

CONFIDENCE INTERVALS UNDER ORDER RESTRICTIONS

Yongseok Park¹, John D. Kalbfleisch² and Jeremy M. G. Taylor²

¹University of Pittsburgh and ²University of Michigan

Abstract: In this paper, we consider the problem of constructing confidence intervals (CIs) for G independent normal population means subject to linear ordering constraints. For this problem, CIs based on asymptotic distributions, likelihood ratio tests, and bootstraps do not have good properties, particularly when some of the population means are close to each other. We propose a new method based on defining intermediate random variables that are related to the original observations and using the CIs of the means of these intermediate random variables to restrict the original CIs from the separate groups. The coverage rates of the intervals are shown to exceed, but be close to, the nominal level for two groups, when the ratio of the variances is assumed known. Simulation studies show that the proposed CIs have coverage rates close to nominal levels with reduced average widths. Data on half-lives of an antibiotic are analyzed to illustrate the method.

Key words and phrases: Convex combination, elliptical unimodal distribution, linear ordering, normal distribution, restricted confidence interval.

1. Introduction

Consider a G -sample problem where the observations X_{gi} , $g = 1, \dots, G$, $i = 1, \dots, n_g$ are independent random variables with distribution function $F_g(x; \mu_g)$. When estimating $\boldsymbol{\mu} = (\mu_1, \dots, \mu_G)$, there often exists information about inequality orderings of these parameters. For example, if μ_g is the average height of children of age g or μ_g is the toxicity rate of a drug for dose level g in a clinical trial, the parameters should satisfy the restriction:

$$\mu_1 \leq \mu_2 \leq \dots \leq \mu_G. \quad (1.1)$$

This is called *simple ordering* or *linear ordering*. The natural estimator for order restricted parameters is the restricted maximum likelihood estimator (MLE). For the case where $F_g(x; \mu_g)$, $g = 1, \dots, G$, is normal with mean μ_g and variance σ^2 , the MLE of μ_g under restriction (1.1) is the isotonic regression estimator (Barlow et al. (1972); Robertson, Wright, and Dykstra (1988)). The restricted MLE $\hat{\mu}_g$ has been shown to dominate the unrestricted MLE \bar{X}_g in the sense that

$$P(|\hat{\mu}_g - \mu_g| \leq c) \geq P(|\bar{X}_g - \mu_g| \leq c), g = 1, \dots, G, \quad (1.2)$$

for all $c > 0$ (Kelly (1989); Lee (1981)). In this paper, we focus on constructing confidence intervals (CIs) for the parameters μ_g under the linear ordering constraints.

Estimation problems in a restricted parameter space have been studied since the 1950s. Marchand and Strawderman (2004) and van Eeden (2006) reviewed estimation methods that have been developed in the past and discussed such “good” properties of restricted estimators as dominance, minimaxity and admissibility. Cohen and Sackrowitz (2004) discussed some inference issues and pointed out that traditional inference methods, such as the likelihood based method, can lead to some undesirable properties in restricted parameter problems. Andrews (2000), Hwang (1995), and Peddada (1997) pointed out that the bootstrap method, which has been very useful for constructing CIs of complicated parameters, fails when a parameter is on the boundary or close to the boundary of the parameter space. Thus, it is of interest to develop an inference procedure without depending on traditional inference methods.

Specialized methods for constructing CIs under order restrictions have been suggested. Schoenfeld (1986) proposed a method for one-sided intervals based on inverting the likelihood ratio test for the ordered means from a normal distribution. Hwang and Peddada (1994) proposed a constant length CI, in which the CI, derived without the order restriction assumption, is shifted and centered at an improved estimator, e.g. centered at the restricted MLE in the linear ordering case. According to dominance as at (1.2), coverage rates of these restricted methods exceed the nominal levels obtained from unrestricted intervals. Bayesian, bootstrap, and other resampling methods are also discussed by Dunson and Neelon (2003), Peddada (1997) and Li, Taylor, and Nan (2010).

We propose a novel method to construct CIs under a linear ordering constraint. In Section 2, we consider a two-sample case of ordered normal means with known variances and obtain some theoretical results about the coverage rate and width of the CI. In Section 3, we show how to adapt the methods to the case when the population variances are unknown. We extend the methods to the case with three or more samples in Section 4. In Section 5, we describe some other CIs that have been proposed in the literature. In Section 6, we illustrate the method using data on half-lives of an antibiotic in an animal study, and in Section 7, we conduct simulation studies to compare those CIs with ours.

2. Confidence Intervals for μ_1 and μ_2 with Known Variances

2.1. Family of confidence intervals

Let $X_g \sim N(\mu_g, \sigma_g^2)$, $g = 1, 2$, where σ_g^2 is known. Our goal is to construct $1 - \alpha$ CIs for μ_1 and μ_2 when it is known that $\mu_1 \leq \mu_2$. Let $X(\gamma) = \gamma X_1 + (1 -$

$\gamma)X_2$, where $\gamma \in [0, 1]$. The mean and variance of $X(\gamma)$ are $\mu(\gamma) = EX(\gamma) = \gamma\mu_1 + (1 - \gamma)\mu_2$ and $\sigma^2(\gamma) = \text{var}\{X(\gamma)\} = \gamma^2\sigma_1^2 + (1 - \gamma)^2\sigma_2^2$. Let $z_{1-\alpha/2}$ be the upper $\alpha/2$ percentile of a standard normal distribution and let $t_{1-\alpha/2,\nu}$ be the upper $\alpha/2$ percentile of a standard t distribution with degree of freedom ν , which we denote for convenience by z and t_ν . The unrestricted CIs for μ_1, μ_2 , and $\mu(\gamma)$ are $\mu_g \in [X_g - z\sigma_g, X_g + z\sigma_g], g = 1, 2$, and $\mu(\gamma) \in [X(\gamma) - z\sigma(\gamma), X(\gamma) + z\sigma(\gamma)]$. Since $\mu_1 \leq \mu(\gamma) \leq \mu_2$, it is sensible to consider modifying the limits of the CIs for μ_1 and μ_2 , based on the limits of the CI for $\mu(\gamma)$. We propose a family of CI $[L_1(\gamma), U_1(\gamma)]$ for μ_1 and $[L_2(\gamma), U_2(\gamma)]$ for μ_2 as

$$\begin{aligned} L_1(\gamma) &= \min\{X_1 - z\sigma_1, X(\gamma) - z\sigma(\gamma)\}, \\ U_1(\gamma) &= \min\{X_1 + z\sigma_1, X(\gamma) + z\sigma(\gamma)\}, \\ L_2(\gamma) &= \max\{X_2 - z\sigma_2, X(\gamma) - z\sigma(\gamma)\}, \\ U_2(\gamma) &= \max\{X_2 + z\sigma_2, X(\gamma) + z\sigma(\gamma)\}. \end{aligned} \tag{2.1}$$

In Sections 2 and 3, we consider upper and lower limits for μ_1 ; those for μ_2 are of the same form except for changing min to max and changing the subscript from 1 to 2. In this section, we develop the method and theory for the case of one observation per group; the results apply to multiple observations per group by replacing X_g by the group mean and replacing σ_g^2 by σ_g^2/n_g .

Definition 1. $\mathbf{X} = (X_1, \dots, X_k)$ has an elliptical unimodal distribution with location $\boldsymbol{\mu}$ and positive-definite matrix $\boldsymbol{\Sigma}$ if its probability density function is

$$f(\mathbf{x}) = C h\{(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\},$$

where $h(t)$ is a nonincreasing function in t .

Theorem 1. Suppose $\mathbf{Y} = (Y_1, Y_2)^T$ has a bivariate elliptical unimodal distribution with location $\boldsymbol{\mu} = (0, \Delta)$ and $\boldsymbol{\Sigma} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$, where $\Delta \geq 0$. If c_α satisfies $P(|Y_1| \leq c_\alpha) = 1 - \alpha$, then

$$Q = P\{\min(Y_1 - c_\alpha, Y_2 - c_\alpha) \leq 0 \leq \min(Y_1 + c_\alpha, Y_2 + c_\alpha)\} \geq 1 - \alpha. \tag{2.2}$$

Proof. The joint probability density function of \mathbf{Y} is

$$f(y_1, y_2) = C h\{y_1^2 + (y_2 - \Delta)^2 - 2\rho y_1(y_2 - \Delta)\}.$$

From (2.2),

$$\begin{aligned} Q &= P(Y_1 \geq -c_\alpha, Y_2 \geq -c_\alpha) - P(Y_1 \geq c_\alpha, Y_2 \geq c_\alpha) \\ &= \int_{\mathbb{D} \cup \mathbb{E} \cup \mathbb{H}} f(y_1, y_2) dy_1 dy_2 \\ &= \int_{\mathbb{D} \cup \mathbb{E} \cup \mathbb{F}} f(y_1, y_2) dy_1 dy_2 + \int_{\mathbb{H}} f(y_1, y_2) dy_1 dy_2 - \int_{\mathbb{F}} f(y_1, y_2) dy_1 dy_2, \end{aligned} \tag{2.3}$$




Figure 1. Sample space of (y_1, y_2) .

where $\mathbb{D}, \mathbb{E}, \mathbb{F}$ and \mathbb{H} are defined in Figure 1. It is clear that

$$\int_{\mathbb{D} \cup \mathbb{E} \cup \mathbb{F}} f(y_1, y_2) dy_1 dy_2 = P\{|Y_1| \leq c_\alpha\} = 1 - \alpha. \tag{2.4}$$

If $\mathbf{y}_H = (y_1, y_2)$ is a point in \mathbb{H} with $\mathbf{y}_F = (-y_2, -y_1)$ a corresponding point in \mathbb{F} , it can be seen that

$$(\mathbf{y}_H - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y}_H - \boldsymbol{\mu}) - (\mathbf{y}_F - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y}_F - \boldsymbol{\mu}) = -2\Delta(1 - \rho)(y_1 + y_2) \leq 0$$

for all $\Delta \geq 0$ and $(y_1, y_2) \in \mathbb{H}$. Since $h(t)$ is nonincreasing, we have $f(y_1, y_2) \geq f(-y_2, -y_1)$ if $\Delta \geq 0$ and $(y_1, y_2) \in \mathbb{H}$. Thus the density at each point in \mathbb{H} is greater than or equal to the density at the corresponding point in \mathbb{F} . Since $(-y_2, -y_1)$ varies over all of \mathbb{F} as (y_1, y_2) varies over \mathbb{H} ,

$$\int_{\mathbb{H}} f(y_1, y_2) dy_1 dy_2 - \int_{\mathbb{F}} f(y_1, y_2) dy_1 dy_2 \geq 0. \tag{2.5}$$

Thus (2.2) follows from (2.3), (2.4), and (2.5).

Corollary 1. *If $X_1 \sim N(\mu_1, \sigma_1^2)$ and $X_2 \sim N(\mu_2, \sigma_2^2)$ are independent with $\mu_1 \leq \mu_2$, then $P\{L_1(\gamma) \leq \mu_1 \leq U_1(\gamma)\} \geq 1 - \alpha$ for all $\gamma \in [0, 1]$.*

Proof. This follows by setting $Y_1 = (X_1 - \mu_1)/\sigma_1$, $Y_2 = \{X(\gamma) - \mu_1\}/\sigma(\gamma)$ and $h(t) = \exp(-t/2)$ in Theorem 1. The corresponding $\Delta = \{\mu(\gamma) - \mu_1\}/\sigma(\gamma) =$

$(1 - \gamma)(\mu_2 - \mu_1)/\sigma(\gamma) \geq 0$, $\rho = \{1 + (1 - \gamma)^2\sigma_2^2/(\gamma^2\sigma_1^2)\}^{-1/2}$, and $C = (2\pi)^{-1}(1 - \rho^2)^{-1/2}$.

This corollary shows that the coverage rate of the interval $[L_1(\gamma), U_1(\gamma)]$ always exceeds the nominal level when the variances of X_1 and X_2 are known.

2.2. Proposed confidence intervals and their properties

We seek a value of γ to make the width of the CIs for μ_1 at (2.1) as small as possible. One choice of γ is the value that minimizes $W_1(\gamma) = E\{U_1(\gamma) - L_1(\gamma)\}$, but $W_1(\gamma)$ depends on the unknown mean difference $\mu_2 - \mu_1$. There is a $\gamma \in (0, 1)$ for which $\sigma^2(\gamma) < \sigma_1^2$, and it can be seen that when $\sigma^2(\gamma) < \sigma_1^2$, $U_1(\gamma) - L_1(\gamma) \leq 2z\sigma_1$ for any observations X_1 and X_2 . Thus the width of the interval can be reduced by suitable choice of γ . Another intuitive choice of γ is the value that minimizes $\sigma^2(\gamma)$. It is easy to see that $\sigma^2(\gamma)$ is minimized at $\gamma_0 = \sigma_2^2/(\sigma_1^2 + \sigma_2^2)$.

Theorem 2. $W_1(\gamma)$ is minimized at $\gamma = \gamma_0$ if $\mu_1 = \mu_2$.

Proof. Consider the case $\gamma \geq 1 - 2\sigma_1^2/(\sigma_1^2 + \sigma_2^2)$, for which $\sigma^2(\gamma) \leq \sigma_1^2$. Let $c(\gamma) = z\{\sigma_1 - \sigma(\gamma)\}/(1 - \gamma)$. Then we have

$$U_1(\gamma) - L_1(\gamma) = \begin{cases} 2z\sigma_1 & \text{if } X_2 - X_1 > c(\gamma), \\ (1 - \gamma)(X_2 - X_1) + z\sigma_1 + z\sigma(\gamma) & \text{if } -c(\gamma) < X_2 - X_1 \leq c(\gamma), \\ 2z\sigma(\gamma) & \text{if } X_2 - X_1 \leq -c(\gamma). \end{cases}$$

So,

$$W_1(\gamma) = \int_{c(\gamma)}^{\infty} 2z\sigma_1 f_{X_2 - X_1}(x) dx + \int_{-c(\gamma)}^{c(\gamma)} \{(1 - \gamma)x + z\sigma_1 + z\sigma(\gamma)\} f_{X_2 - X_1}(x) dx + \int_{-\infty}^{-c(\gamma)} 2z\sigma(\gamma) f_{X_2 - X_1}(x) dx = z\{\sigma_1 + \sigma(\gamma)\},$$

because $X_2 - X_1 \sim N(0, \sigma_1^2 + \sigma_2^2)$, $\int_{-c(\gamma)}^{c(\gamma)} x f_{X_2 - X_1}(x) dx = 0$ and $\int_c^{\infty} f_{X_2 - X_1}(x) dx = \int_{-\infty}^{-c} f_{X_2 - X_1}(x) dx$.

Similarly we can show that if $\gamma < 1 - 2\sigma_1^2/(\sigma_1^2 + \sigma_2^2)$, $W_1(\gamma) = z\{\sigma_1 + \sigma(\gamma)\}$.

Thus minimizing $W_1(\gamma)$ is the same as minimizing $\sigma(\gamma)$, which happens at γ_0 .

Using this γ_0 , the proposed CI for μ_1 is

$$\hat{L}_1 = \min(X_1 - z\sigma_1, \hat{X} - z\sigma_0), \quad \hat{U}_1 = \min(X_1 + z\sigma_1, \hat{X} + z\sigma_0), \tag{2.6}$$

where $\hat{X} = (X_1\sigma_2^2 + X_2\sigma_1^2)/(\sigma_1^2 + \sigma_2^2)$ and $\sigma_0^2 = \sigma_1^2\sigma_2^2/(\sigma_1^2 + \sigma_2^2)$.

Let $\rho = \sigma_2/\sqrt{\sigma_1^2 + \sigma_2^2}$ and let $\Delta = (\mu_2 - \mu_1)(1 - \rho^2)/(\rho\sigma_1)$. Let $Y_1 = (X_1 - \mu_1)/\sigma_1$ and let $Y_2 = (\hat{X} - \mu_1)/\sigma_0$. Then the joint distribution of Y_1 and Y_2 is

Table 1. Theoretical maximum coverage rate of CI for μ_1 in the situations with different ratio of variances.

$1 - \alpha$	σ_2^2/σ_1^2								
	10^{-9}	10^{-3}	0.01	0.1	0.2	0.5	1	2	10
0.95	0.969	0.969	0.968	0.965	0.962	0.959	0.956	0.953	0.950
0.90	0.933	0.932	0.930	0.924	0.920	0.913	0.909	0.905	0.901
0.80	0.852	0.850	0.846	0.835	0.829	0.819	0.812	0.806	0.801
0.70	0.761	0.759	0.753	0.740	0.732	0.721	0.713	0.707	0.701

$$f(y_1, y_2) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left\{-\frac{y_1^2 + (y_2 - \Delta)^2 - 2\rho y_1(y_2 - \Delta)}{2(1-\rho^2)}\right\}.$$

The coverage probability is

$$P = 1 - \alpha + \int_z^\infty \int_{-\infty}^z f(y_1, y_2) dy_2 dy_1 - \int_{-z}^\infty \int_{-\infty}^{-z} f(y_1, y_2) dy_2 dy_1.$$

Setting $dP/d\Delta = 0$, we find that maximum coverage probability of the proposed CI for fixed ρ occurs at $\hat{\Delta}$ that solves the equation

$$\Phi\left\{\frac{z\rho - z - \Delta\rho}{\sqrt{1-\rho^2}}\right\} - \Phi\left\{\frac{z - z\rho - \Delta\rho}{\sqrt{1-\rho^2}}\right\} \exp(-2z\Delta) = 0, \quad (2.7)$$

where Φ is the cumulative distribution function of a standard normal. It can be shown that (2.7) has one and only one positive solution for Δ , for any $0 < \rho < 1$. As can be seen in Table 1 and Figure 3a, the theoretical maximum coverage rate increases as σ_2^2/σ_1^2 (or equivalently ρ) decreases, and approaches 0.969 for nominal level of 95% when σ_2^2/σ_1^2 goes to 0.

As to how much wider this CI is compared to the narrowest possible CI when $\mu_1 \neq \mu_2$, we compute $\min_\gamma W_1(\gamma)$ through numerical minimization over γ for a given $\mu_2 - \mu_1$, and compare it with $W_1(\gamma_0)$. The results are shown in Figure 2. The largest possible average width for the CI using γ_0 , compared to the optimal, occurs when $\sigma_2^2/\sigma_1^2 = 0.063$ for nominal level of 95% and when $\sigma_2^2/\sigma_1^2 = 0.081$ for the nominal level of 99%. Compared to the minimum possible $W_1(\gamma)$, $W_1(\gamma_0)$ is only at most 0.4% wider at the nominal level of 95% and at most 0.8% wider at the nominal level of 99%. This indicates that the CI using γ_0 is almost as efficient as the most efficient CI in this family.

The reduction of average width compared to the unrestricted CI depends on σ_2^2/σ_1^2 and $\mu_2 - \mu_1$, as can be seen in Figure 3b. For the CI of μ_1 , the smaller the σ_2^2/σ_1^2 , the smaller the average width, and the closer the means, the smaller the average width. The average width can be half the width of the unrestricted CI when $\sigma_2^2/\sigma_1^2 \rightarrow 0$ and $\mu_1 = \mu_2$. If the variance of X_1 and X_2 are similar, the average width is about 85% of the unrestricted one when $\mu_1 = \mu_2$.




Figure 2. Comparison of width of the CI using γ_0 and minimum possible width.

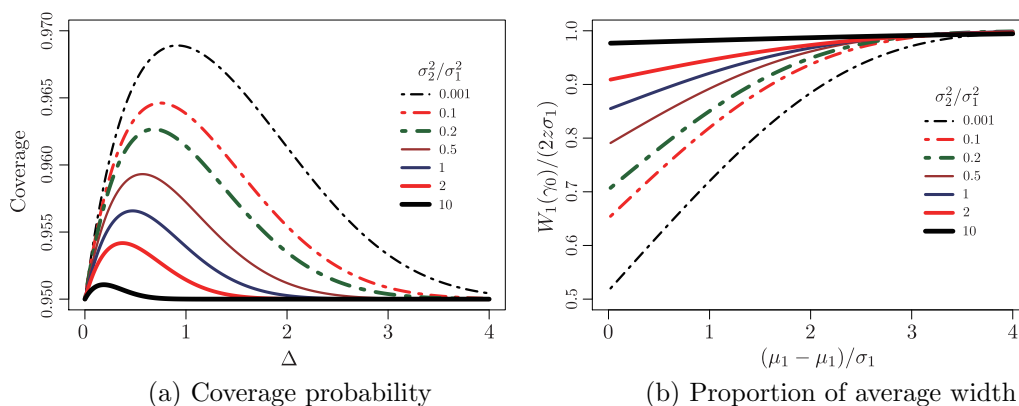


Figure 3. Coverage probability and ratio of average width of restricted CI for μ_1 compared to unrestricted CI for μ_1 .

3. Confidence Intervals for μ_1 and μ_2 with Unknown Variances

3.1. Restricted confidence interval when σ_2^2/σ_1^2 is known

Suppose we observe $X_{gi} \sim N(\mu_g, \sigma_g^2)$, $g = 1, 2, i = 1, \dots, n_g$. Let $\bar{X}_g = \sum_{i=1}^{n_g} X_{gi}/n_g, g = 1, 2$. Let $s_1^2 = \left\{ \sum_{i=1}^{n_1} (X_{1i} - \bar{X}_1)^2 + \sum_{i=1}^{n_2} (X_{2i} - \bar{X}_2)^2/p \right\}/\nu$, and let $\tilde{X} = (\bar{X}_1 n_1 p + \bar{X}_2 n_2)/(n_2 + n_1 p)$, where $p = \sigma_2^2/\sigma_1^2$ and $\nu = n_1 + n_2 - 2$. Then $\tilde{\mu} = E\tilde{X} = (\mu_1 n_1 p + \mu_2 n_2)/(n_2 + n_1 p)$ and $\tilde{\sigma}^2 = \text{var}(\tilde{X}) = \sigma_1^2 p/(n_2 + n_1 p)$. If $\tilde{s}^2 = p s_1^2/(n_2 + n_1 p)$, then $(\tilde{X} - \tilde{\mu})/\tilde{s}$ and $\sqrt{n_1}(\bar{X}_1 - \mu_1)/s_1$ follow standard T distributions with ν degrees of freedom. The unrestricted CIs for μ_1 and $\tilde{\mu}$ are

$$\mu_1 \in \left[\bar{X}_1 - t_\nu \frac{s_1}{\sqrt{n_1}}, \bar{X}_1 + t_\nu \frac{s_1}{\sqrt{n_1}} \right],$$

$$\tilde{\mu} \in [\tilde{X} - t_\nu \tilde{s}, \tilde{X} + t_\nu \tilde{s}].$$

We propose the restricted CI for μ_1 as

$$\begin{aligned} \hat{L}_1 &= \min(\bar{X}_1 - t_\nu \frac{s_1}{\sqrt{n_1}}, \tilde{X} - t_\nu \tilde{s}), \\ \hat{U}_1 &= \min(\bar{X}_1 + t_\nu \frac{s_1}{\sqrt{n_1}}, \tilde{X} + t_\nu \tilde{s}). \end{aligned} \tag{3.1}$$

When σ_2^2/σ_1^2 is known, the pivotal random variables $(X_1 - \mu_1)/s_1$ and $(\tilde{X} - \tilde{\mu})/\tilde{s}$ follow a bivariate T distribution. Since the multivariate T belongs to the elliptical distribution family, the result in Theorem 1 concerning coverage rates of CIs is applicable.

Corollary 2. *The CI at (3.1) satisfies $P(\mu_1 \in [\hat{L}_1, \hat{U}_1]) \geq 1 - \alpha$.*

Proof. This follows by setting $Y_1 = (\bar{X}_1 - \mu_1)/s_1$, $Y_2 = \{\tilde{X} - \mu_1\}/\tilde{s}$ and $h(t) = \{1 + t/\nu\}^{-(\nu+2)/2}$ in Theorem 1. The corresponding $\Delta = (\tilde{\mu} - \mu_1)E(1/\tilde{s}) = (\mu_2 - \mu_1)n_2E(1/\tilde{s})/(n_2 + n_1p) \geq 0$, $\rho = \sqrt{pn_1/(pn_1 + n_2)}$, and $C = 1/(2\pi\sqrt{1 - \rho^2})$.

3.2. Restricted confidence intervals when σ_1^2/σ_2^2 is unknown

Let $s_g^2 = \sum_{i=1}^{n_g} (X_{gi} - \bar{X}_g)^2/(n_g - 1)$, $g = 1, 2$. The unrestricted CI for μ_g is $\bar{X}_g \pm t_{\nu_g} s_g/\sqrt{n_g}$, where $\nu_g = n_g - 1$. We again consider an intermediate random variables \tilde{X} , with mean $\tilde{\mu}$, obtain a CI $[\tilde{L}, \tilde{U}]$ for $\tilde{\mu}$, and define the restricted CI for μ_1 as

$$\hat{L}_1 = \min(\bar{X}_1 - t_{\nu_1} \frac{s_1}{\sqrt{n_1}}, \tilde{L}), \quad \hat{U}_1 = \min(\bar{X}_1 + t_{\nu_1} \frac{s_1}{\sqrt{n_1}}, \tilde{U}). \tag{3.2}$$

In this case, it is not possible to find an intermediate random variable $\tilde{X} = \gamma\bar{X}_1 + (1 - \gamma)\bar{X}_2$ with exactly appropriate properties. Even for a fixed γ , the interval estimation for $\mu(\gamma) = E(\tilde{X})$ is in fact a variant of the Behrens-Fisher problem.

We propose two methods to approximate the distribution of \tilde{X} . The first adjusts the width of the interval by incorporating uncertainty in the estimates of σ_1^2 and σ_2^2 while the second additionally modifies the effective sample size.

Method 1. For $X(\gamma) = \gamma\bar{X}_1 + (1 - \gamma)\bar{X}_2$, $\text{var}\{X(\gamma)\} = \gamma^2\sigma_1^2/n_1 + (1 - \gamma)^2\sigma_2^2/n_2$ is minimized at $\gamma^* = n_1\sigma_2^2/(n_1\sigma_2^2 + n_2\sigma_1^2)$, giving $X^* = (n_1\sigma_2^2\bar{X}_1 + n_2\sigma_1^2\bar{X}_2)/(n_1\sigma_1^2 + n_2\sigma_2^2)$ and the minimum variance $\sigma^{*2} = \sigma_1^2\sigma_2^2/(n_2\sigma_1^2 + n_1\sigma_2^2)$. Replacing σ_1^2 and σ_2^2 with unbiased estimators s_1^2 and s_2^2 , the approximation for X^* is

$$\tilde{X} = \frac{n_1s_2^2\bar{X}_1 + n_2s_1^2\bar{X}_2}{n_1s_2^2 + n_2s_1^2}.$$

However, the coverage rate for $\tilde{\mu} = E(\tilde{X})$ based on $\tilde{X} \pm z\hat{\sigma}^*$, where $\hat{\sigma}^{*2} = s_1^2s_2^2/(n_2s_1^2 + n_1s_2^2)$, is too low, because the estimate $\hat{\sigma}^{*2}$ does not incorporate

the uncertainty in the estimation of σ_1^2 and σ_2^2 . One approach to allow for this is to modify the estimated variance based on thresholds of t_{ν_g} distributions. Since $P(\bar{X}_g - \mu_g > t_{\nu_g} s_g / \sqrt{n_g}) = P[\bar{X}_g - \mu_g > \{s_g t_{\nu_g} / (z \sqrt{n_g})\} z] = \alpha/2$, we approximate the distribution of $\bar{X}_g - \mu_g$ with a $N(0, s_g^2 t_{\nu_g}^2 / (z^2 n_g))$ distribution. This gives exactly the same $1 - \alpha$ CI for $\mu_g, g = 1, 2$ as using a t distribution. Thus, replacing s_g^2 with $s_g^2 t_{\nu_g}^2 / z^2$ in $\hat{\sigma}^{*2}$ gives the estimate

$$\tilde{\sigma}^2 = \frac{t_{\nu_1}^2 t_{\nu_2}^2 s_1^2 s_2^2}{n_2 t_{\nu_1}^2 s_1^2 + n_1 t_{\nu_2}^2 s_2^2} \times \frac{1}{z^2}. \tag{3.3}$$

The approximate CI for $\tilde{\mu}$ is then $\tilde{X} \pm z\tilde{\sigma}$.

Method 2. Since $(\bar{X}_g - \mu_g)\sqrt{n_g}/s_g \sim T_g$, conditional on \bar{X}_g and s_g^2 , $\mu_g \sim_f \bar{X}_g + (s_g/\sqrt{n_g})T_g, g = 1, 2$, where T_g is a standard T random variable with degrees of freedom ν_g and \sim_f represents the fiducial distribution, which is equivalent to a Bayesian posterior distribution under the usual noninformative priors. The variance of $\tilde{\mu}$ is minimized at $\hat{\gamma} = n_1^* s_2^2 / (n_2^* s_1^2 + n_1^* s_2^2)$, where $n_g^* = n_g(n_g - 3)/(n_g - 1), g = 1, 2$. This suggests taking

$$\tilde{X} = \frac{n_1^* s_2^2 \bar{X}_1 + n_2^* s_1^2 \bar{X}_2}{n_1^* s_2^2 + n_2^* s_1^2}. \tag{3.4}$$

We still suggest using the variance estimate in (3.3), giving the CI

$$\tilde{L} = \tilde{X} - z\tilde{\sigma}, \quad \tilde{U} = \tilde{X} + z\tilde{\sigma}. \tag{3.5}$$

The use of $\tilde{\sigma}^2$ is desirable because the CI for μ_1 , derived from $(\tilde{X} - z\tilde{\sigma}, \tilde{X} + z\tilde{\sigma})$ using $\tilde{\sigma}^2$ from (3.3), always gives smaller or at least equal length interval compared to the unrestricted interval $(\bar{X}_1 - t_{n_1-1} s_1, \bar{X}_1 + t_{n_1-1} s_1)$, whereas this does not hold if we replace n_g by n_g^* in (3.3).

Another way to calculate \tilde{L} and \tilde{U} is to use the exact fiducial distribution of $\tilde{\mu}$,

$$\tilde{\mu} \sim \tilde{X} + \hat{\gamma} \left(\frac{s_1}{\sqrt{n_1}} \right) T_1 + (1 - \hat{\gamma}) \left(\frac{s_2}{\sqrt{n_2}} \right) T_2, \tag{3.6}$$

and numerically calculate the percentiles of this distribution. Simulations show that the restricted CI using the exact fiducial distribution of $\tilde{\mu}$ in (3.6) gives similar results to the CI at (3.5) using $\tilde{\sigma}^2$ when comparing average width and coverage rate.

4. Confidence Intervals with Three or More Groups

4.1. Confidence intervals with known variances

Suppose $X_{gi} \sim N(\mu_g, \sigma_g^2), i = 1, \dots, n_g, g = 1, \dots, G$, and assume that $\mu_1 \leq \mu_2 \leq \dots \leq \mu_G$. Let $\bar{X}_g = \sum_{i=1}^{n_g} X_{gi} / n_g$ and $\bar{X}_{\ell,u} = \sum_{g=\ell}^u (\bar{X}_g n_g \sigma_g^{-2}) / \sum_{g=\ell}^u (\sigma_g^{-2})$,




Figure 4. Upper limits of the CI for the middle group in three or more sample case.

$1 \leq \ell \leq u \leq G$, with $\mu_{\ell,u}$ denoting its mean. The unrestricted CI for $\mu_{\ell,u}$ is $\bar{X}_{\ell,u} \pm z\sigma_{\ell,u}$, where $\sigma_{\ell,u}^2 = 1/\sum_{g=\ell}^u(n_g\sigma_g^{-2})$, and $\bar{L}_{\ell,u}$ and $\bar{U}_{\ell,u}$ denote lower and upper limits. Note that $\mu_{g,g} = \mu_g$ and $\sigma_{g,g}^2 = \sigma_g^2/n_g, g = 1, \dots, G$.

We reduce the problem to one of comparing two groups. For group 1, we can construct the CI for μ_1 based on the comparison of groups 1 and 2. Thus the CI for μ_1 is

$$\hat{L}_1 = \min(\bar{L}_{1,1}, \bar{L}_{1,2}), \quad \hat{U}_1 = \min(\bar{U}_{1,1}, \bar{U}_{1,2}). \tag{4.1}$$

Similarly, the CI for μ_G is based on the comparison of groups $G-1$ and G , and is

$$\hat{L}_G = \max(\bar{L}_{G,G}, \bar{L}_{G-1,G}), \quad \hat{U}_G = \max(\bar{U}_{G,G}, \bar{U}_{G-1,G}). \tag{4.2}$$

We consider the CI for μ_g , where $1 < g < G$. This involves two two-sample problems, groups $g-1$ and g , and groups g and $g+1$, which would each give respectively lower and upper bound $(\bar{L}_{g-1,g}, \bar{U}_{g-1,g})$ and $(\bar{L}_{g,g+1}, \bar{U}_{g,g+1})$. For the upper bound of the CI for μ_g , if $\bar{U}_{g,g+1} \geq \bar{U}_{g-1,g}$, then $\min\{\bar{U}_{g,g+1}, \max(\bar{U}_{g,g}, \bar{U}_{g-1,g})\}$ and $\max\{\bar{U}_{g-1,g}, \min(\bar{U}_{g,g}, \bar{U}_{g,g+1})\}$ are both possible upper bounds, but both are equal to $\text{median}\{\bar{U}_{g-1,g}, \bar{U}_{g,g}, \bar{U}_{g,g+1}\}$ (see Figure 4(a)). If $\bar{U}_{g,g+1} < \bar{U}_{g-1,g}$, it is not clear how to pick a value for \hat{U}_g . In the two-sample case for groups $g-1$ and g , $\hat{U}_g = \max(\bar{U}_{g,g}, \bar{U}_{g-1,g})$ implies that $\hat{U}_g \geq \bar{U}_{g-1,g}$, while in the two-sample case for groups g and $g+1$, $\hat{U}_g = \min(\bar{U}_{g,g}, \bar{U}_{g,g+1})$ implies that $\hat{U}_g \leq \bar{U}_{g,g+1}$. Since $\bar{U}_{g,g+1} < \bar{U}_{g-1,g}$, a good value for \hat{U}_g should be between $\bar{U}_{g,g+1}$ and $\bar{U}_{g-1,g}$. Note that the true means are ordered as $\mu_{g-1,g} \leq \mu_{g-1,g+1} \leq \mu_{g,g+1}$, though $\bar{U}_{g-1,g+1}$ may not be between $\bar{U}_{g,g+1}$ and $\bar{U}_{g-1,g}$. We propose $\hat{U}_g = \text{median}(\bar{U}_{g-1,g}, \bar{U}_{g-1,g+1}, \bar{U}_{g,g+1})$ (see Figure 4(b)), and note that $\hat{U}_g = \bar{U}_{g-1,g+1}$

in most, but not all cases. The proposed restricted CI for μ_g is

$$\hat{U}_g = \begin{cases} \text{median}(\bar{U}_{g-1,g}, \bar{U}_{g,g}, \bar{U}_{g,g+1}) & \text{if } \bar{U}_{g,g+1} \geq \bar{U}_{g-1,g}, \\ \text{median}(\bar{U}_{g-1,g}, \bar{U}_{g-1,g+1}, \bar{U}_{g,g+1}) & \text{otherwise,} \end{cases} \tag{4.3}$$

$$\hat{L}_g = \begin{cases} \text{median}(\bar{L}_{g-1,g}, \bar{L}_{g,g}, \bar{L}_{g,g+1}) & \text{if } \bar{L}_{g,g+1} \geq \bar{L}_{g-1,g}, \\ \text{median}(\bar{L}_{g-1,g}, \bar{L}_{g-1,g+1}, \bar{L}_{g,g+1}) & \text{otherwise.} \end{cases}$$

4.2. Confidence intervals with unknown variances

Here the restricted CIs can also be defined using (4.1)–(4.3). We discuss how to define the limits of unrestricted CI $(\bar{L}_{\ell,u}, \bar{U}_{\ell,u}), 1 \leq \ell \leq u \leq G$.

When $w_g = \sigma^2/\sigma_g^2, g = 1, \dots, G$, is known, for some unknown σ^2 , take the intermediate random variable $\bar{X}_{\ell,u} = \sum_{g=\ell}^u (\bar{X}_g n_g w_g) / \sum_{g=\ell}^u (n_g w_g)$ with corresponding mean $\mu_{\ell,u}$. The unrestricted CI for $\mu_{\ell,u}$ is $(\bar{L}_{\ell,u}, \bar{U}_{\ell,u}) = \bar{X}_{\ell,u} \pm t_\nu s (\sum_{g=\ell}^u n_g w_g)^{-1/2}$, where $s^2 = \sum_{g=1}^G \{w_g \sum_{i=1}^{n_g} (X_{gi} - \bar{X}_g)^2\} / (\sum_{g=1}^G n_g - G)$ and $\nu = \sum_{g=1}^G n_g - G$.

If $\sigma_g^2, g = 1, \dots, G$, needs to be estimated separately, let $s_g^2 = \sum_{i=1}^{n_g} (X_{gi} - \bar{X}_g)^2 / (n_g - 1)$. In this situation, we use Method 2 of Section 3.2 to obtain the means and CIs for the combined groups. With $n_g^* = n_g(n_g - 3) / (n_g - 1)$, the mean estimate for groups from ℓ to u is

$$\bar{X}_{\ell,u} = \left(\sum_{g=\ell}^u n_g^* s_g^{-2} \bar{X}_g \right) \left(\sum_{g=\ell}^u n_g^* s_g^{-2} \right)^{-1}.$$

The variance approximation for this mean is

$$\bar{\sigma}_{\ell,u}^2 = \left(\sum_{g=\ell}^u n_g s_g^{-2} t_{n_g-1}^{-2} z^2 \right)^{-1}.$$

With $\mu_{\ell,u} = E\bar{X}_{\ell,u}$, the approximation of the unrestricted CI for $\mu_{\ell,u}$ is $(\bar{L}_{\ell,u}, \bar{U}_{\ell,u}) = \bar{X}_{\ell,u} \pm z\bar{\sigma}_{\ell,u}$.

5. Other Restricted Confidence Intervals

There are other approaches to construct restricted CIs, including bootstrap based CIs and constant length CIs (Hwang and Peddada (1994)).

Two sampling schemes for the bootstrap based on pivotal distributions are considered. The first is based on the unrestricted MLE, in which $X_{gi}^b, b = 1, \dots, B, i = 1, \dots, n_g, g = 1, \dots, G$, is sampled from $N(\bar{X}_g, \sigma_g^2)$ if σ_g^2 is known, X_{gi}^b is sampled from $\bar{X}_g + s T_g / \sqrt{w_g}$, where T_g is standard T random variable

Table 2. Half life of an antibiotic in rats.

Dose (mg/kg)	Data (<i>h</i>)					Mean	Standard Deviation	restricted MLE	
								Same Vars	Diff. Vars
5	1.17	1.12	1.07	0.98	1.04	1.076	0.073	1.076	1.076
10	1.00	1.21	1.24	1.14	1.34	1.186	0.126	1.186	1.186
25	1.55	1.63	1.49	1.53		1.550	0.059	1.524	1.547
50	1.21	1.63	1.37	1.50	1.81	1.504	0.231	1.524	1.547
200	1.78	1.93	1.80	2.07	1.70	1.856	0.145	1.856	1.856

with degrees of freedom $\sum_{g=1}^G(n_g - 1)$ if w_g is known, or X_{gi}^b is sampled from $\bar{X}_g + s_g T_g$, where T_g is standard T random variable with degrees of freedom $n_g - 1$ if σ_g^2 is estimated. The second scheme is based on the restricted MLE $\hat{\mu}_g$, in which X_{gi}^b is sampled in the three different ways described except that the mean is $\hat{\mu}_g$ instead of \bar{X}_g , where $\hat{\mu}_g$ is the restricted MLE of μ_g . For each bootstrap sample, a bootstrap estimate $\hat{\mu}_g^b$, $b = 1, \dots, B$, is obtained by applying the restricted maximum likelihood method. CIs are based on the percentiles of the bootstrap distribution of $\hat{\mu}_g^b$.

Hwang and Peddada (1994) proposed a constant length CI in which the center of the CI is shifted from the unrestricted MLE to the restricted MLE. They showed that, under fairly general conditions, the coverage probability of the CI centered on the restricted MLE exceeds the nominal level. In our setting, the constant length CI is $\hat{\mu}_g \pm z\sigma_g/\sqrt{n_g}$ if σ_g^2 is known, or is $\hat{\mu}_g \pm t_v s_g/\sqrt{n_g}$ if σ_g^2 is unknown, where $v = \sum_{g=1}^G(n_g - 1)$ for known w_g or $v = n_g - 1$ for the case when σ_g^2 s are estimated separately.

6. Example

The half-life of a drug is the time needed to halve the concentration of the drug in the body of a human or an animal. The half-life may vary with the concentration of the drug, and usually is longer for higher concentration levels. Table 2 has data from Hirotsu (2005). It shows the half-lives in hours of an antibiotic at four different doses injected into rats. The higher dose level should result in a higher concentration and hence it is reasonable to assume the half life is shorter for the lower dose level.

The analysis is based on two scenarios. First, we assume the observations are normal with means that depend on the doses but with the same variance. Second, we assume that the variances of different dose levels may not be equal. The results are shown in Table 3. Some CIs are unchanged while some become narrower. The CIs for doses of 25 mg/kg and 50 mg/kg are narrower for both scenarios, whereas the sample means of the half-lives do not satisfy the constraint. The most noticeable reduction in the width of the restricted CI is for the dose

Table 3. Estimates and CIs of half lives in the example in Section 6.

Dose (mg/kg)	Assuming Equal variances				Assuming Different variances		
	MLE	RMLE	standard CI	restricted CI	RMLE	standard CI	restricted CI
5	1.08	1.08	(0.94, 1.21)	(0.94, 1.21)	1.08	(0.99, 1.17)	(0.99, 1.17)
10	1.19	1.19	(1.05, 1.32)	(1.05, 1.32)	1.19	(1.03, 1.34)	(1.03, 1.34)
25	1.55	1.52	(1.40, 1.70)	(1.40, 1.62)	1.55	(1.46, 1.64)	(1.46, 1.63)
50	1.50	1.52	(1.37, 1.64)	(1.42, 1.64)	1.55	(1.22, 1.79)	(1.46, 1.79)
200	1.86	1.86	(1.72, 1.99)	(1.72, 1.99)	1.86	(1.68, 2.04)	(1.68, 2.04)

of 50 mg/kg when we estimate the variances separately, where the width of the restricted CI is about 58% of the unrestricted one.

7. Simulation Study

We have undertaken numerous simulation studies for the two-sample case. We considered many different scenarios by varying $\mu_1, \mu_2, \sigma_1^2, \sigma_2^2, n_1$, and n_2 . We found that the proposed approaches give excellent coverage rates close to the nominal level even in small sample sizes, and that the widths are narrower than those of unrestricted intervals and can be substantially narrower. The two methods in 3.2 give similar results with very slightly better properties for Method 2. We present results only for the more interesting and challenging three-sample case.

Let the population means of the three groups be ordered as $\mu_1 \leq \mu_2 \leq \mu_3$. Coverage probabilities and the average width of CIs were calculated based on 10,000 simulated datasets and each bootstrap CI was based on 1,999 bootstrap estimates. The distributions and sample sizes for the simulations are listed in Table 4. We included in the comparison the CI based on the unrestricted estimates (Unres), the shifted constant length CI (Shifted Const), the bootstrap methods, and the method of Section 4. The three parametric bootstrap methods are, the completely unrestricted (Bootstrap Unres), the method where the bootstrap samples are simulated from a distribution centered at \bar{X}_g and the restricted MLE is estimated (Bootstrap RMLE), and the method where the bootstrap samples are simulated from a distribution centered at $\hat{\mu}_g$ and the restricted MLE is estimated (Bootstrap-R RMLE).

We present the results for coverage rates and average CI widths in Table 5 for known ratios of variances, and in Table 6 for the case where all variances are estimated.

As expected, the shifted constant length CI centered on the restricted MLE has higher coverage probability than the nominal level; however, the CI can be extremely conservative for μ_2 when all three population means are close to one another and the sample size for group 2 is not large (cases (a), (b), (f) in Table

Table 4. Different combinations of population means, variances and sample sizes used in simulation studies.

	μ	σ^2	n
(a)	-0.1, 0, 0.1	10, 10, 10	5, 5, 5
(b)	-0.1, 0, 0.1	10, 10, 10	5, 5, 50
(c)	-0.1, 0, 0.1	10, 10, 10	5, 50, 10
(d)	-1.0, 0, 1.0	10, 10, 10	10, 10, 10
(e)	-0.1, 0, 2.0	10, 10, 10	10, 50, 10
(f)	0, 0, 0	10, 10, 10	10, 10, 10

Table 5. Empirical coverage rate and average width of 95% CI for μ_1 , μ_2 and μ_3 when the ratios of variances are known.

		Method					
		Unres	Shifted Const	Bootstrap Unres	Bootstrap RMLE	Bootstrap-R RMLE	New method
(a)	μ_1	95.3(6.06)	97.5(6.06)	95.1(6.02)	92.1(4.58)	91.7(4.83)	95.8(5.20)
	μ_2	95.2(6.06)	99.2(6.06)	95.0(6.02)	95.9(4.07)	97.5(4.25)	96.5(4.37)
	μ_3	95.2(6.06)	97.3(6.06)	95.3(6.02)	91.1(4.57)	90.5(4.82)	95.3(5.20)
(b)	μ_1	95.1(5.64)	97.3(5.64)	94.9(5.61)	89.6(3.66)	88.3(4.02)	95.0(4.84)
	μ_2	95.0(5.64)	99.6(5.64)	95.0(5.61)	95.7(2.65)	97.2(2.98)	97.0(2.93)
	μ_3	95.2(1.78)	95.6(1.78)	94.7(1.78)	93.5(1.69)	93.5(1.70)	95.0(1.74)
(c)	μ_1	94.8(5.62)	97.5(5.62)	94.6(5.60)	90.7(3.48)	89.9(3.89)	95.7(3.75)
	μ_2	94.7(1.78)	96.5(1.78)	94.4(1.77)	94.7(1.63)	95.5(1.65)	95.5(1.67)
	μ_3	94.6(3.98)	97.2(3.98)	94.5(3.96)	91.5(2.68)	90.9(2.93)	95.4(2.88)
(d)	μ_1	94.7(4.07)	96.4(4.07)	94.5(4.04)	94.6(3.48)	95.2(3.54)	95.1(3.75)
	μ_2	95.4(4.07)	98.3(4.07)	94.9(4.05)	95.7(3.13)	97.2(3.19)	96.2(3.43)
	μ_3	94.8(4.07)	96.7(4.07)	94.8(4.05)	94.9(3.49)	95.8(3.53)	95.3(3.75)
(e)	μ_1	94.7(3.97)	97.5(3.97)	94.5(3.96)	93.3(2.70)	93.1(2.95)	95.7(2.85)
	μ_2	95.4(1.78)	96.4(1.78)	95.2(1.77)	95.1(1.66)	94.9(1.67)	95.5(1.70)
	μ_3	94.9(3.97)	96.0(3.97)	94.6(3.96)	94.7(3.60)	95.9(3.61)	94.9(3.80)
(f)	μ_1	94.5(4.06)	97.2(4.06)	94.2(4.04)	90.6(3.04)	90.3(3.23)	94.9(3.46)
	μ_2	95.2(4.06)	99.5(4.06)	95.2(4.04)	95.5(2.72)	97.5(2.86)	96.3(2.89)
	μ_3	95.6(4.06)	97.8(4.06)	95.2(4.04)	91.0(3.04)	89.8(3.23)	95.6(3.47)

5 and 6. Even though the bootstrap method works well when all the population means are well separated (case (d)), the coverage rates for some population means can be well below the nominal level in some situations (μ_1 for cases (a), (b), (c) and (f)). There are no noticeable improvements from using the bootstrap method with sampling centered on the restricted MLE compared to the bootstrap method with sampling centered on the unrestricted MLE. Our method gives fairly accurate coverage rate with reduced width of the intervals in all the situations considered in this study.

The findings regarding the coverage rate and length of CI's for the three-group case apply to situations with more than three groups, since the CI for

Table 6. Empirical coverage rate and average width of 95% CI for μ_1 , μ_2 and μ_3 when the variances are estimated separately.

		Method					
		Unres	Shifted Const	Bootstrap Unres	Bootstrap RMLE	Bootstrap-R RMLE	New method
(a)	μ_1	95.2(7.45)	96.6(7.45)	95.1(7.41)	90.6(5.53)	88.2(5.63)	95.1(6.58)
	μ_2	95.7(7.39)	97.1(7.39)	95.7(7.35)	93.4(4.65)	92.5(4.59)	94.2(4.89)
	μ_3	95.1(7.39)	96.3(7.39)	95.0(7.35)	89.9(5.49)	86.9(5.59)	94.6(6.53)
(b)	μ_1	95.0(7.35)	97.0(7.35)	94.7(7.31)	88.5(4.61)	87.2(4.98)	95.2(6.53)
	μ_2	94.8(7.37)	98.2(7.37)	94.6(7.34)	94.4(2.98)	95.2(3.21)	95.2(3.35)
	μ_3	95.3(1.79)	95.4(1.79)	95.0(1.78)	93.7(1.70)	92.2(1.69)	95.1(1.78)
(c)	μ_1	95.5(7.39)	97.6(7.39)	95.4(7.36)	90.9(4.37)	90.3(4.78)	96.2(5.12)
	μ_2	95.1(1.79)	96.3(1.79)	94.7(1.78)	94.3(1.65)	94.5(1.65)	94.3(1.68)
	μ_3	94.7(4.39)	97.4(4.39)	94.7(4.37)	91.3(2.90)	90.8(3.14)	96.0(3.40)
(d)	μ_1	95.0(4.40)	96.1(4.40)	94.6(4.37)	94.1(3.69)	94.6(3.75)	95.8(4.09)
	μ_2	95.1(4.39)	97.2(4.39)	94.9(4.37)	95.1(3.25)	96.2(3.30)	95.4(3.61)
	μ_3	95.0(4.42)	96.3(4.42)	94.9(4.40)	94.9(3.71)	95.2(3.76)	96.0(4.11)
(e)	μ_1	94.4(4.39)	97.4(4.39)	94.2(4.37)	93.2(2.92)	93.3(3.17)	96.1(3.36)
	μ_2	95.3(1.79)	96.4(1.79)	95.2(1.77)	94.9(1.67)	94.5(1.67)	95.4(1.71)
	μ_3	94.8(4.42)	95.5(4.42)	94.6(4.40)	94.6(3.91)	95.4(3.91)	95.5(4.16)
(f)	μ_1	94.9(4.43)	97.2(4.43)	94.5(4.41)	89.5(3.27)	88.4(3.42)	95.7(4.01)
	μ_2	95.4(4.40)	98.5(4.40)	95.1(4.38)	93.8(2.84)	95.2(2.93)	94.9(3.01)
	μ_3	95.1(4.36)	97.3(4.36)	95.0(4.34)	89.9(3.24)	88.6(3.40)	95.6(3.97)

a group only depends on the data for that group and its immediately adjacent neighbors, and not on the data from more distant groups.

The results for known σ_g^2 are uniformly excellent for the new method and are not presented.

8. Discussion

The CI developed in Section 3 utilizes a CI for an intermediate variable, where that CI is centered at a weighted average of X_1 and X_2 . An alternative is to define the center of the CI for the intermediate variable as the restricted MLE under the assumption that the mean of X_1 equals the mean of X_2 . In simulation studies this method gave CI's with similar properties to that of the second method in Section 3.2.

The methods developed in Sections 2, 3 and 4 are applicable to normal observations. On account of the Central Limit Theorem, we expect the coverage rates of the restricted CIs for the means to be close to the nominal level in nonnormal populations if the sample size is fairly large. We found this to be empirically true in simulations (not shown), except when the distribution is highly skewed and the sample size is relatively small. Even in this case, the coverage rates were not

substantially below the nominal level, and showed much better coverage rates than the alternative bootstrap CIs.

The method proposed in this paper can be generalized to other distributions by using transformations. For example, if T is a monotone transformation and $T(\hat{\mu}_g)$ is approximately normally distributed, then it is possible to apply our method to estimate the CI for $T(\mu_g)$ and then apply T^{-1} to obtain the CI for μ_g . For example, for the binomial case with success probability μ_g , the variance stabilizing transformation $\sin^{-1}(\sqrt{\mu_g})$ could be used.

The strategy for obtaining a confidence interval for group g that we developed can be broadly described using two stages: in the first stage, for some sets of neighboring combined groups ℓ, \dots, u , an unrestricted CI, $(\bar{L}_{\ell,u}, \bar{U}_{\ell,u})$, for $\mu_{\ell,u}$ is obtained using a normal or t distribution; in the second stage, the CI for μ_g is modified using the bounds $\bar{L}_{\ell,u}$ and $\bar{U}_{\ell,u}$ based on the order restrictions using, for example, (4.1), (4.2), and (4.3). An alternative is to use the bootstrap in the first stages to obtain $(\bar{L}_{\ell,u}, \bar{U}_{\ell,u})$ instead of the normal or t distribution.

The method for three or more samples discussed in Section 4 could potentially be made more efficient by combining groups rather than just considering the closest group. Specifically when obtaining the upper bound for μ_g , consider a combination of groups $g + 1, \dots, u$, rather than just group $g + 1$, then apply the obvious generalizations of (4.1), (4.2), and (4.3). Whether it is beneficial to combine groups depends on the closeness of the means of neighboring groups and the decision of whether to combine groups could be based on either prior knowledge, or potentially a pre-test from the available data. For example, one approach for the three group situation would be to test $H_0 : \mu_1 \leq \mu_2 = \mu_3$ vs $H_a : \mu_1 \leq \mu_2 < \mu_3$ at a certain significance level to decide whether μ_2 and μ_3 are close to each other, and so to decide whether or not to combine groups 2 and 3 to construct the restricted CIs for μ_1 .

References

- Andrews, D. W. K. (2000). Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica* **68**, 399-405.
- Barlow, R. E., Bartholomew, D. J., Bremner, J. M. and Brunk, H. D. (1972). *Statistical Inference under Order Restrictions*. Wiley, New York.
- Cohen, A. and Sackrowitz, H. B. (2004). A discussion of some inference issues in order restricted models. *Canad. J. Statist.* **32**, 199-205.
- Dunson, D. B. and Neelon, B. (2003). Bayesian inference on order-constrained parameters in generalized linear models. *Biometrics* **59**, 286-295.
- Hirotsu, C. (2005). Isotonic inference. In *Encyclopedia of Biostatistics* (Edited by P. Armitage and T. Colton), 2nd edition. Wiley, New York.
- Hwang, J. T. G. (1995). Fieller's problems and resampling techniques. *Statist. Sinica* **5**, 161-171.

- Hwang, J. T. G. and Peddada, S. D. (1994). Confidence interval estimation subject to order restrictions. *Ann. Statist.* **22**, 67-93.
- Kelly, R. E. (1989). Stochastic reduction of loss in estimating normal means by isotonic regression. *Ann. Statist.* **17**, 937-940.
- Lee, C.-I. C. (1981). The quadratic loss of isotonic regression under normality. *Ann. Statist.* **9**, 686-688.
- Li, Z., Taylor, J. M. G. and Nan, B. (2010). Construction of confidence intervals and regions for ordered binomial probabilities. *Amer. Statist.* **64**, 291-298.
- Marchand, É. and Strawderman, W. E. (2004). Estimation in restricted parameter spaces: A review. In A. DasGupta, ed., *Lecture notes-monograph series*, **45**. Institute of Mathematical Statistics, Beachwood, OH, 21-44.
- Peddada, S. D. (1997). Confidence interval estimation of population means subject to order restrictions using resampling procedures. *Statist. Probab. Lett.* **31**, 255-265.
- Robertson, T., Wright, F. T. and Dykstra, R. L. (1988). *Order Restricted Statistical Inference*. Wiley, New York.
- Schoenfeld, D. A. (1986). Confidence bounds for normal means under order restrictions, with application to dose-response curves, toxicology experiments, and low-dose extrapolation. *J. Amer. Statist. Assoc.* **81**, 186-195.
- van Eeden, C. (2006). *Restricted Parameter Space Estimation Problems: Admissibility and Minimality Properties*. Springer, New York.

Department of Biostatistics, University of Pittsburgh, Pittsburgh, PA 15261, USA.

E-mail: yongpark@pitt.edu

Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109-2029, USA.

E-mail: jdkalbf@umich.edu

Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109-2029, USA.

E-mail: jmgt@umich.edu

(Received January 2012; accepted January 2013)