A Bayesian Analysis of the Effect of Polyaromatic Hydrocarbons on Hormone Levels During the Human Menstrual Cycle

Fletcher Christensen

Department of Mathematics & Statistics
University of New Mexico

March 22, 2018

Outline of Presentation

- An Introduction to the Problem
- Modeling and the Analysis Pipeline
- Complication #1 – Missing Data
- Complication #2 – Expensive Data
- Complication #3 – Comparing Models
- Conclusions and Discussion

Hormones, Ovulation, and the Human Menstrual Cycle

- Before we get started, it'll be helpful to make sure everyone's on the same page biologically.
- I'm going to assume everyone knows what ovulation and the menstrual cycle are.
- Some terminology you might not be familiar with:
  - The **follicular phase** is the phase of the menstrual cycle leading up to ovulation. During this phase, one or more *ovarian follicles* release their eggs.
  - The **luteal phase** refers to the time after ovulation and before the start of menses. During this phase, the follicles that released eggs develop into *corpora lutea* that decay at the end of the month if the egg is not fertilized.

Luteinizing Hormone

- One of the two hormones we will consider in this study is luteinizing hormone (LH), which controls when ovulation occurs and the post-ovulatory conversion of a follicle into a corpus luteum.
- The most distinctive feature of LH during the menstrual cycle is the "LH surge", a spike in LH levels that marks the beginning of ovulation and the transition from the follicular to the luteal phase.
- Daily monitoring of LH levels is a standard method used to detect ovulation.

Luteinizing Hormone

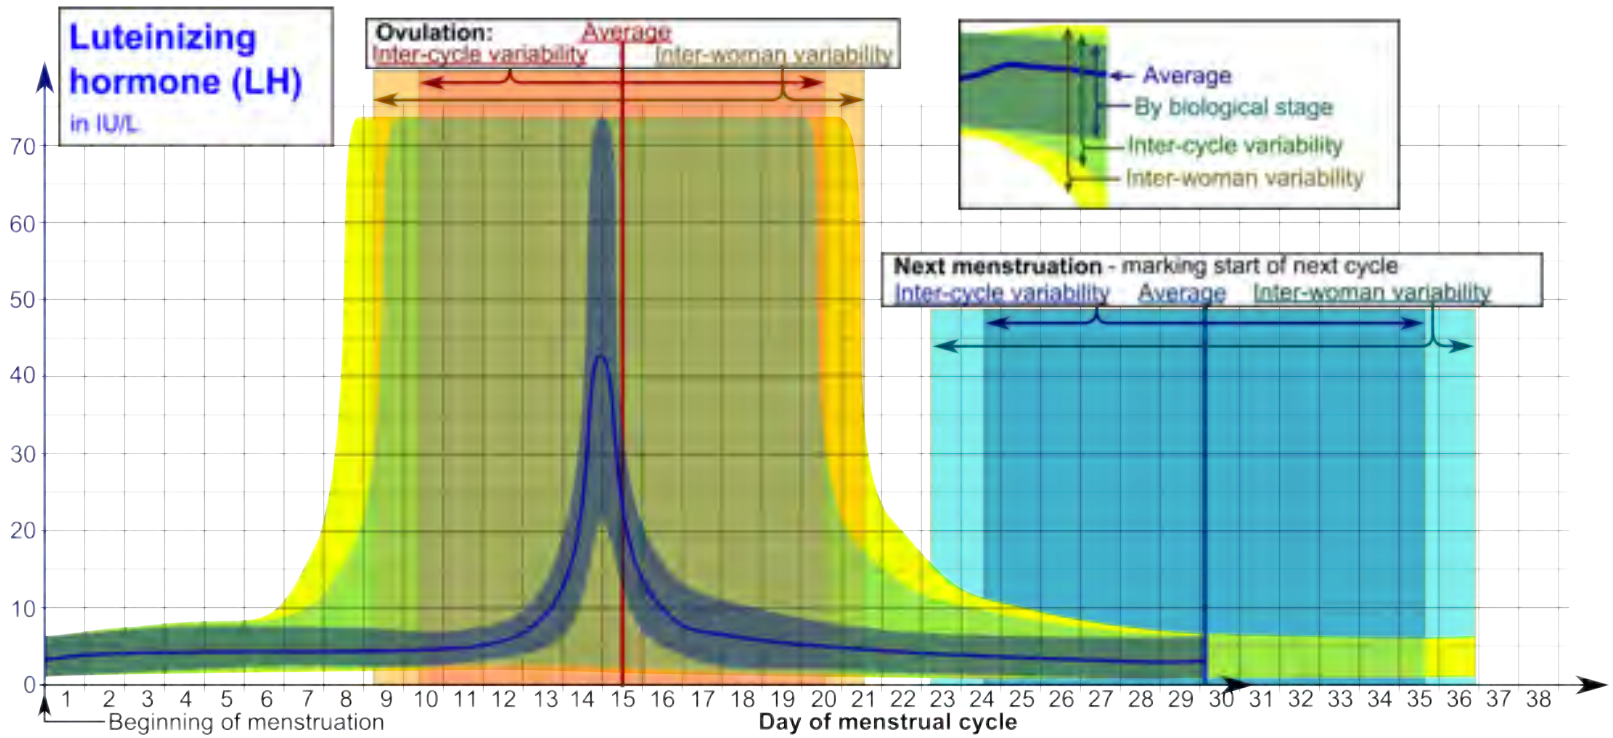

Levels of luteinizing hormone throughout the menstrual cycle.[1]

[1] Häggström, 2014

Estradiol

- The other hormone we will consider is estrone 3-glucuronide ($E_1$3G), which is used as a marker for estradiol levels.
- In the follicular phase, estradiol rises as the oocyte within a follicle matures into a viable egg. When estradiol levels rise sufficiently, the hormone kicks off a chain of events that lead to the LH surge and ovulation.
- In the luteal phase, estradiol (with progesterone) is critical to the development of the endometrium, the lining of the uterus in which a fertilized egg would implant itself.

Estradiol

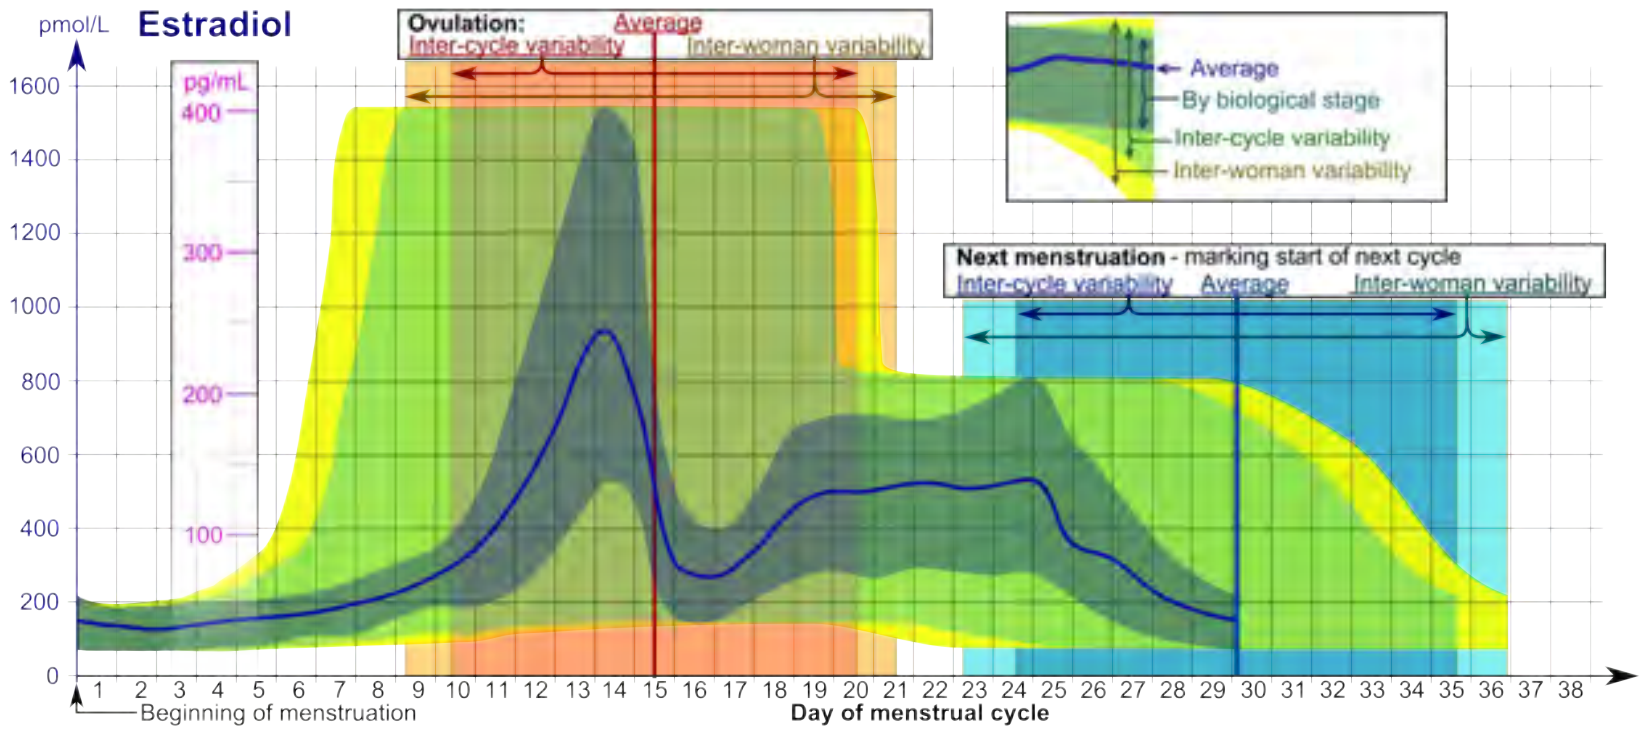

Levels of estradiol throughout the menstrual cycle.[2]

[2] Häggström, 2014

Polycyclic Aromatic Hydrocarbons – What They Are

- Polycyclic aromatic hydrocarbons (PAHs) are organic compounds containing only hydrogen and carbon in which atoms are arranged in multiple connected rings.
- PAHs can be considered a type of environmental pollutant. They are often produced through incomplete combustion of organic matter (e.g. oil, coal, and tar).
- We frequently encounter them through grilled meats, cigarette smoke, engine exhaust, etc.

Polycyclic Aromatic Hydrocarbons – What We Know

- PAHs build up in the human body, and exposure levels can be monitored by looking at PAH concentration in urine samples and stool samples.
- Epidemiologial studies have linked PAH exposure *in utero* with low birth weight[3], reduced immune function[4], and impaired neurological development[5].
- Animal studies have shown that certain PAHs cause significant follicular damage in female mice and rats[6].

[3] Šrám et al, 2005
[4] Winans et al, 2011
[5] Suades-González et al, 2015
[6] Takizawa et al, 1984; Borman et al, 2000

Polycyclic Aromatic Hydrocarbons – What We Use

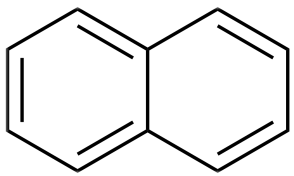

Naphthalene

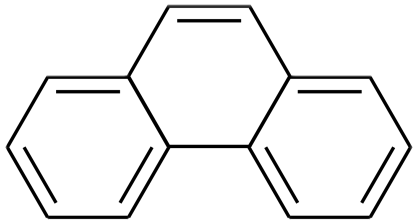

Phenanthrene

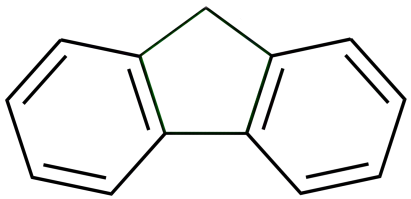

Fluorene

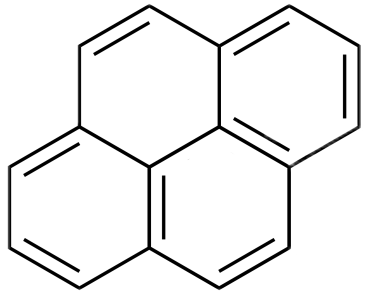

Pyrene

Polycyclic Aromatic Hydrocarbons – Those Crazy Numbers

- I'll be talking about a lot of numbered PAH compounds, like 2-hydroxyl naphthalene (2NAP) or 9-hydroxyl fluorene (9FLUO).
- These are *hydroxylated PAHs* – PAHs that have had a hydroxyl group (-OH) added to them.
- Hydroxylation is a detoxification process by which organisms break down organic compounds into compounds that are more easily excreted.
- The numbering associated with these OH-PAHs refers to the bind point in the PAH where the hydroxyl group connects.

An Analysis Hierarchy

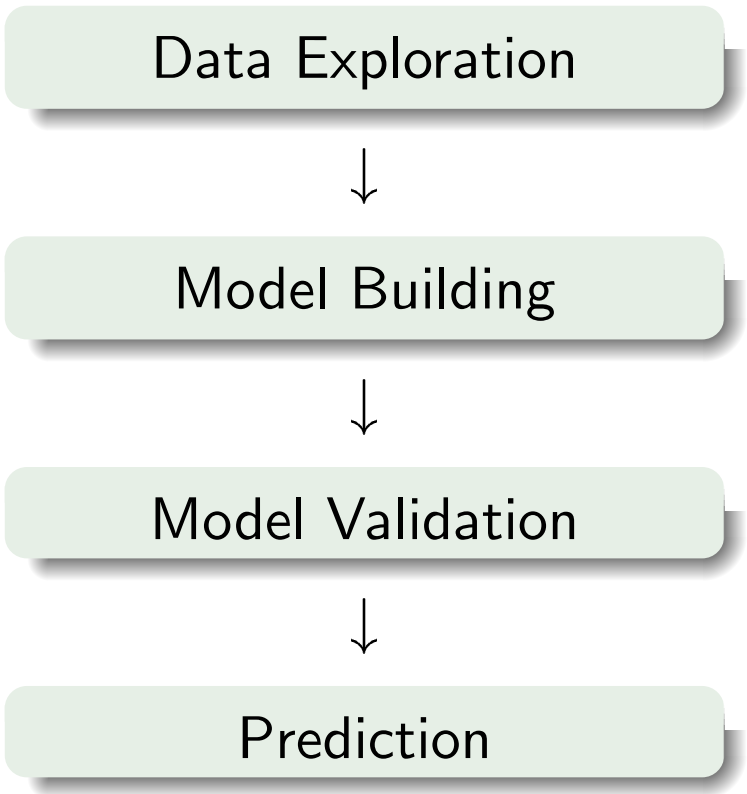

Our Question of Interest

Research Question:

Do OH-PAHs have an effect on hormone levels and ovulation during the human menstrual cycle?

Our Place in the Hierarchy

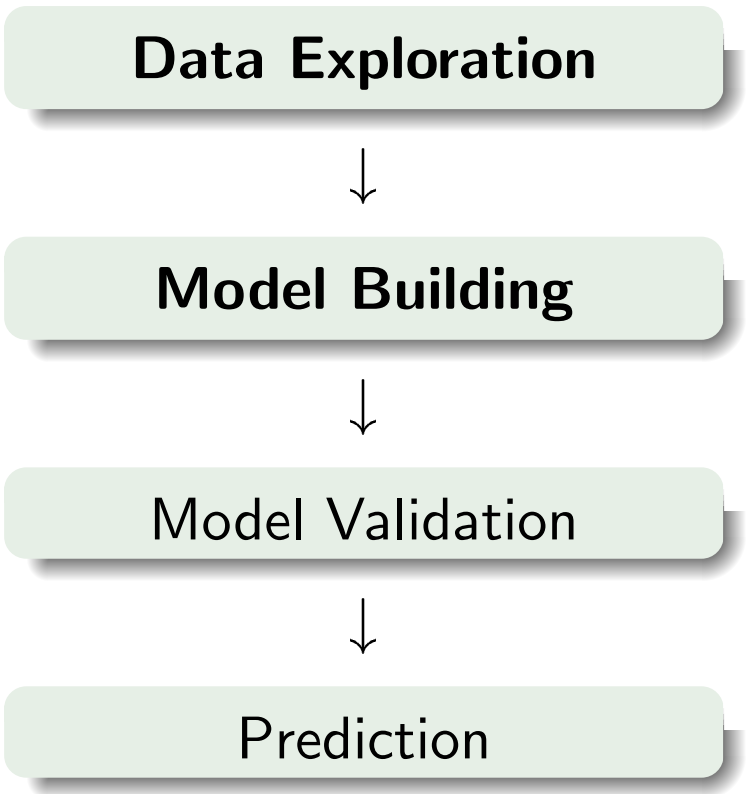

- Although previous work linking PAH to fertility has been done in animals, this is the first major study examining the link in humans.
- We don't really know what to expect.
- We want to establish whether PAHs are important to this topic and begin looking at their role.
- We're not ready to validate a model or predict outcomes.

Response Variables – General

- **Menstrual Cycle Length** – What it says on the tin, this is just how many days a woman's menstrual cycle lasts. We remove all cycles lasting longer than 35 days, which are generally anovulatory cycles and thus members of a fundamentally different population than the ¡35 day cycles.
- **Follicular Phase Length** – How many days pass between menses and the LH surge. We need to detect an LH surge to record this.
- **Ovulation Status** – Whether ovulation occurred or not (as measured by whether we can detect an LH surge). This turns out to be a complicated diagnostic testing problem because of missing data issues, and we won't be considering it here.

Response Variables – LH-based

- **Average Follicular Phase LH Level** – Recall that the hormone profile for LH during the menstrual cycle shows very low baseline levels everywhere except at the LH surge. This is a measure of ground-state LH levels.
- **Highest Observed LH Level** – When we believe a surge occurred, this is the highest LH level we recorded.
- **Peak Level at LH Surge** – When we have complete data for the days surrounding the surge, this is the level we observe at the peak. (*Peak LH* and *Highest LH* are identical when both are recorded; *Peak LH* has some missing values where *Highest LH* is recorded, though.

Response Variables – $E_1$3G-based

- **Average Follicular Phase $E_1$3G Level** – This is the ground state for estradiol observed early in the follicular phase.
- **Slope of Pre-Ovulatory $E_1$3G Rise** – Recall that estradiol increases as the follicles mature during the cycle, and then after a certain threshold is reached, this begins a chain reaction resulting in the LH surge. *$E_1$3G Slope* is the rate of increase in estradiol during this process.
- **Peri-ovulatiory $E_1$3G Level** – Recall also that estradiol drops off again during the LH surge. This is a measure of estradiol levels at that time.

Predictors of Interest

- We have a large group of predictors we can draw from in this problem.
- However, we're only interested in whether PAHs have an effect on hormone levels.
- We want to make the most effective use of our predictors in answering our key question.

Baseline covariates:

Race	Stress	Education
Age	Caffeine Intake	
BMI	Alcohol Consumption	
Walking Minutes per Week		
Vigorous Activity per Week		

PAH covariates:

2Fluo	3Fluo	9Fluo
1Phen	2Phen	3Phen
1Nap	2Nap	1Pyr

Correlated Data

- Data were collected across multiple menstrual cycles for a total of 51 women.
- The number of cycles observed for each woman varied from 2 to 13, with 6 cycles collected per woman on average.
- We need to account for correlation within observations on a single subject.

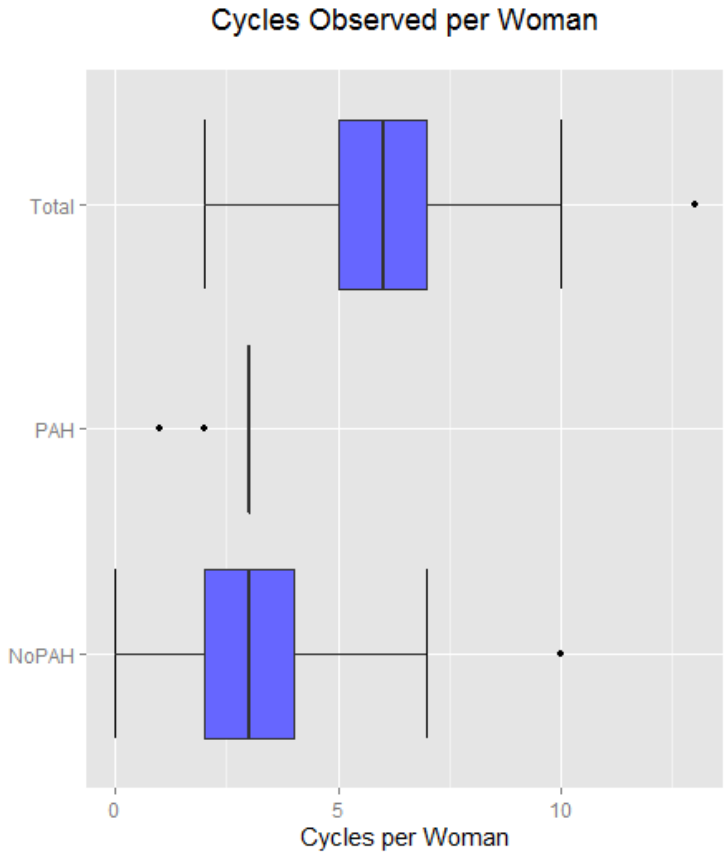

Bayesian vs. Frequentist

- With complicated modeling situations, the Bayesian approach is often easier to construct and understand.
- Construction – the missing data situations in our problem make a Bayesian approach much more tractable here.
- Understanding – as you'll see in a second, the model we come up with is pretty complicated. It's a lot easier to keep track of the various parts of that model, and to make sure we're making reasonable assumptions about those parts, when we put everything into a Bayesian framework.

The Analysis Pipeline

- Recall that our primary interest is in detecting whether OH-PAHs have an effect on our response variables.
- More than that, we're interested in whether they have a meaningful effect—that is, do OH-PAHs tell us something about our response variables that we can't learn more easily with other predictors?
- We will consider two models for each endpoint:
  1. A baseline model using only our non-PAH predictors.
  2. A final model using both non-PAH predictors and our OH-PAH concentrations.
- We will compare these two models to see which one is better.

The Model Structure

- We begin with the basic linear mixed model:

$$Y_{ij} = X_{ij}\beta + Z_{ij}\gamma_i + \varepsilon_{ij}.$$

- In an easy setting, we'd assume something like:

$$\varepsilon_{ij} \overset{\text{iid}}{\sim} N(0, \sigma^2), \qquad \gamma_i \overset{\text{iid}}{\sim} N(0, \tau^2).$$

- We'll have to do things a little differently to deal with missing data—specifically, our ε_{ij}'s aren't going to be identically distributed.

How We Get Our Models

- We will use a backward stepwise procedure to get our models.
- For the baseline model:
  1. We use MCMC to fit a model with all non-PAH predictors included.
  2. For each β in the model, we calculate the proportion of iterates that were above zero. Call this p_β.
  3. We drop the covariate for which p_β is closest to 50%. If the covariate is a factor variable, we consider only the p_β farthest from 0.5 among the β's for all its levels.
  4. We refit the model until all remaining covariates have $p_\beta > 0.85$ or $p_\beta < 0.15$ (i.e. until all remaining coefficients are consistently positive or negative).
  5. Covariates for Age and Race are forced into the model because we believe *a priori* that they are important to modeling our response variables.

How We Get Our Models

- We will use a backward stepwise procedure to get our models.
- For the model including OH-PAH covariates:
  1. Our initial model includes all the non-PAH predictors from the baseline model, as well as all the OH-PAH predictors.
  2. We conduct the same stepwise procedure using p_β, removing whichever covariate has p_β closest to 50%.
  3. Again, we refit the model after each covariate is removed until all remaining covariates have $p_\beta > 0.85$ or $p_\beta < 0.15$.
  4. We do allow the model to drop non-PAH covariates if their p_β's fall back below the retention threshold after OH-PAH covariates are included.

Traditional Missing Data – Predictors

- Our study has a lot of self-report data on demographic predictors that we're interested in using.
- Unsurprisingly, though, when you give a lot of surveys to a lot of people, some of them won't complete every survey, or will provide data that isn't easy to categorize.
- For BMI, Stress, Education, Caffeine Intake, and Alcohol Consumption, we are missing data for 1-3 women (out of a total sample of $k = 51$).
- We will assume these data are missing at random – that there are no systematic differences between the women for whom these data are missing and the women for whom these data are present.
- Thankfully, we have very little missing predictor data, so our analysis is probably robust to this assumption.

Traditional Missing Data – Responses

- The missing data problem on our response variables is trickier.
- Most of the responses we're using are calculated from raw, daily data obtained by an electronic device that monitors hormone concentrations in urine. Again unsurprisingly, people differ in their commitment to peeing on a stick every morning in the pursuit of science.
- It is probably unreasonable to assume that there are no important differences between women for whom we have more consistent data and women for whom we lack consistent data...
- But we can still make the MAR assumption—that *conditional on what we know about these women through our model,* there is no systematic differences between cycles where we have consistent data and cycles where we don't.

Simple Bayesian Imputation

- Missing data is a common complication in data analysis, but the Bayesian approach makes dealing with it very easy.
- On the next two slides, I'll give two examples of WinBUGS code. The first is a complete data formulation; the second is a missing data formulation.
- The change here is relatively minor—though it will have some knock-on effects on our ability to compare models down the road.

A Complete Data Example

```
Model {
  for( i in 1:n2 ){
    folli_length[i] ~ dnorm(mu[i], sigma)
    mu[i] <- b_0 + b_stress * stress[SID[i]] + b_rand[SID[i]]
  }
  for( k in 1:n_s ){
    b_rand[k] ~ dnorm(0, tau)
  }
  sigma ~ dgamma(0.001, 0.001)
  sigma_sd <- pow(sigma,-0.5)
  tau_sd ~ dunif(0,50)
  tau <- pow(tau_sd,-2)
  b_0 ~ dnorm(0,.01)
  b_stress ~ dnorm(0,.01)
}
```

A Missing Data Example

```
Model {
  for( i in 1:n2 ){
    folli_length[i] ~ dnorm(mu[i], sigma)
    mu[i] <- b_0 + b_stress * stress[SID[i]] + b_rand[SID[i]]
  }
  for( k in 1:n_s ){
    b_rand[k] ~ dnorm(0, tau)
    stress[k] ~ dnorm(mu_stress,sigma_stress)
  }
  sigma ~ dgamma(0.001, 0.001)
  sigma_sd <- pow(sigma,-0.5)
  tau_sd ~ dunif(0,50)
  tau <- pow(tau_sd,-2)
  b_0 ~ dnorm(0,.01)
  b_stress ~ dnorm(0,.01)
  mu_stress ~ dnorm(0,1)
  sigma_stress ~ dgamma(2,2)
}
```

- In the BUGS setting, all we have to do is add statements to indicate that our covariate will be probabilistically modeled with some given prior distribution.
- Whenever data are available, the MCMC will use the actual data for the variable.
- When data aren't available, the MCMC will impute the variable automatically, given its role in the model.
- In the example above, if information on Stress is missing, the MCMC will tend toward a weighted average of a b_{rand} term and a *stress* term according to their respective variances and the *stress* coefficient.
- If we think about what this means in terms of the model, that's exactly what we want.

OH-PAH Data Costs Money

- We've got a different type of missing data situation going on with our OH-PAH predictors.
- OH-PAH data are collected once per cycle for each woman, on the 10^{th} day after menses.
- OH-PAH concentrations are measured from a urine sample (collected independently of daily hormone measurements).
- We pay the CA Dept. of Public Health's Environmental Health Lab to calculate the OH-PAH concentrations based on the urine samples we've collected.
- Unfortunately, we run out of money to pay the lab before we run out of samples to test.

What We Can Access

- Recall the figure I showed earlier, with boxplots of the number of cycles recorded per woman.
- We have OH-PAH data for 1-3 cycles on every woman (with fewer than 3 cycles iff we have 3 or fewer cycles total for the woman).
- Testing priority was given for samples corresponding to cycles with more complete daily hormone data (c.f. our issue with missing data on the responses).

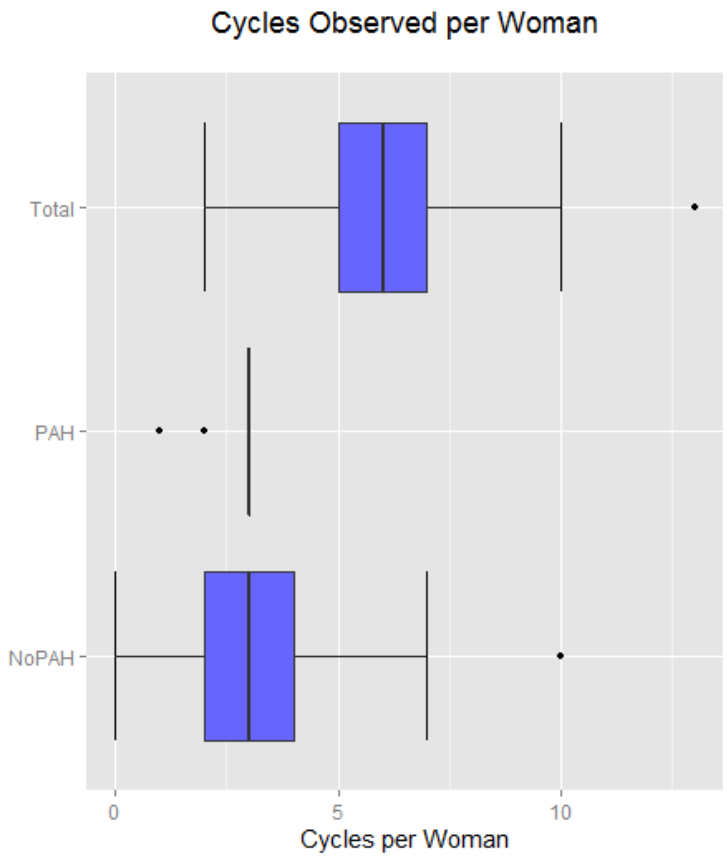

How Does OH-PAH Missingness Affect Our Analysis

- As always, we should worry about the MCAR/MAR/MNAR assumptions. We have a process to choose which samples to test. Are samples we test likely to be structurally different from samples we don't test, conditional on the other information we have?
- It is likely that missingness in response data is related to individual subjects – and this is important because it's part of our algorithm for choosing which OH-PAH samples to test.
- It is *unlikely* that missingness is related to individual *cycles*, conditional on subject.
- Since we are making sure to get 3 OH-PAH samples from every subject based on the cycles with her most complete response data, we should be able to make the MAR assumption.

How Does OH-PAH Missingness Affect Our Analysis

- We have another serious problem here, though. Remember that the whole point of our study is to identify whether OH-PAH's have an effect on hormones levels during, and characteristics of, the menstrual cycle.
- We only have data on OH-PAH concentrations for a bit under half (48%) of cycles we studied.
- Do we have to throw out half of our data because it doesn't talk about our research question, or is there something else we can do?

Shared-Parameter Modeling

Think about the role our classes of variables play here:

- Response – For which we want to build a model.
- OH-PAH Predictor – For which we want to determine whether they have a role in that model.
- Non-PAH Predictor – For which we want to control, so we can see whether the OH-PAH predictors' role is practically meaningful.

Shared-Parameter Modeling

- We would like to use the information on the PAH-missing data to help us understand the role the non-PAH covariates play in the model.
- We can do this by assuming that we're modeling the same parameters when we have PAHs and when we don't have PAHs.
- Let's take a look at what this means in terms of WinBUGS code:

The PAH-Missing Model

```
Model {
  for( i in 1:n2 ){
    folli_length[i] ~ dnorm(mu[i], sigma)
    mu[i] <- b_0 + b_stress * stress[SID[i]] + b_rand[SID[i]]
  }
  for( k in 1:n_s ){
    b_rand[k] ~ dnorm(0, tau)
    stress[k] ~ dnorm(mu_stress,sigma_stress)
  }
  ...
}
```

The PAH-Included Model

```
Model {
  for( i in 1:n1 ){
    folli_length[i] ~ dnorm(mu[i], sigma_1 )
    mu[i] <- b_0 + b_stress * stress[SID[i]] + b_rand[SID[i]]
  }
  for( i in (n1+1):n2 ){
    folli_length[i] ~ dnorm(mu[i], sigma_2)
    mu[i] <- b_0 + b_stress * stress[SID[i]] + b_rand[SID[i]] +
             b_pah[1] * Cr_1PYR[i] + b_pah[2] * Cr_3FLUO
  }
  for( k in 1:n_s ){
    b_rand[k] ~ dnorm(0, tau)
    stress[k] ~ dnorm(mu_stress,sigma_stress)
  }
  ...
}
```

What Are We Doing?

- We fit a model without OH-PAHs on the PAH-missing data, and then we fit a second model on the PAH-included data.
- These two models share parameters for the intercept and for each non-PAH covariate.
- These two models also share random effects terms, because the random effect value is subject-specific and subjects have observations in both models.
- The **error variance** on the models is different, because we expect that we are accounting for more response variability in the PAH-included data.

Places to Be Cautious

- By structuring these two models with shared parameters, we assume that these parameters will be the same across models.
- This is stricter than simply assuming the parameters should be related between models, and involves a complicated set of assumptions.
- In particular, I worry about:
  - Whether the OH-PAH covariates have a biasing effect. We deal with this by standardizing the OH-PAH values so they have mean 0. (I think) we could also let the intercept term in the two models differ.
  - Whether there is collinearity between OH-PAH concentrations and the non-PAH covariates. There is, which means that our estimates of the non-PAH parameters in the PAH-included model will probably shrink towards zero, relative to our estimates from the PAH-missing model.

How We Select Our Models

- Recall that we obtain our models through a pair of backward stepwise procedures.
- We choose to remove and retain covariates based on p_β, a measure of the proportion of iterates from the posterior distribution for β that fall below 0.
- Selection by p_β is ad-hoc, and while it seems reasonable for the stepwise procedure where we're using it, we would prefer to have additional criteria for evaluating the relative quality of our models.

Traditional Model Selection Criteria

- If we weren't being Bayesians, we might choose something like AIC with its ties to the Kullback-Leibler divergence.
- As Bayesians, though, we have three traditional measures of fit we might consider:
  - Bayes Factor (BF) – Analogous to a likelihood ratio; looks at the ratio of the probability of the data arising under the two models being compared.
  - Log Pseudo-marginal Likelihood (LPML) – Like the Bayes Factor, but involves looking at how well the model fits the data when that data point is not used in fitting the model; analogous to a leave-one-out approach to the BF.
  - Deviance Information Criterion (DIC) – Computationally "easy", a weighted measure of the difference between the mean posterior (Bayesian) deviance under a model, and the (Bayesian) deviance evaluated at the posterior mean.

Different Types of DIC

- In the original development of DIC[7], the authors discuss the idea of **focal parameters** and suggest marginalizing out non-focal parameters before computing the DIC.
- This is difficult to do in practice. The great advantage of DIC is its computability—but its computability relies on having the joint distribution of all parameters in the model.
- In the random effects and missing data settings, variants of the traditional DIC have been proposed[8] to deal with mismatch between the focal parameters and the parameters actually modeled.

[7] Spiegelhalter et al, 2002
[8] Celeux et al, 2006

What WinBUGS Provides

- WinBUGS calculates a DIC score automatically.

"This deviance is defined as -2 * log(likelihood): 'likelihood' is defined as $p(y|\theta)$, where y comprises all stochastic nodes given values (i.e. data), and θ comprises the stochastic parents of y – 'stochastic parents' are the stochastic nodes upon which the distribution of y depends, when collapsing over all logical relationships."

- Unfortunately, this is not the deviance we want. It treats both the random effects and the parameters used to model our missing covariates as focal parameters.
- Take-away message: don't trust naive DIC calculations on complicated models.

The Marginal DIC

- Consider a model involving some data Y, some model parameters θ, and some augmentation Z.
- We can think of Z in a number of different ways: as random effects, as missing data, or as augmented data.
- What WinBUGS gives us is based on

$$D(\theta, Z|Y) = -2 \ln L(\theta, Z|Y) = -2 \ln f(Y|\theta, Z).$$

- What we want is

$$D(\theta|Y) = -2 \ln \int_{supp(Z)} f(Y|\theta, Z)$$

The Marginal DIC

- We can get $D(\theta|Y)$ directly in this problem, by using the multivariate normal distribution with covariance matrix corresponding to the random effects model.
- We ignore the effect of missing data on the DIC, because the missing data relates to the observed data but not the model parameters themselves.
- The current solution is computationally intensive (it takes about 10m to calculate the DIC for each model), but it is tractable, and it gives us the correct DIC.

Areas Where PAHs Seem Important

- Basically everywhere.
- We show a DIC reduction of 3 or more[9] in our models for cycle length, follicular phase length, follicular LH level, highest observed LH level, peak LH level, and slope of pre-ovulatory estradiol rise.
- Only models for follicular estradiol level and periovulatory estradiol level don't appear to show significant improvement with the inclusion of OH-PAH concentrations.

[9] A difference 3 is suggested by Spiegelhalter et al (2002) as sufficient indication that one model is preferred to another.

Which PAHs Seem Important

- Pyrene plays a role in cycle length, follicular phase length, and follicular phase LH levels.
- Fluorene is indicated in models for cycle length, follicular phase length, and surge LH levels.
- Naphthalene appears in all LH models, as well as $E_1 3G\ Slope$.
- Phenanthrene appears in models for all response variables except follicular phase length.

The Role of PAH Profiles

- OH-PAH concentrations tend to be highly collinear, especially among different isomers of the same compound.
- Naphthalene is an exception to this – 1- and 2-naphthalene are known to arise from exposure to different pollutants.
- Because of collinearity, we model fluorene and phenanthrene with both an "average concentration" term and terms for the degree to which individual isomers deviate from the average.
- We find that the profile of isomers being created appears just as influential as the overall concentration of the OH-PAHs.

From the Data Analysis

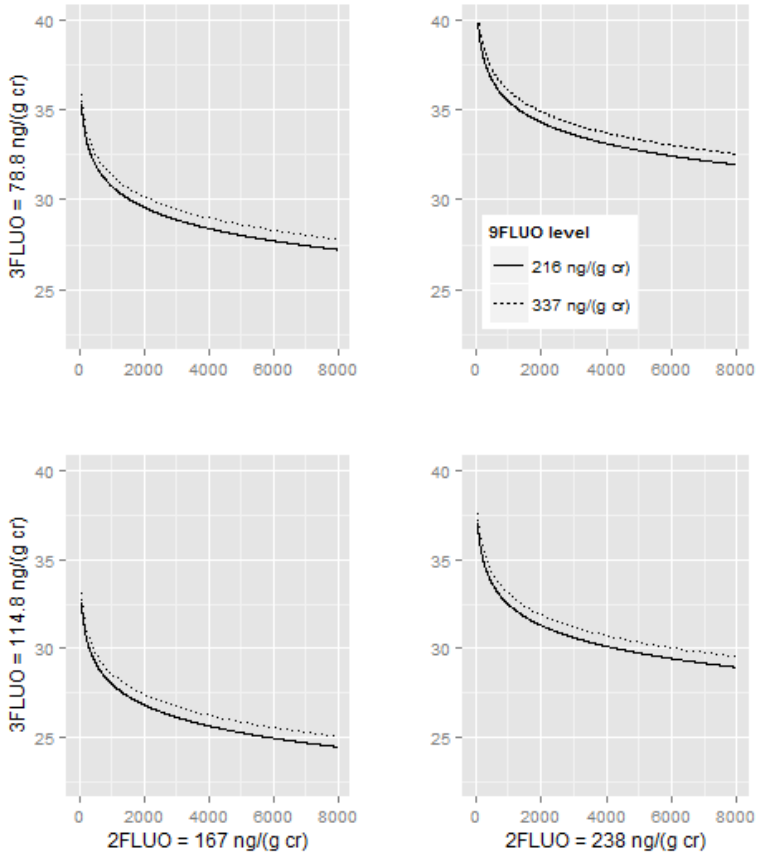

Plots of 1-naphthalene vs. highest observed LH for various concentrations of fluorene isomers.

Revisiting the Analysis Hierarchy

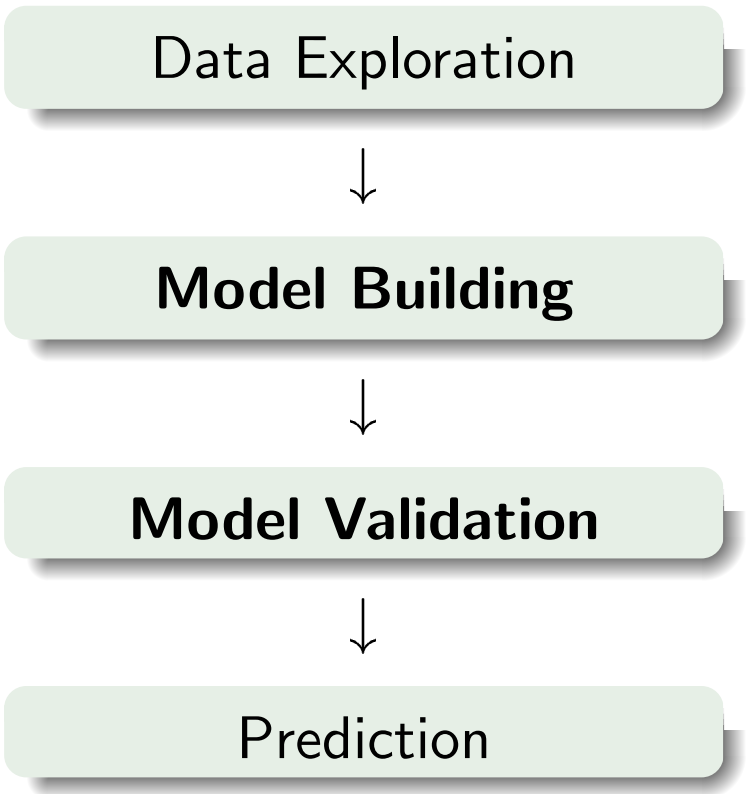

- We've verified that PAH concentrations are predictive of hormone levels during the human menstrual cycle.
- We've begun to examine possible models of their action.
- Further research should consider biologic pathways through which they may act, and should seek to verify and expand upon the results we've found.

The Ovulation Model Problem

- I mentioned early on that we wanted to look at whether ovulation occurred as a response variable. It's actually the response variable we're most interested in studying with these data.
- Unfortunately, the missing data issues described above become even more complicated for ovulation, because MAR is no longer a valid assumption.
- The problem we encounter is that our proxy for whether ovulation occurred is our detection of an LH surge. If we see a surge, we say ovulation occurred.
- If we *don't* see a surge, it may be because no surge occurred, or because we don't have data for the 2-3 days during which we would observe heightened LH levels.

The Ovulation Model Problem

- What we end up with is a division of observations into three categories:
  - (A) Observations where we know an LH surge occurred (i.e. we saw a surge).
  - (B) Observations where we know no LH surge occurred (i.e. we have complete data and can see nothing happened).
  - (C) Observations where we're unsure whether a surge occurred.
- The problem here is that our criteria for sorting observations into groups (A) and (B) aren't equally stringent. We are much more likely to see a surge than to see enough observations to be certain no surge occurred.
- Because our missingness mechanism is directly tied to what we want to predict, we need a different approach (e.g. the diagnostic testing paradigm).

Bibliography

- SM Borman, PJ Christian, IG Sipes, & PB Hoyer. (2000) "Ovotoxicity in female Fischer rats and B6 mice induced by low-dose exposure to three polycyclic aromatic hydrocarbons: Comparison through calculation of an ovotoxic index." *Toxicology and Applied Pharmacology* **167** 191-198.
- G Celeux, F Forbes, CP Robert, & DM Titterington. (2006) "Deviance information criteria for missing data models." *Bayesian Analysis* **1**, 651-674.
- M Häggström. (2014) "Reference ranges for estradiol, progesterone, luteinizing hormone and follicle-stimulating hormone during the menstrual cycle." *Wikiversity Journal of Medicine* **1** (1).
- RJ Šrám, B Binková, J Dejmek, & M Bobak. (2005) "Ambient air pollution and pregnancy outcomes: a review of the literature." *Environmental Health Perspectives* **113**, 375-382.

Bibliography

- E Suades-González, M Gascon, M Guxens, & J Sunyer. (2015) "Air pollution and neuropsychological development: a review of the latest evidence." *Endocrinology* **156**, 3473-3482.
- DJ Spiegelhalter, NG Best, BP Carlin, & A van der Linde. (2002) "Bayesian measures of model complexity and fit." *J.R. Statist. Soc. B* **64**, 583-639.
- K Takizawa, H Yagi, DM Jerina, & DR Mattison. (1984) "Murine strain differences in ovotoxicity following intraovarian injection with benzo(a)pyrene, (+)-(7R,8S)-oxide, (-)-(7R,8R)-dihydrodiol, or "(+)-(7R,8S)-diol-(9S,10R)-epoxide-2." *Cancer Research* **44**, 2571-2576.
- B Winans, MC Humble, & BP Lawrence. (2011) "Environmental toxicants and the developing immune system: A missing link in the global battle against infectious disease?" *Reproductive Toxicology* **31**, 327336.