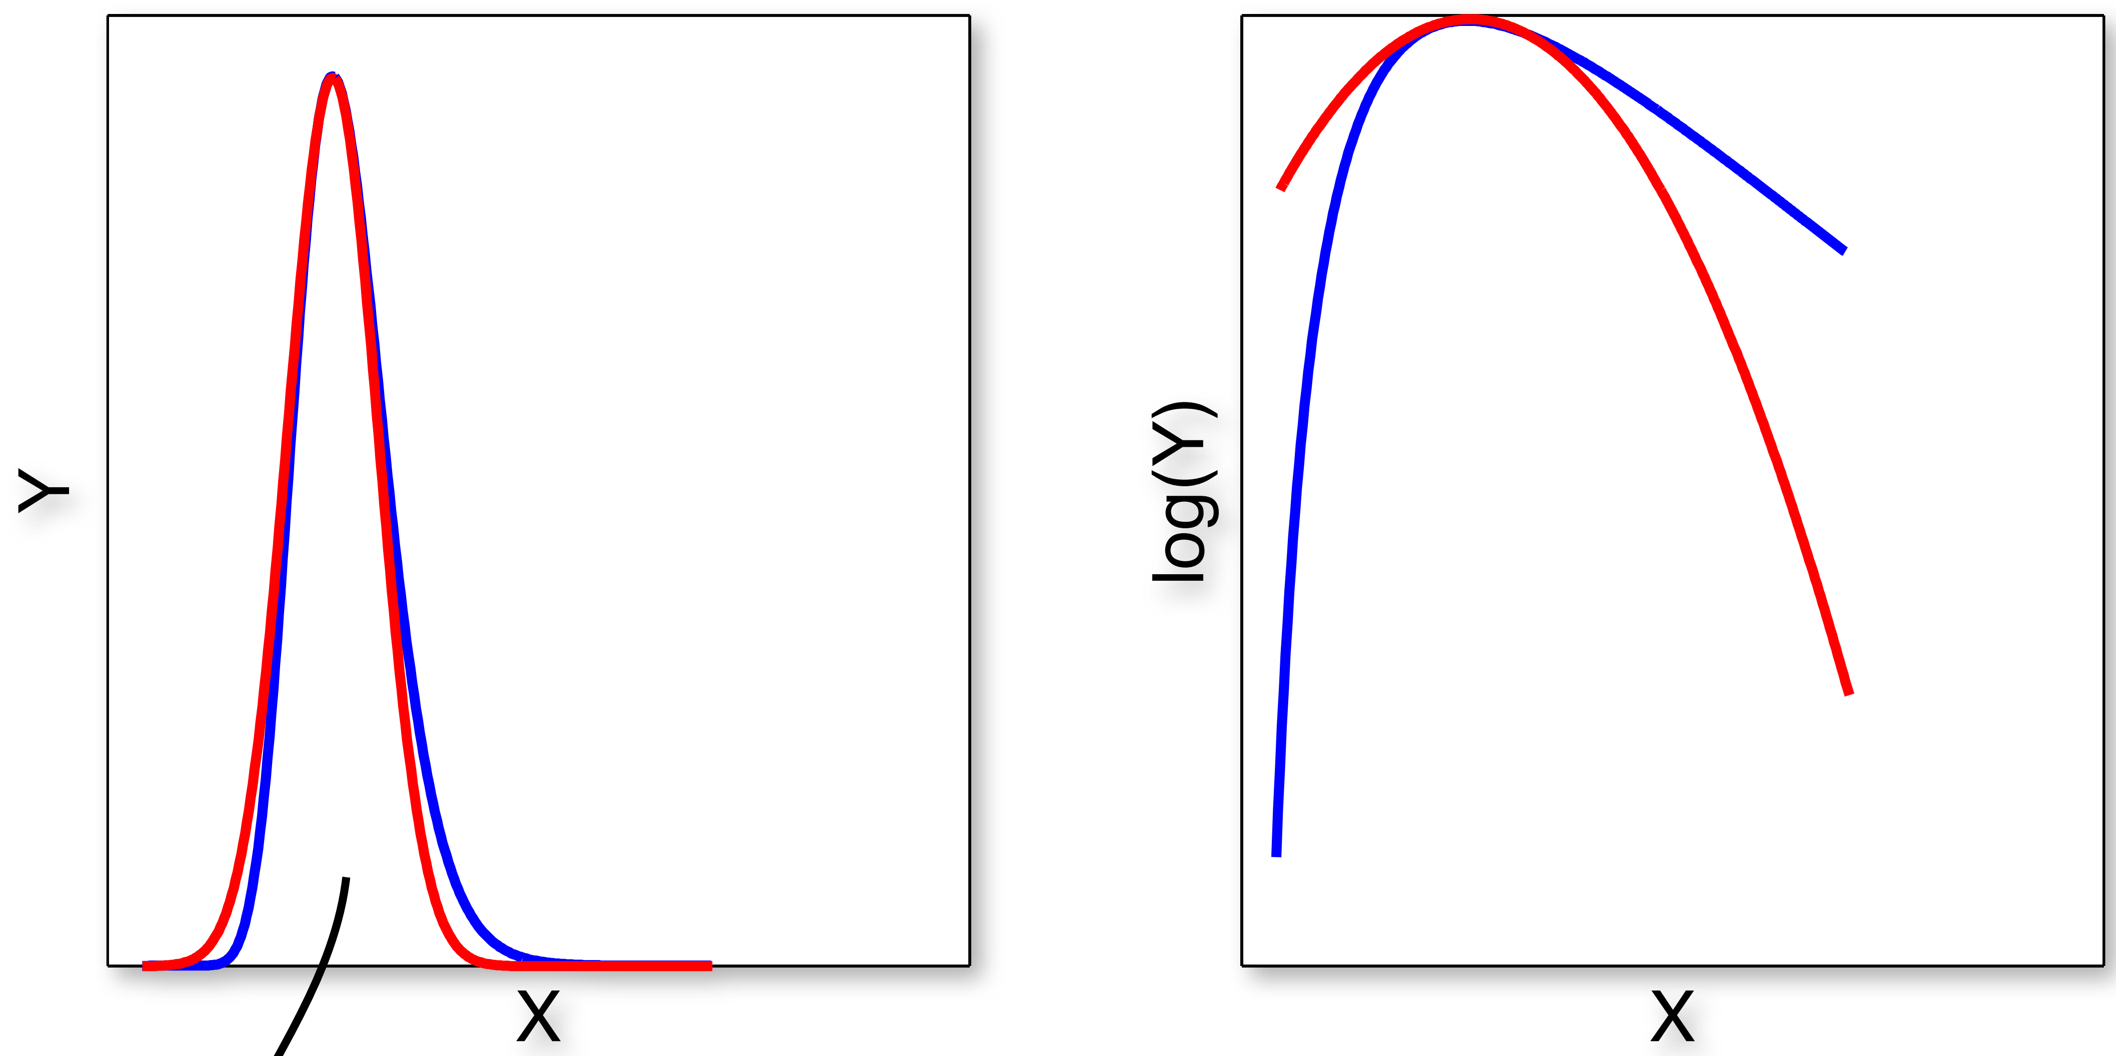

Just a Gaussian

$$\int dx\, f(x) \approx f^*(x_0)\, \sqrt{2\pi\sigma^2}$$

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$$\begin{aligned} p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M}) \\ &\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \end{aligned}$$

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$p(\mathcal{A}|\zeta,\mathcal{M})\; p(\zeta|\mathcal{M})$

$$\begin{aligned} p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\; p(\zeta|\mathcal{M}) \\ &\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M}) p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N |\Sigma|} \end{aligned}$$

Wednesday, 11 May 2011

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M})$
is propto Gaussian

$$\begin{aligned} p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M}) \\ &\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \end{aligned}$$

Evaluating p(A|M)

- **Two integrals**
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- **Step by step: approximating levels separately**
  - Top level first:

$p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M})$
is propto Gaussian

$$\begin{aligned} p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M}) \\ &\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \end{aligned}$$

Model doesn't prefer particular ζ

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M})$
is propto Gaussian

$$\begin{aligned}
p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M}) \\
&\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \\
\log p(\mathcal{A}|\mathcal{M}) &\approx \log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) + \frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)
\end{aligned}$$

Model doesn't prefer particular ζ

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M})$
is propto Gaussian

$$\begin{aligned}
p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M}) \\
&\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \\
\log p(\mathcal{A}|\mathcal{M}) &\approx \log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) \boxed{+\frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)}
\end{aligned}$$

Model doesn't prefer particular ζ

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M})$
is propto Gaussian

Model doesn't prefer particular ζ

$$\begin{aligned}
p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M}) \\
&\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \\
\log p(\mathcal{A}|\mathcal{M}) &\approx \log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) \boxed{+\frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)}
\end{aligned}$$

$\approx \quad -N \qquad$ Akaike Information Criterion (AIC)

$\approx \quad -\frac{N}{2}\log(KT) \qquad$ Bayesian Information Criterion (BIC)

Evaluating p(A|M)

- Two integrals
  - tricky

$$p(\mathcal{A}|\mathcal{M}) = \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M}) \int d\zeta\, p(\theta|\zeta)\, p(\zeta|\mathcal{M})$$

- Step by step: approximating levels separately
  - Top level first:

$p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M})$
is propto Gaussian

$$p(\mathcal{A}|\mathcal{M}) = \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\ p(\zeta|\mathcal{M})$$

$$\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|}$$

Model doesn't prefer particular ζ

$$\log p(\mathcal{A}|\mathcal{M}) \approx \log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) \boxed{+ \frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)}$$

$\approx \quad -N \quad$ Akaike Information Criterion (AIC)

$\approx \quad -\frac{N}{2}\log(KT) \quad$ Bayesian Information Criterion (BIC)

Model fit vs Model complexity

Approximating level 1

- Still leaves the first level:
  - Approximate integral by sampling, e.g. importance sampling:

$$\begin{aligned}\log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) &= \log \int d\theta\, p(\mathcal{A}|\theta)\, p(\theta|\zeta^{ML}) \\ &\approx \log \frac{1}{B}\sum_{b=1}^{B} p(\mathcal{A}|\theta^b) \\ \theta^b &\sim p(\theta|\zeta^{ML})\end{aligned}$$

Group-level BIC

$$
\begin{aligned}
\log p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta)\, p(\zeta|\mathcal{M}) \\
&\approx -\frac{1}{2}\mathsf{BIC}_{\mathsf{int}} \\
&= \log \hat{p}(\mathcal{A}|\hat{\zeta}^{ML}) - \frac{1}{2}|\mathcal{M}|\log(|\mathcal{A}|)
\end{aligned}
$$

- So:

Group-level BIC

$$
\begin{aligned}
\log p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta)\, p(\zeta|\mathcal{M}) \\
&\approx -\frac{1}{2}\mathsf{BIC}_{\mathsf{int}} \\
&= \log \hat{p}(\mathcal{A}|\hat{\zeta}^{ML}) - \frac{1}{2}|\mathcal{M}|\log(|\mathcal{A}|)
\end{aligned}
$$

- So:

$\boxed{\text{Model fit vs Model complexity}}$

Individual level BIC

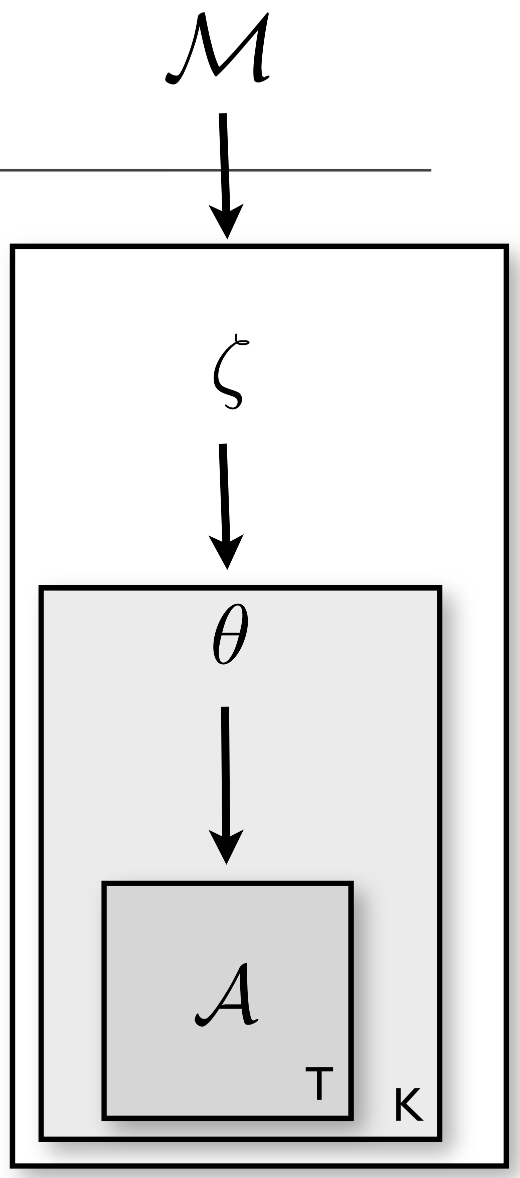

- No group level

$$
\begin{aligned}
p(\mathcal{A}|\mathcal{M}) &= \int d\zeta \, p(\mathcal{A}|\zeta,\mathcal{M}) \; p(\zeta|\mathcal{M}) \\
&\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M}) p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N |\Sigma|} \\
\log p(\mathcal{A}|\mathcal{M}) &\approx \log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) + \frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)
\end{aligned}
$$

- Model comparison for each subject individually
  - Treat them as data points -> do classical pairwise tests

Individual level BIC

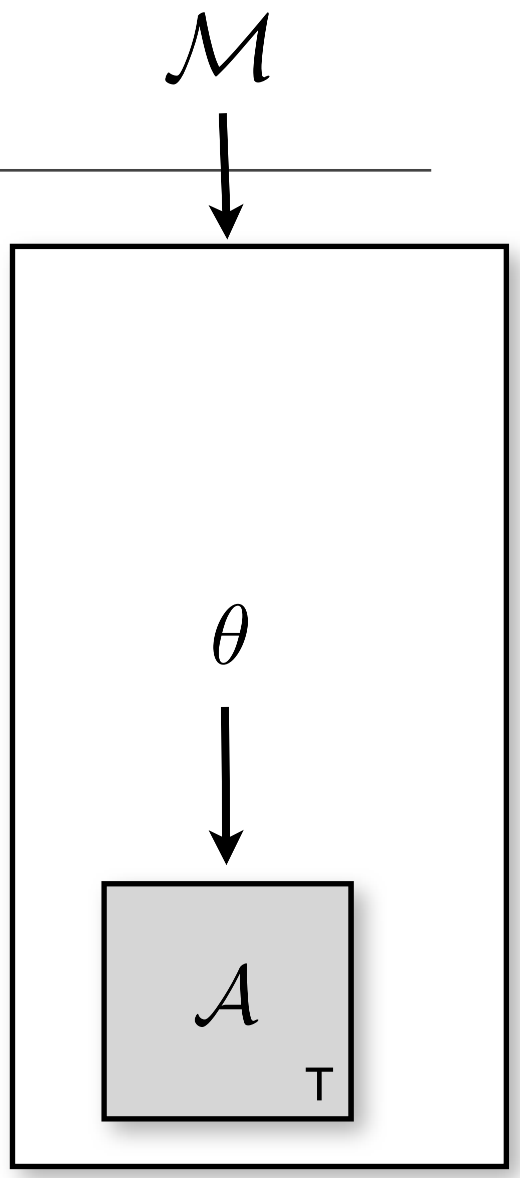

- No group level

$$
\begin{aligned}
p(\mathcal{A}|\mathcal{M}) &= \int d\zeta\, p(\mathcal{A}|\zeta,\mathcal{M})\; p(\zeta|\mathcal{M}) \\
&\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \\
\log p(\mathcal{A}|\mathcal{M}) &\approx \log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) + \frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)
\end{aligned}
$$

- Model comparison for each subject individually
  - Treat them as data points -> do classical pairwise tests

Individual level BIC

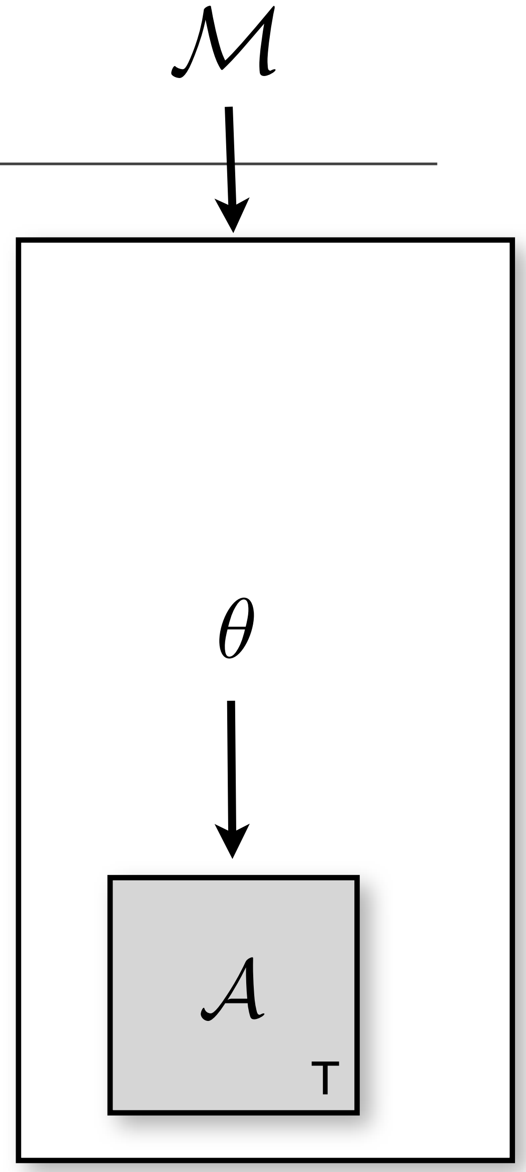

- No group level

$$
\begin{aligned}
p(\mathcal{A}|\mathcal{M}) &= \int d\zeta \, p(\mathcal{A}|\zeta,\mathcal{M}) \; p(\zeta|\mathcal{M}) \\
&\approx p(\mathcal{A}|\zeta^{ML},\mathcal{M})p(\zeta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \\
\log p(\mathcal{A}|\mathcal{M}) &\approx \log p(\mathcal{A}|\zeta^{ML},\mathcal{M}) + \frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)
\end{aligned}
$$

known!

- Model comparison for each subject individually
  - Treat them as data points -> do classical pairwise tests

Individual level BIC

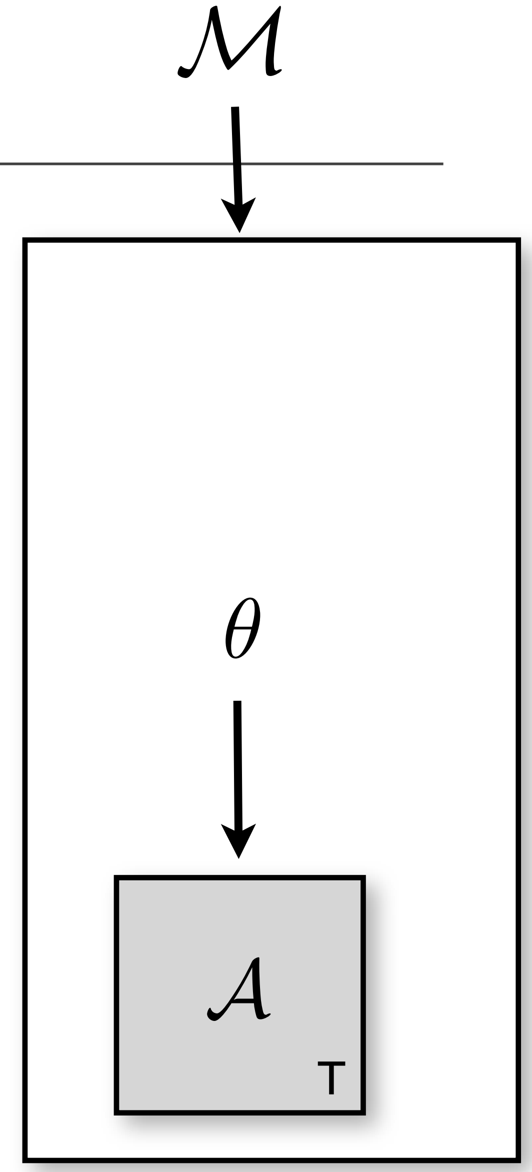

- No group level

$$
\begin{aligned}
p(\mathcal{A}|\mathcal{M}) &= \int d\theta\, p(\mathcal{A}|\theta,\mathcal{M})\; p(\theta|\mathcal{M}) \\
&\approx p(\mathcal{A}|\theta^{ML},\mathcal{M})p(\theta^{ML}|\mathcal{M}) \times \sqrt{(2\pi)^N|\Sigma|} \\
\log p(\mathcal{A}|\mathcal{M}) &\approx \log p(\mathcal{A}|\theta^{ML},\mathcal{M}) + \frac{1}{2}\log(|\Sigma|) + \frac{N}{2}\log(2\pi)
\end{aligned}
$$

known!

- Model comparison for each subject individually
    - Treat them as data points -> do classical pairwise tests

How does it do?

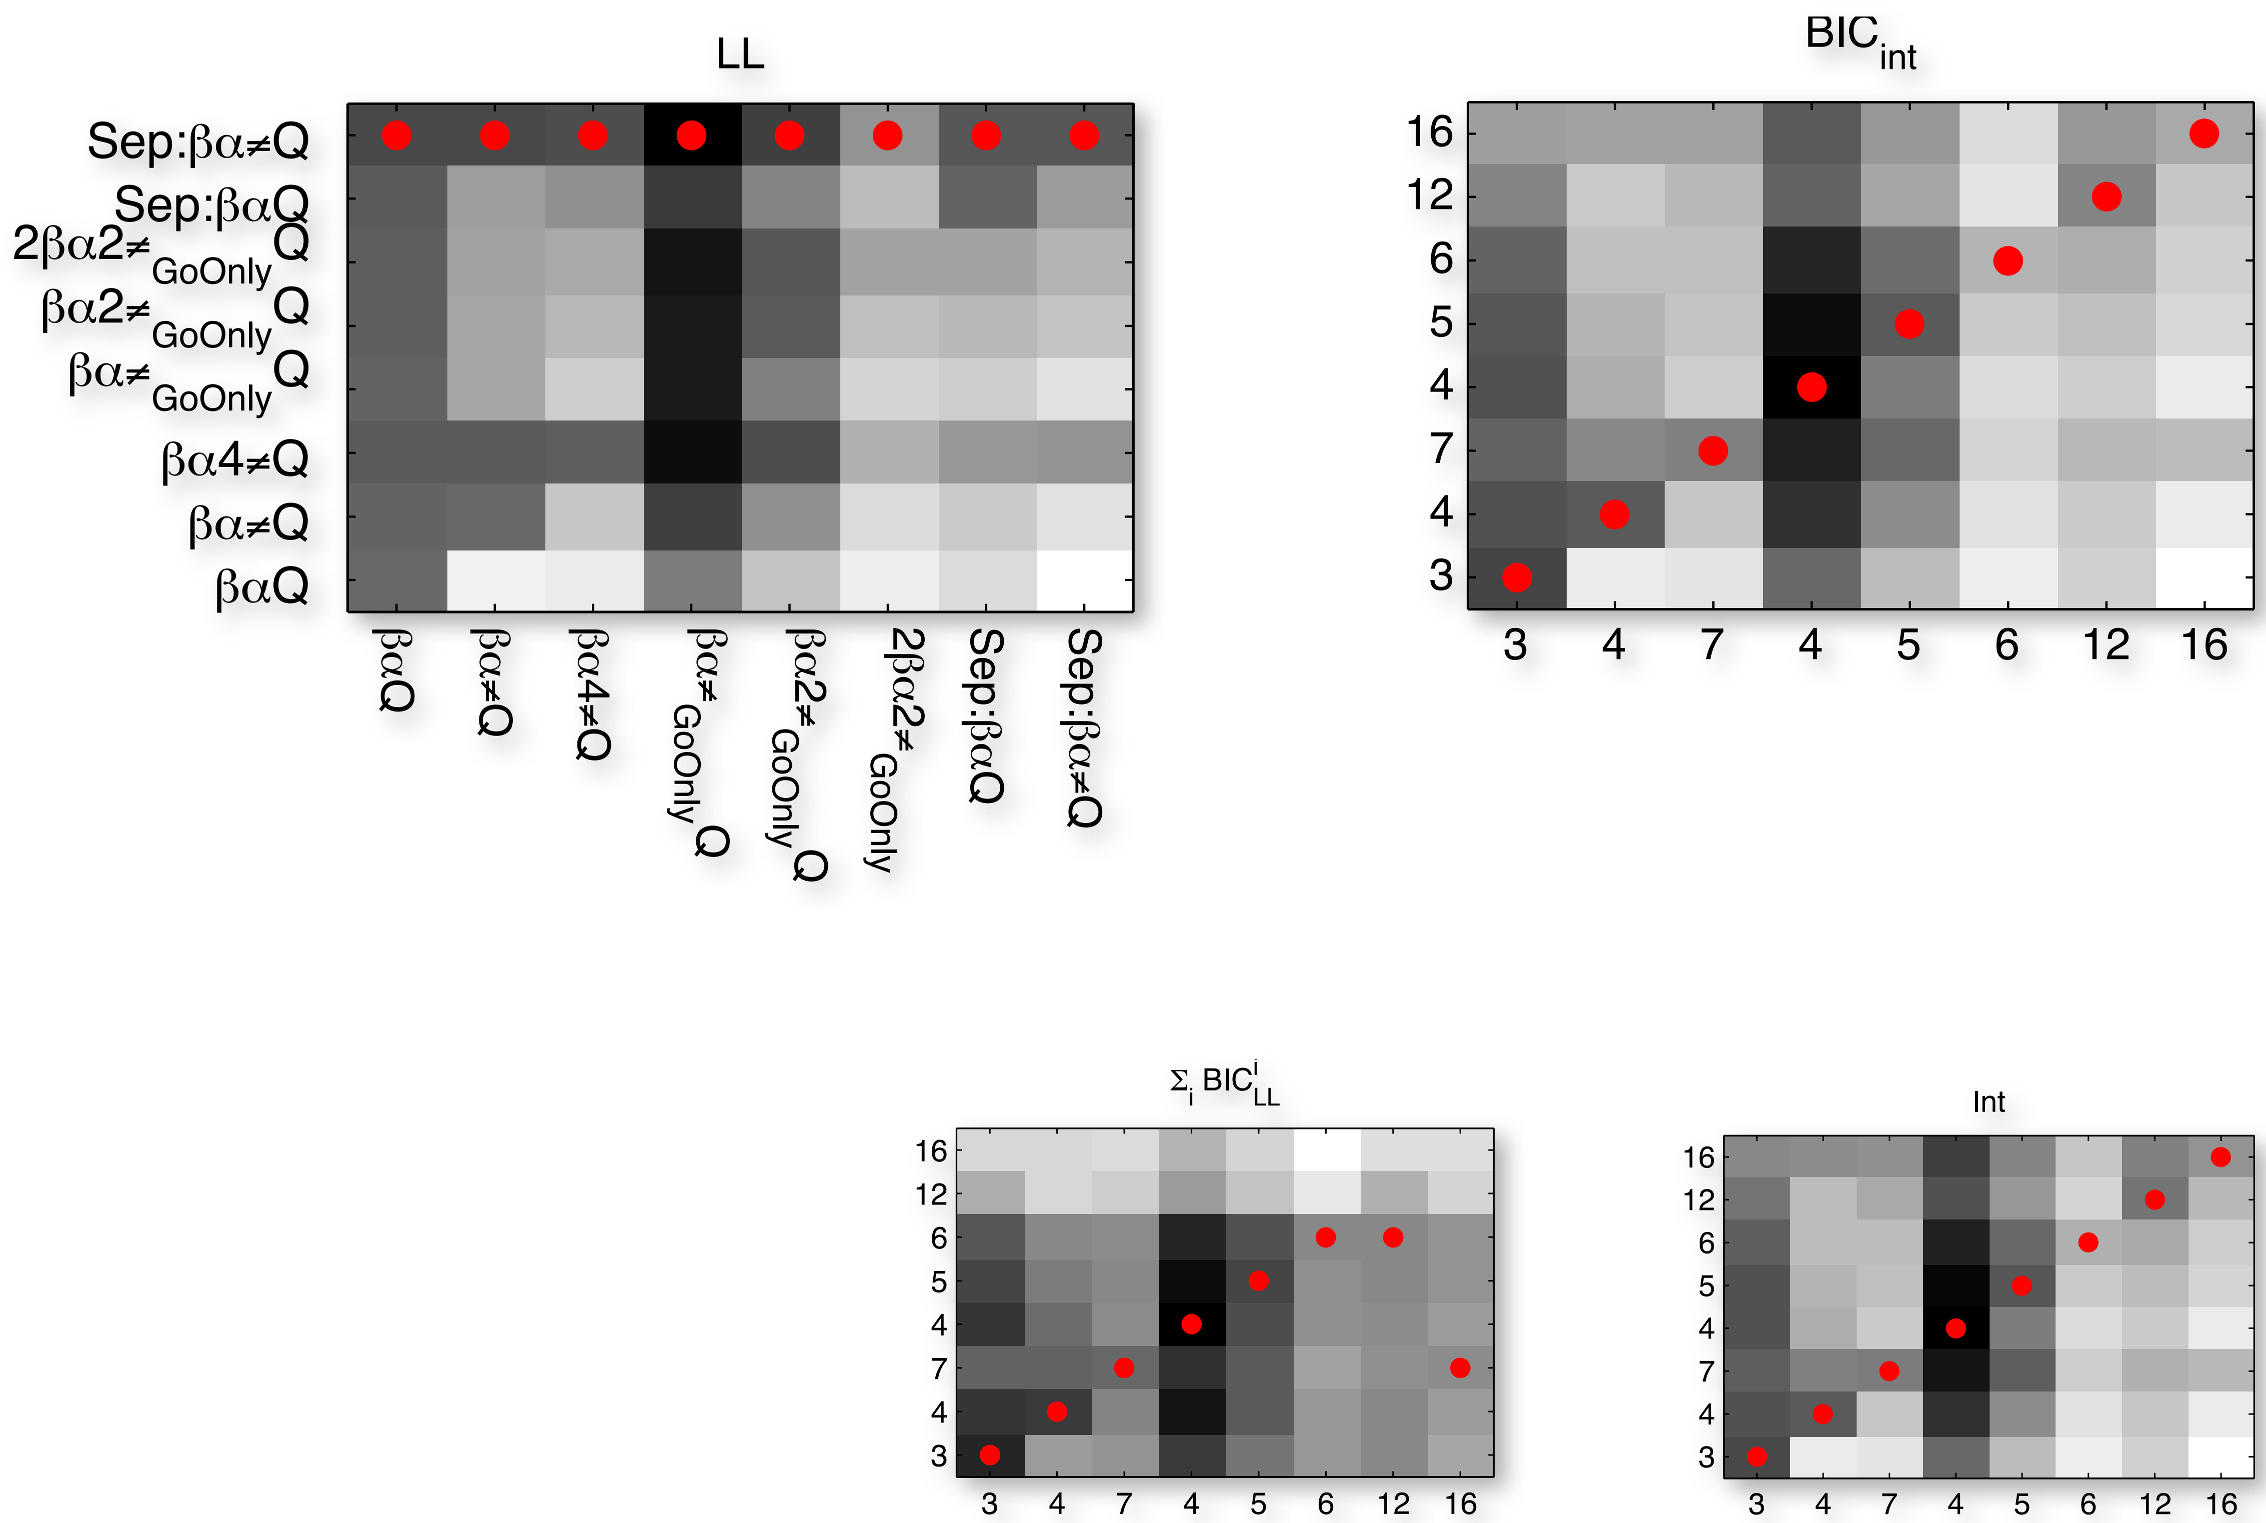

How does it do?

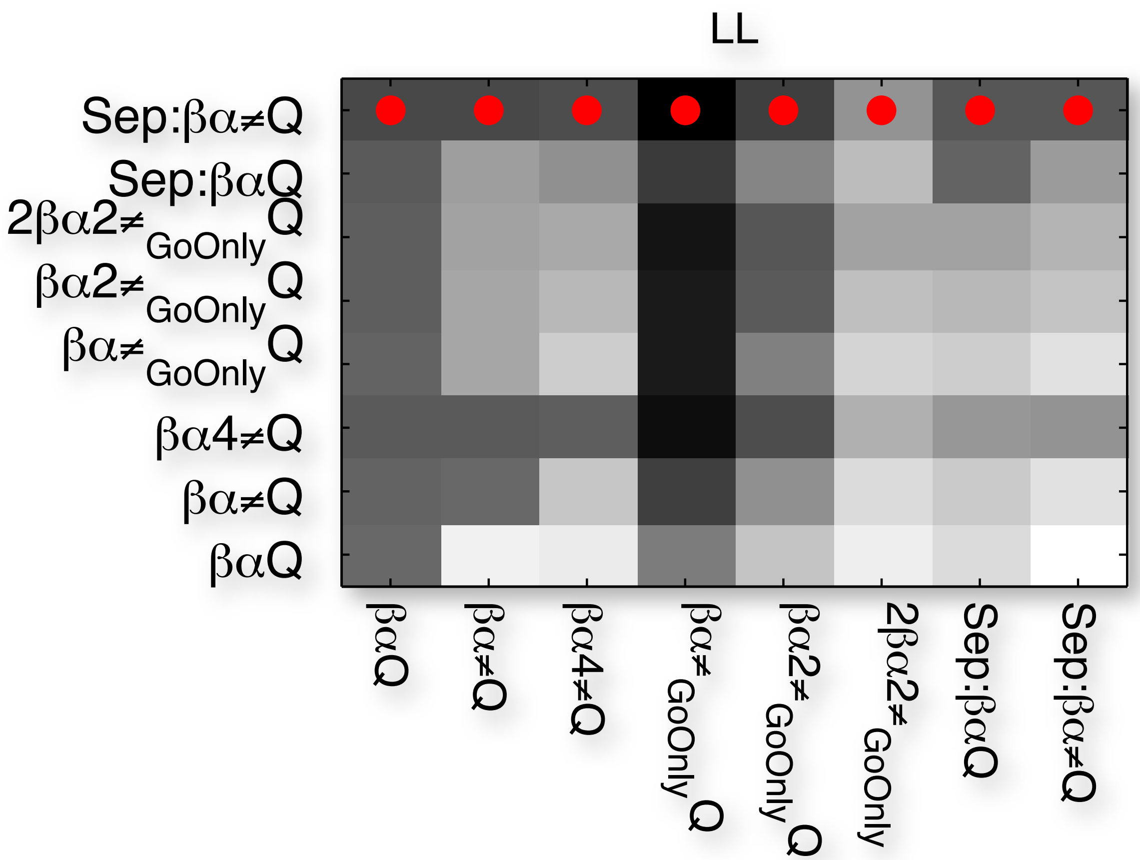

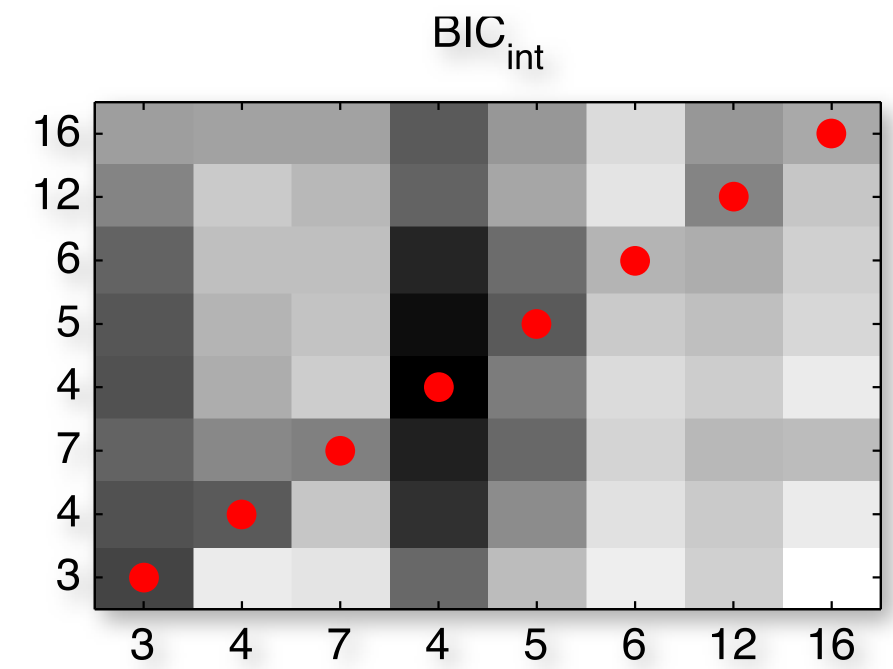

Fitted by EM...
too nice?

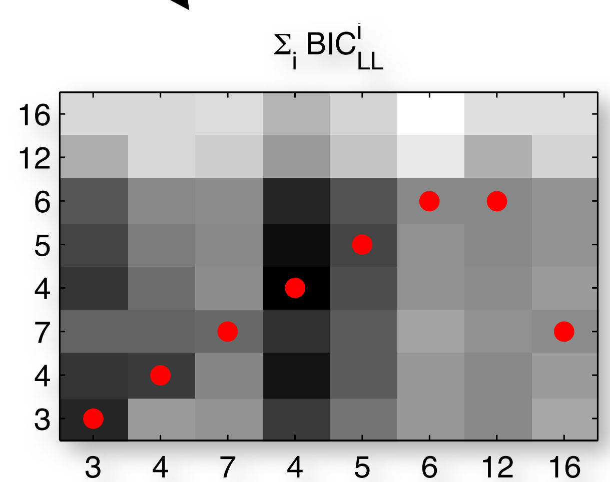

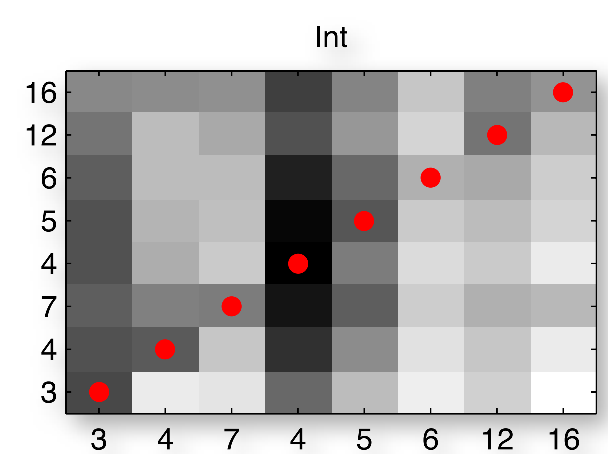

Top-level Laplacian approximation

- Estimating the top-level determinant
  - using 2nd order finite differences

$$\frac{d^2}{dh_{ij}^2} p(\mathcal{A}|\zeta)\Big|_{\zeta=\hat{\zeta}^{ML}} \approx \frac{1}{\delta^2}\Big[p(\mathcal{A}|\hat{\boldsymbol{\zeta}}^{ML} + \delta \mathbf{e}_i) - 2p(\mathcal{A}|\hat{\boldsymbol{\zeta}}^{ML}) + p(\mathcal{A}|\hat{\boldsymbol{\zeta}}^{ML} - \delta \mathbf{e}_j)\Big]$$

  - the shifted likelihoods can be evaluated by shifting the samples.

Group comparisons

- Are two groups similar in parameter x?
  - ANOVA: compare likelihood of two means to likelihood of one global mean. Take degrees of freedom into account.
  - But: this tries to account for the **parameters** with one or two groups, not for the **data**
- Need to:
  - 1: Compare models with separate or joint parameter & prior:

Model 1	ε	β_1, β_2
Model 2	ε	β

  - 2: **IF** Model 1 > Model 2, then can do classical test on parameters, as splitting does not say that group means should be significantly different, or which direction.

Priors and 2nd level analysis

- Posterior parameter estimates
  - do classical second level analyses
  - can use Hessians as weights

$$\text{E step:} \quad q_k(\theta) = \mathcal{N}(\mathbf{m}_k, \mathbf{S}_k)$$

$$\mathbf{m}_k \leftarrow \underset{\theta}{\operatorname{argmax}}\, p(\mathbf{a}_k|\theta)p(\theta|\zeta^{(i)})$$

$$\boxed{\mathbf{S}_k^{-1} \leftarrow \left.\frac{\partial^2 p(\mathbf{a}^k|\theta)p(\theta|\zeta^{(i)})}{\partial\theta^2}\right|_{\theta=\mathbf{m}_k}}$$

matlab: `[m,L,,,S]=fminunc(…)`

Regression

- Standard regression analysis:

$$\mathbf{m}_i = \mathbf{C}\mathbf{r}_i + \Sigma^{1/2}\boldsymbol{\eta} \qquad \forall i$$

- Including uncertainty about each subject's inferred parameters

$$\mathbf{m}_i = \mathbf{C}\mathbf{r}_i + (\Sigma^{1/2} + \mathbf{S}_i^{1/2})\boldsymbol{\eta} \qquad \forall i$$

Overview

- Formulate probabilistic model for choices
  - model fit: predictive probability
- ML / MAP
  - parameter inference
  - prior inferred from all joint data
- Empirical prior
  - Infer with approximate EM
  - second level analysis:
    - priors
    - individual posterior parameters

RL models

- Are no panacea
  - statistics about specific aspects of decision machinery
  - only account for part of the variance
- Model needs to match experiment
  - ensure subjects actually do the task the way you wrote it in the model
  - model comparison
- Model = Quantitative hypothesis
  - strong test
  - need to compare **models**, not **parameters**
  - includes all consequences of a hypothesis for choice