

Efficient Bayesian Experimental Design for **Benchmark Dose Estimation**:

Effects of Dose Group Allocation with **Overdispersed** Toxicological Data

Trevor Wroblewski

Bayesian Biostatistics

Leiden, NL

October 24, 2025

Benchmark dose (BMD) estimation underpins regulatory toxicology:

- ▶ Determines **acceptable exposure limits** for chemicals, pharmaceuticals, environmental contaminants
- ▶ Informs **risk assessment** and public health policy worldwide
- ▶ Replaces traditional NOAEL/LOAEL with continuous dose-response modeling

*Accurate, precise BMD estimates require getting both the **experimental design** and **statistical model** right.*

Explore to reduce model bias.

- ▶ Use many dose groups (G)
- ▶ Better captures the curve's true shape
- ▶ Guards against model misspecification

VS.

Replicate to reduce sampling variance.

- ▶ Use more total subjects (N)
- ▶ Increases within-group precision
- ▶ Crucial when data are noisy (overdispersed)

Given fixed resources, how do we resolve this trade-off?


Specifically, what combination of dose placement strategy, total subjects (N), and number of groups (G) best minimizes BMD error in the presence of **overdispersion** while retaining coverage?

Overdispersion as a Noise Structure


In an animal bioassay, individual subjects are not independent.

- ▶ **Litter Effects** → clustering or intra-litter correlation
 - ▶ Biological and genetic heterogeneity
 - ▶ Unmeasured environmental factors (handling, placement)
-

Step 1: Draw a probability p_i



Step 2: Generate count $Y_i \sim \text{Bin}(n, p_i)$



Overdispersion is structural

Beta–Binomial likelihood per group:

$$Y_i \sim \text{Beta-Binomial}(n_i, \phi p_i, \phi(1 - p_i)), \quad \mathbb{E}[Y_i/n_i] = p_i$$

Intraclass correlation (ICC): $\rho = \frac{1}{\phi+1}$.

$$\phi = 10 \Rightarrow \rho \approx 0.09 \quad (\text{typical } 0.03\text{--}0.15)$$

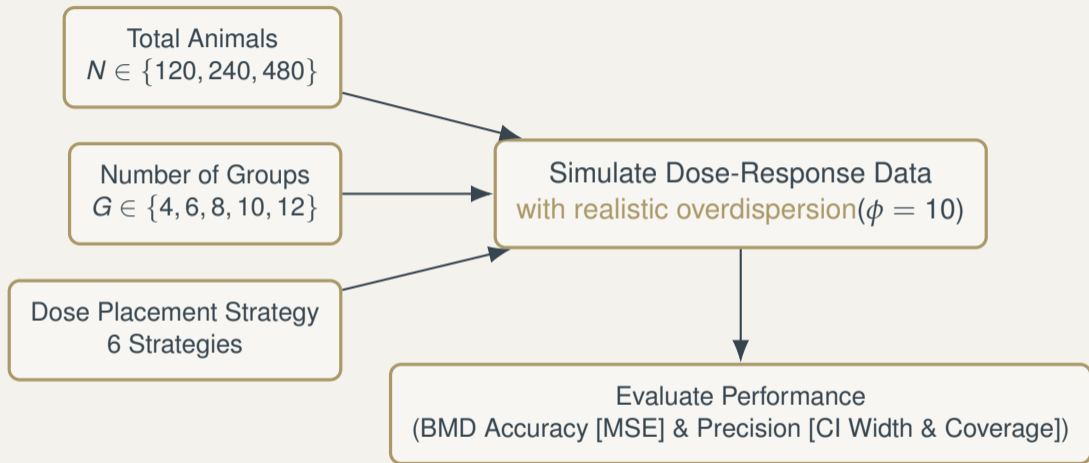
Ignoring overdispersion narrows CIs and can bias functionals; modeling it restores coverage.

The Unaddressed Complications

Models often assumes simple binomial variance for dichotomous outcomes.

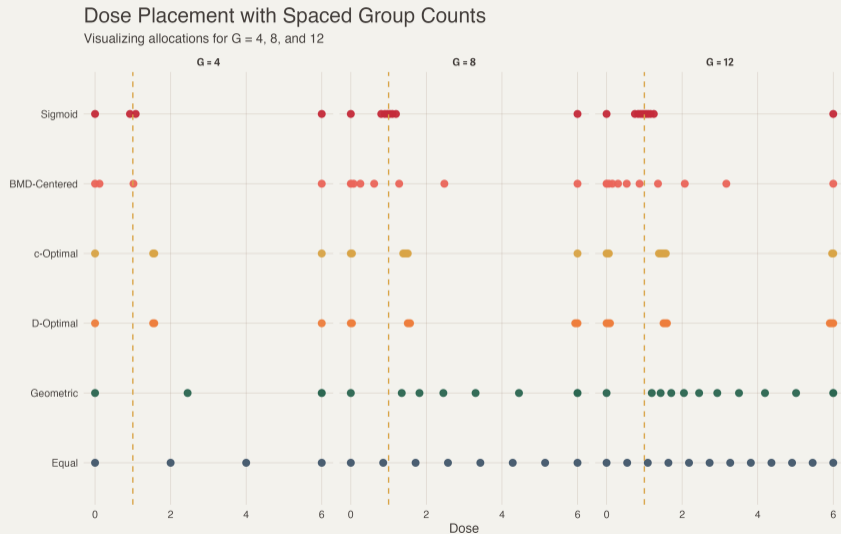
Toxicological reality is rarely so simple. Data frequently exhibit significant noise that can undermine theoretically optimal plans.

The Key Idea: A Simulated Contest




90 designs (6 strategies \times 5 group sizes \times 3 sample sizes) \times 30 replicates = 2,700.

Results: Strategies place dose groups very differently



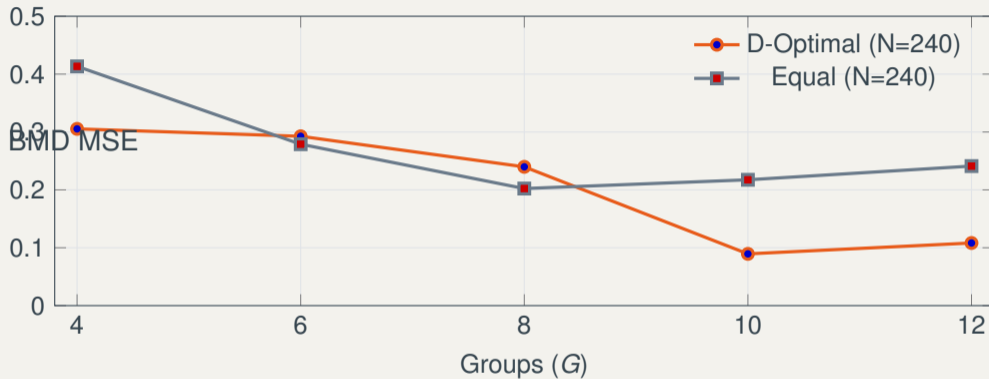
Dashed line: true BMD = 1.00.

Results: Resource Allocation Can Outweigh Total Sample Size



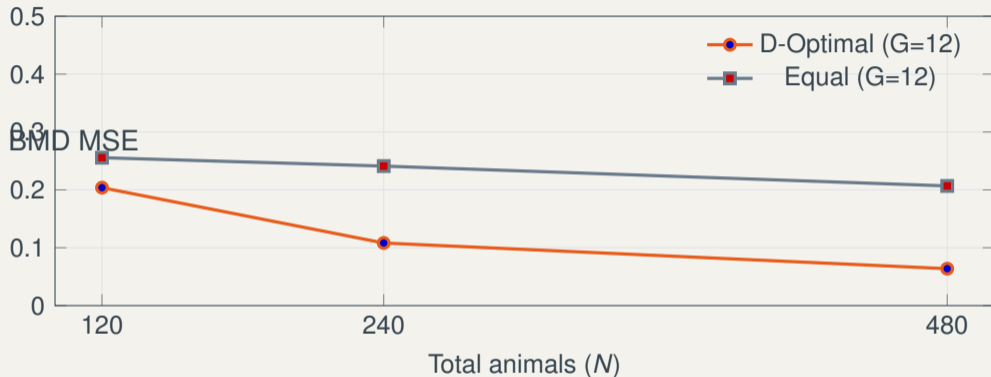
~5x lower MSE with half the animals (Beta-Binomial $\phi=10$, ICC ≈ 0.09).

Results: With better dose placement, more groups help



With suboptimal spacing, increasing N or G can sharpen the wrong curve; with optimal placement, extra groups extract more structure.

Results: With Fixed Groups, More Subjects Offer Marginal Benefits




With fixed $G=12$, increasing N reduces BMD MSE across strategies, but limited effect; optimal placement remains uniformly better than equal spacing.

Results: Full comparison for D-Optimal Allocation

BMD MSE for D-Optimal


Lower values indicate better performance



BMD MSE (lower is better)

BMD CI width for D-Optimal

Lower values indicate better precision



BMD 95% CI width (lower is better)

We've shown **dose placement strategy** is an important driver of BMD accuracy. More **subjects** and **dose groups** also can improve estimation.

But even with **optimal design**, you need the right inference approach:

- ▶ Likelihood choice (Beta-Binomial vs Binomial)
- ▶ Dose handling
- ▶ Prior specification

Thank you

Questions or suggestions?

Email t.wrobleski@lse.ac.uk

Affiliation London School of Economics, Department of Statistics

Dichotomous Hill:

$$p(d) = \gamma_0 + (1 - \gamma_0) \frac{d^k}{d^k + \theta^k}, \quad d \geq 0$$

BMD at BMR = $\rho = 0.10$:

$$\text{BMD} = \theta \left(\frac{\rho}{1 - \rho} \right)^{1/k}$$

Priors (estimation):

$$\gamma_0 \sim \text{Beta}(1, 1), \quad \theta \sim \mathcal{N}^+(3, 3), \quad k \sim \mathcal{N}^+(2, 2), \quad \phi \sim \text{Gamma}(2, 0.1)$$

Simulation: $\phi = 10$.

Appendix — Overdispersion rationale

Beta-Binomial ICC: $\rho = 1 / (\phi + 1)$. With $\phi = 10$, $\rho \approx 0.09$ (typical 0.03–0.15).

- ▶ Reflects litter/cluster effects common in toxicology.
- ▶ Ignoring overdispersion underestimates uncertainty; design should account for it.

Appendix — Optimal design (high level)

Fisher Information $M(\zeta; \Theta)$ under DH; criteria:

- ▶ **D-optimal**: maximize $\det M$ (global parameter precision).
- ▶ **c-optimal**: minimize $\text{Var}(\text{BMD}) \approx \mathbf{c}^\top M^{-1} \mathbf{c}$ (delta method).


Practical:

- ▶ Fedorov-style exchange; discretize to equal group sizes.
- ▶ Local optimality uses nominal (γ_0, θ, k) .

Appendix — Results: Full comparison at N=240

BMD MSE comparison (N = 240)


Lower values indicate better performance



BMD MSE (lower is better)

BMD credible interval width (N = 240)

Lower values indicate better precision




BMD 95% CI width (lower is better)

Appendix — Results: Full Comparison MSE

BMD Mean Squared Error (MSE) vs N and G

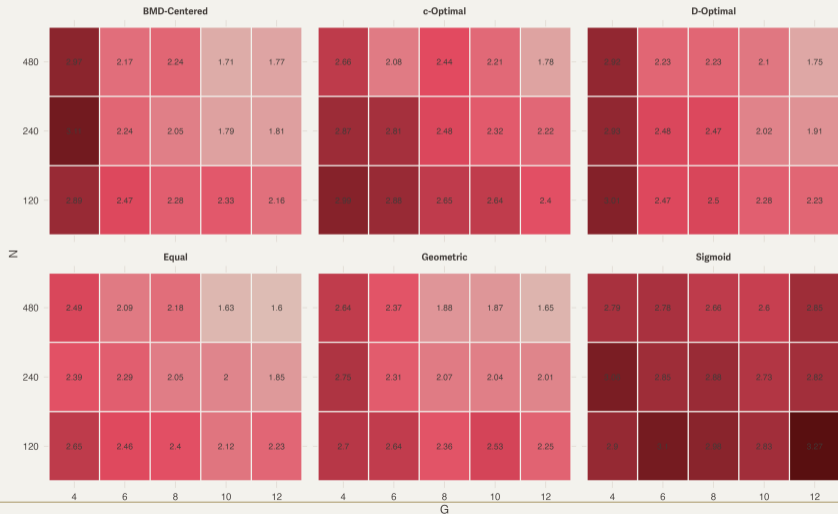
Faceted by strategy



Appendix — Results: Full Comparison Credible Interval Width


BMD credible interval width vs N and G

Faceted by strategy



Appendix — Coverage probability (N=240)


Coverage probability (N = 240)



Appendix — Results: Full comparison for Equal Allocation

BMD MSE for Equal


Lower values indicate better performance



BMD MSE (lower is better)

BMD CI width for Equal

Lower values indicate better precision



BMD 95% CI width (lower is better)

Appendix — Convergence summary

Threshold	# Designs	% of total
Perfect (100%)	38	42.2%
Excellent ($\geq 95\%$)	66	73.3%
Good ($\geq 90\%$)	84	93.3%
Lower ($< 90\%$)	6	6.7%

Lowest coverage noted: 80% for D-Optimal at $N=120$, $G=8$.