Bayesian (stable isotope) mixing models: MixSIAR

BRIAN STOCK

MAY 22, 2017

Calculate source % to a mixture

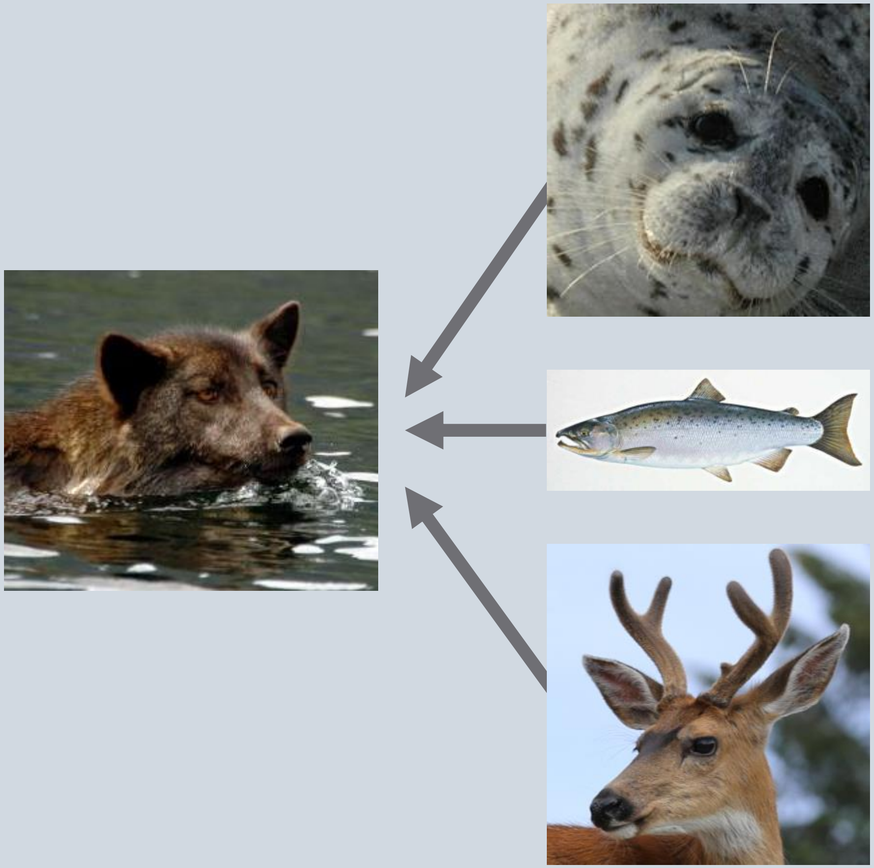

Calculate **prey** % to a **diet**

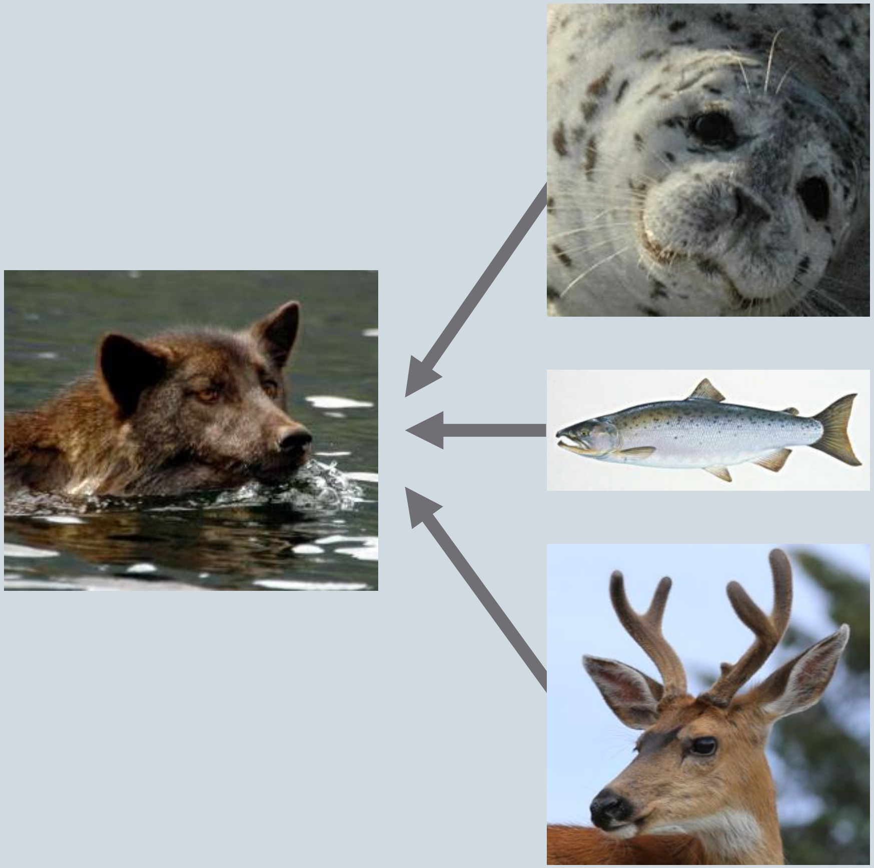

Semmens et al. (2009)

Calculate **colony** % to a **bird**

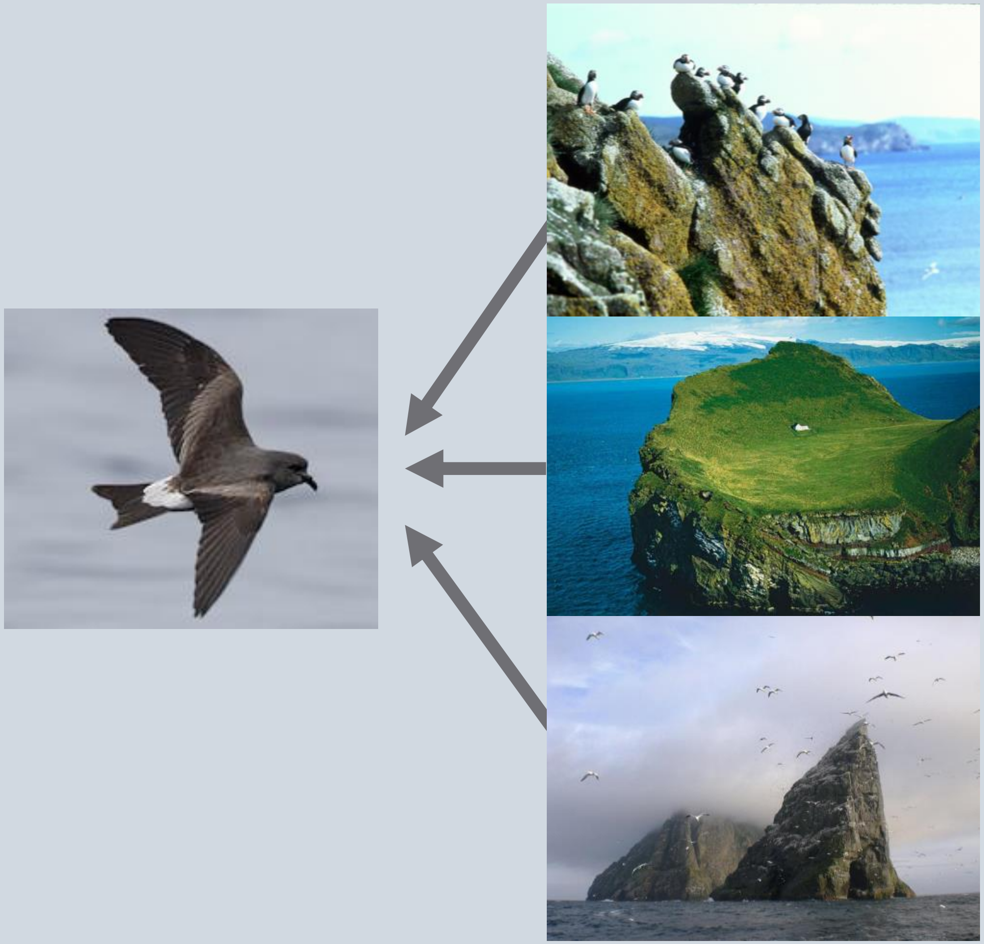

Bicknell et al. (2014)

Calculate **soil** % to a **sediment**

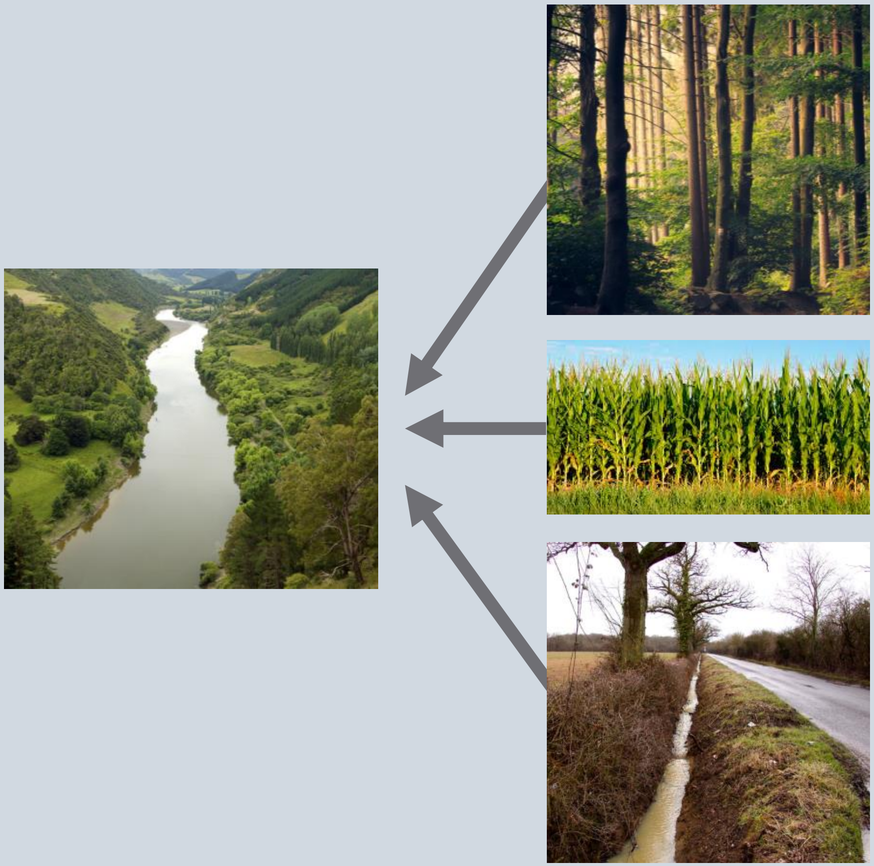

Scientists use mixing models a lot

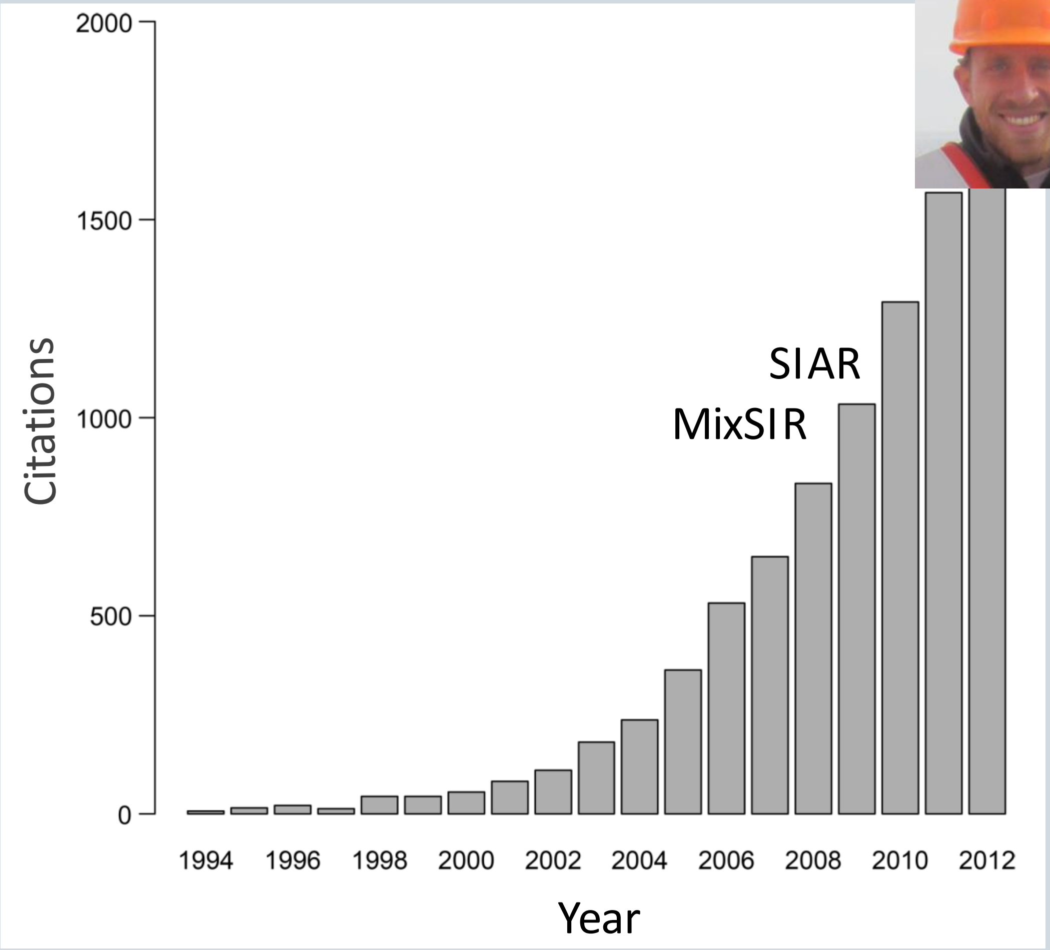

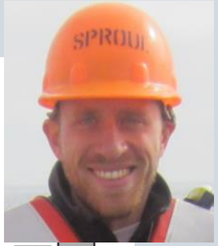

Phillips et al. (2014)

Calculate **prey** % to a **diet**

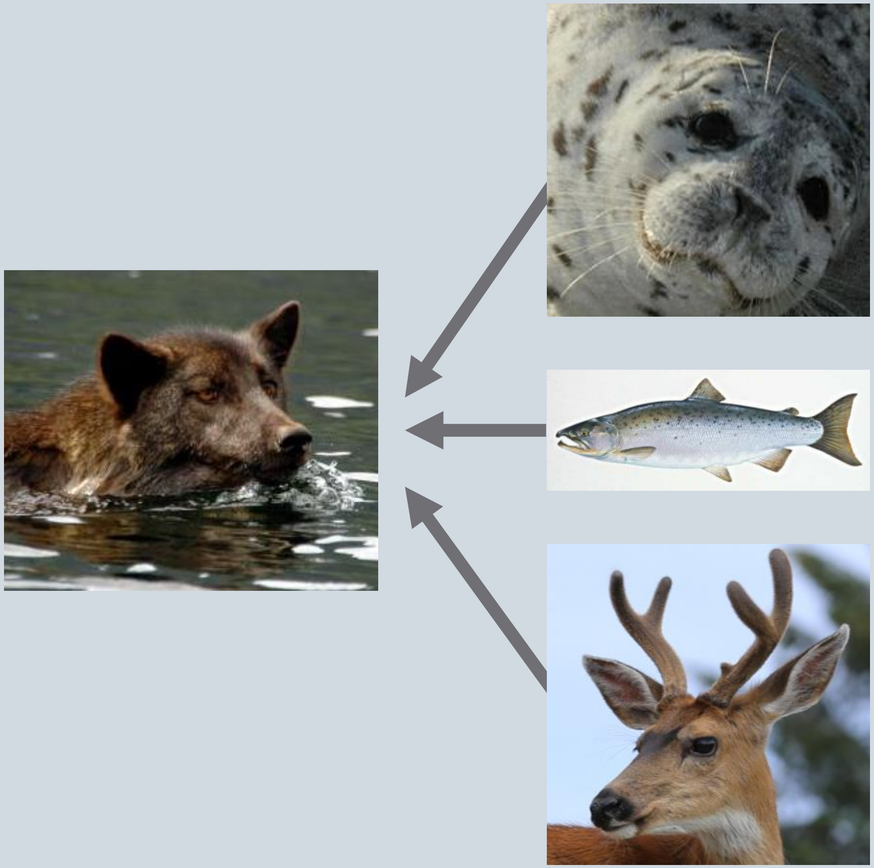

Calculate **prey** % to a **diet**

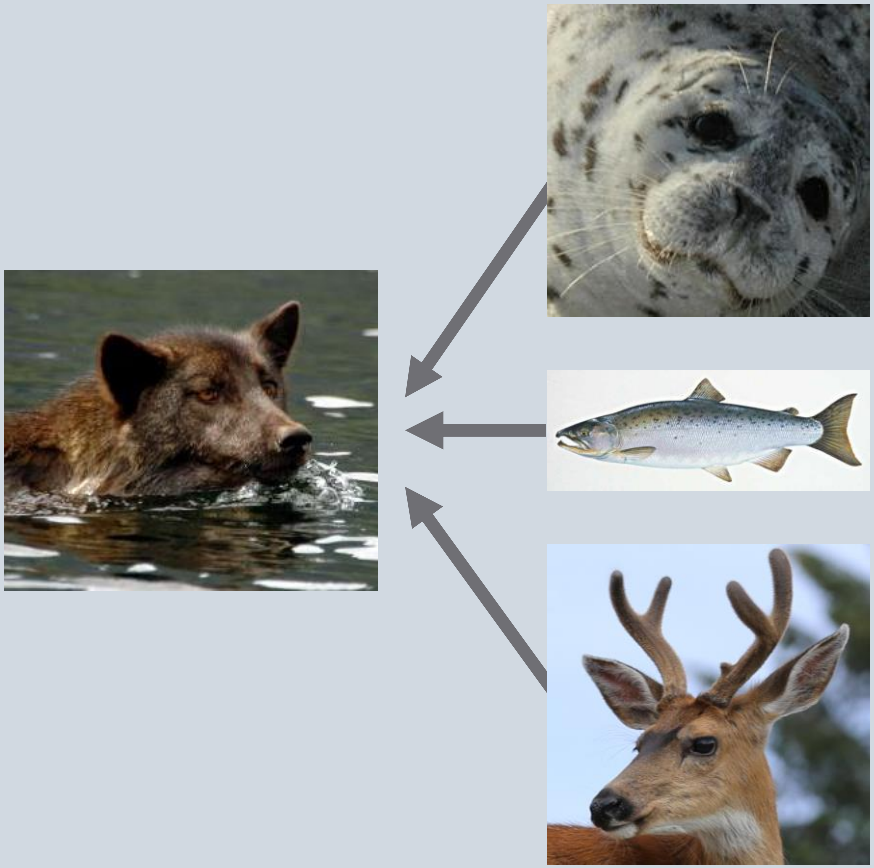

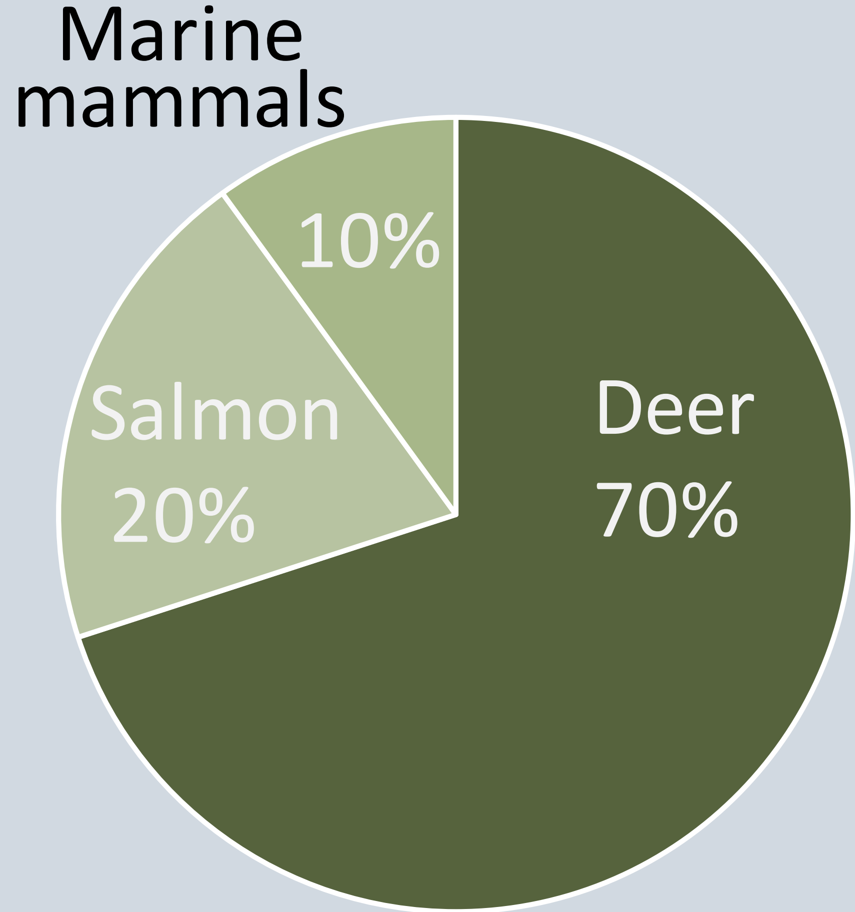

Calculate **prey** % to a **diet**

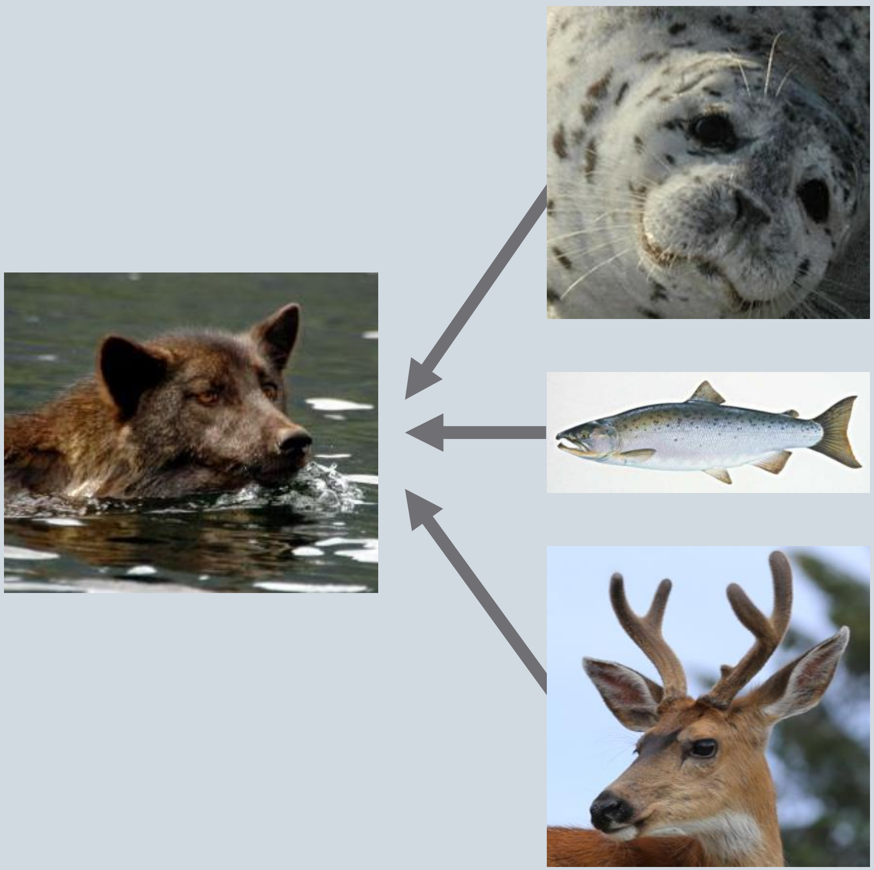

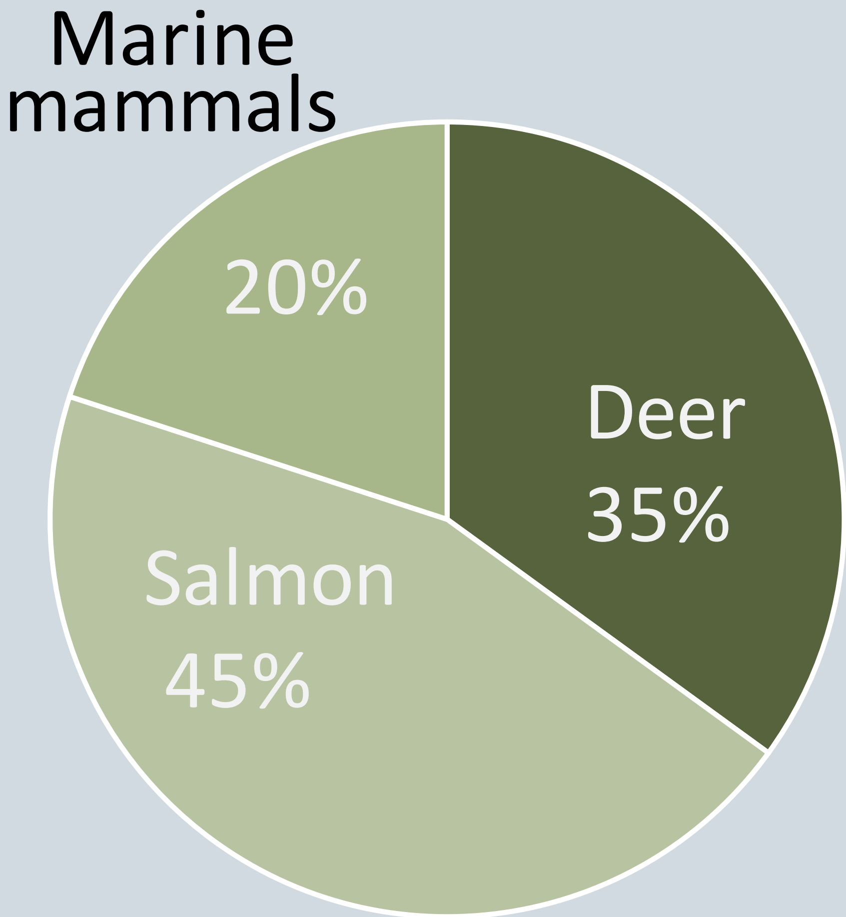

Using stable isotope data

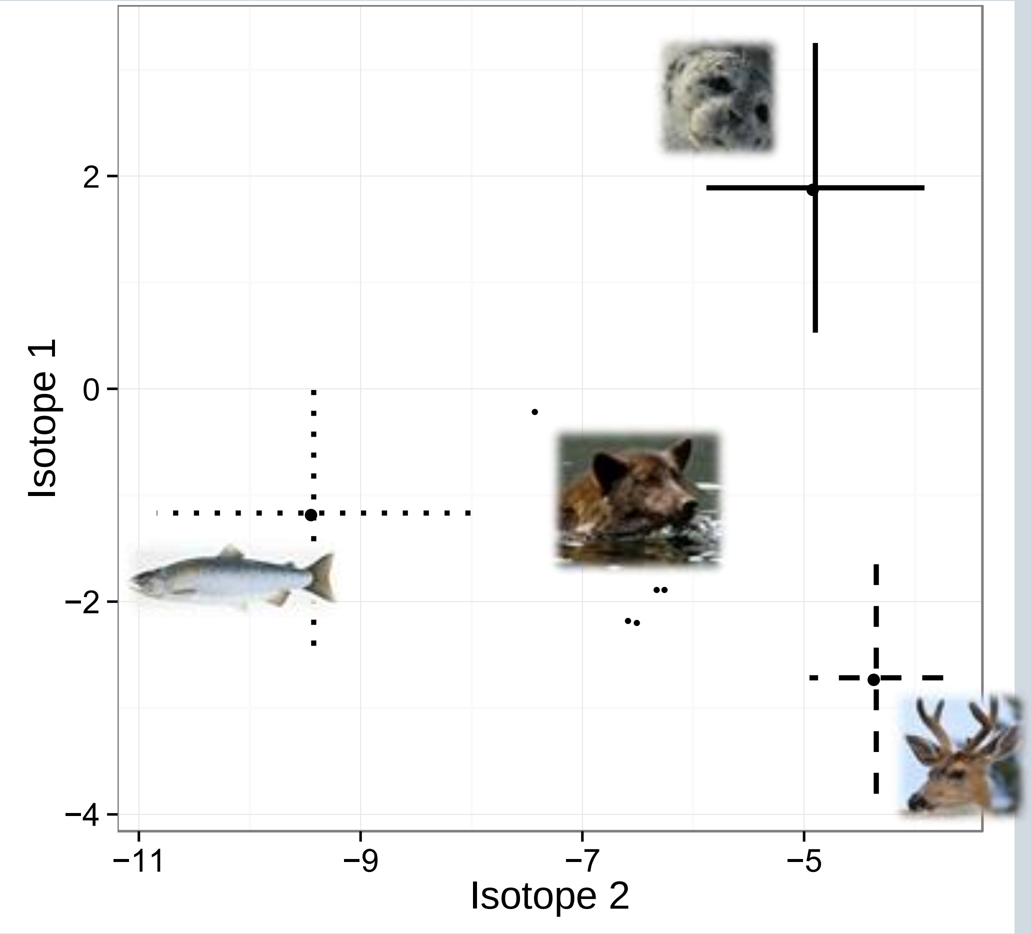

Using stable isotope data

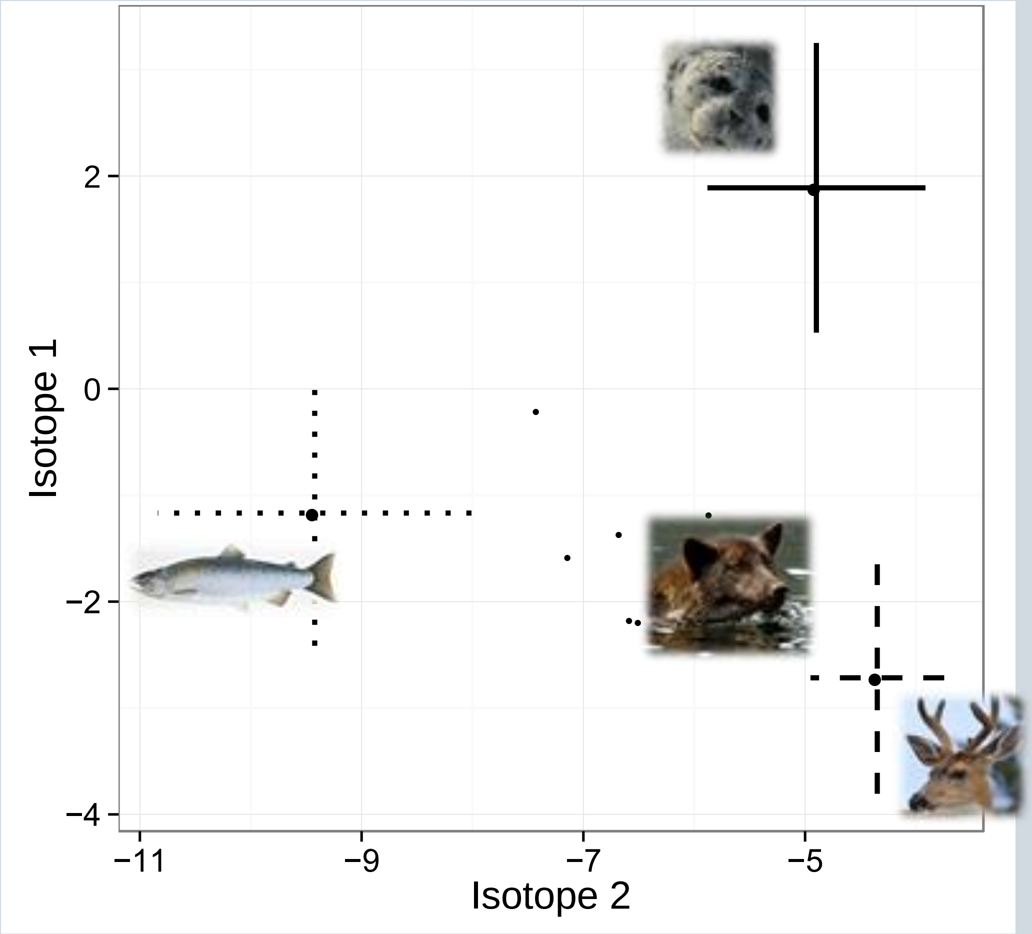

How mixing models work

Linear mixing model:

Diet = ?

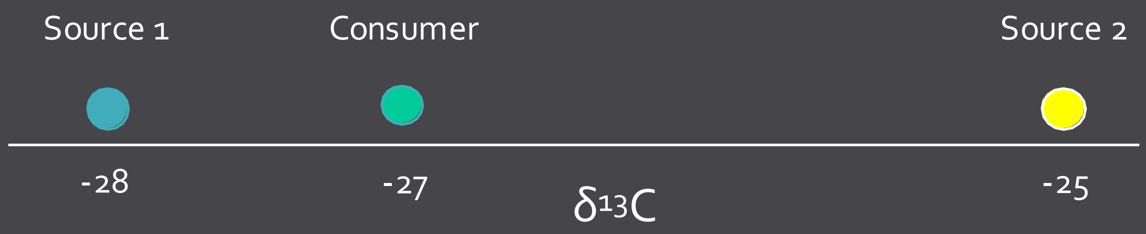

How mixing models work

Linear mixing model:

$$\text{Consumer} = p_1 * s_1 + p_2 * s_2 \qquad (p_1 + p_2 = 1)$$

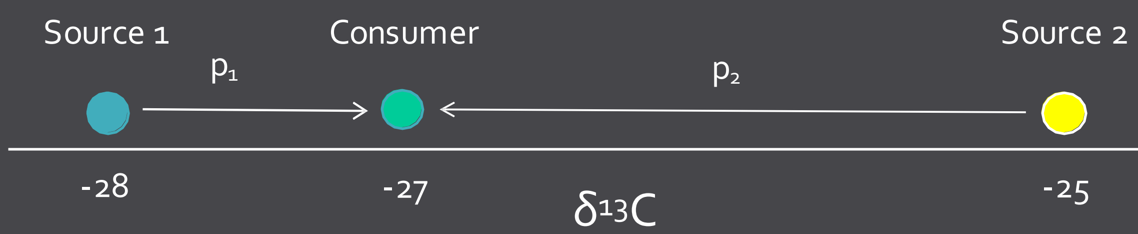

How mixing models work

Linear mixing model:

$$\text{Consumer} = 2/3\ s_1 + 1/3\ s_2$$

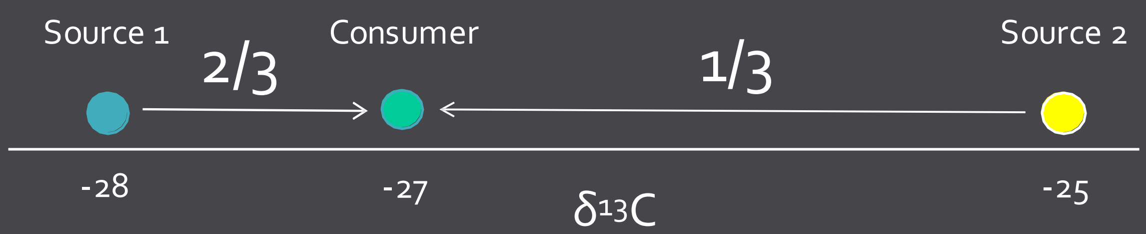

How mixing models work

Linear mixing model:

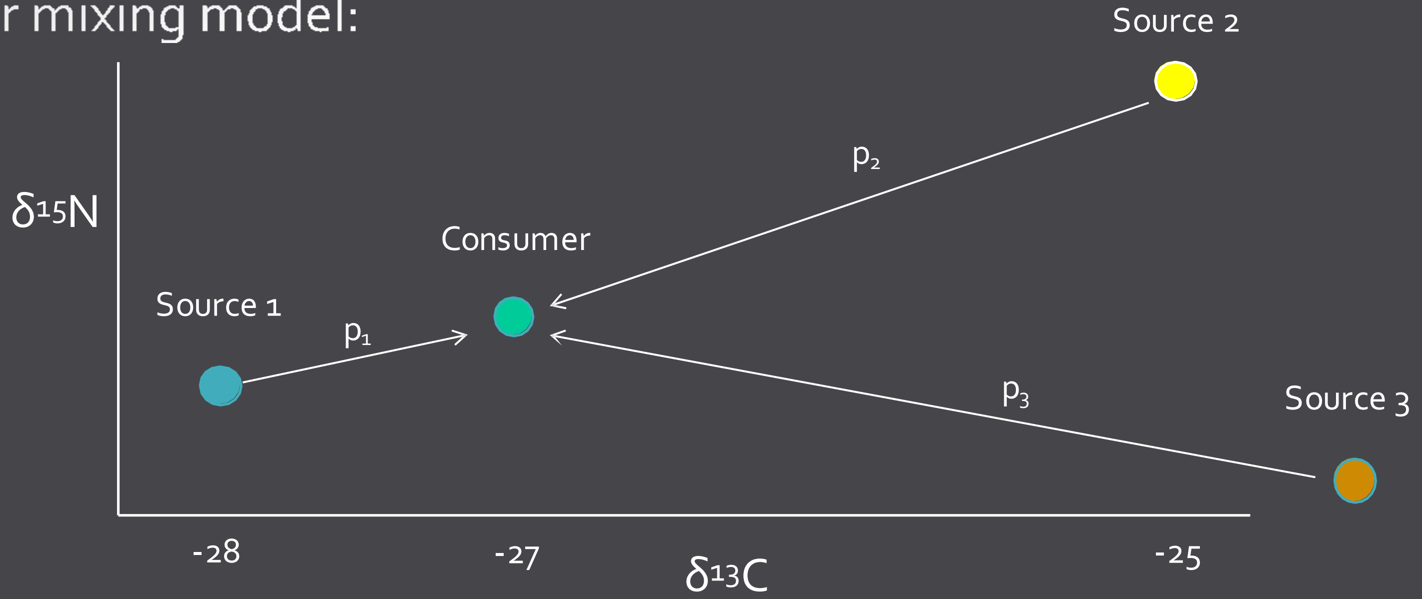

$$Consumer_C = p_1 s_{1C} + p_2 s_{2C} + p_3 s_{3C}$$

$$Consumer_N = p_1 s_{1N} + p_2 s_{2N} + p_3 s_{3N}$$

$$p_1 + p_2 + p_3 = 1$$

How mixing models work

~~Linear mixing model:~~

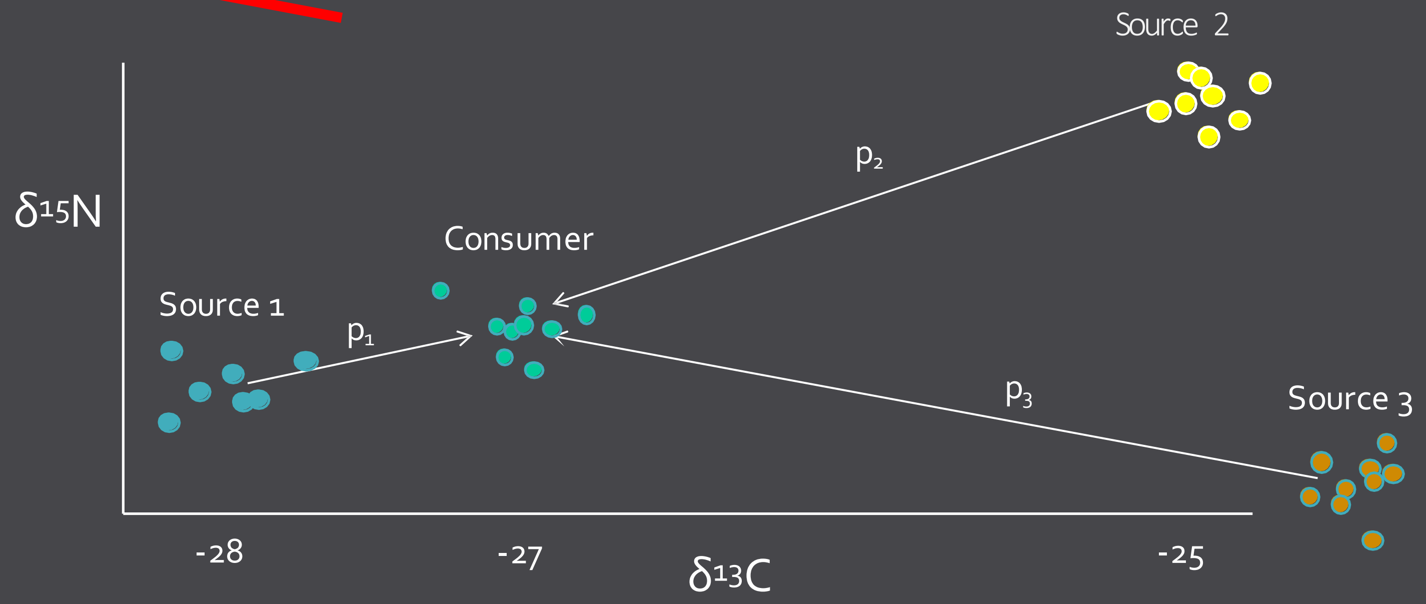

How mixing models work

~~Linear mixing model:~~

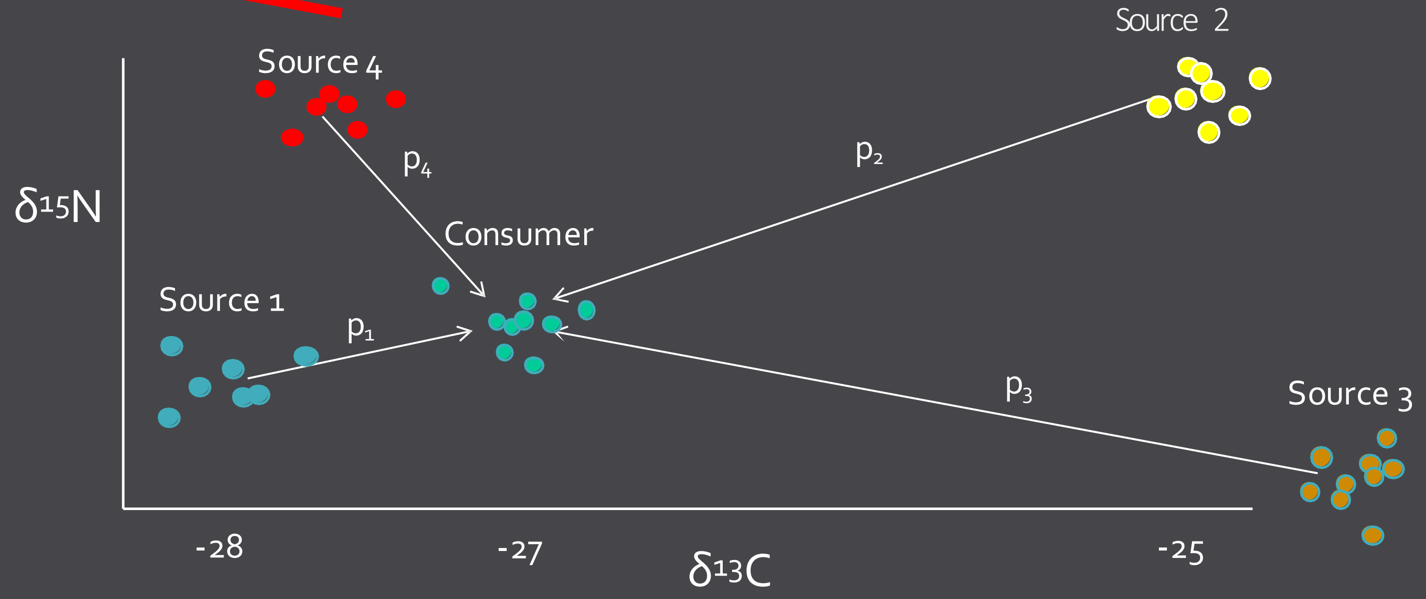

How mixing models work

~~Linear mixing model:~~

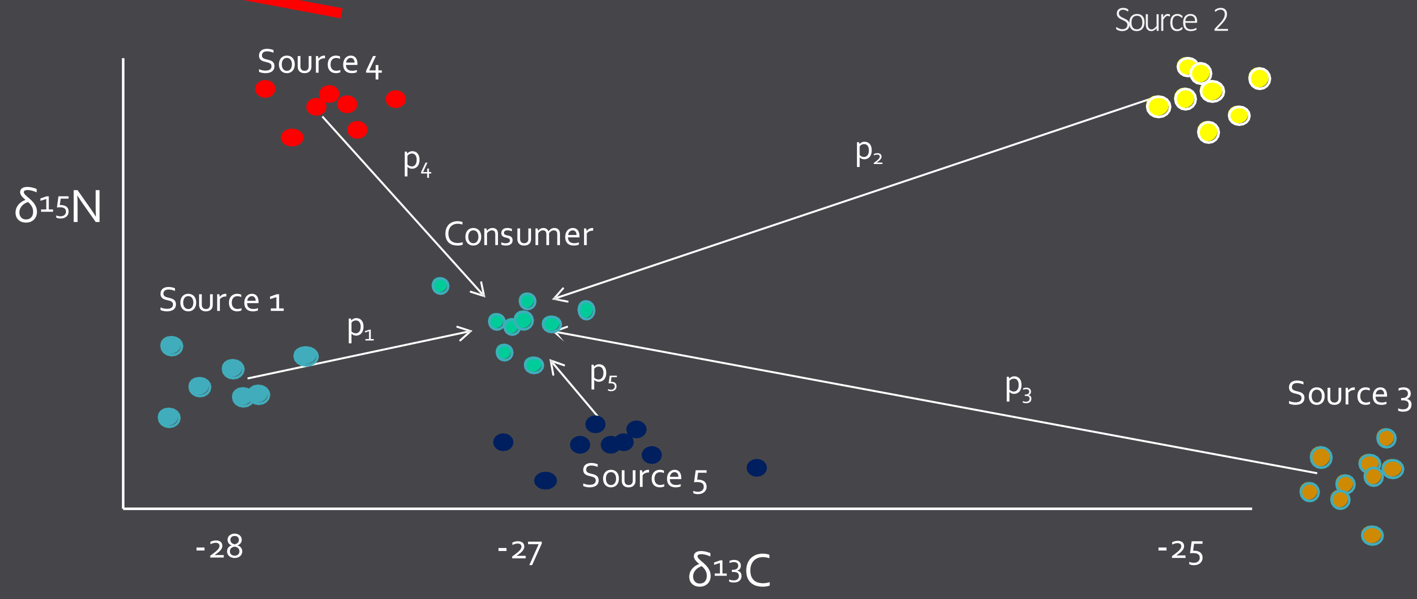

How mixing models work

Bayesian mixing model:

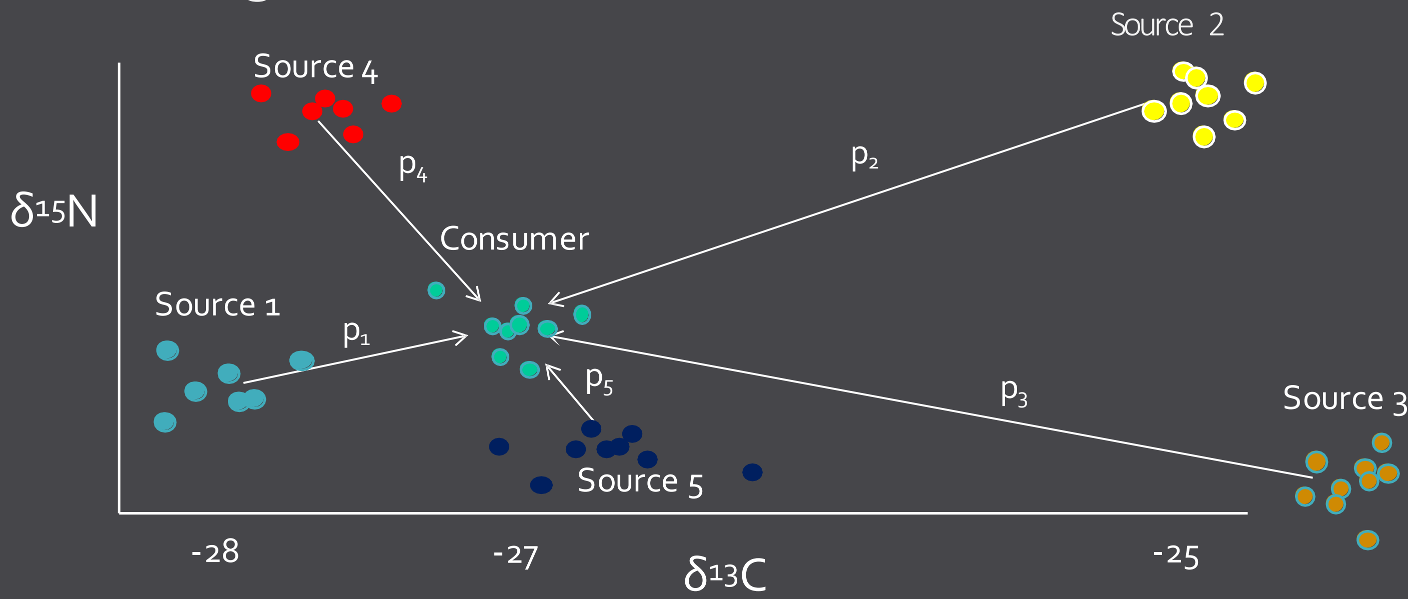

$$\mathbf{p} \sim Dir(\boldsymbol{\alpha})$$

$$Y_{ij} \sim MVN\left(\sum_k p_k \mu^s_{jk}, \Sigma\right)$$

$$Y^s_{jk} \sim MVN\left(\mu^s_{jk}, \Sigma^s_k\right)$$

$$\mu^s_{jk} \sim \mathcal{N}(0, .001)$$

$$\tau_{jk} \sim gamma(.001, .001)$$

How mixing models work

Bayesian mixing model:

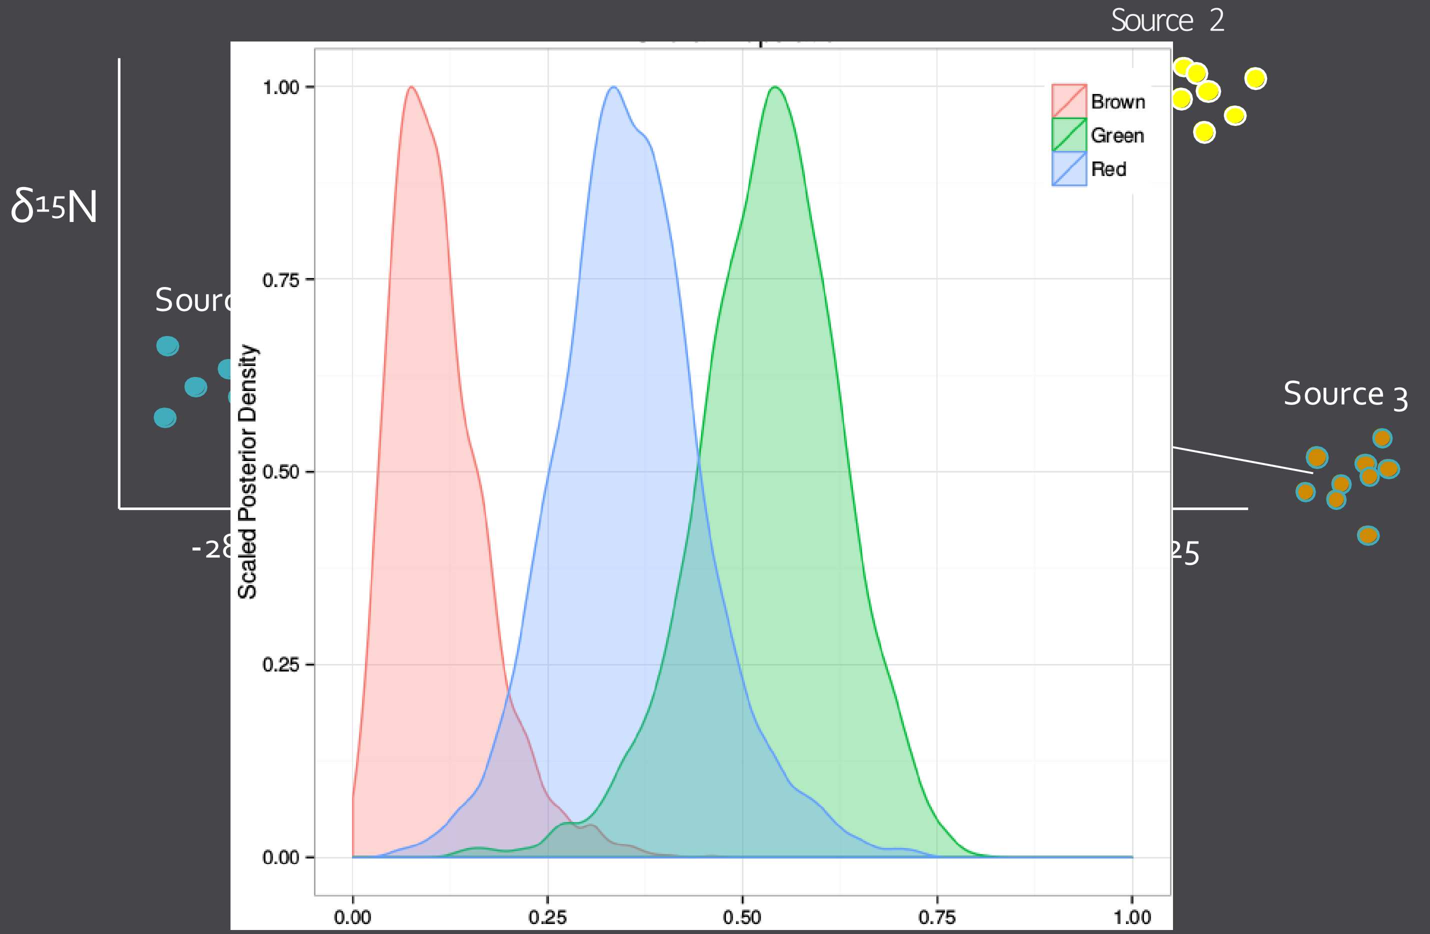

Bayesian models are better

1. Account for uncertainty in data (mix, source, TDF)
2. Solid statistical basis (likelihood)
3. Include additional info as priors
   - stomach/fecal contents
   - prey abundance

Bayesian models are better

1. Account for uncertainty in data (mix, source, TDF)
2. Solid statistical basis (likelihood)
3. Include additional info as priors
   - stomach/fecal contents
   - prey abundance

LETTER

Incorporating uncertainty and prior information into stable isotope mixing models

MixSIR 2008

Source Partitioning Using Stable Isotopes: Coping with Too Much Variation

Andrew C. Parnell[1], Richard Inger[2], Stuart Bearhop[2], Andrew L. Jackson[3*]

SIAR 2010

MixSIAR adds further improvements

1. Graphical User Interface (GUI)

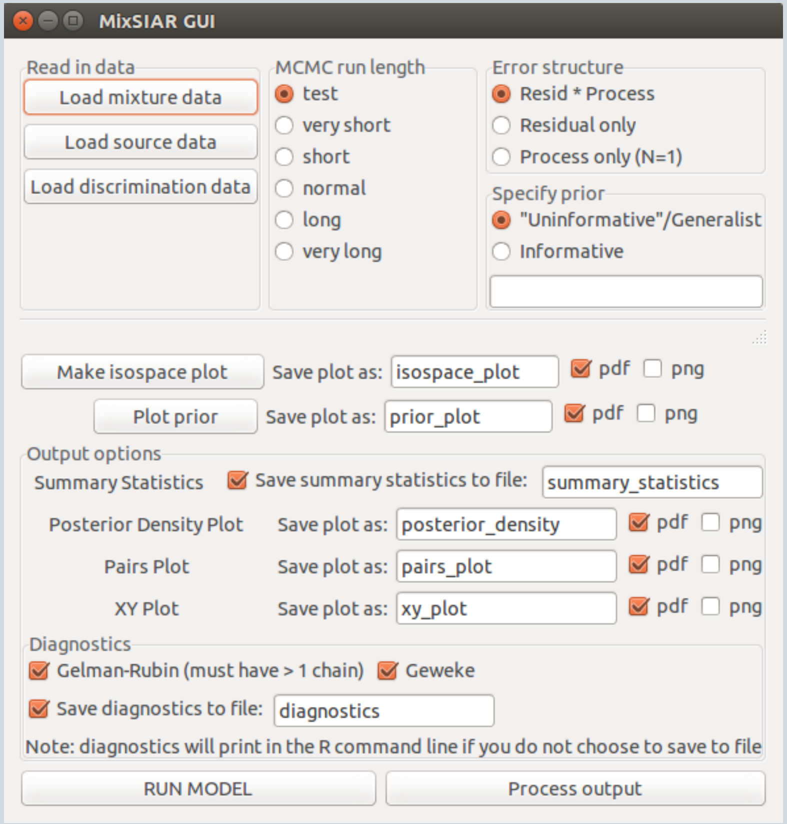

MixSIAR adds further improvements

1. Graphical User Interface (GUI)
2. Covariate effects

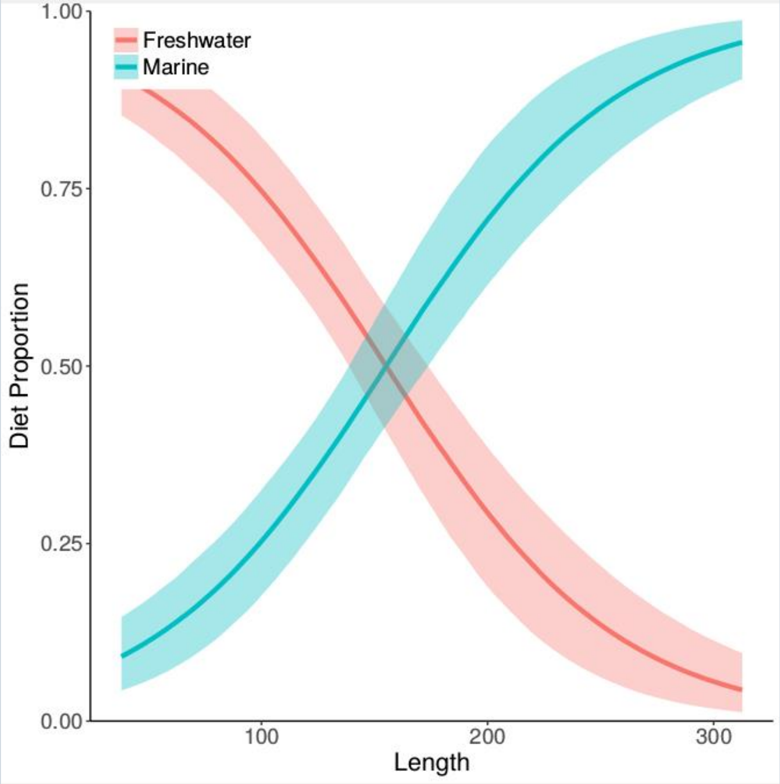

MixSIAR adds further improvements

1. Graphical User Interface (GUI)
2. Covariate effects
3. Fit source data within model

$$Y_{jk}^{s} \sim MVN\left(\mu_{jk}^{s}, \Sigma_{k}^{s}\right)$$

$$\mu_{jk}^{s} \sim \mathcal{N}(0, .001)$$

$$\tau_{jk} \sim gamma(.001, .001)$$

MixSIAR adds further improvements

1. Graphical User Interface (GUI)
2. Covariate effects
3. Fit source data within model
4. Better error structure(s)

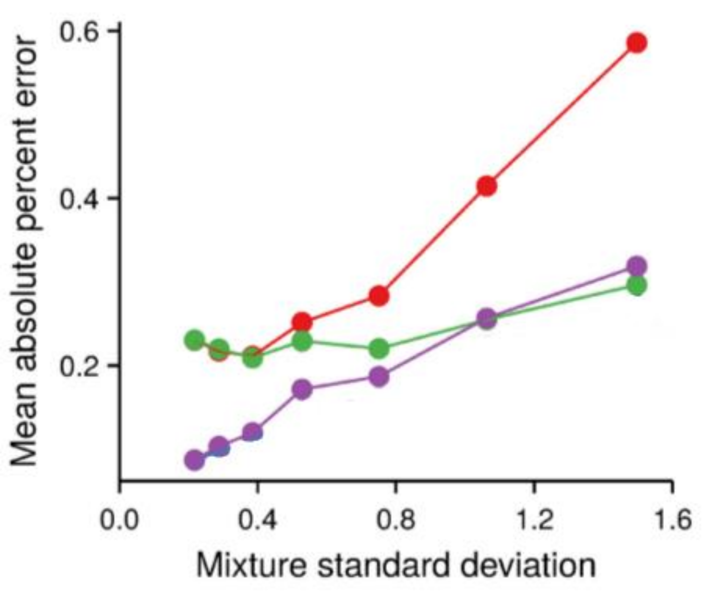

MixSIAR adds further improvements

1. Graphical User Interface (GUI)
2. Covariate effects
3. Fit source data within model
4. Better error structure(s)
5. Plot/modify your prior

Covariate effects in MixSIAR

No covariate effects…

$$\boldsymbol{p} = [20\%, 50\%, 20\%\ 10\%]$$

Assumes that all consumers have the *same diet*

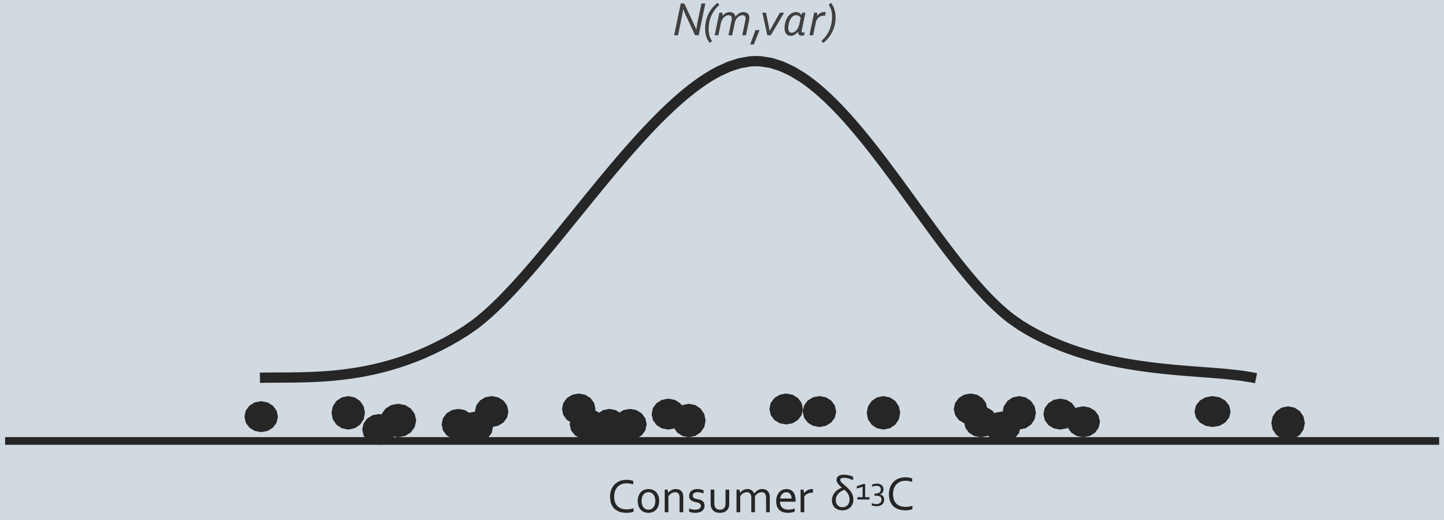

Covariate effects in MixSIAR

Transform p's

Linear regression in ILR-space

$$\mathbf{p} \sim Dir(\boldsymbol{\alpha})$$

$$p_{ik} = inverseILR(ilr.global_k + ilr.fac1_{mk} + ilr.cont1_k Cont1_i)$$

Covariate effects in MixSIAR

Transform p's

Linear regression in ILR-space

$$\mathbf{p} \sim Dir(\boldsymbol{\alpha})$$

$$p_{ik} = inverseILR(\underbrace{ilr.global_k}_{\text{Intercept/mean}} + \underbrace{ilr.fac1_{mk}}_{\text{Fixed/random effect}} + \underbrace{ilr.cont1_k}_{\text{Continuous effect ("slope")}} Cont1_i)$$

Fixed effects

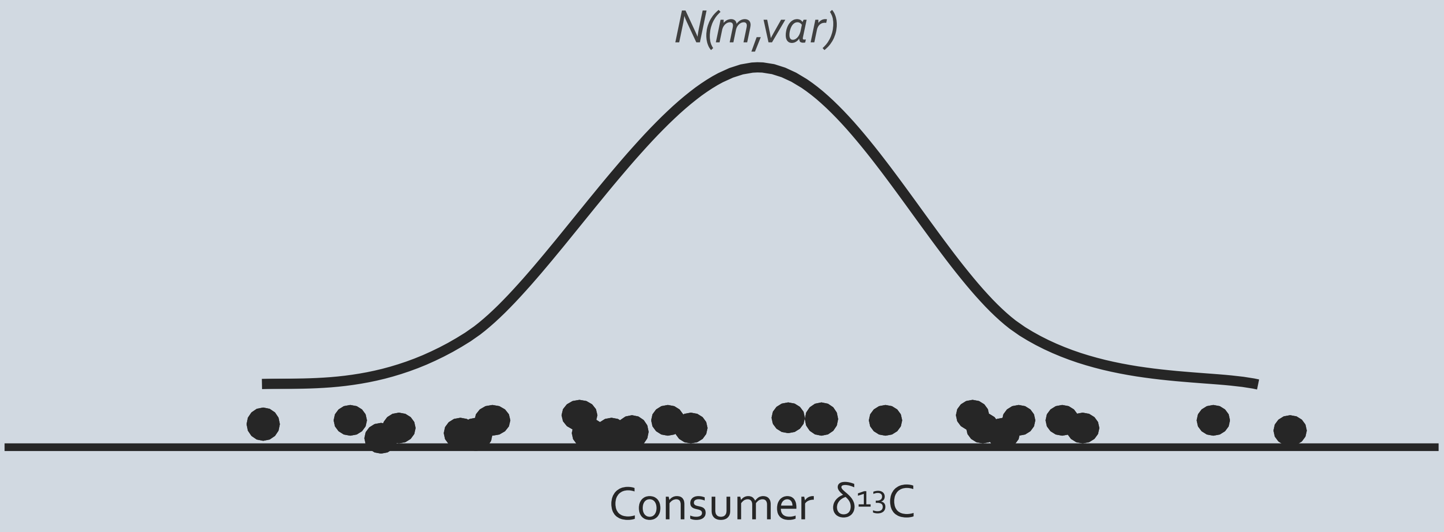

1. Covariate effects

Fixed effects

Simplest = estimate mean for different groups independently

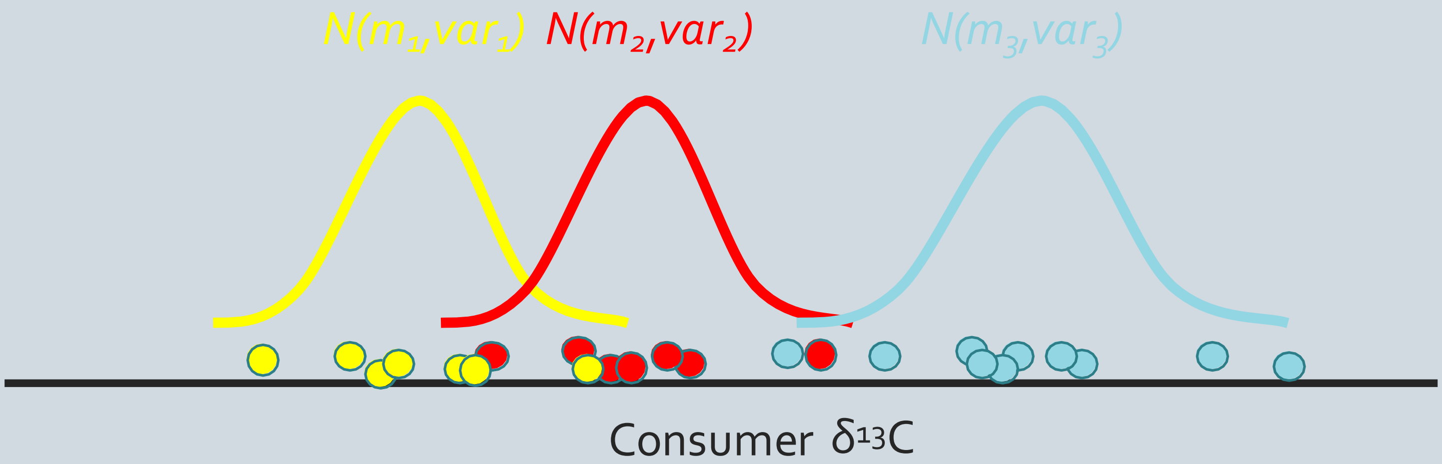

1. Covariate effects

Random effects

More complex

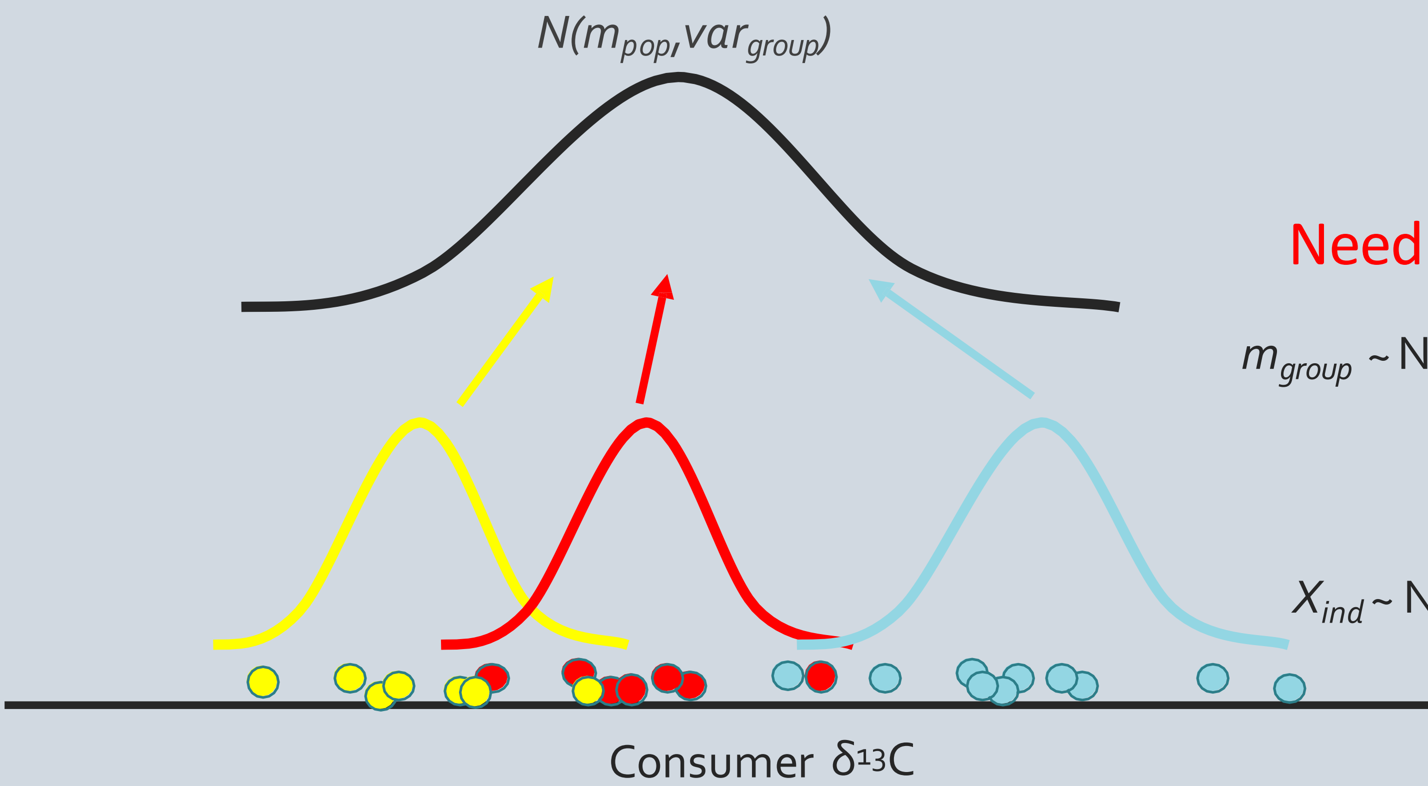

Need 2+ groups

$m_{group} \sim N(m_{pop}, \boxed{var_{group}})$

$X_{ind} \sim N(m_{group}, var_{ind})$

1. Covariate effects

(Nested) Random effects: Wolves Ex

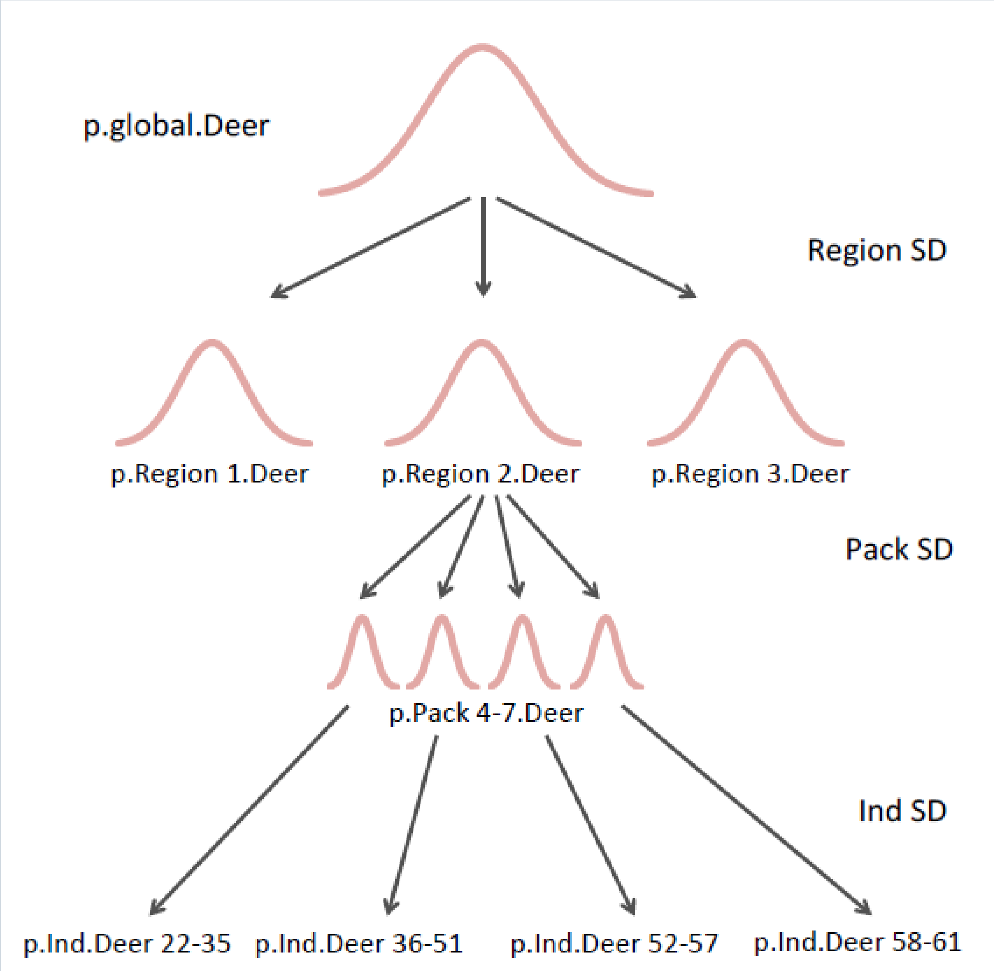

3 Regions

8 Packs

64 Wolves

Semmens et al. (2009)

(Nested) Random effects: Wolves Ex

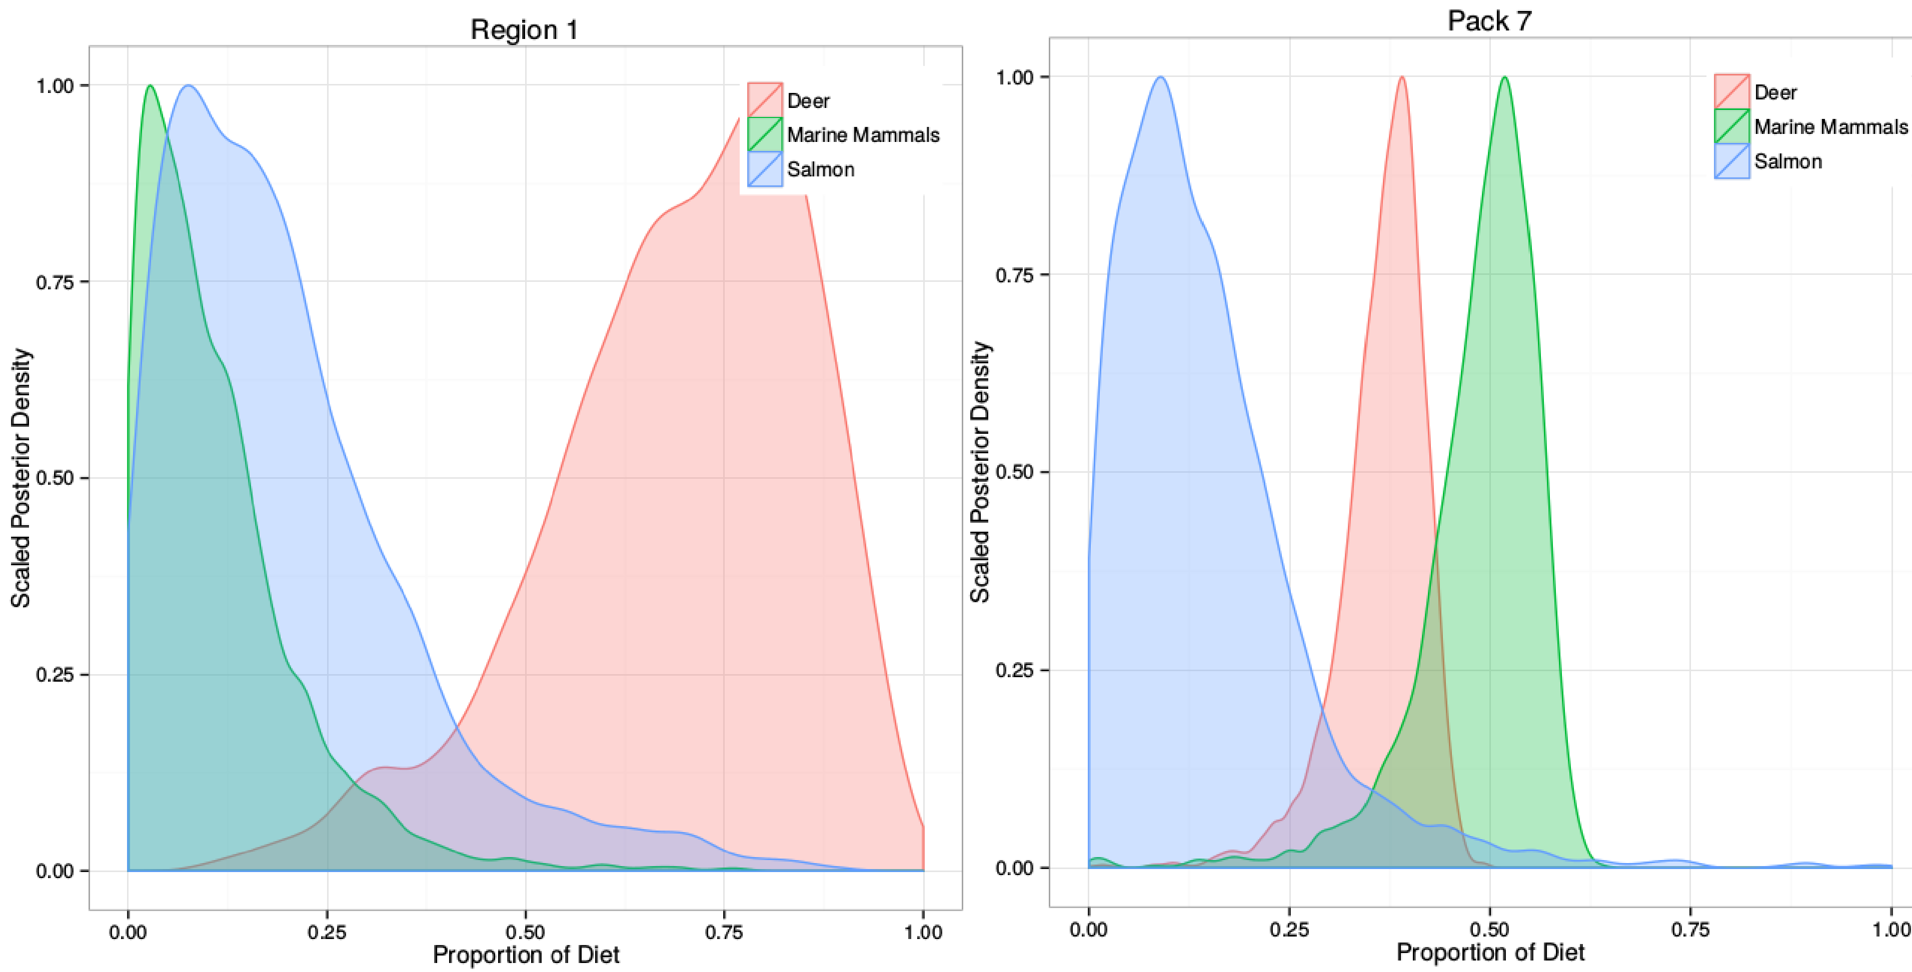

1. Covariate effects

Semmens et al. (2009)

Continuous effect: Alligator length

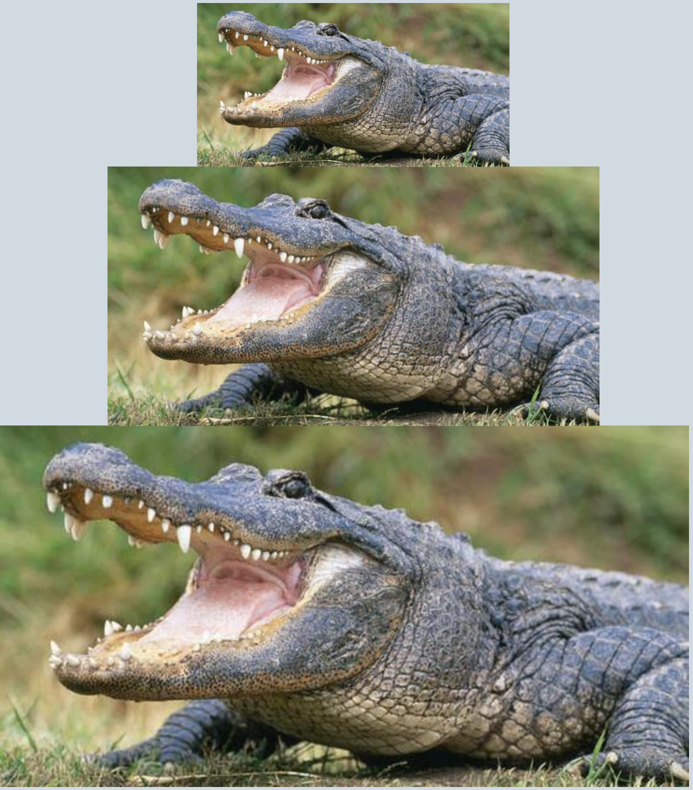

Freshwater

Marine

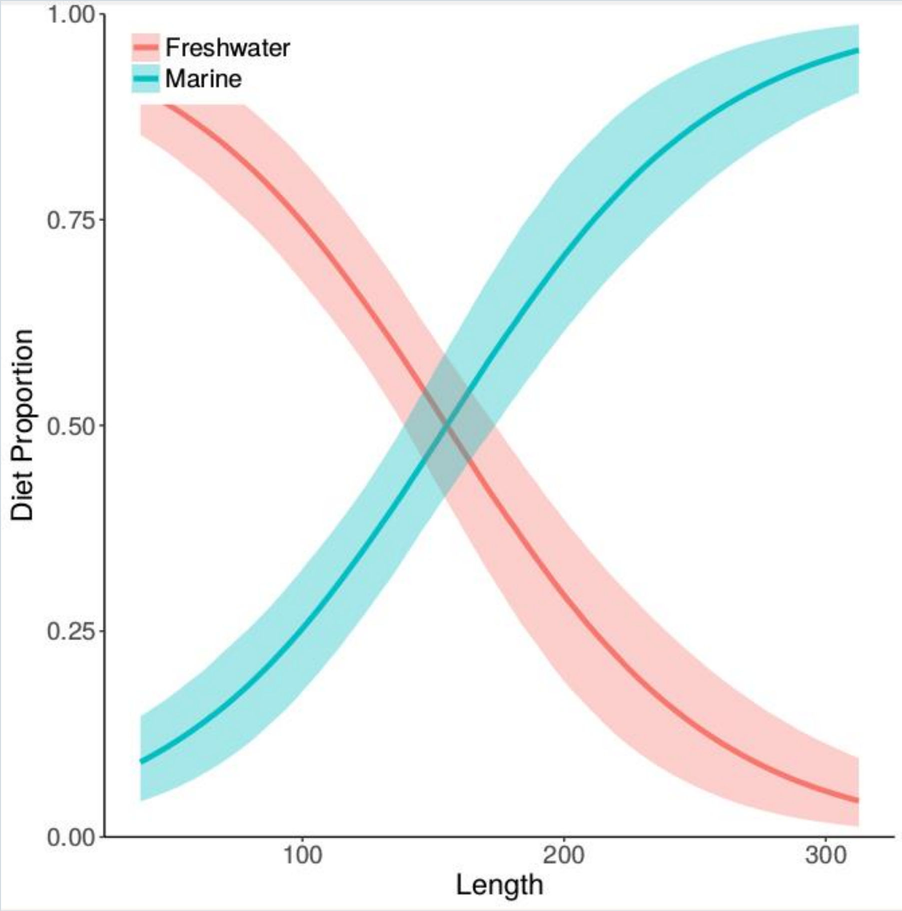

Nifong et al. (2015)

Fit source data within model

Ask me later (>1 way to do this)

Boring to talk about… but reduces error

$$Y^s_{jk} \sim MVN\left(\mu^s_{jk}, \Sigma^s_k\right)$$

$$\mu_{jk} \sim \mathcal{N}\left(m_{jk}, n_k/s^2_{jk}\right),$$

$$tmp.X_{jk} \sim \chi^2(n_k),$$

$$\tau_{jk} = \frac{tmp.X_{jk}}{s^2_{jk}\left(n_k - 1\right)},$$

$$\Sigma^s_k = diag\left(\frac{1}{\tau_{jk}}\right),$$

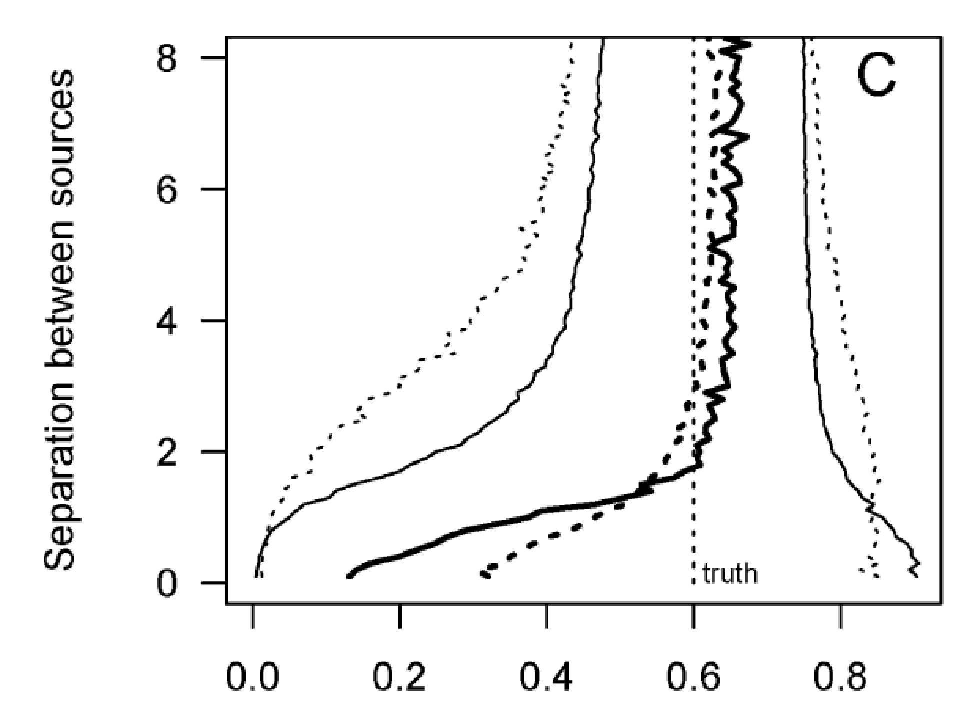

Previous models unrealistic

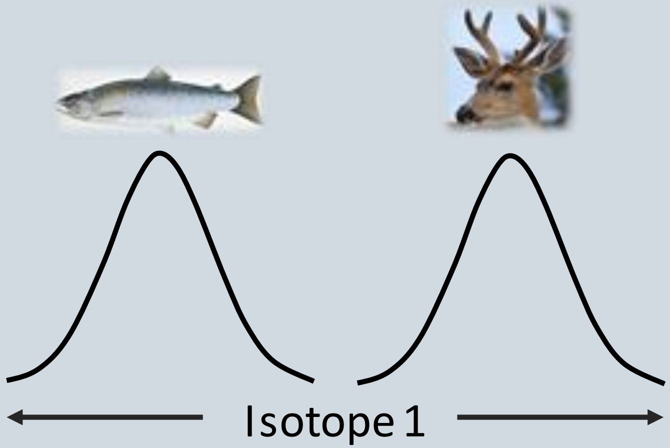

3. Better error structures

Stock and Semmens (2016)

Previous models unrealistic

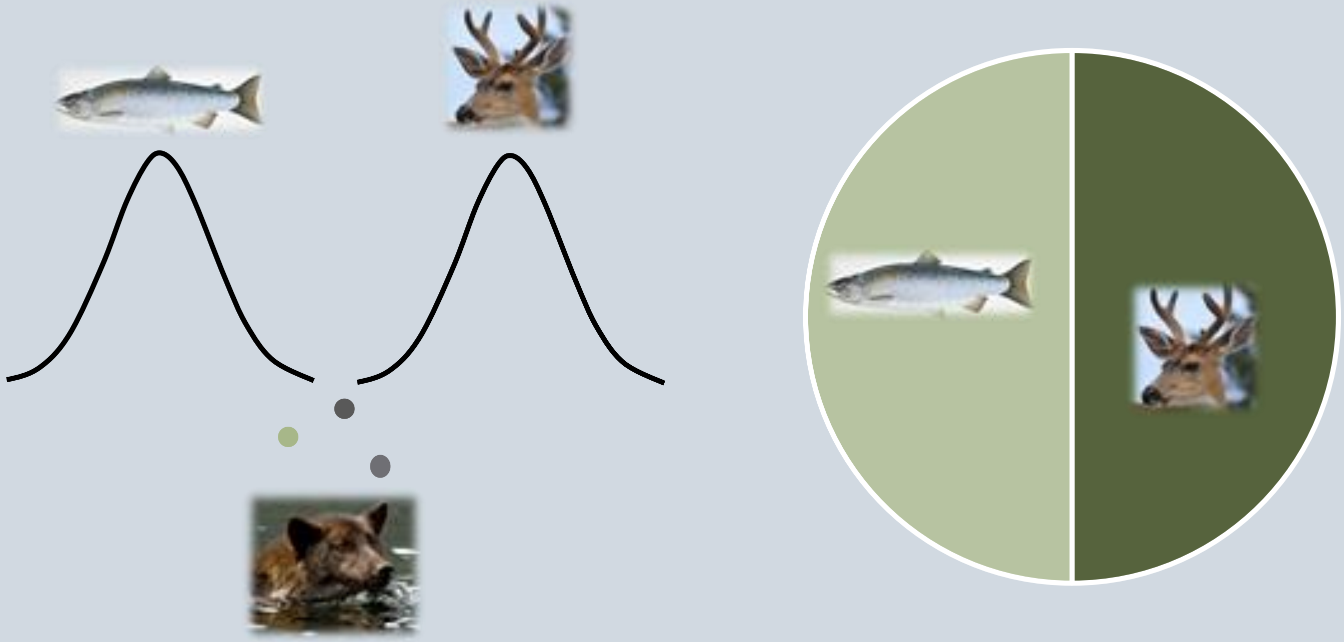

3. Better error structures

Stock and Semmens (2016)

Previous models unrealistic

$n = 1$

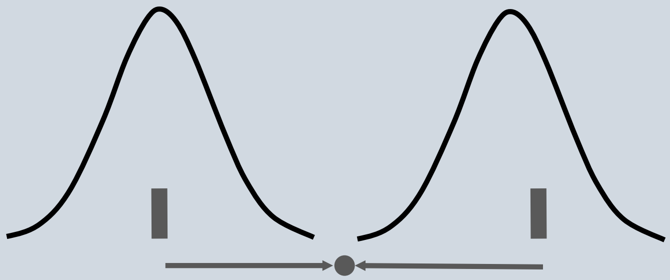

Stock and Semmens (2016)

Previous models unrealistic

$n = 1$

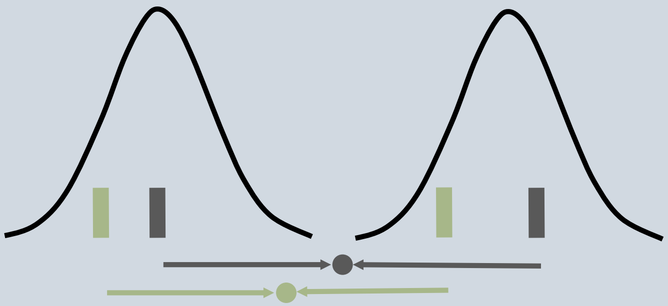

Stock and Semmens (2016)

Previous models unrealistic

$n = 1$

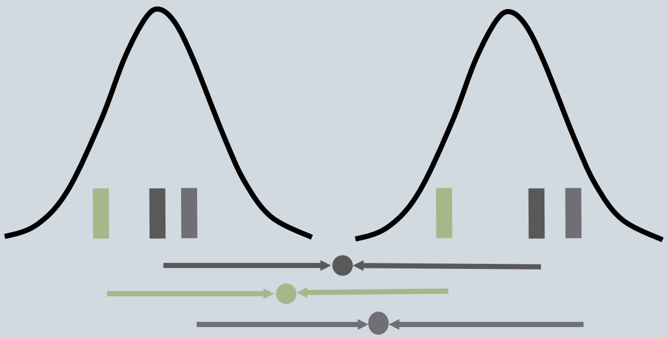

Stock and Semmens (2016)

Previous models unrealistic

$n = 1$

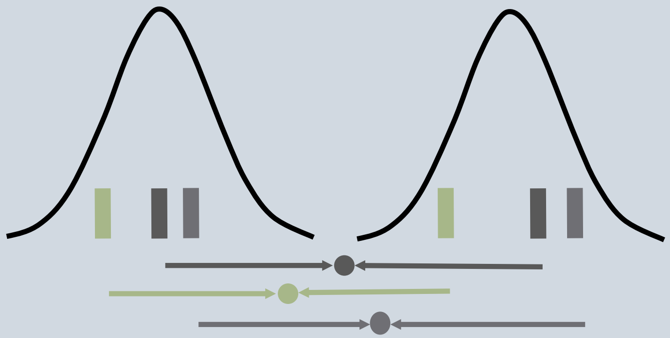

$n = \infty$

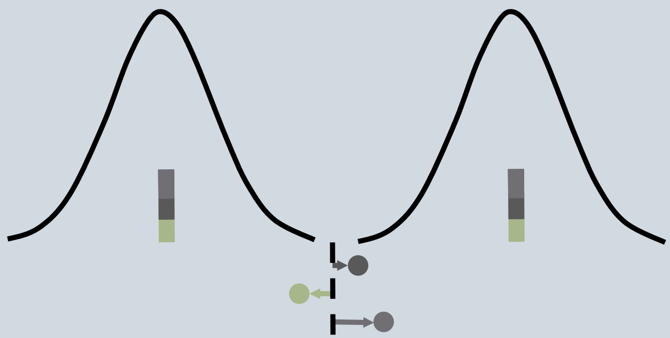

Stock and Semmens (2016)

Previous models unrealistic

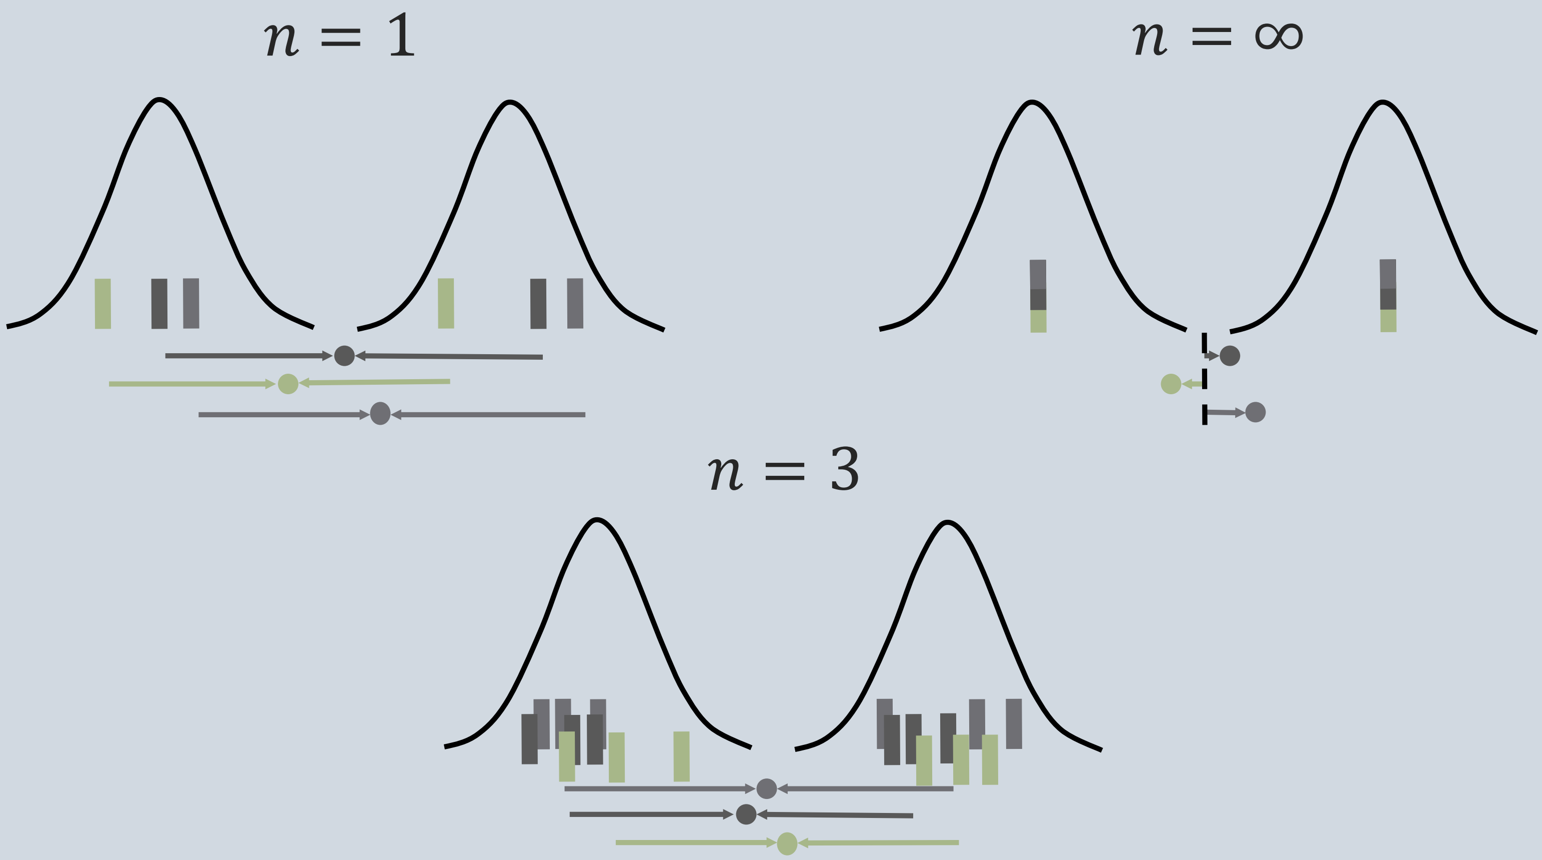

Stock and Semmens (2016)

Previous models unrealistic

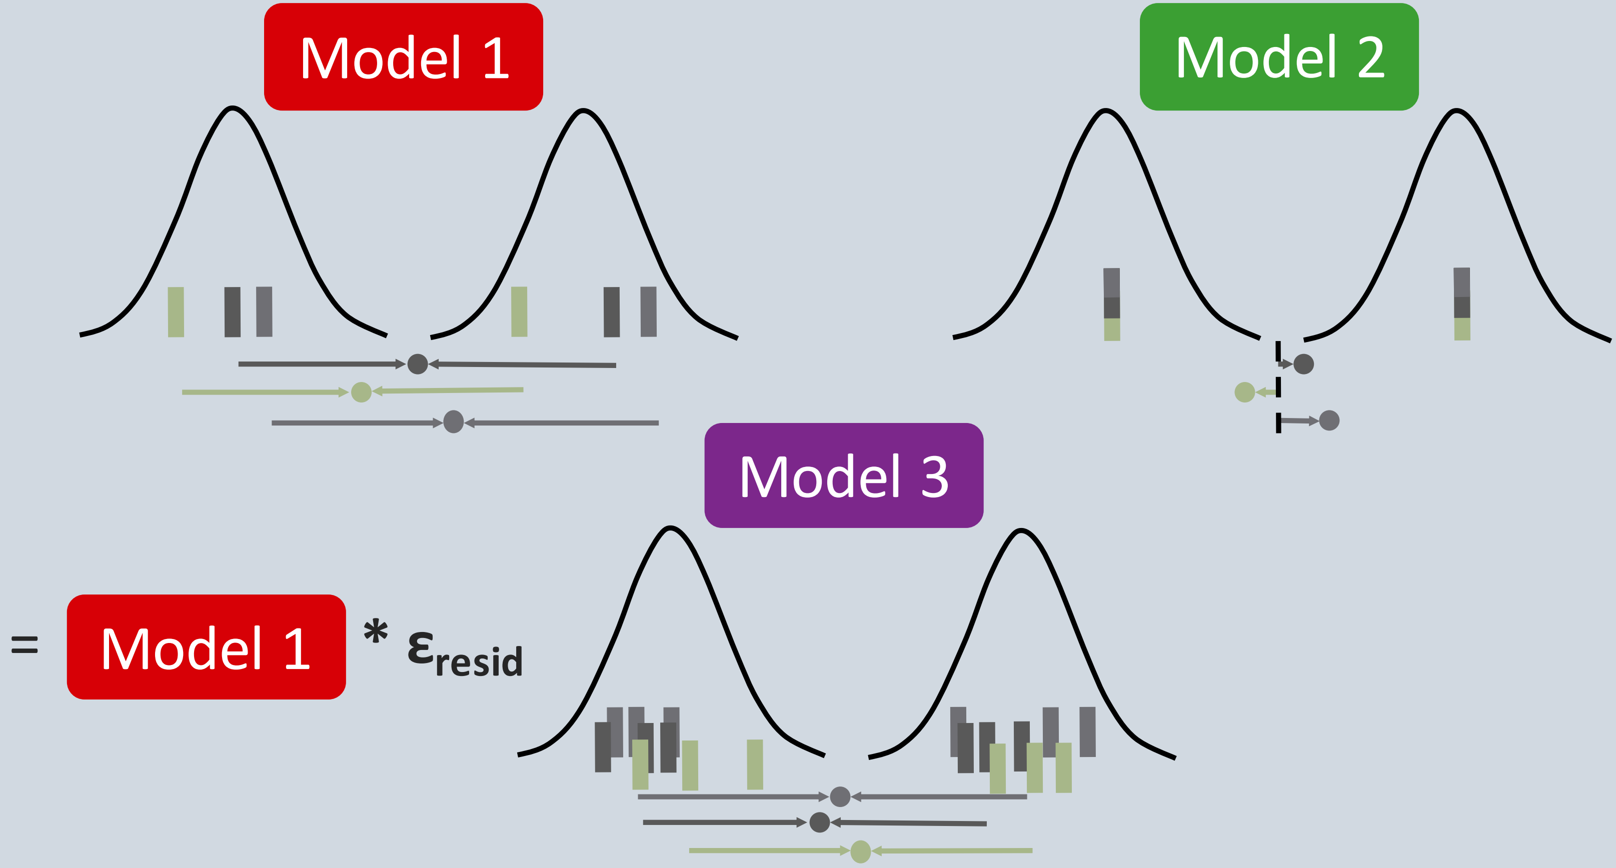

3. Better error structures

Stock and Semmens (2016)

Ecologically meaningful

$$\frac{1}{\varepsilon_{\text{resid}}} \propto \text{consumption}$$

Stock and Semmens (2016)

Ecologically meaningful

$$\frac{1}{\varepsilon_{resid}} \propto \text{consumption}$$

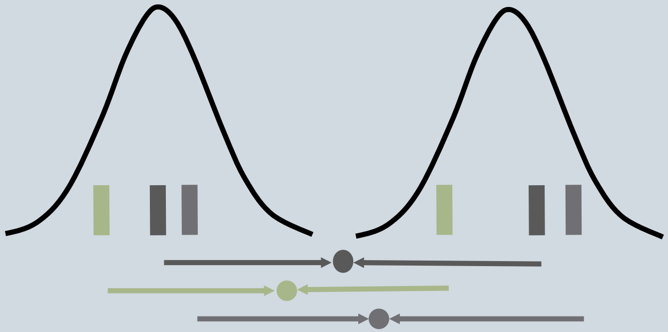

C = 1

3. Better error structures

Stock and Semmens (2016)

Ecologically meaningful

$$\frac{1}{\varepsilon_{resid}} \propto \text{consumption}$$

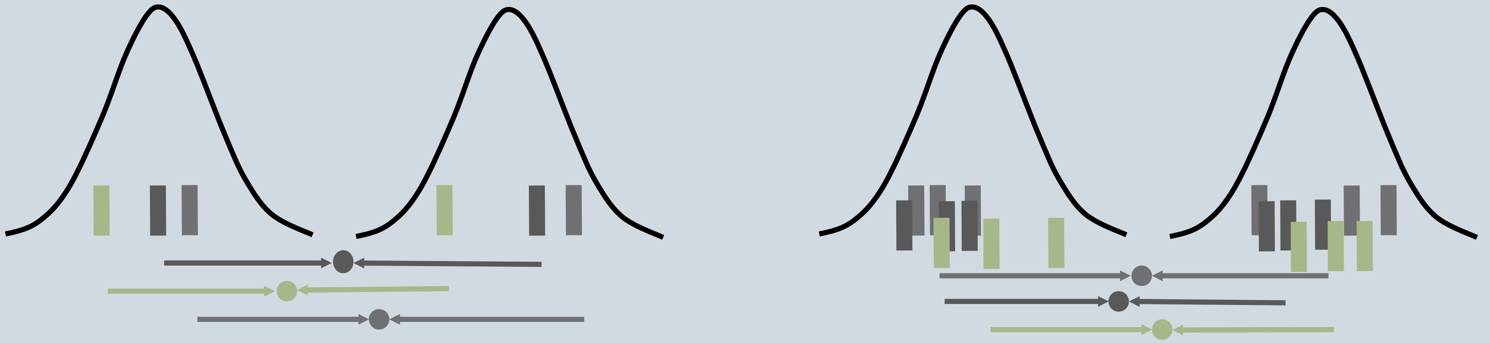

C = 1

C = 3

Stock and Semmens (2016)

Ecologically meaningful

$$\frac{1}{\varepsilon_{resid}} \propto \text{consumption}$$

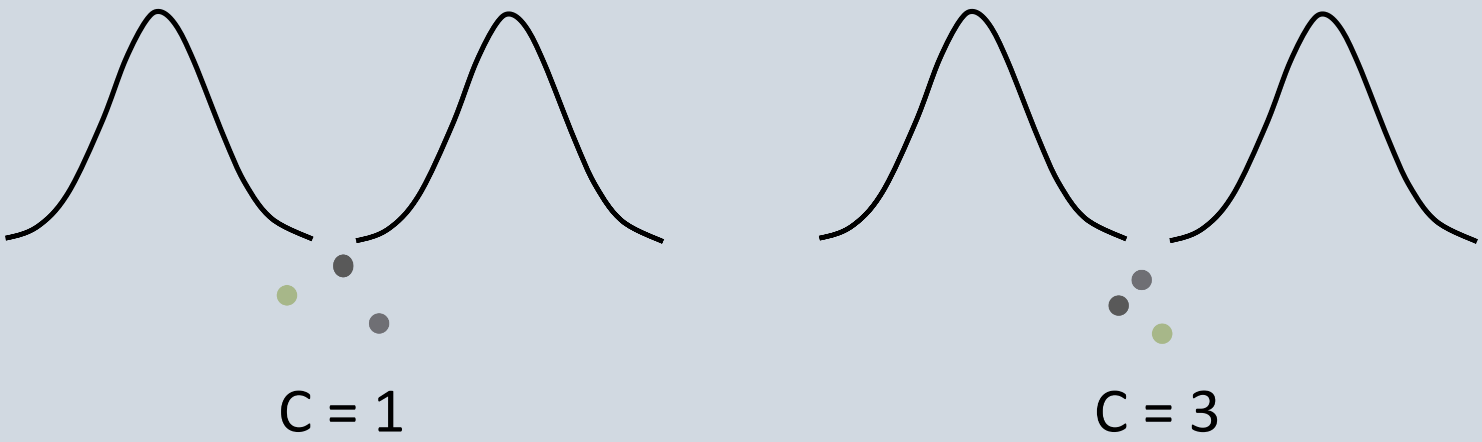

Stock and Semmens (2016)

Ecologically meaningful

$$\frac{1}{\varepsilon_{resid}} \propto \text{consumption}$$

		Tissue	ε_{avg}
	Oyster	Muscle	0.23
	Human	Bone	1.78

Stock and Semmens (2016)

Confounding of ε_{resid}

Depends on:

- Inclusion of covariates that explain variability
- TDF variance (rarely known)

Confounding of ε_{resid}

Depends on:

- **Inclusion of covariates that explain consumer variability**
- TDF variance (rarely known)

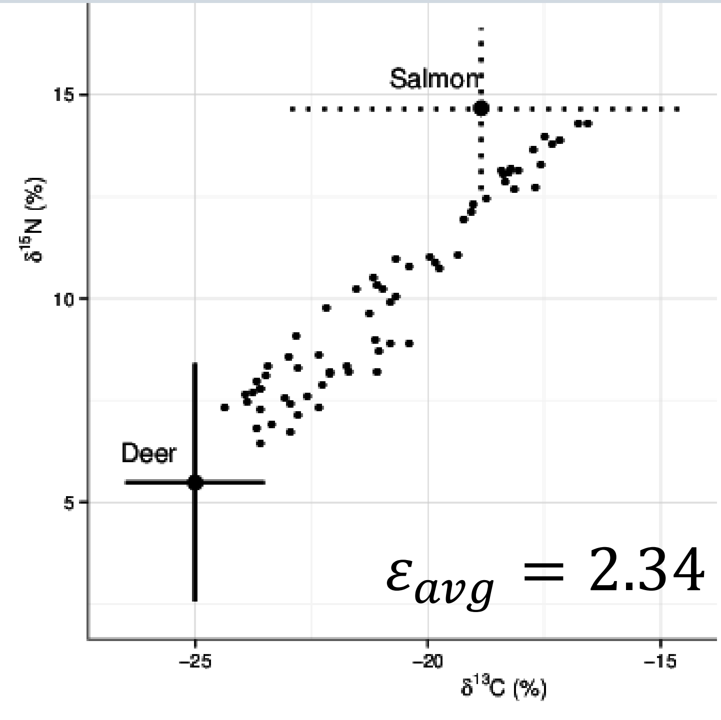

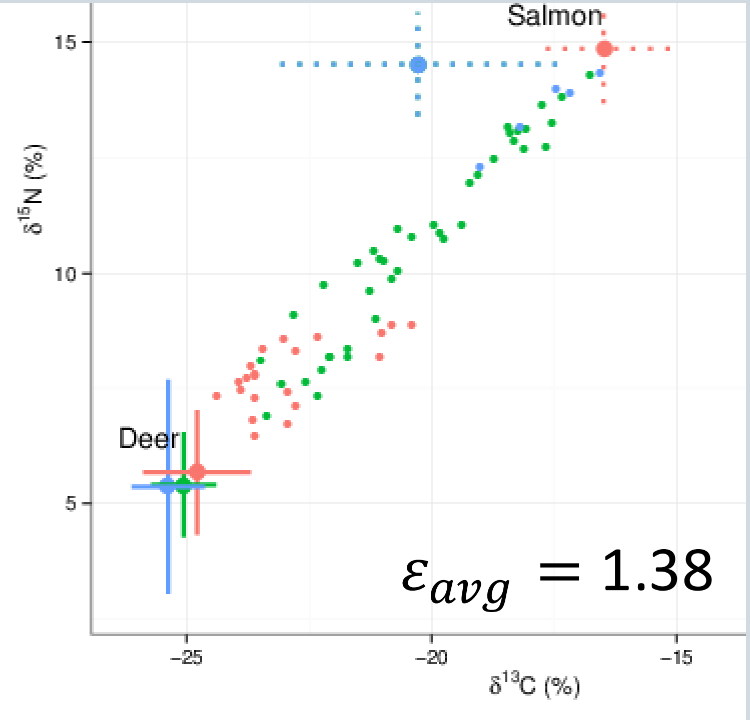

3. Better error structures

Confounding of ε_{resid}

Depends on:

- **Inclusion of covariates that explain consumer variability**
- TDF variance (rarely known)

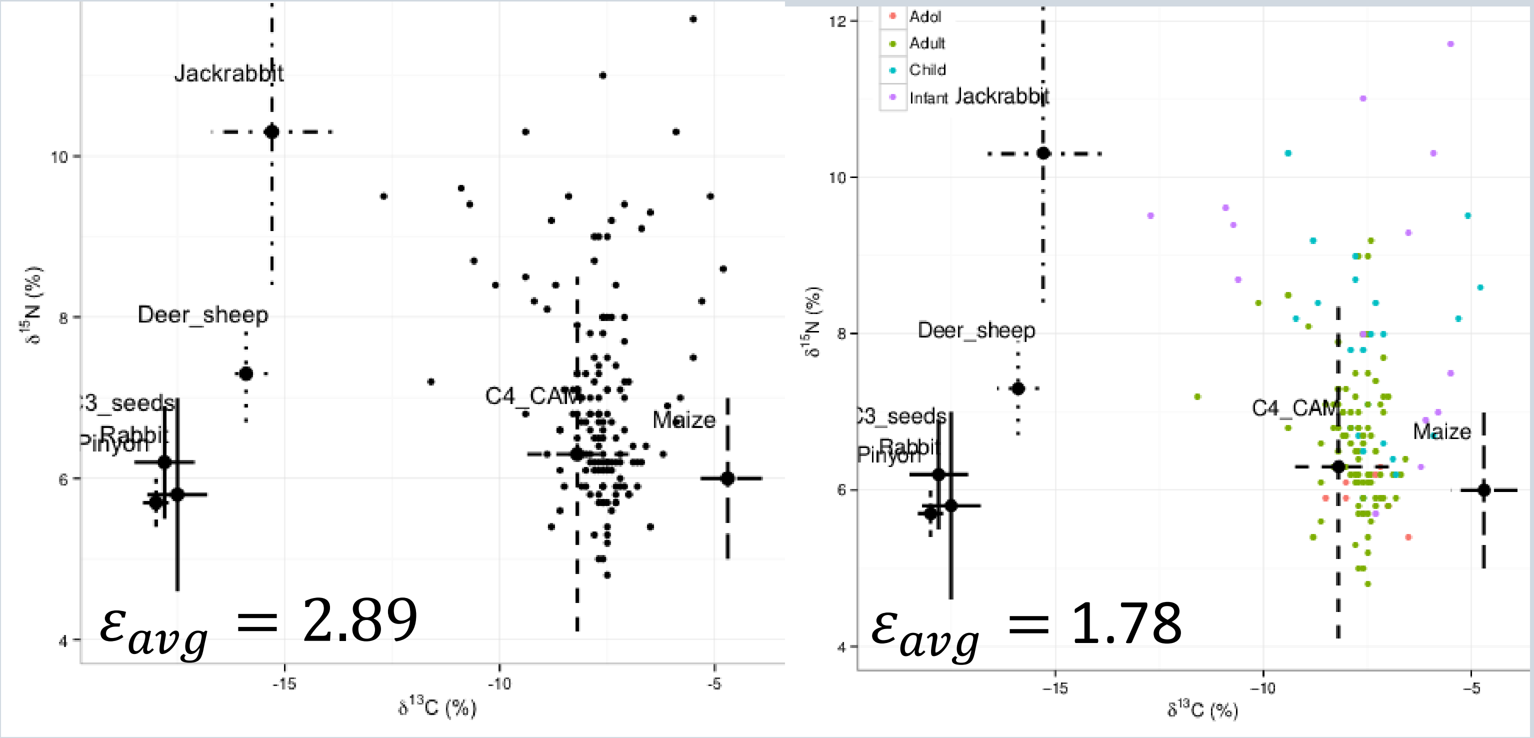

3. Better error structures

More accurate

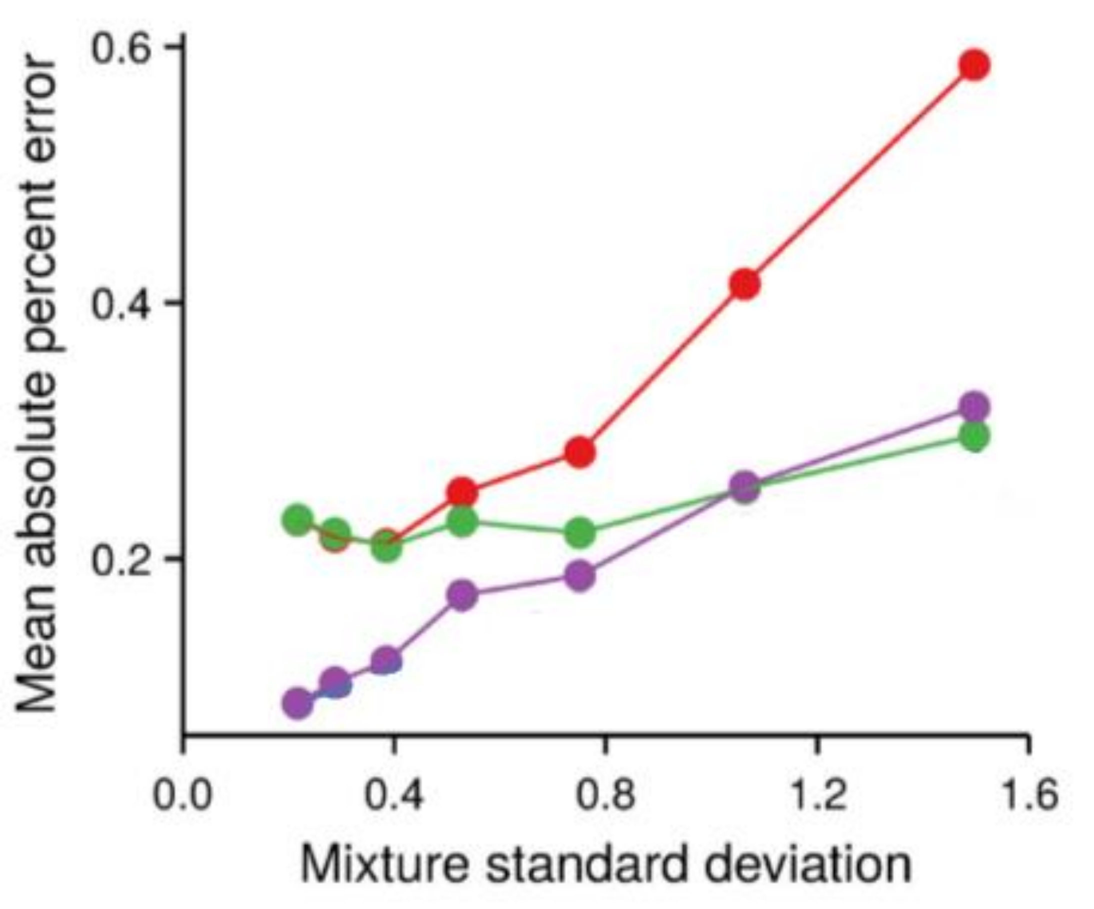

MixSIR

SIAR

MixSIAR ★

3. Better error structures

Stock and Semmens (2016)

4. Effect of priors/ “Bayesian mixing models are biased”

0. What is a *prior*?

$$Pr(\theta|data) \propto \boldsymbol{Pr(\theta)} * Pr(data|\theta)$$

Posterior **Prior** Likelihood

"From a Bayesian perspective, the principle of unbiasedness is reasonable in the limit of large samples, but otherwise it is potentially misleading."

Gelman et al. (1995)

1. There is no “uninformative” prior

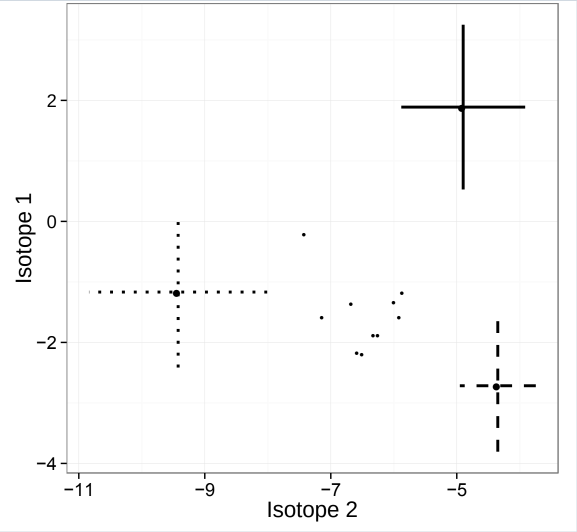

1. There is no "uninformative" prior

Problem: proportions are not independent!

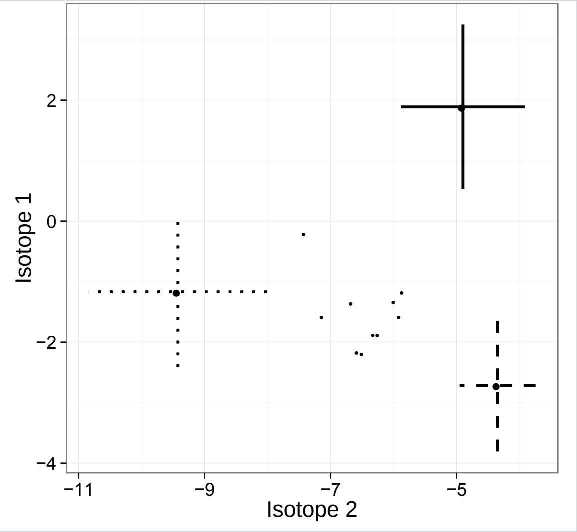

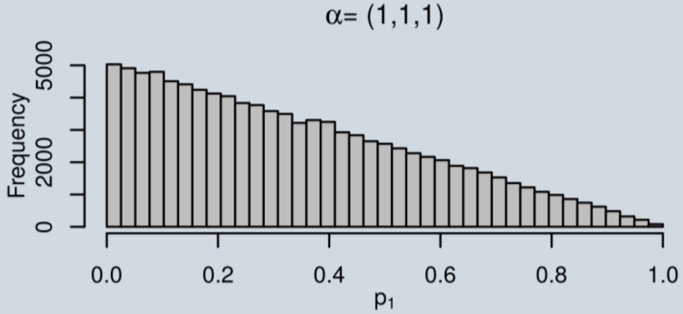

1. There is no “uninformative” prior

Problem: proportions are not independent!

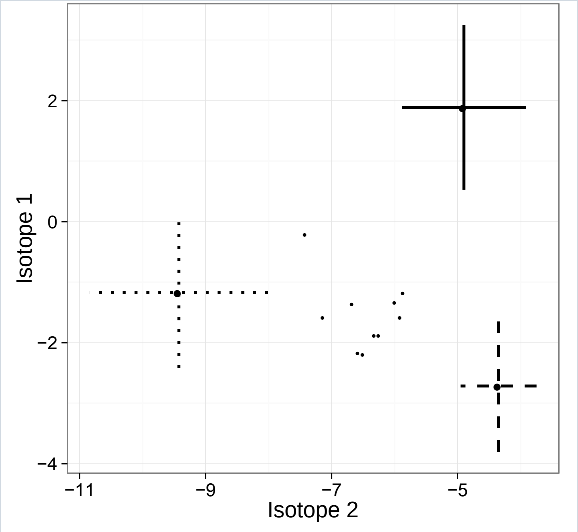

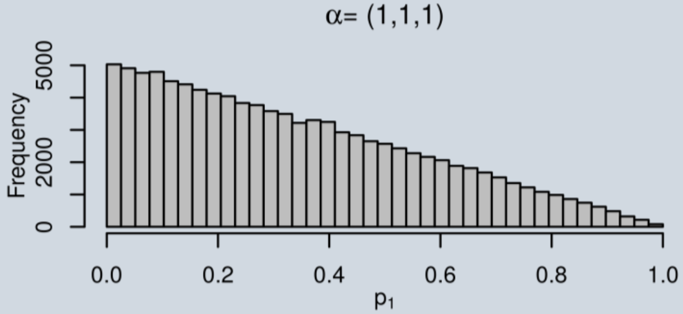

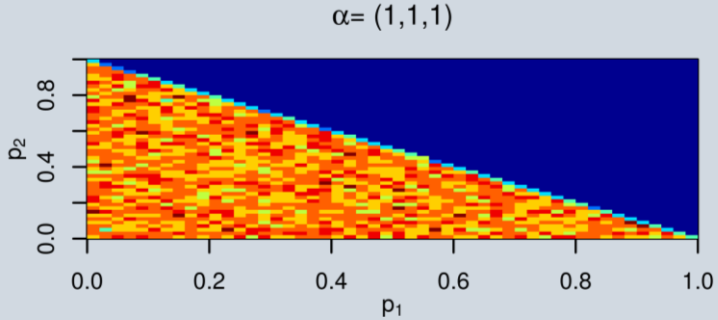

1. There is no "uninformative" prior

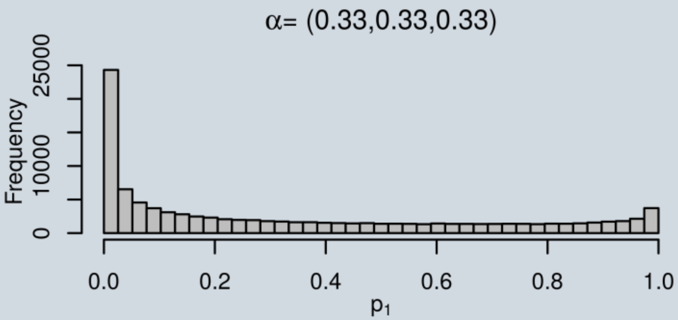

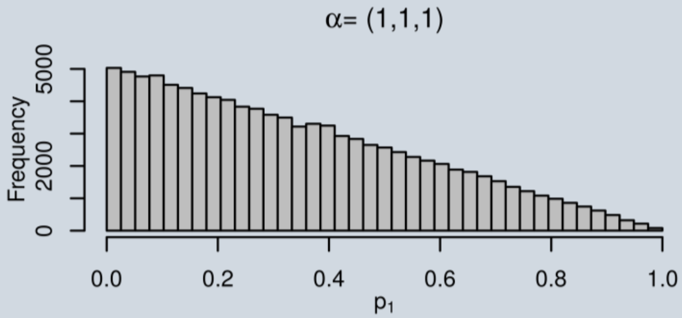

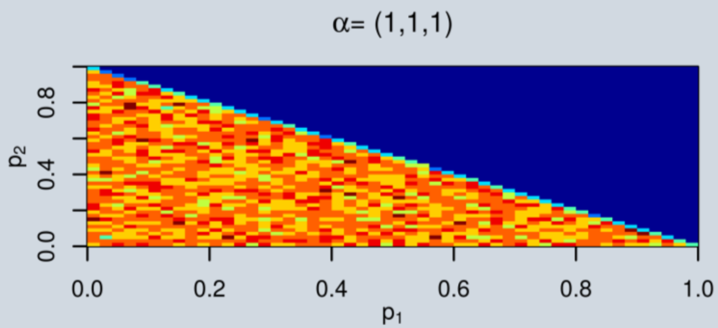

1. There is no “uninformative” prior

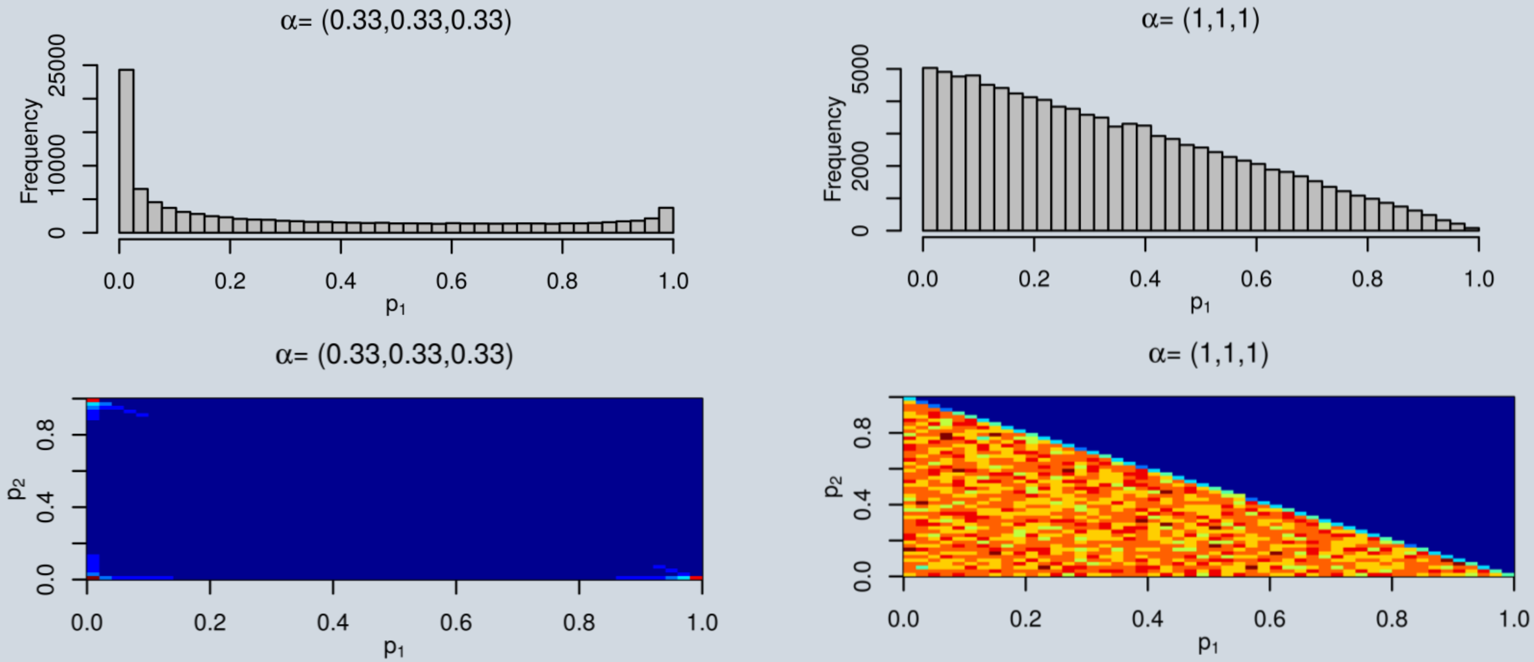

2. Effect of the "uninformative" prior

1. How good is your data?

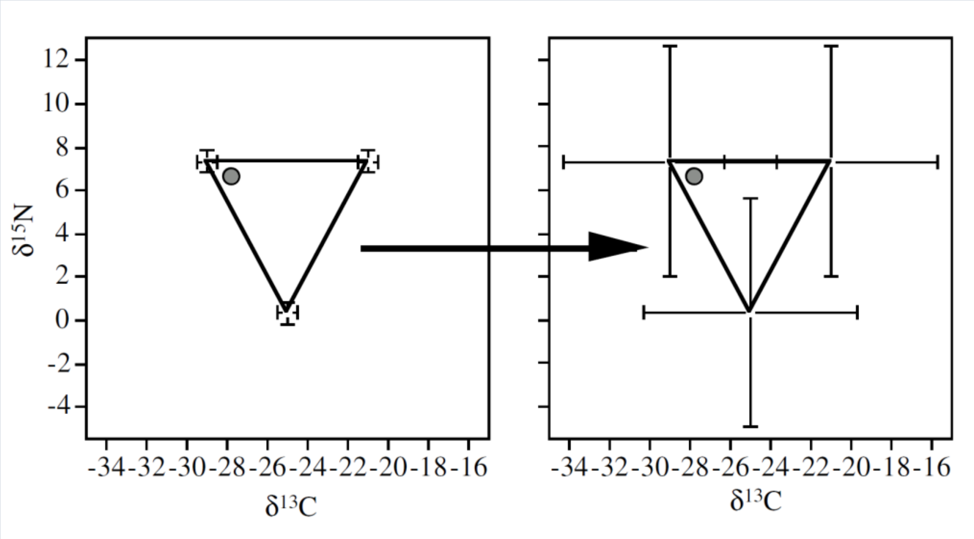

2. Effect of the "uninformative" prior

1. How good is your data?

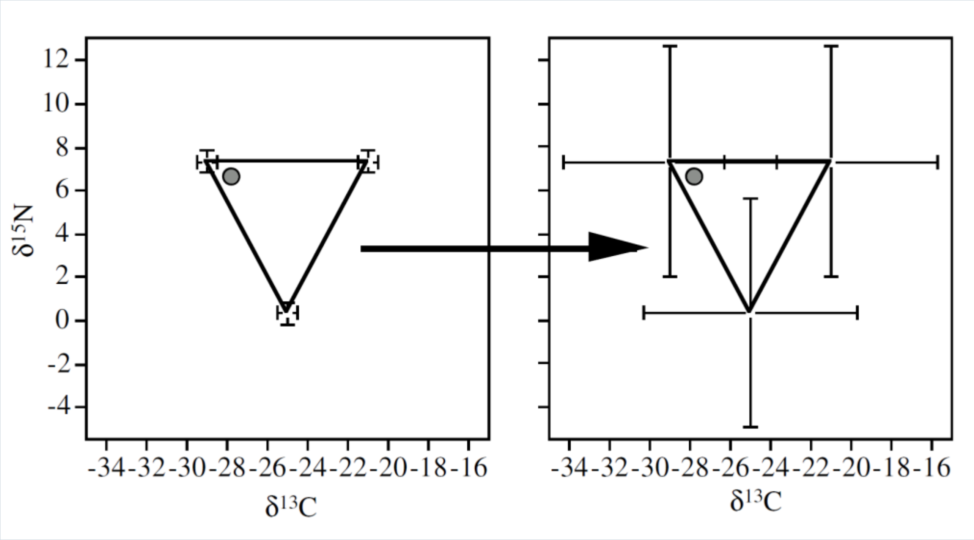

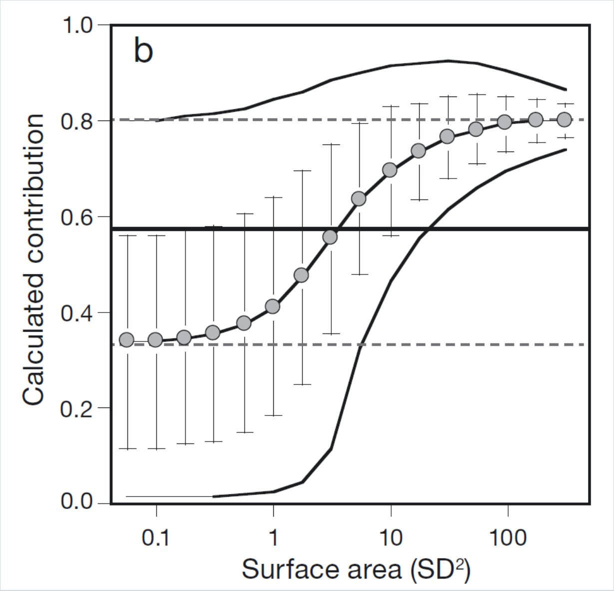

2. Effect of the "uninformative" prior

1. How good is your data?

2. How much data do you have?

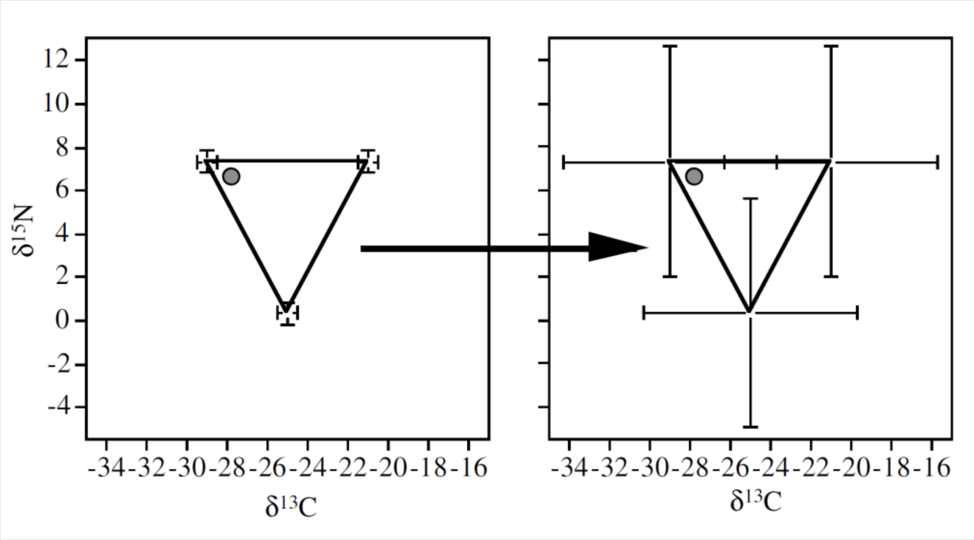

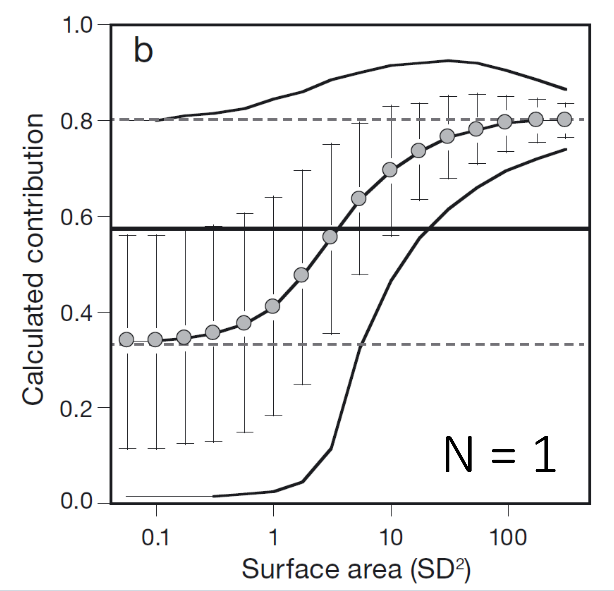

3. Constructing informative priors

You control the mean proportions AND the variance ("informativeness")

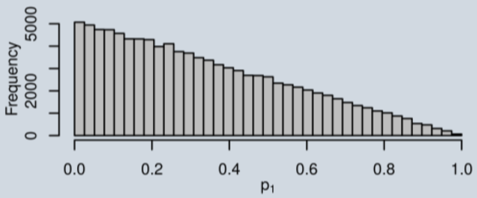

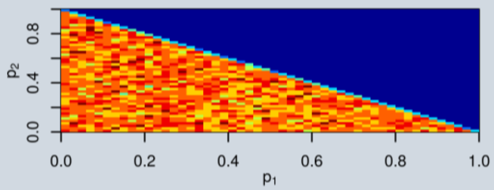

$\alpha = (1, 1, 1)$

3. Constructing informative priors

You control the mean proportions AND the variance (“informativeness”)

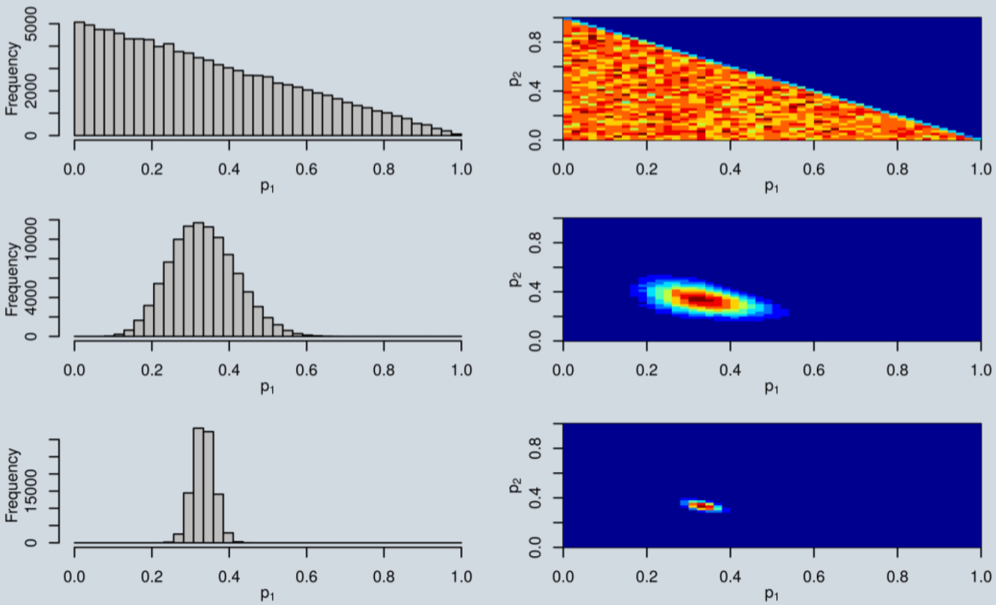

$\alpha = (1, 1, 1)$

$\alpha = (10, 10, 10)$

$\alpha = (100, 100, 100)$

3. Constructing informative priors

You control the mean proportions AND the variance ("informativeness")

30 

8 

25 

3. Constructing informative priors

You control the mean proportions AND the variance ("informativeness")

30 

8 

25 

$\boldsymbol{\alpha} = (\mathbf{30}, \mathbf{8}, \mathbf{25})$

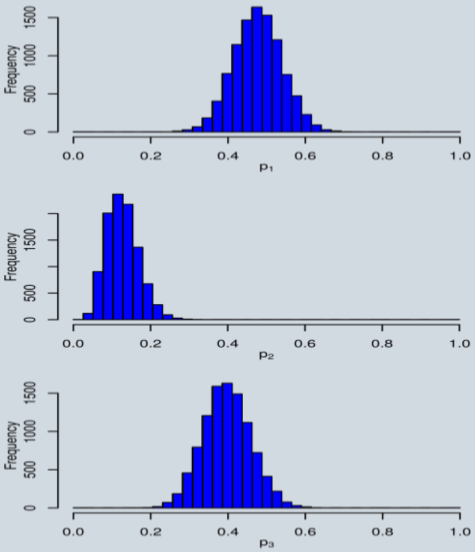

3. Constructing informative priors

You control the mean proportions AND the variance ("informativeness")

30 

8 

25 

$\boldsymbol{\alpha = (30, 8, 25)}$

$\boldsymbol{\alpha = \dfrac{3 * (30, 8, 25)}{63}}$

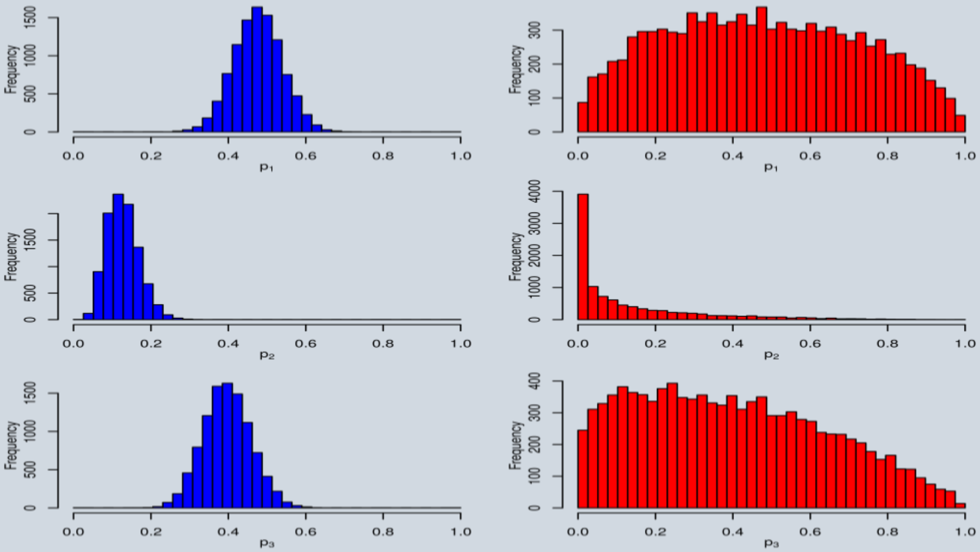

4. Effect of priors/ “Bayesian mixing models are biased”

REDUCE THE INFLUENCE OF THE GENERALIST PRIOR:

1. COLLECT MORE DATA (SOURCE AND CONSUMER)
2. SPECIFY A NON-GENERALIST PRIOR

Great! Where do I get MixSIAR?

CRAN (few months ago)

1. Download and install/update R
2. Download and install JAGS
3. Open R and run:

```
install.packages("MixSIAR")
library(MixSIAR)
```

GitHub (latest)

1. Download and install/update R
2. Download and install JAGS
3. Open R and run:

```
library(devtools)
install_github("brianstock/MixSIAR")
```

Great! Where do I get MixSIAR?

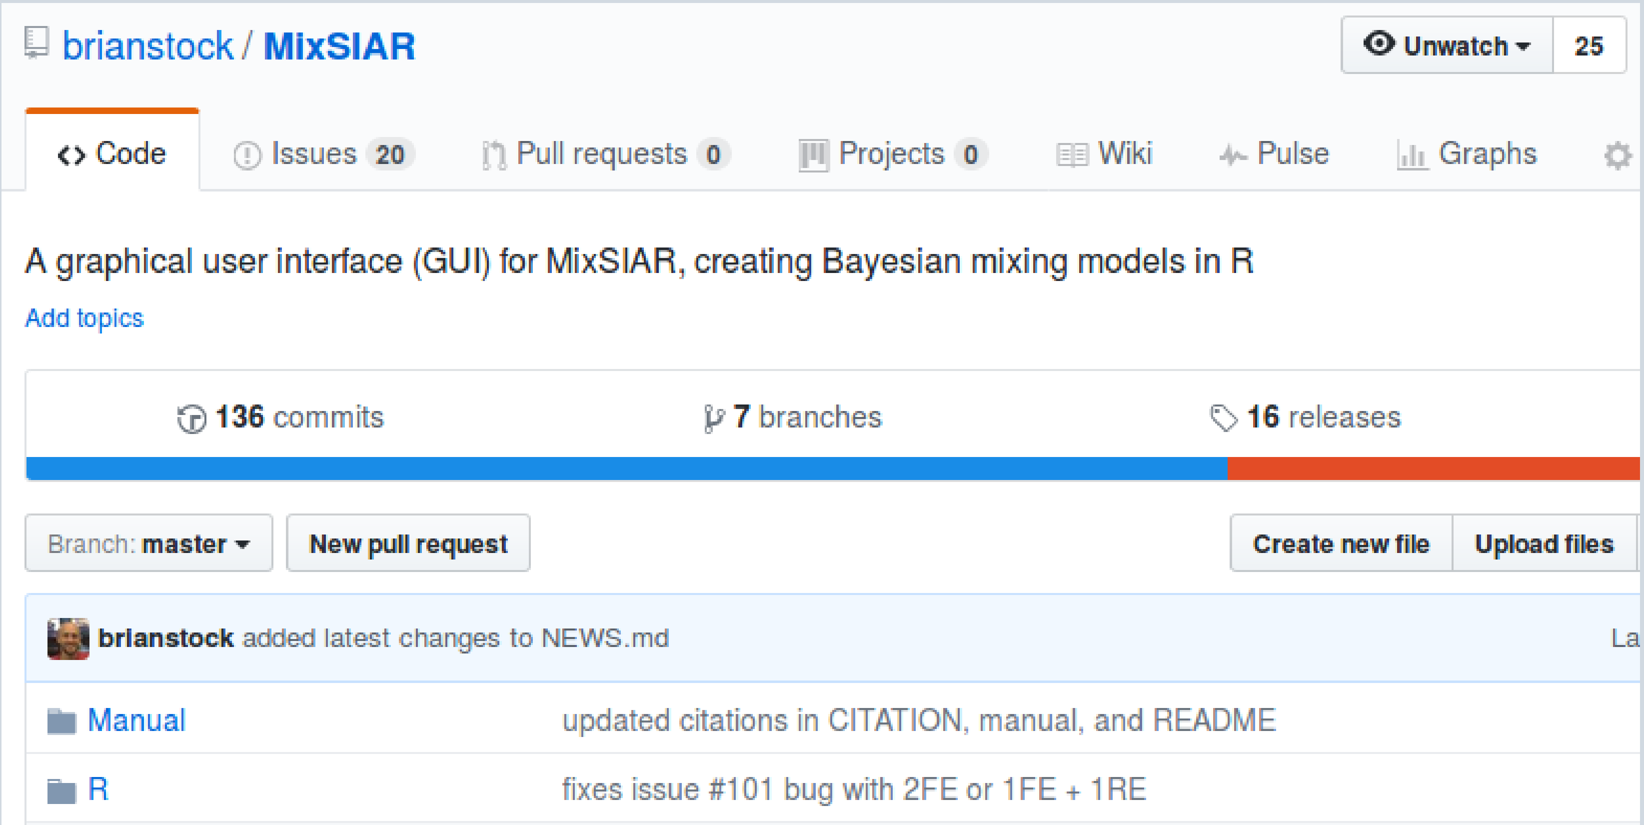

Great! Where do I get MixSIAR?

brianstock added the **bug** label 19 days ago

brianstock added a commit that referenced this issue 18 days ago

`fixes issue #101 bug with 2FE or 1FE + 1RE` `91cf880`

brianstock referenced this issue 18 days ago

Accessing posterior chains with attach.jags function #80 — Closed

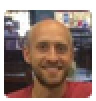

brianstock commented 18 days ago — Owner

Fixed now.

Data files and script: #100
With 1 fixed + 1 random effect, the proper proportions to use are created by combining the global intercept, the factor1 offset, and the factor2 offset:

1. `ilr.both[,f1,f2,src] = ilr.global[,src] + ilr.fac1[,f1,src] + ilr.fac2[,f2,src]`