Bivariate Subgroup Analysis for Benefit-Harm Assessment

Ravi Varadhan & Nicholas Henderson

Division of Oncology Biostatistics
Johns Hopkins University
AND
University of Michigan
Department of Biostatistics

Subgroup Analysis and Heterogeneity in Benefit/Harm

- In clinical trials, subgroup analyses are regularly performed to investigate the consistency of treatment effect across patient subgroups.
- While subgroup analyses are frequently used for looking at heterogeneity in treatment effectiveness (HTE), heterogeneity in treatment safety is seldom examined.
- Even when heterogeneity in treatment-related adverse events (HTAE) is addressed, the subgroup analysis for safety is typically performed separately from the HTE analysis.

Subgroup Analysis and Heterogeneity in Risk/Benefit

- From a patient-centered perspective, separate subgroup analyses of HTAE and HTE ignore potentially important relationships between primary and safety outcomes.
- For example, suppose we have a binary primary event (PE) and binary adverse event (AE) whose joint distribution is given by

	Treatment 1		Treatment 2	
	AE	No AE	AE	No AE
PE	0.1	0	0	0.1
No PE	0.3	0.6	0.4	0.5

- $Pr(PE|Trt = 1) = 0.1$ $\quad$ $Pr(PE|Trt = 2) = 0.1$
- $Pr(AE|Trt = 1) = 0.4$ $\quad$ $Pr(AE|Trt = 2) = 0.4$

Subgroup Analysis and Heterogeneity in Risk/Benefit

- Only comparing the marginal distribution of a PE and AE can hide important risk-benefit considerations.
- For greater relevance to patients, subgroups analysis should assess variation in changes to patients risk-benefit profiles.
- A "truly" bivariate subgroup analyses would allow us to explore joint patient outcomes and heterogeneity of treatment impact.

SPRINT Trial

- The systolic blood pressure intervention (SPRINT) trial (N=9,361) investigated the effect of using a more stringent blood pressure target: $\leq$ 120 mm Hg (intensive) versus $\leq$140 mm Hg (standard).
- At the conclusion of the trial, 243 PEs were observed in the intensive treatment arm and 319 PEs were observed in the standard treatment arm.

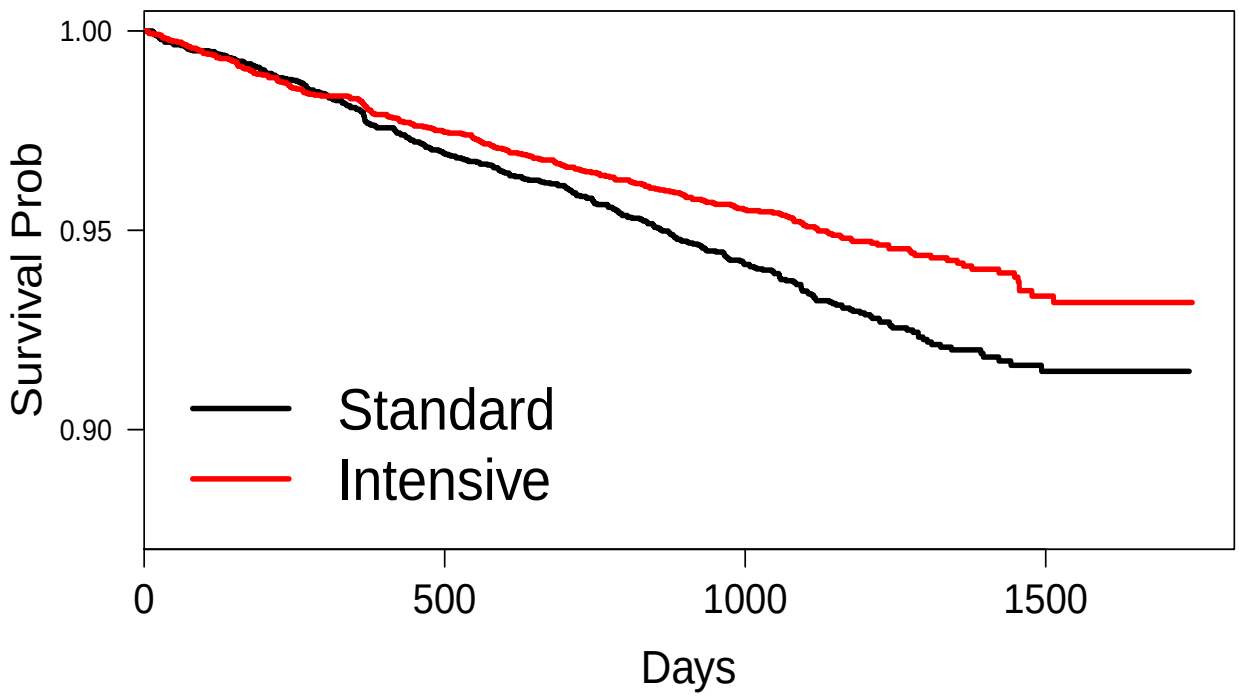

SPRINT Trial

- Joint counts of PEs and treatment-related serious adverse events (SAEs) in the SPRINT trial.

	Standard Treatment		Intensive Treatment	
	SAE	No SAE	SAE	No SAE
PE	18	301	30	213
No PE	100	4264	190	4245

Data for a Bivariate Subgroup Analysis

- **Primary Event:** T_i - time to the primary event
- **Safety Event:** W_i - a binary outcome (an indicator of whether or not patient experienced at least one AE)
- $Y_i = \min\{T_i, C_i\}$: duration of follow-up
  $\delta_i = \mathbf{1}\{T_i \leq C_i\}$: event indicator
- A_i - treatment arm assignment
- G_i - indicator of subgroup membership
- Subgroup memberships G_i are for a "fully stratified" subgroup analysis as opposed to the more typical univariate "one variable at-a-time" subgroup analysis.

Joint model for survival and binary outcomes

- **Model Parameters:**

$$\alpha, \{\lambda_{awg}\}, \{p_{ag}\}; \quad a = 0, 1; \quad w = 0, 1; \quad g = 1, \ldots, G.$$

- We want to specify the joint distribution of (T_i, W_i) conditional on treatment arm assignment A_i and subgroup membership G_i.

- This is done by assuming that

$$\begin{aligned} T_i | A_i = a, W_i = w, G_i = g &\sim \text{Weibull}(\alpha, \lambda_{awg}) \\ W_i | A_i = a, G_i = g &\sim \text{Bernoulli}(p_{ag}). \end{aligned}$$

- The joint distribution $(T_i, W_i)|A_i = a, G_i = g$ depends on the parameters $(\alpha, \lambda_{a0g}, \lambda_{a1g}, p_{ag})$.

Summary Statistics for the Exponential Model

- When $\alpha = 1$, the time-to-event T_i follows an exponential distribution.
- Moreover, when $\alpha = 1$, the likelihood only depends on the following summary statistics

$$
\begin{aligned}
D_{awg} &= \sum_{i=1}^{n} \delta_i I(A_i = a) I(W_i = w) I(G_i = g) \\
U_{awg} &= \sum_{i=1}^{n} Y_i I(A_i = a) I(W_i = w) I(G_i = g) \\
V_{ag} &= \sum_{i=1}^{n} W_i I(A_i = a) I(G_i = g)
\end{aligned}
$$

Summary Statistics for the SPRINT Trial

Subgroup			Standard Treatment			Intensive Treatment			
CKD	Age	Sex	(D_{01g}, U_{01g})	(D_{00g}, U_{00g})	V_{0g}	(D_{11g}, U_{11g})	(D_{10g}, U_{10g})	V_{1g}	N_g
No	< 75	Male	(3, 92.0)	(96, 5546.1)	29	(10, 174.1)	(54, 5446.5)	61	3528
Yes	< 75	Male	(3, 48.4)	(28, 1364.7)	16	(5, 103.2)	(25, 1227.7)	32	858
No	≥ 75	Male	(1, 19.6)	(46, 1277.5)	8	(1, 57.7)	(26, 1265.9)	19	913
Yes	≥ 75	Male	(6, 40.0)	(47, 977.7)	15	(7, 88.3)	(38, 974.8)	31	730
No	< 75	Female	(0, 40.1)	(31, 2641.9)	13	(0, 67.5)	(25, 2705.4)	20	1706
Yes	< 75	Female	(2, 42.4)	(12, 948.7)	14	(4, 50.2)	(16, 1000.3)	16	617
No	≥ 75	Female	(0, 37.8)	(16, 835.1)	12	(1, 60.9)	(18, 778.2)	20	568
Yes	≥ 75	Female	(3, 30.8)	(25, 612.1)	11	(2, 66.4)	(11, 634.9)	21	441

Modeling Subgroup Parameters (Saturated Model)

- The distribution of summary statistics $(D_{awg}, U_{awg}, V_{ag})$ depends on hazard rate parameters λ_{awg} and AE probabilities p_{ag}.
- Assume that

$$\log(\lambda_{awg}) = \mathbf{x}_g^T \boldsymbol{\beta}_{aw} \qquad \text{and} \qquad \text{logit}(p_{ag}) = \mathbf{z}_g^T \boldsymbol{\gamma}_a$$

- For example, in a **saturated model**, we have

$$\log(\lambda_{awg}) = \beta_{aw,g} \qquad \text{and} \qquad \text{logit}(p_{ag}) = \gamma_{a,g},$$

where $\boldsymbol{\beta}_{aw} = (\beta_{aw,1}, \ldots, \beta_{aw,G})^T$ and $\boldsymbol{\gamma}_a = (\gamma_{a,1}, \ldots, \gamma_{a,G})^T$.

Modeling Subgroup Parameters (Additive Model)

- In the saturated model, the λ_{awg} and p_{ag} are treated separately with no additional information used to indicate relationships among the subgroups.
- Subgroups that share much of their characteristics are treated the same as subgroups that are quite different.
- Some regression structure linking the parameters λ_{awg}, p_{ag} can induce more sensible correlation.

Modeling Subgroup Parameters (Additive Model)

- In the **additive model**, λ_{awg} and p_{ag} are determined additively from the variables comprising subgroup g.

- For example, if we have four subgroups arising from each combination of the variables age (young/old) and smoking behavior (smoker/non-smoker)

$$
\begin{aligned}
(\text{Smoker/Young}) \quad & \log(\lambda_{aw1}) = \beta_{aw,1} \\
(\text{Smoker/Old}) \quad & \log(\lambda_{aw2}) = \beta_{aw,1} + \beta_{aw,2} \\
(\text{Non-Smoker/Young}) \quad & \log(\lambda_{aw3}) = \beta_{aw,1} + \beta_{aw,3} \\
(\text{Non-Smoker/Old}) \quad & \log(\lambda_{aw4}) = \beta_{aw,1} + \beta_{aw,2} + \beta_{aw,3}
\end{aligned}
$$

- We could also use a regression that includes higher-order interactions or a model that makes additional assumptions about how hazards can vary across subgroups:
  $\lambda_{a1g}/\lambda_{a0g} = \phi$ (proportional hazards)

Specifying the Prior (Saturated Model)

- In the saturated model, we assumed

$$\log(\lambda_{awg}) = \beta_{aw,g} \qquad \text{for } g = 1, \ldots, G.$$

- For each treatment separately, assume the following for the $\beta_{aw,g}$

$$\begin{bmatrix} \beta_{a0,1} \\ \beta_{a1,1} \end{bmatrix}, \ldots, \begin{bmatrix} \beta_{a0,G} \\ \beta_{a1,G} \end{bmatrix} \Bigg| \boldsymbol{\mu}_a, \boldsymbol{\tau}_a \sim \text{Normal}\left(\begin{bmatrix} \mu_{a0} \\ \mu_{a1} \end{bmatrix}, \begin{bmatrix} \tau^2_{a0} & 0 \\ 0 & \tau^2_{a1} \end{bmatrix} \right)$$

- Place a proper, but "vague" prior on the joint distribution of $\boldsymbol{\mu}_a$ while allowing for user-specified prior correlation.

Specifying the Prior (Saturated Model)

$$\log(\boldsymbol{\tau}_a) \sim \text{Normal}\left(\begin{bmatrix} \log(1/2) \\ \log(1/2) \end{bmatrix}, \begin{bmatrix} \sigma^2_{\tau,a} & \sigma^2_{\tau,a}\rho_{\tau,a} \\ \sigma^2_{\tau,a}\rho_{\tau,a} & \sigma^2_{\tau,a} \end{bmatrix} \right)$$

- For variance components $\boldsymbol{\tau}_a = (\tau_{a0}, \tau_{a1})$, use a "quasi informative" or "weakly informative" prior as a default.
- Place most prior mass on plausible variation across subgroups.
- Consider the hazard ratio $\lambda_{awj}/\lambda_{awk}$ between two subgroups.

$$Pr\left\{ \frac{1}{4} \leq \frac{\lambda_{awj}}{\lambda_{awk}} \leq 4 \middle| \boldsymbol{\tau}_a \right\} \geq 0.95 \quad \text{whenever} \quad \tau_{aw} \leq 1/2$$

- Prior median of τ_{a0} and τ_{a1} is $1/2$. Choose $\sigma^2_{\tau,a}$ so that

$$Pr\{\tau_a \leq 2\} \approx 0.95$$

- Prior for correlation: $\rho_{\tau,a} \sim \text{Uniform}(-1, 1)$.

Targets of Inference

- For each subgroup, we want to target some parameter (or a collection of paramaters) which captures important changes in the joint distribution (T_i, W_i) from treatment $A_i = 0$ to $A_i = 1$.
- With our Bayesian setup, this is easy for any chosen target because we can just transform the posterior draws of parameters λ_{awg} and p_{ag} as needed.
- In our implementation, we consider the following targets:

  (1) Heterogeneity in joint binary outcomes

  (2) Heterogeneity in utility gained

  (3) Heterogeneity in probability of outcome improvement

Heterogeneity in Utility Gain/Loss

- Think of the composite score for patient i

$$H_i = b_1 W_i \min\{T_i, \tau\} + b_2(1 - W_i)\min\{T_i, \tau\}$$

  for weights $b_2 > b_1 > 0$.

- Patient i receives a "score" of $b_1 T_i$ if surviving to time $T_i < \tau$ while experience an AE some time in $(0, T_i)$.

- Patient i receives a "score" of $b_2 T_i$ if surviving to time $T_i < \tau$ while never experiencing an AE.

- For each subgroup g, the parameter of interest is the expected difference in the composite score

$$\eta_g = E[H_i|A_i = 1, G_i = g] - E[H_i|A_i = 0, G_i = g]$$

SPRINT Trial: Heterogeneity in Utility Gain/Loss ($b_2 = 1$)

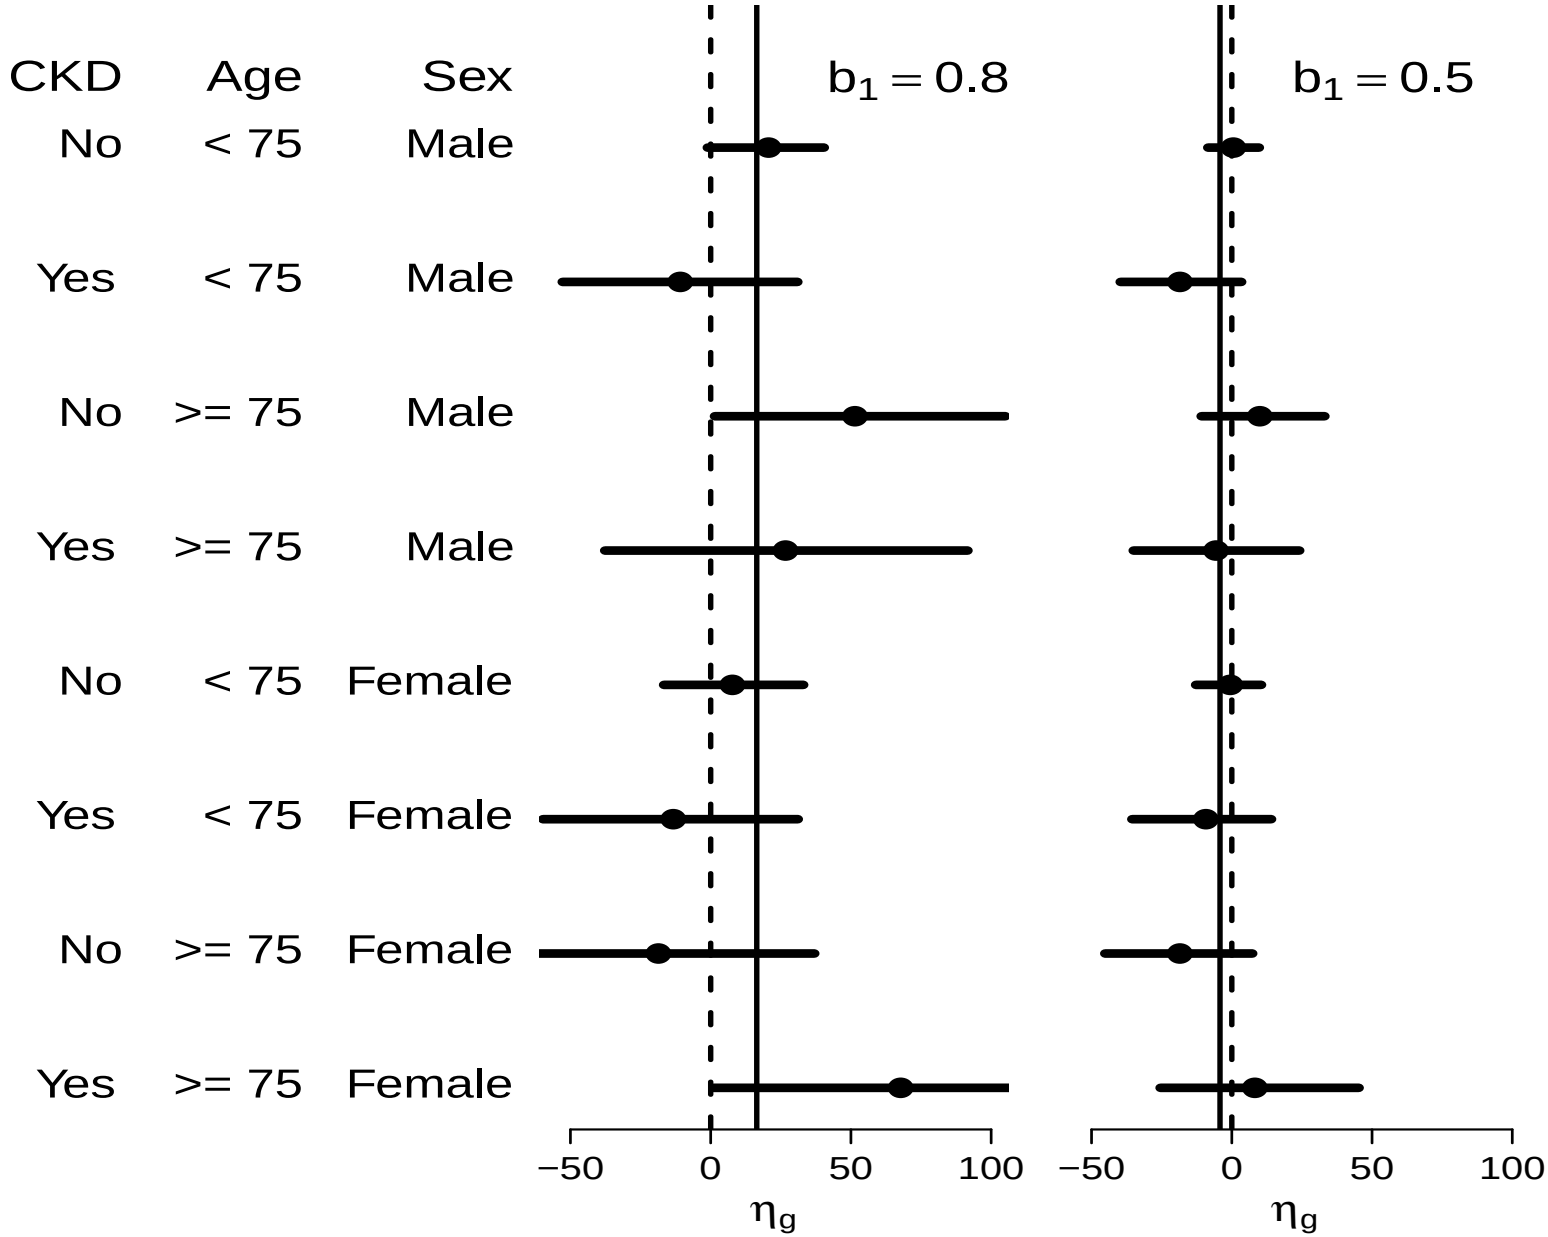

Heterogeneity in Probability of Outcome Improvement (Assuming $A_i = 1$ and $A_j = 0$)

Outcome		Preferred Treatment
$T_i > T_j(1+\delta)$	$W_i = 1, W_j = 0$	$A = 1$
$T_i \leq T_j(1+\delta)$	$W_i = 1, W_j = 0$	$A = 0$
$T_j > T_i(1+\delta)$	$W_i = 0, W_j = 1$	$A = 0$
$T_j \leq T_i(1+\delta)$	$W_i = 0, W_j = 1$	$A = 1$
$T_i > T_j$	$W_i = 1, W_j = 1$	$A = 1$
$T_i > T_j$	$W_i = 0, W_j = 0$	$A = 1$
$T_i \leq T_j$	$W_i = 1, W_j = 1$	$A = 0$
$T_i \leq T_j$	$W_i = 0, W_j = 0$	$A = 0$

The subgroup-specific parameters of interest are

$$\phi_g = 2 \times \Pr\Big\{\text{outcome i} > \text{outcome j} \Big| A_i = 1, A_j = 0, G_i = g, G_j = g\Big\} - 1$$

SPRINT Trial: Outcome Improvement Measure

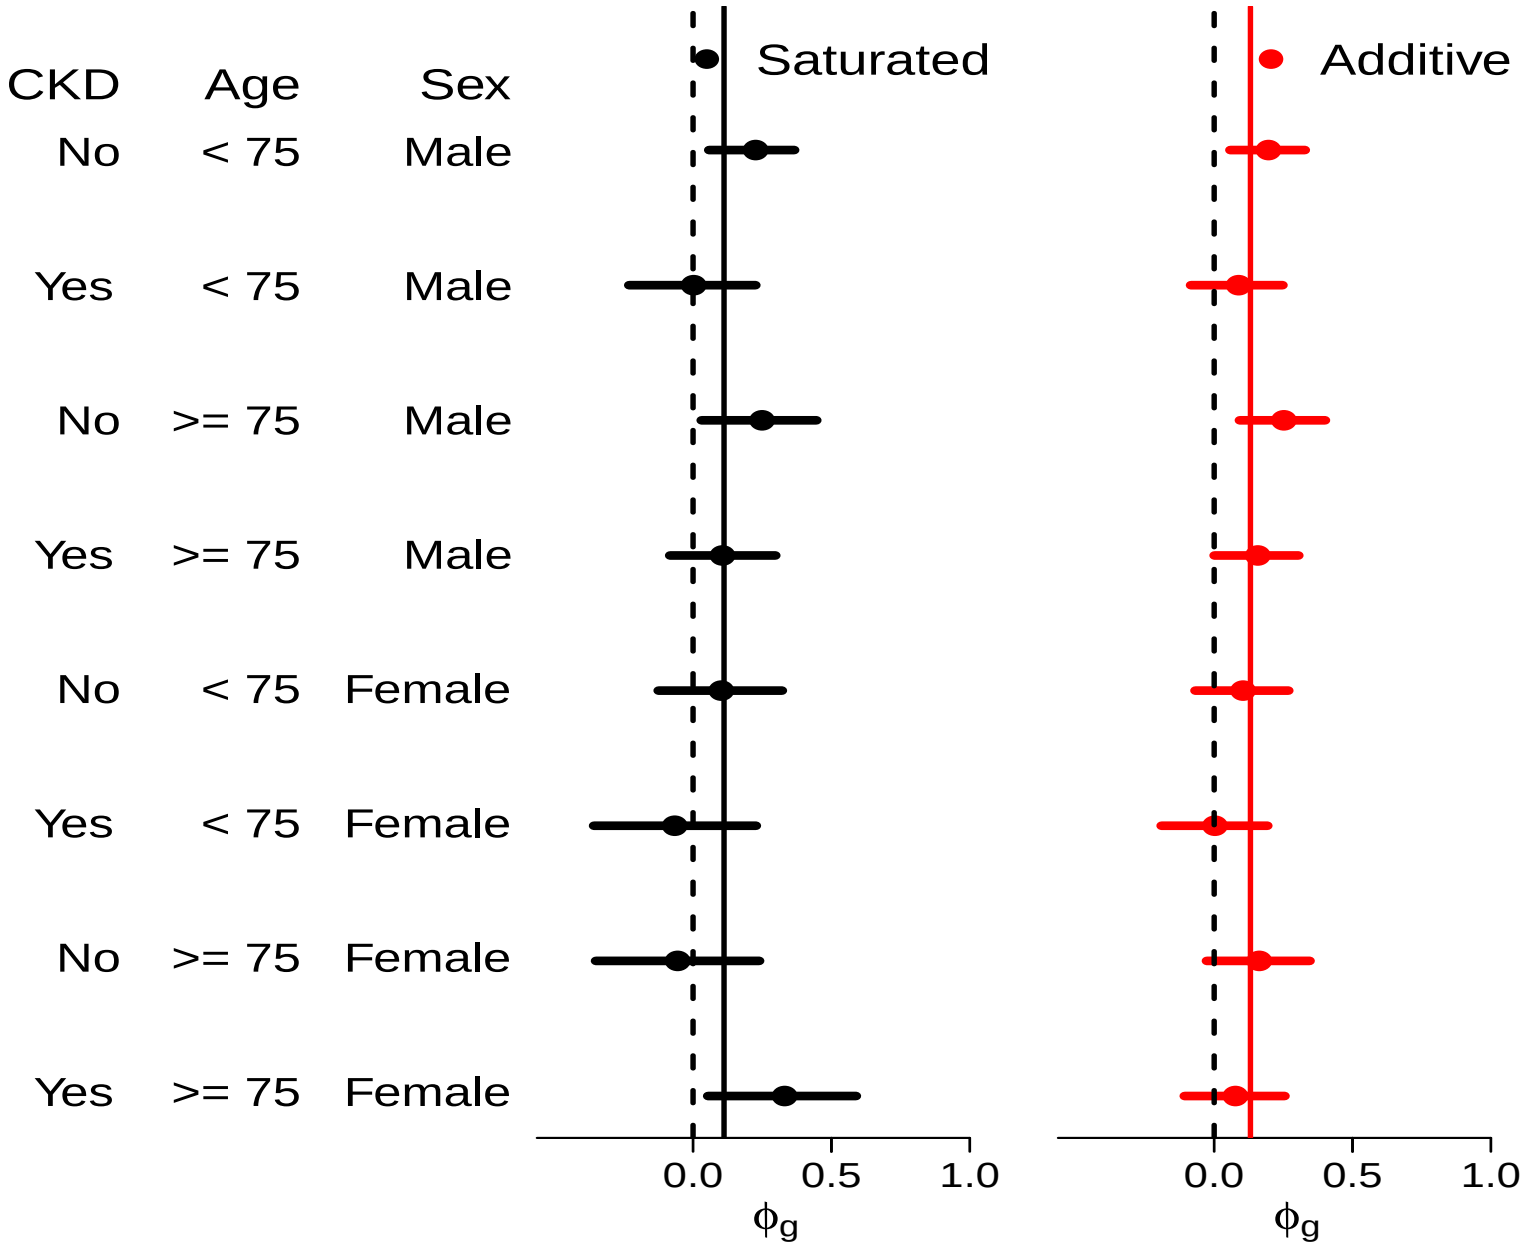

Summary

- Bayesian methods, such as the models proposed here, allow us to undertake patient-centered "joint" benefit-harm assessments
- Patient-level data is not required - only summaries are required
- Software implementing the discussed bivariate subgroup analyses is available at `http://hteguru.com/index.php/bbsga/`
- Software allows one to perform posterior predictive checks and model comparisons.

References

1. Henderson, N.C. and R. Varadhan (2018). "Bayesian Bivariate Subgroup Analysis for Risk-Benefit Evaluation" *Health Services and Outcomes Research Methodology*, **18**(4), 244-264.

2. The SPRINT Research Group (2015), "A randomized trial of intensive versus standard blood-pressure control", *The New England Journal of Medicine*, **373**(22), 21032116.

3. Jones, H. E., Ohlssen, D. I., Neuenschwander, B., Racine, A. and Branson, M. (2011), "Bayesian models for subgroup analysis in clinical trials, *Clinical Trials*, **8**, 129143.

4. Evans, S. R. and Follmann, D. (2016), "Using outcomes to analyze patients rather than patients to analyze outcomes: A step toward pragmatism in benefit:risk evaluation", *Statistics in Biopharmaceutical Research*, **8**(4), 386393.