

Data combination in seamless Phase II/III designs

Christopher Jennison and Lisa Hampson

Department of Mathematical Sciences,

University of Bath, UK

<http://people.bath.ac.uk/mascj>

2009 Quality & Productivity Research Conference

IBM T. J. Watson Research Center,

Yorktown Heights, June 2009

Outline of talk

1. Seamless Phase II/III clinical trials
2. Example: Formulation of the selection and testing problem
3. Four testing procedures
4. Power curves for the four tests
5. Understanding tests' performance via an "optimal" test
6. Assessments under other configurations of the treatment effects
7. Comments and conclusions

1. Phase II and III clinical trials

Phase IIa

A “proof of concept” study intended to show evidence of drug activity at some dose.

Phase IIb

A trial to compare several doses or other variants of a treatment against control, with the aim of selecting a treatment for testing in a “confirmatory” Phase III trial.

Phase III

A trial intended to demonstrate the treatment selected in Phase II is superior to the control.

Seamless Phase II/III trial

A “seamless” design combines Phase IIb and Phase III, selecting a treatment and continuing to test this against control in a single study.

Seamless Phase II/III trials

Joint planning

Combining Phases II and III promises a more rapid final decision since the “white space” between phases is eliminated.

This requires substantial preparation to anticipate all the eventualities at the end of Phase II and, in each case, document how the study should continue.

Combining data

We shall discuss the issue of combining information in Phase II and Phase III data.

This option is of interest when the patient population and primary endpoint remain the same over both phases.

But, the final hypothesis test will need proper consideration of the selection process in Phase II, which generated the hypothesis now being tested.

2. Example

Jennison & Turnbull (*J. Biopharm. Statistics*, 2007, Sec. 5) consider:

Phase II

Four treatments and a control are compared, with $m_1 = 100$ observations on each.

Estimated treatment effects are $\hat{\theta}_{1,i}$, $i = 1, \dots, 4$.

The treatment i^* with highest $\hat{\theta}_{1,i}$ is selected for Phase III.

Phase III

Treatment i^* is compared against control, with $m_2 = 500$ observations on each.

Estimated treatment effect is $\hat{\theta}_{2,i^*}$.

Conclusion

A final decision is made, based on $\hat{\theta}_{1,1}, \dots, \hat{\theta}_{1,4}$ and $\hat{\theta}_{2,i^*}$.

Example: Requirements

There are four null hypotheses, $H_i: \theta_i \leq 0, i = 1, \dots, 4$.

If dose i^* is selected for Phase III, we focus on testing $H_{i^*}: \theta_{i^*} \leq 0$.

Family-wise error

We wish to control **family-wise error**, so, for all vectors $\theta = (\theta_1, \dots, \theta_4)$,

$$Pr\{\text{Reject any true } H_i\} \leq \alpha.$$

Then, the probability of falsely claiming significance for the selected i^* is at most α .

Power

When some of the θ_i are greater than zero, we wish to have a high probability of selecting an effective treatment and rejecting the associated null hypothesis.

Formally, define power as

$$Pr\{\text{Select the treatment } j \text{ with maximum } \theta_i \text{ and reject } H_j: \theta_j \leq 0\}.$$

Example: Further details

Phase II

Responses follow distributions

$$N(\mu_i, \sigma^2), \text{ on treatments 1 to 4,}$$

$$N(\mu_c, \sigma^2), \text{ on control,}$$

where $\sigma^2 = 25$.

Treatment effects are $\theta_i = \mu_i - \mu_c$.

The estimated effects are

$$\hat{\theta}_{1,i} = \hat{\mu}_{1,i} - \hat{\mu}_{1,c} \sim N(\theta_i, \frac{2\sigma^2}{m_1}),$$

with correlation 0.5 between each pair.

Early stopping

If $\hat{\theta}_{1,i^*} = \max_i(\hat{\theta}_{1,i}) < 0$, stop for futility, otherwise proceed to Phase III.

Example: Further details

Phase III

The estimated effect of treatment i^* in Phase III is

$$\hat{\theta}_{2,i^*} = \hat{\mu}_{2,i^*} - \hat{\mu}_{2,c} \sim N(\theta_{i^*}, \frac{2\sigma^2}{m_2}).$$

Question: How should one make the final decision based on

$$\hat{\theta}_{1,1}, \dots, \hat{\theta}_{1,4} \quad \text{and} \quad \hat{\theta}_{2,i^*} ?$$

— How to adjust for multiplicity? How to weight data from the two phases?

Jennison & Turnbull (*J. Biopharm. Statistics*, 2007) consider several methods:

Conventional: Final decision is based on $\hat{\theta}_{2,i^*}$ alone.

Bauer and Köhne: Applying combination tests to data from the two phases.

TSE: A design of Thall, Simon & Ellenberg (*Biometrika*, 1988).

3.1. Conventional: Final test based on Phase III data alone

It is not necessary to adjust for data-driven selection of hypothesis H_{i^*} since this is to be tested using only the *new* data from Phase III.

So, we can define

$$Z = \frac{\hat{\theta}_{2,i^*}}{\sqrt{(2\sigma^2/m_2)}}$$

and reject $H_{i^*}: \theta_{i^*} \leq 0$ in favour of $\hat{\theta}_{i^*} > 0$ if

$$Z > \Phi^{-1}(1 - 0.025) = 1.96.$$

With stopping for futility if $\hat{\theta}_{1,i^*} < 0$, the overall type I error rate is only 0.020.

For type I error rate 0.025 when $m_1 = 100$ and $m_2 = 500$, reject H_{i^*} if

$$Z > 1.86.$$

3.2 Bauer & Köhne (*Biometrics*, 1994): Combination tests

Testing a null hypothesis H_0 against a one-sided alternative, we have p -values:

P_1 from stage 1 (Phase II), P_2 from stage 2 (Phase III).

Under H_0 , P_1 and P_2 have independent $U(0, 1)$ distributions.

a) Inverse χ^2 test

$$-\ln(P_1 P_2) \sim \frac{1}{2} \chi_4^2.$$

Hence, a size α test is obtained by rejecting H_0 if

$$-\ln(P_1 P_2) > \frac{1}{2} \chi_{4, 1-\alpha}^2.$$

This χ^2 test was originally proposed for combining results of several studies by R. A. Fisher (1932) *Statistical Methods for Research Workers*.

Bauer & Köhne: Combination tests

b) Inverse normal test

Let $Z_1 = \Phi^{-1}(1 - P_1)$ and $Z_2 = \Phi^{-1}(1 - P_2)$.

Pre-specify weights w_1 and w_2 , where $w_1^2 + w_2^2 = 1$.

Then under H_0 , $w_1 Z_1 + w_2 Z_2 \sim N(0, 1)$.

Hence, a size α test is obtained by rejecting H_0 if

$$w_1 Z_1 + w_2 Z_2 > \Phi^{-1}(1 - \alpha).$$

Multiple comparisons

Since the choice of the null hypothesis H_{i^*} is based on the data that generated P_1 and Z_1 , we need to allow for multiple comparisons if we are to control the family-wise type I error rate.

Closed testing procedures (Marcus et al, *Biometrika*, 1976)

We have 4 null hypotheses, $H_i: \theta_i \leq 0$, for $i = 1, \dots, 4$. For each subset I of $\{1, \dots, 4\}$, we define the intersection hypothesis

$$H_I = \bigcap_{i \in I} H_i.$$

We construct a level α test of each intersection hypothesis H_I : this test rejects H_I with probability at most α whenever all hypotheses specified in H_I are true.

Closed testing procedure

The simple hypothesis $H_j: \theta_j \leq 0$ is rejected overall if, and only if, H_I is rejected for every set I containing index j .

Proof of strong control of family-wise error rate

Let \tilde{I} be the set of indices of all true hypotheses H_i . For a family-wise error to be committed, $H_{\tilde{I}}$ must be rejected. Since $H_{\tilde{I}}$ is true, $Pr\{\text{Reject } H_{\tilde{I}}\} = \alpha$ and, thus, the probability of a family-wise error is no greater than α .

Closed testing procedures using combination tests

(Bretz et al. and Schmidli et al, *Biometrical Journal*, 2006)

To test the intersection hypothesis H_I , we combine p -values $P_{1,I}$ from stage 1 and $P_{2,I}$ from stage 2.

Stage 1 (Phase II)

We have p -values $P_{1,i}$ for the individual null hypotheses H_i .

If H_I is the intersection of m simple hypotheses, let $P_{1,(k)}$, $k = 1, \dots, m$, denote the m p -values for these hypotheses in increasing order.

Using the method of Simes (*Biometrika*, 1986), the p -value for testing H_I is

$$P_{1,I} = \min_{k=1, \dots, m} (m P_{1,(k)} / k).$$

Stage 2 (Phase III)

We are only interested in hypotheses H_I where I contains i^* and we set the p -value for each of these to be $P_{2,I} = P_{2,i^*}$.

Closed testing procedures using combination tests

At the end of Phase III, we reject H_{i^*} if the specified combination test rejects every H_I where I contains i^* .

a) Inverse χ^2 test

The inverse χ^2 test rejects H_I if

$$-\ln(P_{1,I} P_{2,I}) > \frac{1}{2} \chi_{4, 1-0.025}^2 = 5.572.$$

The overall type I error rate under $\theta = (0, 0, 0, 0)$ is 0.021, less than 0.025 because of stopping for futility at Phase II and the conservatism of Simes' test.

For $m_1 = 100$ and $m_2 = 500$, rejecting each H_I if

$$-\ln(P_{1,I} P_{2,I}) > 5.376$$

gives an overall type I error rate under $\theta = (0, 0, 0, 0)$ of 0.025.

Closed testing procedures using combination tests

b) Inverse normal test

We take $w_1 = \sqrt{m_1/(m_1 + m_2)}$ and $w_2 = \sqrt{m_2/(m_1 + m_2)}$, where m_1 and m_2 are the Phase II and Phase III sample sizes per treatment.

The inverse normal test rejects H_I if

$$w_1 Z_{1,I} + w_2 Z_{2,I} > \Phi^{-1}(1 - 0.025) = 1.96.$$

The overall type I error rate for this procedure under $\theta = (0, 0, 0, 0)$ is 0.020, due to stopping for futility at Phase II and the conservatism of Simes' test.

For $m_1 = 100$ and $m_2 = 500$, rejecting each H_I if

$$w_1 Z_{1,I} + w_2 Z_{2,I} > 1.86$$

gives an overall type I error rate under $\theta = (0, 0, 0, 0)$ of 0.025.

3.3 Thall, Simon & Ellenberg (*Biometrika*, 1988)

Define

$$Z_{1,i^*} = \frac{\hat{\theta}_{1,i^*}}{\sqrt{(2\sigma^2/m_1)}} \quad \text{and} \quad Z_{2,i^*} = \frac{\hat{\theta}_{2,i^*}}{\sqrt{(2\sigma^2/m_2)}}.$$

The hypothesis H_{i^*} is rejected if

$$w_1 Z_{1,i^*} + w_2 Z_{2,i^*} > 2.20.$$

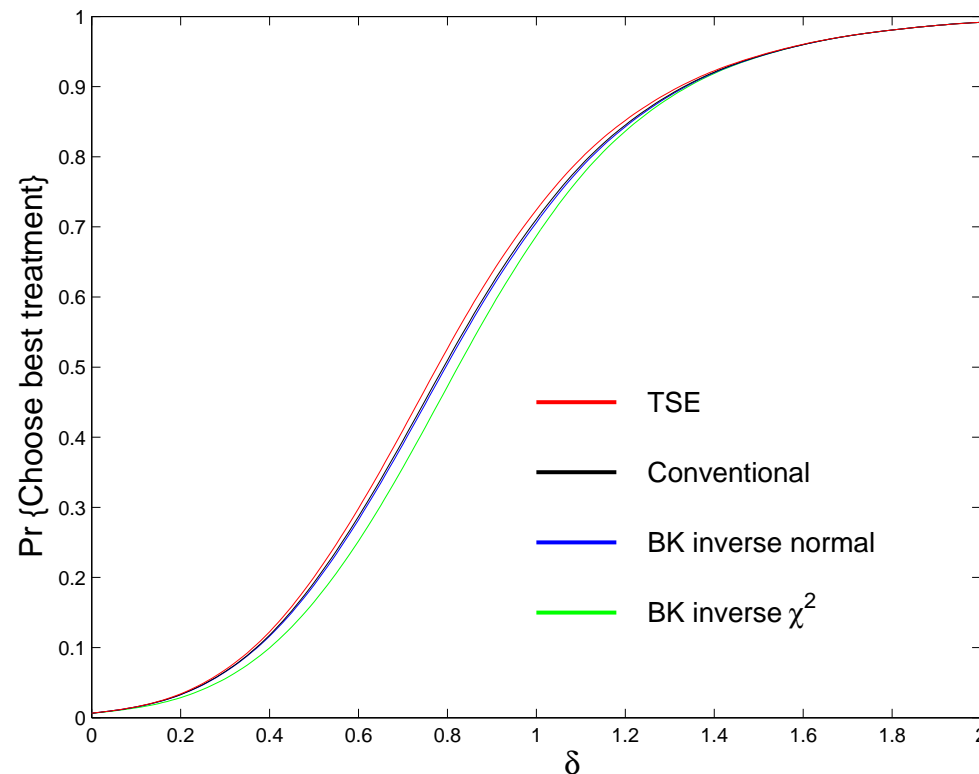
where $w_i = \sqrt{(m_i/(m_1 + m_2))}$, $i = 1, 2$, as before.

The critical value 2.20 is chosen to give overall type I error probability 0.025 if $\theta = (0, 0, 0, 0)$ and this guarantees a maximum family-wise error rate of 0.025.

It is easy to check the final decision is based on the difference in mean responses on treatment i^* and the control, pooled across Phases II and III — a very natural way in which to combine the relevant data.


4. Properties of the four tests

Power of four 2-stage selection/testing procedures when $\theta = (0, 0, 0, \delta)$,
i.e., three are ineffective and the other has effect size δ :



NB: The sample size distribution is the same for all four methods.

Properties of the four tests



The conventional procedure — with no data combination — is surprisingly efficient: only slightly worse than the Thall, Simon & Ellenberg design, and superior to inverse χ^2 and inverse normal combination tests.

Differences in power equate to differences in sample size of around 4% to 8%.

5. Understanding the power curves

We shall seek an “optimal” procedure for this situation, examine its form, and compare with the four tests.

Requirements:

$$Pr\{\text{Type I error}\} = 0.025 \text{ when } \theta = (0, 0, 0, 0),$$

Maximum possible power when $\theta = (0, 0, 0, \delta)$ or a permutation of this.

To find the optimal procedure:

Consider a Bayes decision problem with:

Prior:

$$0.2 \text{ on } \theta = (0, 0, 0, 0),$$

$$0.2 \text{ on each permutation of } \theta = (0, 0, 0, \delta)$$

Costs:

$$c_1 \text{ for rejecting any } H_i \text{ when } \theta = (0, 0, 0, 0),$$

$$-c_2 \text{ for rejecting } H_j \text{ when } \theta_j = \delta.$$

An “optimal” procedure

The Bayes rule minimises

$$0.2 \times c_1 \times \{\text{Type I error probability under } \theta = (0, 0, 0, 0)\} \\ - 4 \times 0.2 \times c_2 \times \{\text{Power at } \theta = (0, 0, 0, \delta)\}.$$

Hence, it maximises power at $\theta = (0, 0, 0, \delta)$ among all procedures with the same type I error probability under $\theta = (0, 0, 0, 0)$.

Choosing c_1 and c_2 so that the procedure has total type I error rate 0.025 under $\theta = (0, 0, 0, 0)$ will give the solution to our stated problem.

It may seem strange not to penalise, say, selecting treatment 2 and rejecting H_2 when $\theta = (0, 0, 0, \delta)$. However, family-wise error is protected under this θ .

Most importantly, this formulation is what we need to solve the problem as posed.

An “optimal” procedure

Consider the final analysis after selecting treatment i^* and observing Phase III data.

Without loss of generality, suppose $i^* = 4$.

Possible costs are:

If $\theta = (0, 0, 0, 0)$ Reject $H_{i^*} \Rightarrow$ cost c_1

If $\theta = (0, 0, 0, \delta)$ Reject $H_{i^*} \Rightarrow$ cost $-c_2$

Other values of θ No cost, whether H_{i^*} is rejected or accepted.

It follows that the Bayes optimal decision rule depends on the likelihood ratio of the Phase II and Phase III data under $\theta = (0, 0, 0, \delta)$ and $\theta = (0, 0, 0, 0)$.

To evaluate this likelihood ratio, recall

$$\hat{\theta}_{1,i} \sim N\left(\theta_i, \frac{2\sigma^2}{m_1}\right), \quad i = 1, \dots, 4, \quad \hat{\theta}_{2,i^*} \sim N\left(\theta_{i^*}, \frac{2\sigma^2}{m_2}\right)$$

and each pair $(\hat{\theta}_{1,i}, \hat{\theta}_{1,j})$ has correlation 0.5.

An “optimal” procedure

Some algebra shows the optimal procedure rejects H_{i^*} if

$$\frac{8 m_1}{m_1 + m_2} \{ \hat{\theta}_{1,i^*} - \sum_{j \neq i^*} \frac{1}{4} \hat{\theta}_{1,j} \} + \frac{5 m_2}{m_1 + m_2} \hat{\theta}_{2,i^*} > k$$

for some constant k .

The terms $\hat{\theta}_{1,j}$ have negative weights due to their positive correlations with $\hat{\theta}_{1,i^*}$.

Remember that for all the values of θ in our stated problem, three of the θ_j are equal to zero. Positive values of $\hat{\theta}_{1,j}$ for $j \neq i^*$ could be attributable to unusually negative responses on the common control arm, and this detracts from the significance of a high positive θ_{i^*} .

An “optimal” procedure

We have seen that the “optimal” procedure gives a negative weight to the Phase II estimates $\hat{\theta}_{1,j}$ $j \neq i^*$.

The Bauer & Köhne methods use closed testing procedures and these combine p -values by Simes rule to test intersection hypotheses. Higher values of $\hat{\theta}_{1,j}$ $j \neq i^*$ contribute positively to rejection of H_{i^*} .

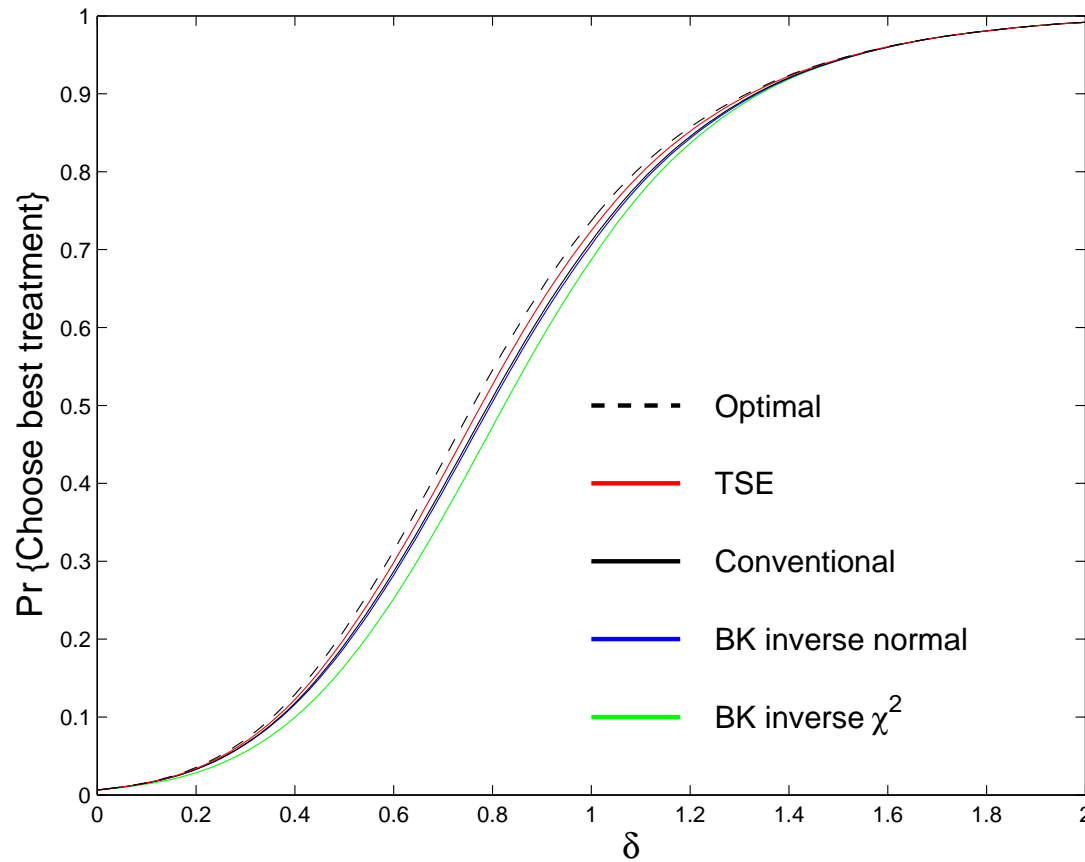
This is a natural way to “borrow strength” when one expects related treatments to have similar effects, but it is counter-productive when the other treatment effects are, in fact, all zero.

The TSE procedure gives zero weight to the other $\hat{\theta}_{1,j}$ s, and so is closer to the “optimal” decision rule.


The conventional rule ignores all phase II estimates and at least it does not suffer from including terms $\hat{\theta}_{1,j}$ with the wrong sign.

Properties of four tests plus the “optimal”

Power of four 2-stage selection/testing procedures when $\theta = (0, 0, 0, \delta)$, plus the “optimal” procedure as a benchmark.



Properties of four tests plus the “optimal”



Differences in power equate to differences in sample size of around 5% to 10%.

But, there is a problem with the “optimal” procedure

The problem with the “optimal” procedure

This procedure rejects H_{i^*} if

$$\frac{8 m_1}{m_1 + m_2} \{ \hat{\theta}_{1,i^*} - \sum_{j \neq i^*} \frac{1}{4} \hat{\theta}_{1,j} \} + \frac{5 m_2}{m_1 + m_2} \hat{\theta}_{2,i^*} > k.$$

Here, k is chosen so type I error probability is 0.025 when $\theta = (0, 0, 0, 0)$.

Negative weights for $\hat{\theta}_{1,j}$, $j \neq i^*$, imply high type I error probability for, say, $\theta = (-\lambda, -\lambda, -\lambda, 0)$ where $-\lambda$ is large and negative. So, the “optimal” test does not satisfy the family-wise error condition.

Question: What procedure is truly optimal, maximising power at $\theta = (0, 0, 0, \delta)$ and permutations of this while protecting the family-wise error rate?

Conjecture: Increasing the weights of $\hat{\theta}_{1,j}$, $j \neq i^*$, so they are non-negative and weighting $\hat{\theta}_{1,i^*}$ and $\hat{\theta}_{2,i^*}$ by the inverse of their variances gives the TSE rule — suggesting this should be close to optimal for this problem.

6. Performance of tests under other configurations of θ

Example 2:

Suppose θ is a permutation of $(\gamma\delta, \gamma\delta, \gamma\delta, \delta)$, where $0 \leq \gamma < 1$.

Aim:

High power to select the treatment i^* for which $\theta_{i^*} = \delta$ and then reject H_{i^*} .

Optimal rules

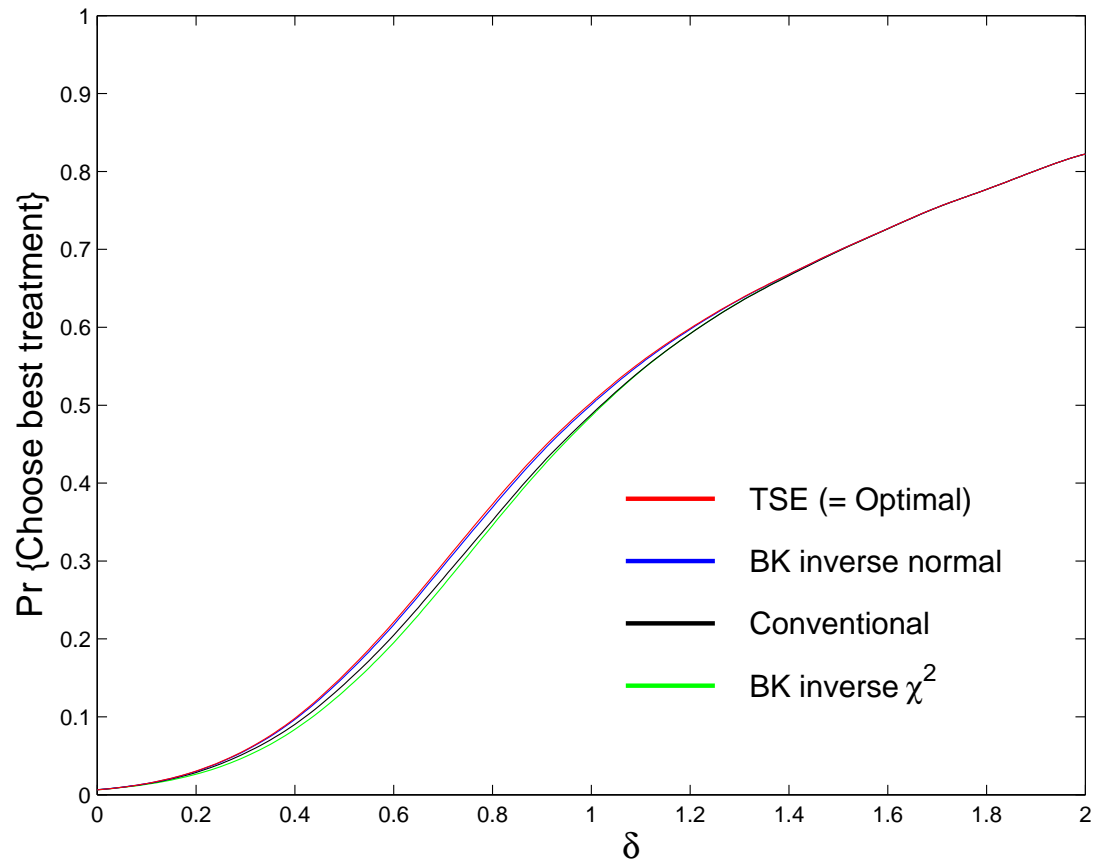
We can follow the same approach of solving a suitably formulated Bayes decision problem to derive an “optimal” rule.

This rule gives non-negative weights to all the $\hat{\theta}_{1,j}$ s when $\gamma \geq 0.5$.

The optimal rule for $\gamma = 0.5$ is the TSE decision rule.


Example 2 with $\gamma = 0.5$

Power of four procedures when $\theta = (\gamma, \gamma, \gamma, 1) \delta$ with $\gamma = 0.5$.




Example 2 with $\gamma = 0.5$

Power of four procedures when $\theta = (\gamma, \gamma, \gamma, 1) \delta$ with $\gamma = 0.5$.



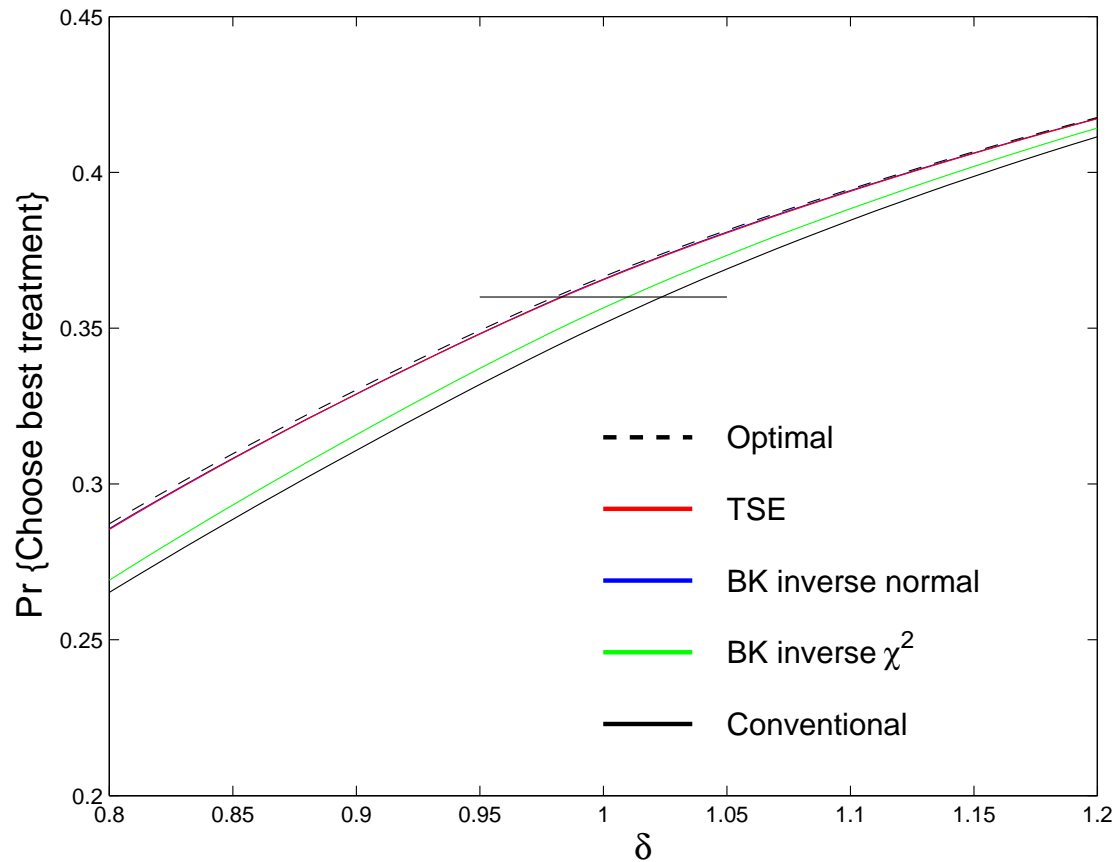
Example 2 with $\gamma = 0.75$

Power of four procedures when $\theta = (\gamma, \gamma, \gamma, 1) \delta$ with $\gamma = 0.75$, plus the “optimal” procedure for this case as a benchmark.



Example 2 with $\gamma = 0.75$

Power of four procedures when $\theta = (\gamma, \gamma, \gamma, 1) \delta$ with $\gamma = 0.75$, plus the “optimal” procedure for this case as a benchmark.



Performance of tests under other configurations of θ

Example 3:

Suppose θ is a permutation of $(\gamma_1, \gamma_2, \gamma_3, 1) \delta$, for $0 \leq \gamma_1 \leq \gamma_2 \leq \gamma_3 < 1$.

Aim:

High power to select the treatment i^* for which $\theta_{i^*} = \delta$ and then reject H_{i^*} .

Optimal rules

Again, we can follow a suitably formulated Bayes decision problem to derive an “optimal” rule, which now involves a combination of likelihood ratios.

For sufficiently high values of $\gamma_1, \dots, \gamma_3$, the optimal rule gives non-negative weights to all the $\hat{\theta}_{1,j}$ s and family-wise type I error is properly protected.

We find the TSE rule remains highly efficient relative to the optimal rule.

Performance of tests under other configurations of θ

Example 4:

Suppose θ is a permutation of $(\gamma_1, \gamma_2, \gamma_3, 1) \delta$, for $0 \leq \gamma_1 \leq \gamma_2 \leq \gamma_3 < 1$.

Suppose also it is thought likely that the θ_i increase monotonically with i — this could be a reasonable view when treatments represent increasing dose levels.

Aim:

High power to select the treatment i^* for which $\theta_{i^*} = \delta$ and then reject H_{i^*} .

Optimal rules

Here, it is of interest to find the “optimal” rule based on the assumption that the θ_i increase monotonically with i .

Once again, a suitably formulated Bayes decision problem leads to an “optimal” rule.

Performance of tests under other configurations of θ

Example 4:

Treatment effects are a permutation of $\theta = (\gamma_1, \gamma_2, \gamma_3, 1) \delta$ and the θ_i are expected to increase monotonically with i .

Calculations show that the additional knowledge about the monotonicity of the θ_i s provides only a very slight efficiency gain.

Thus, the TSE rule remains highly efficient relative to this specialised optimal rule.

This leads us to the conclusion that using a good dose-response model does not greatly improve the final decision to accept or reject H_i^* .

Bretz, Pinheiro & Branson (*Biometrics*, 2005) proposed adaptive choice of a dose-response model to define the best contrast of estimated effects. Our results suggest this may not be necessary!

However, modelling may well be useful in adaptive allocation of treatments (doses) during the Phase II stage.

Relative efficiencies of four decision rules

<i>Decision rule</i>	<i>Configuration of θ</i>				
	A	B	C	D	E
TSE	100	100	99	99	98
BK inverse normal	96	99	99	99	99
BK inverse χ^2	92	94	95	93	93
Phase III data only	97	95	93	90	89

$$m_1 = 100, m_2 = 500$$

$$\text{A: } \theta = (0, 0, 0, \delta)$$

$$\text{B: } \theta = (\gamma, \gamma, \gamma, 1) \delta, \gamma = 0.5$$

$$\text{C: } \theta = (\gamma, \gamma, \gamma, 1) \delta, \gamma = 0.75$$

$$\text{D: } \theta = (\gamma_1, \gamma_2, \gamma_3, 1) \delta, \gamma_1 = 0.75$$

$$\text{E: } \theta = (\gamma_1, \gamma_2, \gamma_3, 1) \delta, \gamma_1 = 0.75, \theta \text{ assumed monotone.}$$

Relative efficiencies of four decision rules

<i>Decision rule</i>	<i>Configuration of θ</i>				
	A	B	C	D	E
TSE	100	100	98	98	97
BK inverse normal	92	99	99	99	98
BK inverse χ^2	93	96	96	96	95
Phase III data only	88	87	82	81	81

$$m_1 = 200, m_2 = 400$$

$$\text{A: } \theta = (0, 0, 0, \delta)$$

$$\text{B: } \theta = (\gamma, \gamma, \gamma, 1) \delta, \gamma = 0.5$$

$$\text{C: } \theta = (\gamma, \gamma, \gamma, 1) \delta, \gamma = 0.75$$

$$\text{D: } \theta = (\gamma_1, \gamma_2, \gamma_3, 1) \delta, \gamma_1 = 0.75$$

$$\text{E: } \theta = (\gamma_1, \gamma_2, \gamma_3, 1) \delta, \gamma_1 = 0.75, \theta \text{ assumed monotone.}$$

7. Comments and conclusions

Modest efficiency gains can be made by combining Phase II and Phase III data (more so for larger Phase II sample sizes).

The TSE decision rule has robust efficiency over a variety of θ configurations, numbers of treatments and sample sizes. These findings apply to qualitatively different treatments and to different doses of a given treatment.

Jennison & Turnbull (*Biometrical Journal*, 2006) showed the TSE rule protects the family-wise error rate.

They also showed TSE is a closed testing procedure. So, like other procedures, it adapts readily to a treatment other than that with highest $\hat{\theta}_{1,i}$ being selected for Phase III (e.g., for safety considerations).

We adjusted critical values for Bauer & Köhne procedures to eliminate conservatism under $\theta = (0, 0, 0, 0)$. This could conceivably increase the family-wise error rate elsewhere — but this only strengthens our recommendation of the TSE rule.

Conclusions

Multiple comparison procedures

The notion of “borrowing strength” implicit in Simes’ rule may not always be desirable for multiple comparisons with a control.

Combining data from Phases II and III

The TSE procedure protects family-wise type I error and can be employed flexibly.

TSE is robustly efficient and even has good efficiency when compared to model based methods which assume the correct model.

TSE only uses observations on the selected treatment and control in the final decision — with appropriate adjustment for data-dependent treatment selection.

Now we have a simple decision rule to apply, we can turn attention to optimal choice of Phase II and Phase III sample sizes — as TSE already did in their 1988 paper.

Appendix 1: Likelihood ratios for a linear model

In comparing k treatments with a control, suppose $\theta_1 = \gamma_1 \delta, \dots, \theta_k = \gamma_k \delta$, so

$$\theta = \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_k \end{pmatrix} = \begin{pmatrix} \gamma_1 \\ \vdots \\ \gamma_k \end{pmatrix} \delta = X\delta$$

for a scalar parameter δ .

In Phase II, we have $\hat{\theta} \sim N(\theta, \Sigma) = N(X\delta, \Sigma)$, where

$$\Sigma = \begin{pmatrix} 2 & 1 & \dots & 1 \\ 1 & 2 & \dots & 1 \\ \vdots & \vdots & & \vdots \\ 1 & 1 & \dots & 2 \end{pmatrix} \frac{\sigma^2}{m_1}.$$

Likelihood ratios for a linear model

The probability density of $\hat{\theta}$ is

$$f_{\delta}(\hat{\theta}) = (2\pi)^{-k/2} (\det \Sigma)^{-1/2} \exp\left\{-\frac{1}{2}(\hat{\theta} - X\delta)^T \Sigma^{-1}(\hat{\theta} - X\delta)\right\}.$$

The likelihood ratio for $\delta = \delta^*$ vs $\delta = 0$ is

$$\frac{f_{\delta^*}}{f_0} = \exp(\delta^* X^T \Sigma^{-1} \hat{\theta} + \text{constant})$$

and the log likelihood ratio is

$$\delta^* X^T \Sigma^{-1} \hat{\theta} + \text{constant}. \tag{1}$$

Now, the maximum likelihood estimate of δ for this normal linear model is

$$\hat{\delta} = (X^T \Sigma^{-1} X)^{-1} X^T \Sigma^{-1} \hat{\theta} \sim N(\delta, (X^T \Sigma^{-1} X)^{-1}).$$

Likelihood ratios for a linear model

So, the log likelihood ratio (1) is a constant plus

$$\delta^* (X^T \Sigma^{-1} X) \hat{\delta} = \delta^* \frac{\hat{\delta}}{\text{Var}(\hat{\delta})}.$$

Denote the Phase II estimate of δ by $\hat{\delta}_1$ so the log likelihood ratio from Phase II data is

$$\delta^* \frac{\hat{\delta}_1}{\text{Var}(\hat{\delta}_1)} + \text{constant}.$$

Similarly (with a simpler model as only $\hat{\theta}_{i^*}$ is observed), the log likelihood ratio from Phase III data is

$$\delta^* \frac{\hat{\delta}_2}{\text{Var}(\hat{\delta}_2)} + \text{constant}.$$

Likelihood ratios for a linear model

Combining across Phase II and Phase III, the log likelihood ratio for $\delta = \delta^*$ vs $\delta = 0$ is a constant plus

$$\delta^* \left\{ \frac{\hat{\delta}_1}{\text{Var}(\hat{\delta}_1)} + \frac{\hat{\delta}_2}{\text{Var}(\hat{\delta}_2)} \right\}.$$

Thus, we reject $\delta = 0$ in favour of $\delta > 0$ for high values of

$$\frac{\hat{\delta}_1}{\text{Var}(\hat{\delta}_1)} + \frac{\hat{\delta}_2}{\text{Var}(\hat{\delta}_2)},$$

the overall MLE of δ from the pooled Phase II and Phase III data.

Appendix 2: Formula for the likelihood ratio from Phase II

The log likelihood ratio is an increasing function of $X^T \Sigma^{-1} \hat{\theta}$, i.e., of

$$\begin{aligned}
 & (\gamma_1, \dots, \gamma_k) \frac{m_1}{\sigma^2(k+1)} \begin{pmatrix} k & -1 & \dots & -1 \\ -1 & k & \dots & -1 \\ \vdots & \vdots & & \vdots \\ -1 & -1 & \dots & k \end{pmatrix} \begin{pmatrix} \hat{\theta}_{1,1} \\ \hat{\theta}_{1,2} \\ \vdots \\ \hat{\theta}_{1,k} \end{pmatrix} \\
 &= \frac{m_1}{\sigma^2} (\gamma_1, \dots, \gamma_k) \left(I_n - \frac{1}{k+1} \mathbf{1}_{n \times n} \right) \begin{pmatrix} \hat{\theta}_{1,1} \\ \vdots \\ \hat{\theta}_{1,k} \end{pmatrix} \\
 &= \frac{m_1}{\sigma^2} \left\{ \sum_{j=1}^k \gamma_j \hat{\theta}_{1,j} - \frac{1}{k+1} \sum_{j=1}^k \gamma_j \sum_{j+1}^k \hat{\theta}_{1,j} \right\}. \tag{2}
 \end{aligned}$$

Special cases of the likelihood ratio from Phase II

Case 1.

With $i^* = k$, consider testing $\theta = (0, \dots, 0)$ vs $\theta = (0, \dots, 0, \delta)$.

We take $(\gamma_1, \dots, \gamma_k) = (0, \dots, 0, 1)$, then (2) becomes

$$\begin{aligned} & \frac{m_1}{\sigma^2} \left(\hat{\theta}_{1,i^*} - \frac{1}{k+1} \sum_{j=1}^k \hat{\theta}_{1,j} \right) \\ &= \frac{m_1}{\sigma^2} \frac{k}{k+1} \left(\hat{\theta}_{1,i^*} - \frac{1}{k} \sum_{j \neq i^*} \hat{\theta}_{1,j} \right). \end{aligned}$$

This lies behind the term

$$\hat{\theta}_{1,i^*} - \frac{1}{4} \sum_{j \neq i^*} \hat{\theta}_{1,j}$$

in the formula on Slide 22 for the “optimal” test with $k = 4$.

Special cases of the likelihood ratio from Phase II

Case 2.

With $i^* = k$, consider testing $\theta = (0, \dots, 0)$ vs $\theta = (\gamma\delta, \dots, \gamma\delta, \delta)$.

We take $(\gamma_1, \dots, \gamma_k) = (\gamma, \dots, \gamma, 1)$, then (2) becomes

$$\begin{aligned} & \frac{m_1}{\sigma^2} \left\{ \hat{\theta}_{1,i^*} + \sum_{j \neq i^*} \gamma \hat{\theta}_{1,j} - \frac{(k-1)\gamma + 1}{k+1} \left(\hat{\theta}_{1,i^*} + \sum_{j \neq i^*} \hat{\theta}_{1,j} \right) \right\} \\ &= \frac{m_1}{\sigma^2(k+1)} \left\{ [k - (k-1)\gamma] \hat{\theta}_{1,i^*} + \sum_{j \neq i^*} (2\gamma - 1) \hat{\theta}_{1,j} \right\}. \end{aligned}$$

The coefficient $(2\gamma - 1)$ of terms $\hat{\theta}_{1,j}$ for $j \neq i^*$ is negative when $\gamma < 0.5$ and positive when $\gamma > 0.5$.

The case $\gamma = 0.5$ gives a coefficient of zero — and the TSE decision rule.

Appendix 3: Alternative derivation of MLEs in Phase II

Our original model assumed independent treatment estimates

$$\hat{\mu}_{1,i} \sim N\left(\mu_i, \frac{\sigma^2}{m_1}\right), \quad i = 1, \dots, k,$$

$$\hat{\mu}_{1,c} \sim N\left(\mu_c, \frac{\sigma^2}{m_1}\right).$$

We defined $\theta_i = \mu_i - \mu_c$, $i = 1, \dots, k$, and worked with the estimates $\hat{\theta}_{1,i}$.

However, we can also work with the $\hat{\mu}_{1,i}$ s and $\hat{\mu}_c$ directly.

Example 1

We can write the case $\theta = (0, \dots, 0, \delta)$ as

$$\begin{aligned} \mu_1 &= \dots = \mu_{k-1} = \mu_c = a, \\ \mu_k &= a + \delta. \end{aligned}$$


Alternative derivation of MLEs in Phase II

So, estimates follow a regression model $E(\hat{\mu}_{1,i}) = a + \delta x$ where

$\hat{\mu}_{1,1}, \dots, \hat{\mu}_{1,k-1}$ and $\hat{\mu}_c$ have $x = 0$,

$\hat{\mu}_{1,k}$ has $x = 1$,

and we are interested in the slope δ .



Clearly,

$$\hat{\delta} = \hat{\mu}_{1,k} - \frac{1}{k} (\hat{\mu}_{1,1} + \dots + \hat{\mu}_{1,k-1} + \hat{\mu}_{1,c}).$$

Alternative derivation of MLEs in Phase II

But, we can write

$$\begin{aligned}\hat{\delta} &= \hat{\mu}_{1,k} - \frac{1}{k} (\hat{\mu}_{1,1} + \dots + \hat{\mu}_{1,k-1} + \hat{\mu}_{1,c}) \\ &= (\hat{\mu}_{1,k} - \hat{\mu}_{1,c}) - \frac{1}{k} \sum_{j=1}^{k-1} (\hat{\mu}_{1,j} - \hat{\mu}_{1,c}) \\ &= \hat{\theta}_{1,k} - \frac{1}{k} \sum_{j=1}^{k-1} \hat{\theta}_{1,j},\end{aligned}$$

in agreement with previous results.

Note: In the estimate of a slope, the sum of the weights of the $\hat{\mu}_{1,i}$ s and $\hat{\mu}_c$ is zero (it is a contrast).

So, it is automatically the case that $\hat{\delta}$ can be expressed in terms of the $\hat{\theta}_{1,i}$ s.

Alternative derivation of MLEs in Phase II

Example 2

We can write the case $\theta = (\gamma, \dots, \gamma, 1) \delta$ as

$$\begin{aligned}\mu_c &= a, \\ \mu_1 = \dots = \mu_{k-1} &= a + \gamma \delta, \\ \mu_k &= a + \delta.\end{aligned}$$

Now, the estimates follow a regression model $E(\hat{\mu}_{1,i}) = a + \delta x$ where


$\hat{\mu}_{1,c}$ has $x = 0$,

$\hat{\mu}_{1,1}, \dots, \hat{\mu}_{1,k-1}$ have $x = \gamma$,

$\hat{\mu}_{1,k}$ has $x = 1$,

and we are interested in the slope δ .

Alternative derivation of MLEs in Phase II



It is easy to see that

for $\gamma = 0.5$, $\hat{\delta} = \hat{\mu}_{1,k} - \hat{\mu}_{1,c} = \hat{\theta}_{1,k}$,

for $\gamma < 0.5$, $\hat{\mu}_{1,1}, \dots, \hat{\mu}_{1,k}$ contribute to $\hat{\delta}$ with negative weights,

for $\gamma > 0.5$, $\hat{\mu}_{1,1}, \dots, \hat{\mu}_{1,k}$ contribute to $\hat{\delta}$ with positive weights.