Supplementary Data

Given only summary statistics, it is necessary to make a distributional assumption. An assumption of normally distributed outcomes appears reasonable from experience with data from other similar experiments. This leads to the likelihood function,

$$p\left(x|\mu,\sigma^2\right) = \left(2\pi\sigma^2\right)^{-\frac{n}{2}} \exp\left(-\frac{\sum_{i=1}^{n}(x_i-\mu)^2}{2\sigma^2}\right),$$

where $x = (x_1, x_2, \ldots, x_n)$, $x_i =$ observed variable (e.g., C_{max}) for individual i. It is mathematically convenient to reparameterize in terms of the precision $\tau = 1/\sigma^2$ (the inverse of the variance), so this becomes,

$$p(x|\mu,\tau) \propto \tau^{\frac{n}{2}} \exp\left(-(\tau/2)\sum_{i=1}^{n}(x_i-\mu)^2\right),$$

where $\propto$ indicates proportionality, that is, "equal to up to a constant of proportionality," so that $x \propto y$ is equivalent to $x = cy$, c a constant.

Under the normality assumption, the sample mean, sample standard deviation, and number of observations are sufficient statistics to fully determine the likelihood, so Bayesian inference is then possible.

The conventional vague, uniform prior for μ and $\log(\tau)$, namely

$$p(\mu,\tau) \propto \frac{1}{\tau},\ -\infty < \mu < \infty,\ \tau > 0,$$

is a standard choice to represent uninformative prior beliefs concerning μ and τ.

This leads to the joint posterior density,

$$p(\mu,\tau|x) \propto \tau^{\frac{n}{2}-1} \exp\left(-(\tau/2)\sum_{i=1}^{n}(x_i-\mu)^2\right).$$

Integrating w.r.t τ leads to the marginal posterior density for μ,

$$p(\mu|x) \propto \left(1 + \frac{n(\mu-\overline{x})^2}{\nu s^2}\right)^{-\frac{n}{2}},$$

which is a Student's t-density with $\nu = n-1$, $\overline{x} = \sum x_i/n$, and $s^2 = \sum_{i=1}^{n}(x_i-\overline{x})^2/\nu$.

Marginalizing w.r.t to μ gives the marginal density for τ (or σ^2 if we substitute for $\tau = 1/\sigma^2$),

$$p(\tau|x) \propto \tau^{\frac{\nu}{2}-1} \exp\left(-\frac{\tau\nu s^2}{2}\right),$$

which is a Gamma density (see, e.g., Lancaster 2004, pp. 120–125).

The marginal density for μ is generally the posterior of interest. Given two sample means we wish to compare, we can generate a large Monte Carlo (MC) pseudosample by obtaining random draws from this Student's t-distribution and then compute the differences to obtain the posterior density of the difference in means for two samples.

With an MC sample from the posterior, the accuracy can be made arbitrarily high by increasing the pseudosample size. The law of large numbers then assures that the expected value of any function of the MC sample converges its true value, that is, for pseudosample of R draws for z, as $R \to \infty$,

$$E(f(z)) \to \frac{1}{R}\sum_{r=1}^{R} f\left(z^{(r)}\right),$$

where $z^{(r)}$ is the rth pseudosample draw.

The algorithm for comparison of means is then as follows.

1. Given sample means $\overline{x}_1$ and $\overline{x}_2$, sample standard deviations s_1 and s_2, and samples sizes n_1 and n_2, randomly draw R values, μ_j^r, $r = 1, 2, \ldots, R$, from each of the marginal posterior distributions $p(\mu_1|x_1)$ and $p(\mu_2|x_2)$, which are Student-t distributions as given above.
2. Compute R differences in means, $\Delta\mu^{(r)} = \mu_1^{(r)} - \mu_2^{(r)}$.
3. Quantiles, means, and so on can then be computed from the sample of differences in means.