Hierarchical modeling and prior information: an example from toxicology

Andrew Gelman
Department of Statistics
Columbia University, New York

Lamont-Doherty Earth Observatory
conference on climate reconstruction, 9 Feb 2011

Outline of 45-minute talk

- Central story: 4-compartment model of toxicokinetics of perchloroethylene
- Bayesian inference combines prior information and data
- Unresolved questions
- How the model all fits together

- Time at the end for discussion and questions

Toxicokinetics of perchloroethylene

- Goal:
  - How much PERC is metabolized at high doses
  - Population distribution
- Experimental data: Expose 6 healthy volunteers to PERC for four hours, then measure concentrations in blood and air for 2 weeks
- 4-compartment model, metabolism in liver
- Our analysis:
  - Simple data-fitting didn't work
  - Use Bayes to combine data and prior info within model

4-compartment model

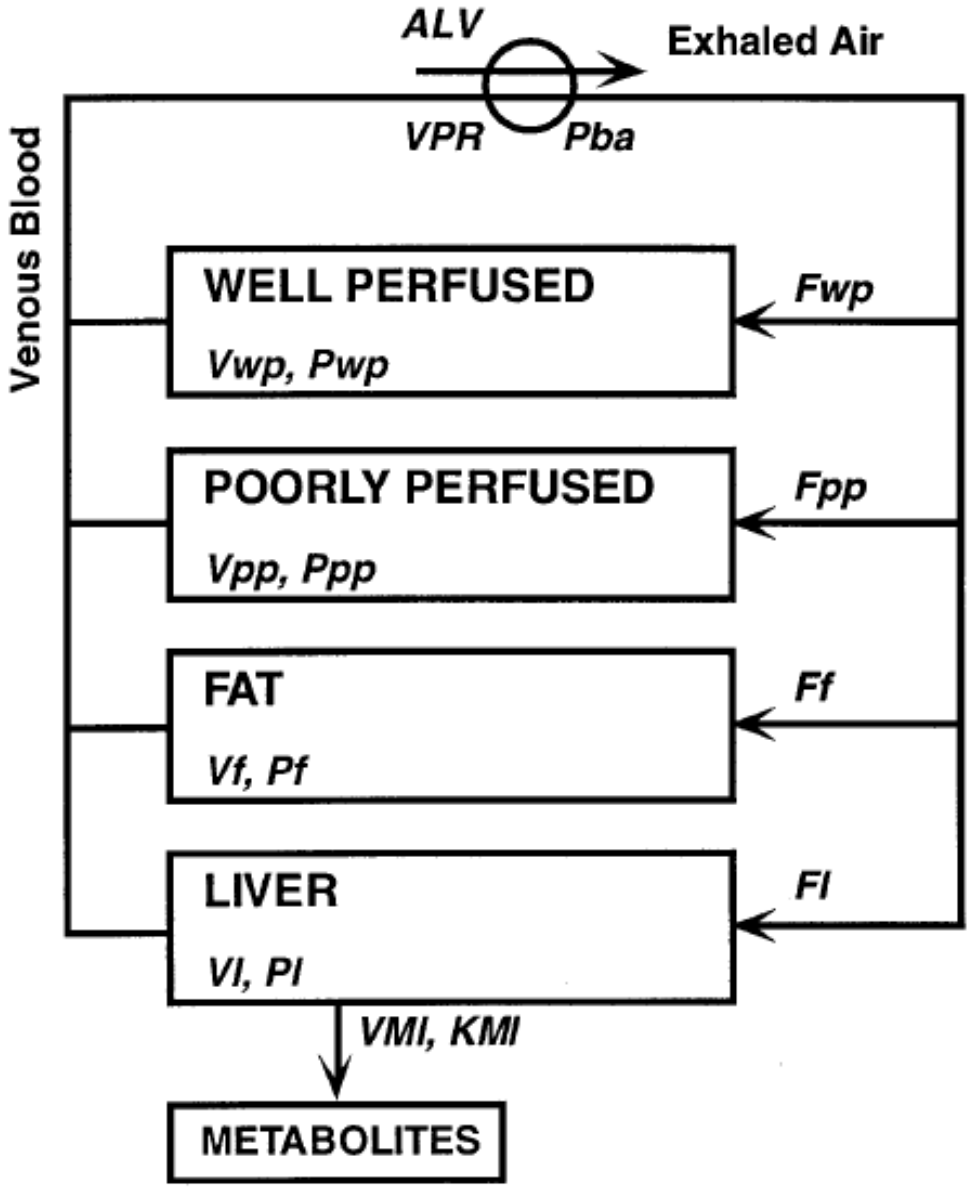

Some data

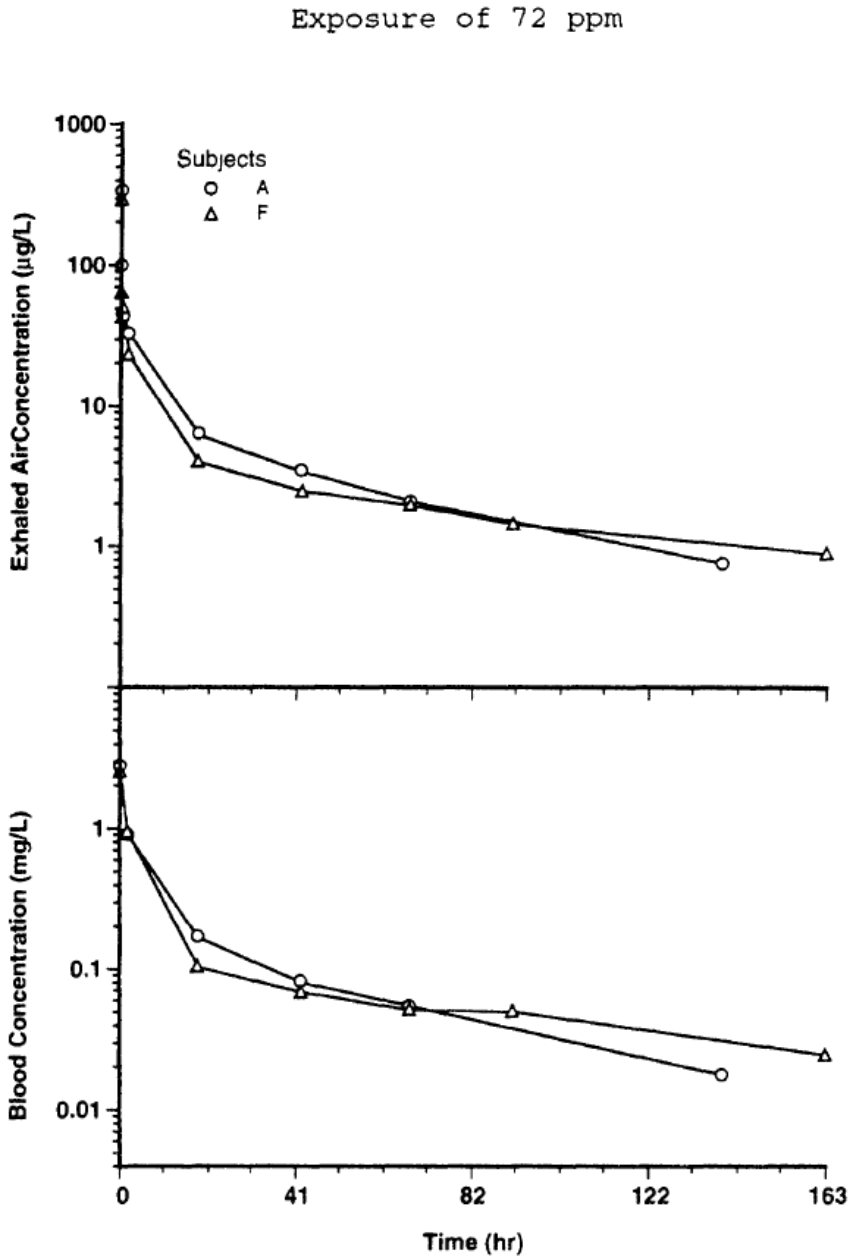
Exposure of 72 ppm
Subjects
A
F
Exhaled AirConcentration (µg/L)
1000
100
10
1
Blood Concentration (mg/L)
1
0.1
0.01
0
41
82
122
163
Time (hr)

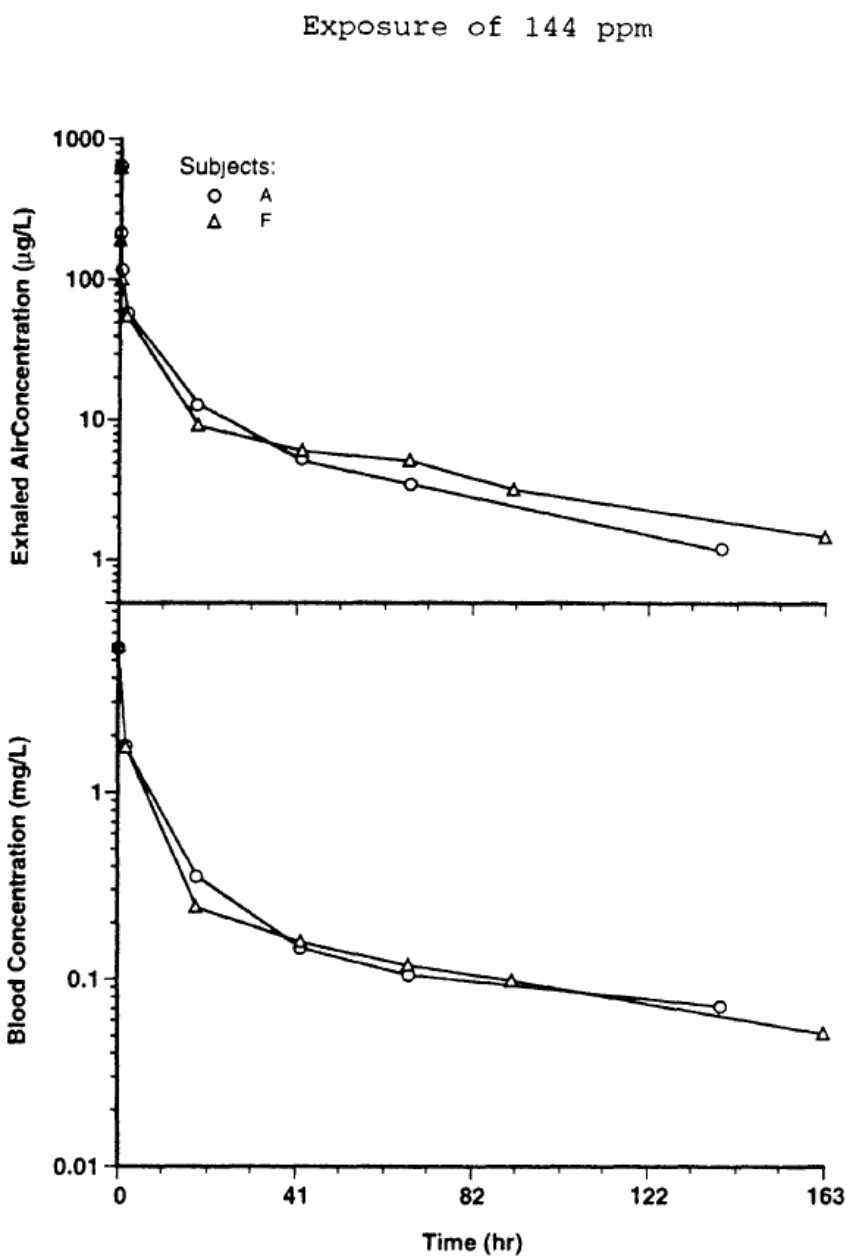
Exposure of 144 ppm
Subjects:
A
F
Exhaled AirConcentration (µg/L)
1000
100
10
1
Blood Concentration (mg/L)
1
0.1
0.01
0
41
82
122
163
Time (hr)

Simple statistical ideas did not work

- Fit 4-compartment model directly to data
- Assisted model fit
- 1 or 2-compartment model
- Simulation from prior distribution

Simple statistical ideas did not work:
Fit 4-compartment model directly to data

- Nonlinear least squares

- Fitting to each person separately:
  - Unstable: approx 30 data points, 15 parameters
  - “8 kg liver”
- Pooling data and estimating parameters for “the standard man”
  - Not useful for our goal of population inference

Simple statistical ideas did not work:
Assisted model fit

- Set some parameters to fixed values (from the pharmacology literature)
- Estimate the other parameters

- Couldn’t fit the data well
- Difficult to get fixed values for PERC-specific parameters such as equilibrium concentration ratios

Simple statistical ideas did not work: 1 or 2-compartment model

- Simpler model can be estimated easily and robustly

- Does not fit the data well
  - Most of the PERC leaves in a few hours, but some stays in the body after a week or more
- Not realistic for low-dose extrapolation

Simple statistical ideas did not work: Simulation from prior distribution

- Get prior information on parameters from pharmacology literature
- Try to fit data within these prior constraints

- Does not fit the data well
- Difficult to get good prior information for PERC-specific parameters such as equilibrium concentration ratios

Bayesian inference

- 4-compartment model
- 15 parameters for each person
- Prior information
  - Strong for some parameters (e.g., volume of liver)
  - Weak for others (e.g., Michaelis-Menten coef)
  - Model includes *uncertainty* and *variation*
- Posterior simulation: random walk through parameter space
- Inference for parameters and predictions
- Model checking

Hierarchical prior distributions

Scaled parameter		Prior on μ		Prior on Σ
	Multiplier	exp (M)	exp (S)	Truncation
Ventilation over perfusion ratio (VPR)	1	1.6	1.3	3
Blood flows				
Well perfused tissues (Fwp)	$0.7 \times MV/VPR$	0.48[e]	1.2[f]	3
Poorly perfused tissues (Fpp)	$0.7 \times MV/VPR$	0.20	1.2[f]	3
Fat (Ff)	$0.7 \times MV/VPR$	0.07	1.2[f]	3
Liver (Fl)	$0.7 \times MV/VPR$	0.25	1.1[f]	3
Volumes				
Well perfused tissues (Vwp)	LBM	0.28[e]	1.2[f]	3
Poorly perfused tissues (Vpp)	LBM	0.56[e]	1.2[f]	3
Liver (Vl)	LBM	0.033	1.1[f]	3
Blood/air partition coefficient (Pba)	1	12	1.5	3
Tissue/blood partition coefficients				
Well perfused tissues (Pwp)	1	4.8	1.5	3
Poorly perfused tissues (Ppp)	1	1.6	1.5	3
Fat (Pf)	1	125	1.5	3
Liver (Pl)	1	4.8	1.5	3
Max. rate of metabolism in liver (VMl)	$LBM^{0.7}$	0.042	10	2
Km in liver (KMl)	1	16	10	2

Fitting and using the model

- Use Gibbs sampler and Metropolis algorithm to take a random walk through parameter space
- Computationally intensive
  - Each step requires evaluation of the numerical differential equation solver
- Check inferences: Do they make sense?
- Re-run the model several times to simulate what would happen under different conditions

Inference for 6 individuals

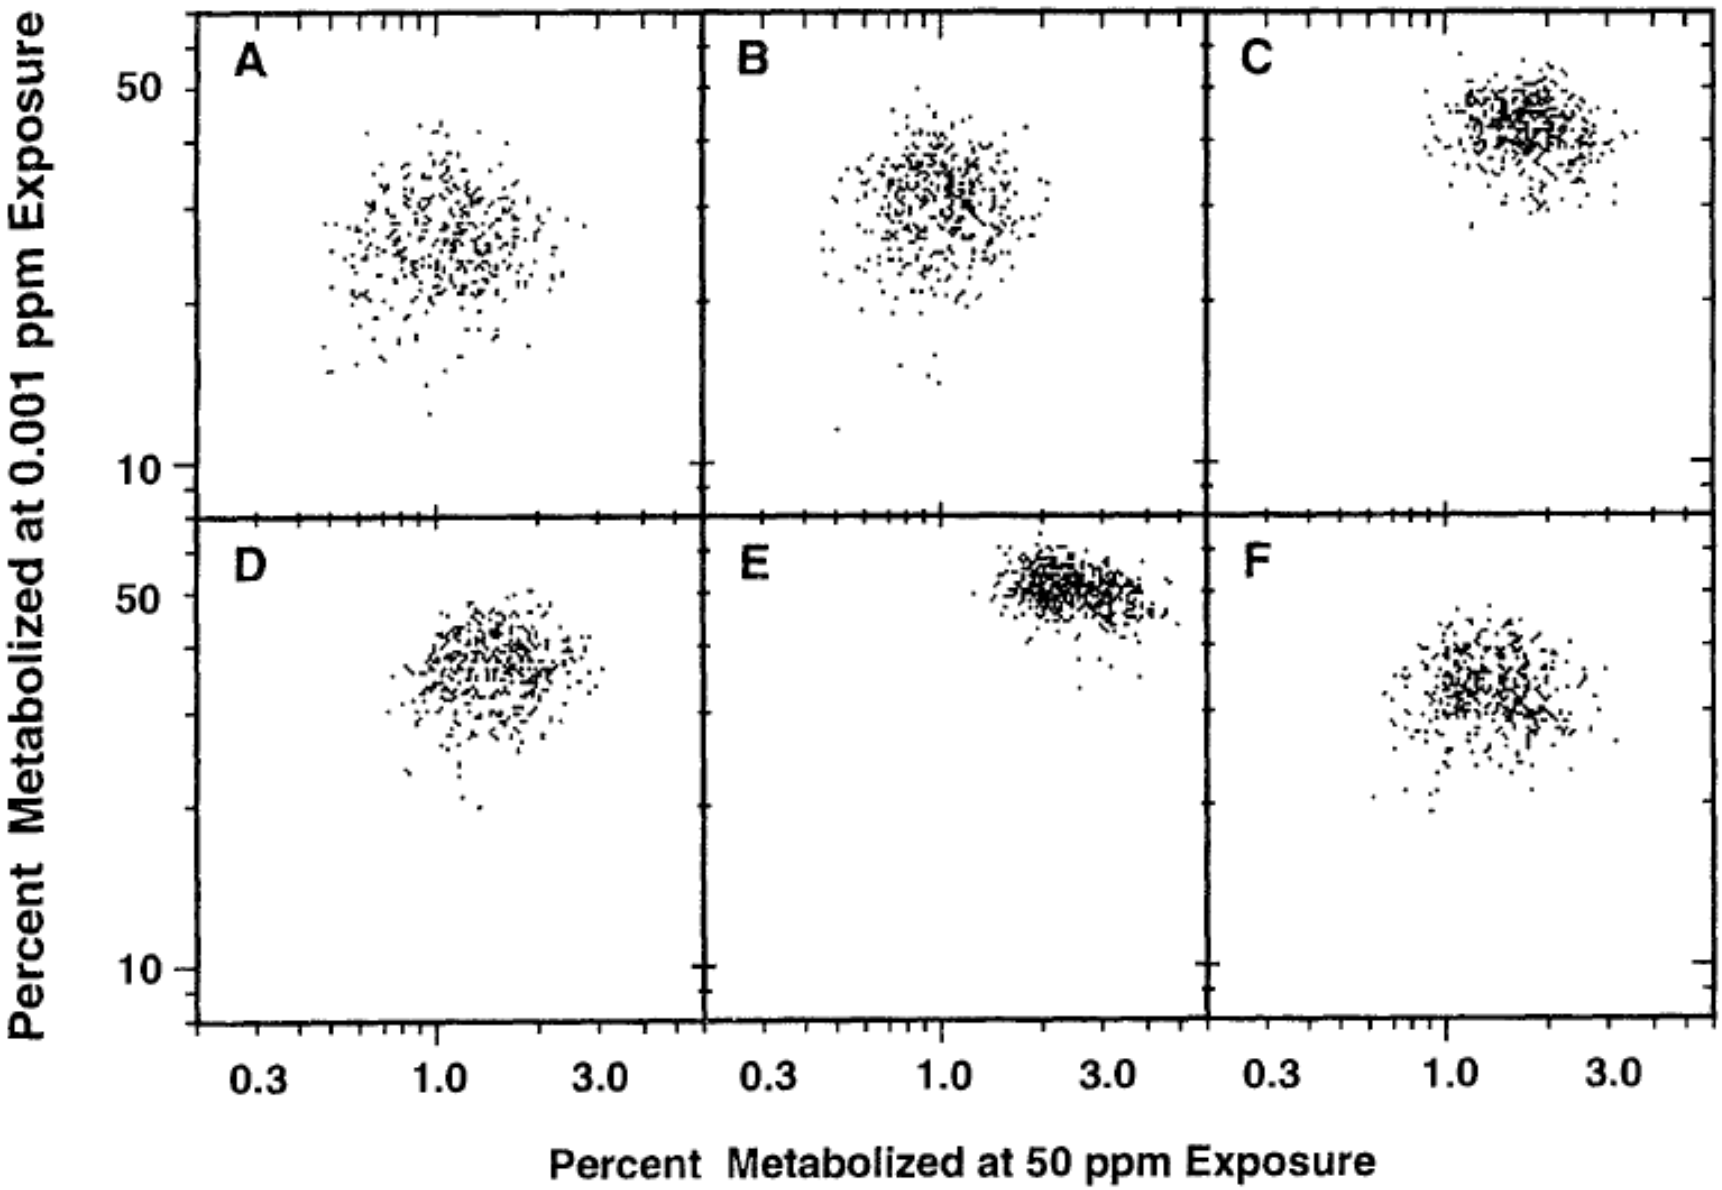

Inference for the population

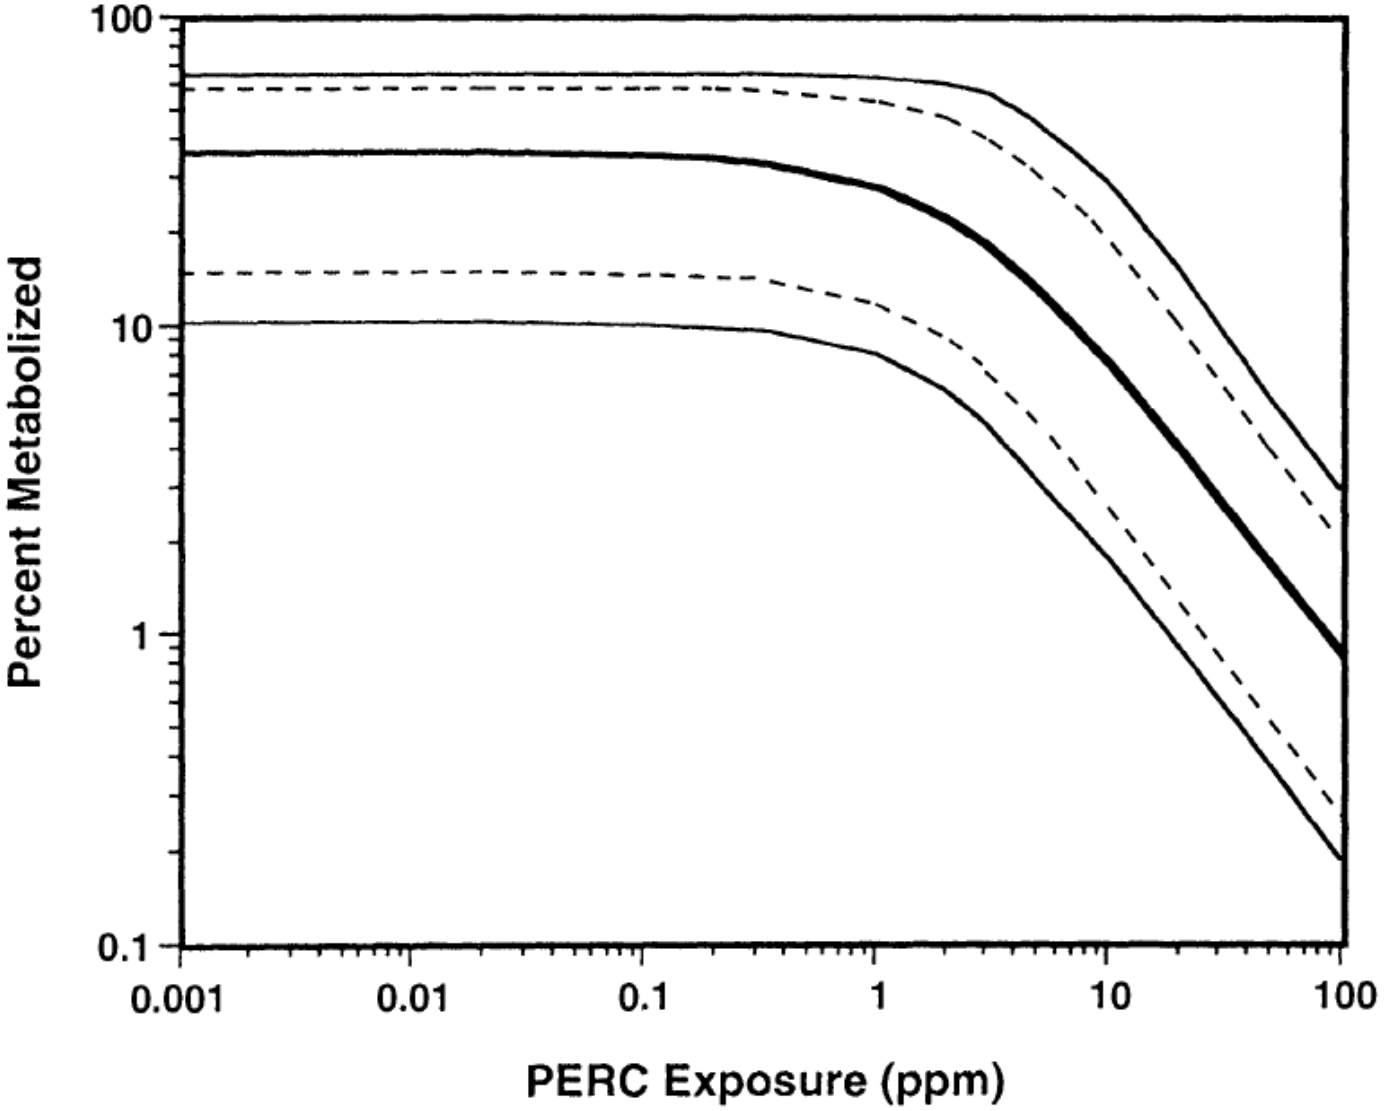

Prediction of data from a new study

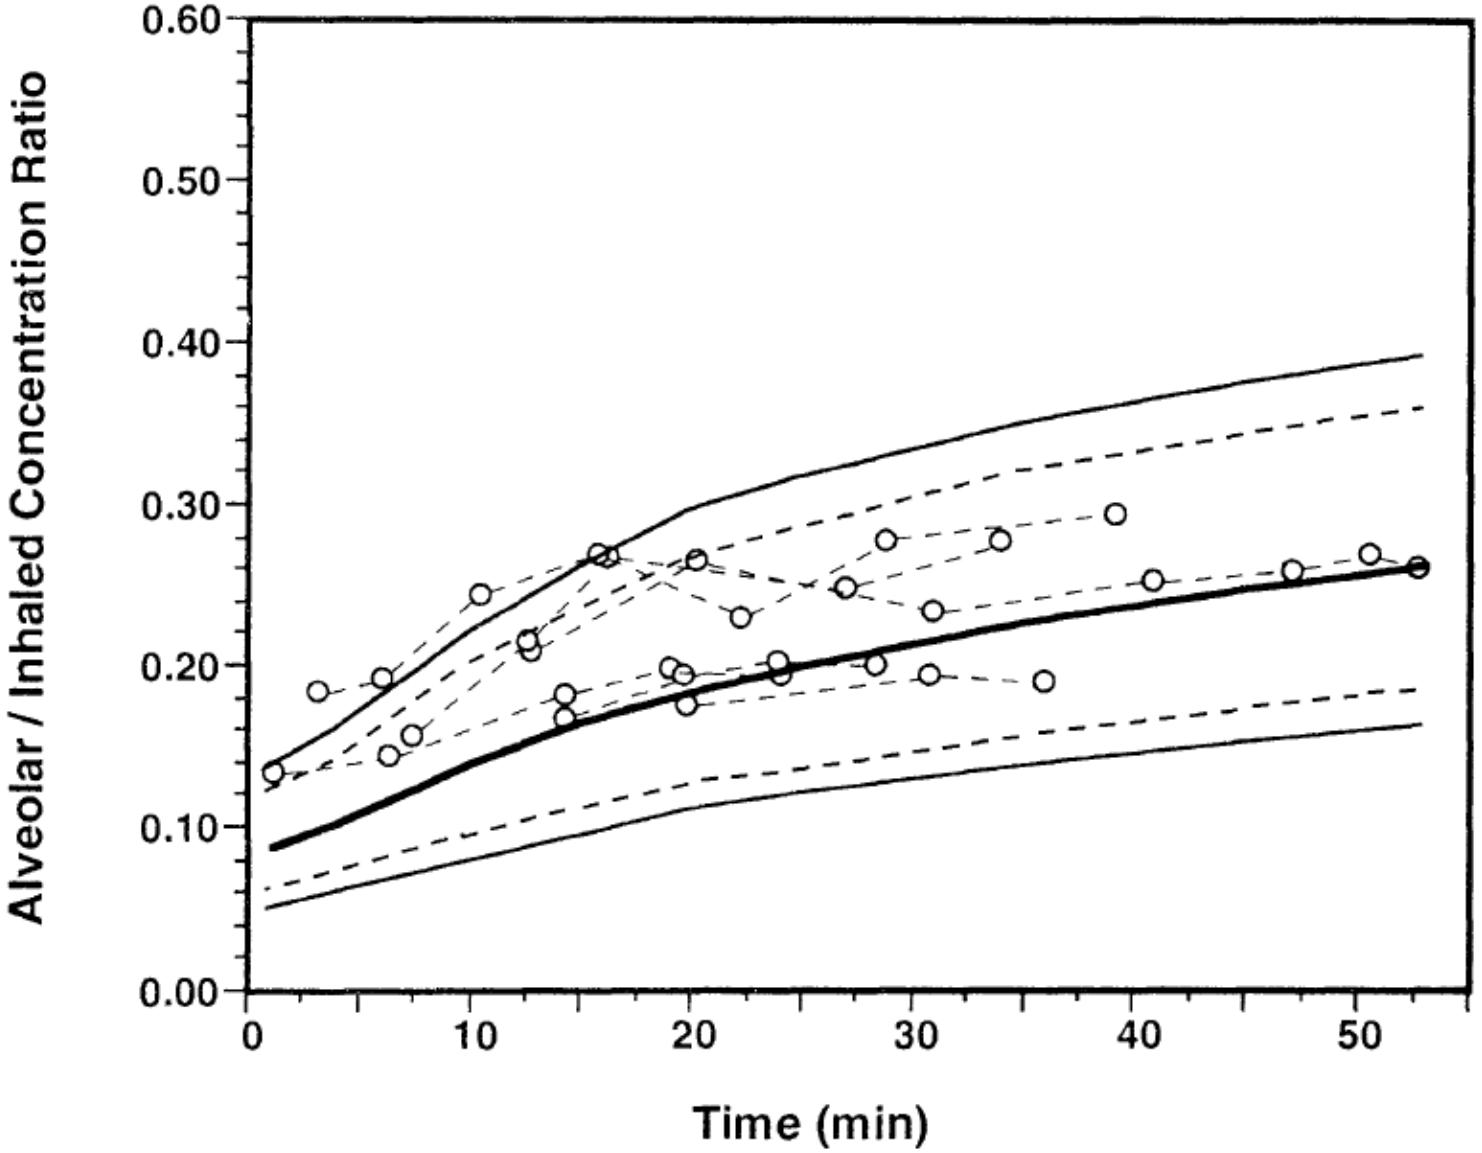

What we did

- Set up a hierarchical prior distribution with *uncertainty* and *population variation* for a 4-compartment model
- Fit the model to data (much computation)
- Checked inferences about parameters to see that they made sense
- Re-ran model under hypothetical low-dose, high-dose exposures

Challenges

- That was 1995; what have we done since?
- Goal of automating the process of inference
- More realistic models
  - More than 4 compartments
  - More happening within each compartment
- Technical challenges in modeling and computation
- Many thousands of drugs and toxins to study

Exchange of ideas between statistics and pharmacology

- Statistics → pharmacology
  - Bayesian inference for combining prior and data
  - Hierarchical models for population variation
- Pharmacology → statistics
  - Models for constrained parameters
  - Hierarchical prior distributions
  - New ideas in understanding and checking models

Putting it all together

1. Physiological pharmacokinetic model
2. Hierarchical population model
3. Prior information
4. Experimental data
5. Bayesian inference
6. Computation
7. Model checking

We need all of these!

1. Physiological pharmacokinetic model

- Without a physiological model, there is no good way to get prior information on the parameters
- We need physiological parameters (not just curve-fitting of the data) to efficiently combine information across different people

2. Hierarchical population model

- Without a population model, there generally are not enough data to estimate the model separately for each individual
- And there is too much variation among bodies (even among healthy young male volunteers) to pool all the data together and estimate common parameters

3. Prior information
4. Experimental data

- We need prior information. Otherwise, our estimates don't make sense (the 8 kg liver)
- We need experimental data to learn about perchloroethylene in particular

5. Bayesian inference

- Using Bayesian inference, we can find parameter that are consistent with both prior information and data, if such agreement is possible
- Automatically includes uncertainty and variability, so inferences can be plugged in directly to risk assessment and decision analysis

6. Computation

- Our models are big. Least squares, maximum likelihood, etc., are not enough
- Old-fashioned differential-equation solver is still sitting inside the model
- Our computers are never fast enough. We want more, more, more!

7. Model checking

- Check inferences about parameters
  - Do they make sense?
  - Are they consistent with prior distributions
- Check fit to data
- Check predictions on new data

Summary of PK example

- Population pharmacokinetic models have many moving parts
- Often, complexity in one place makes it *easier*, not harder, to add information in other places

Using Bayesian ideas to improve existing analyses

- Regularization (for example, avoiding estimates on the boundary of parameter space)
- Accounting for uncertainty (especially for decisions)
- Checking model fit
- Using models to combine different sources of information (partial pooling)
- Better dialogue with subject-matter experts (more windows into the model and data)

Now: Your questions and comments!