# Mid-p Adjustment for exact2x2: Computational Details 

Michael P. Fay
December 8, 2021

## Overview

These notes give details on how the mid-p adjustment is done. Section 1 describes the mid-p adjustment as it is done for the exact $2 \times 2$ and uncondExact $2 \times 2$ functions. Section 2 describes the mid-p adjustment as implemented in the binomMeld.test function.

## 1 Usual Mid-p Adjustment for Two Binomial Distributions

The following is how the usual mid-p adjustment is done (for example in the exact $2 \times 2$ and uncondExact $2 \times 2$ functions). The mid-p value has a long history (see e.g., Lancaster, 1961 or the list of references in Hirji 2006, p. 50).

Let $\mathbf{X}=\left[X_{1}, X_{2}\right]$ with $X_{a} \sim \operatorname{Binom}\left(n_{a}, \theta_{a}\right)$ for $a=1,2$. Suppose we are interested in $\beta=b(\theta)$, where $b(\theta)$ is some function of $\theta_{1}$ and $\theta_{2}$. Common examples are the difference, $\beta_{d}=\theta_{2}-\theta_{1}$, the ratio, $\beta_{r}=\theta_{2} / \theta_{1}$, and the odds ratio, $\beta_{o r}=\left\{\theta_{2}\left(1-\theta_{1}\right)\right\} /\left\{\theta_{1}\left(1-\theta_{2}\right)\right\}$.

Let $T(\mathbf{X})$ be some test statistic, where larger values are most extreme with respect to the null hypothesis. Let $\Theta_{0}$ be the set of all possible values of $\left[\theta_{1}, \theta_{2}\right]$ under the null hypothesis. Then a valid (i.e., exact) p-value is

$$
p\left(\mathbf{x}, \Theta_{0}\right)=\sup _{\theta \in \Theta_{0}} \operatorname{Pr}_{\theta}[T(\mathbf{X}) \geq T(\mathbf{x})]
$$

These exact p-values are necessarily conservative because for most $\theta \in \Theta_{0}$ we have $\operatorname{Pr}_{\theta}\left[p\left(\mathbf{X}, \Theta_{0}\right) \leq \alpha\right]<\alpha$. A less conservative approach, but one that is no longer valid (i.e., no longer exact), is to use a mid-p value. The mid-p value is

$$
p_{m i d}\left(\mathbf{x}, \Theta_{0}\right)=\sup _{\theta \in \Theta_{0}}\left\{\operatorname{Pr}_{\theta}[T(\mathbf{X})>T(\mathbf{x})]+\frac{1}{2} \operatorname{Pr}_{\theta}[T(\mathbf{X})=T(\mathbf{x})]\right\}
$$

It is convenient to write $\Theta_{0}$ in terms of $\beta$. For example,

$$
\Theta_{0}=\left\{\theta: b(\theta)=\beta_{0}\right\}
$$

For this example, instead of writing the null hypothesis as $H_{0}: \theta \in \Theta_{0}$, we write it in terms of $\beta=b(\theta)$ as $H_{0}: \beta=\beta_{0}$. We are generally interested in three classes of hypotheses: two-sided hypotheses,

$$
\begin{array}{ll}
H_{0}: & \beta=\beta_{0} \\
H_{1}: & \beta \neq \beta_{0}
\end{array}
$$

or one of the one-sided hypotheses,

$$
\begin{array}{cl}
\frac{\text { Alternative is Less }}{} & \text { Alternative is Greater } \\
H_{0}: \beta \geq \beta_{0} & H_{0}: \beta \leq \beta_{0} \\
H_{1}: \beta<\beta_{0} & H_{1}: \beta>\beta_{0}
\end{array}
$$

Let $p_{t s}\left(\mathbf{x}, \beta_{0}\right)$ be the p -value for testing the two-sided hypotheses, let $p_{U}\left(\mathbf{x}, \beta_{0}\right)$ be the p -value for testing $H_{0}: \beta \geq \beta_{0}$, and $p_{L}\left(\mathbf{x}, \beta_{0}\right)$ be the p -value for testing $H_{0}: \beta \leq \beta_{0}$.

Then we can create $100(1-\alpha) \%$ confidence regions as the set of $\beta_{0}$ value that fail to reject the associated null hypothesis. For example,

$$
C_{t s}(\mathbf{x}, 1-\alpha)=\left\{\beta: p_{t s}(\mathbf{x}, \beta)>\alpha\right\}
$$

gives a "two-sided" confidence region. The region may not be an interval if the p-value function is not unimodal. This problem occurs with Fisher's exact test (the Fisher-Irwin version, or 'minlike' version). For central confidence regions we take the union of the onesided confidence regions, in other words,

$$
C_{c}(\mathbf{x}, 1-\alpha)=C_{L}(\mathbf{x}, 1-\alpha / 2) \cup C_{U}(\mathbf{x}, 1-\alpha / 2),
$$

where $C_{L}$ and $C_{U}$ are the one-sided confidence regions,

$$
C_{L}(\mathbf{x}, 1-\alpha / 2)=\left\{\beta: p_{L}(\mathbf{x}, \beta)>\alpha / 2\right\}
$$

and

$$
C_{U}(\mathbf{x}, 1-\alpha / 2)=\left\{\beta: p_{U}(\mathbf{x}, \beta)>\alpha / 2\right\} .
$$

If the regions are intervals, and we let $L(\mathbf{x}, 1-\alpha / 2)=\min C_{L}(\mathbf{x}, 1-\alpha / 2)$ and $U(\mathbf{x}, 1-\alpha / 2)=$ $\max C_{U}(\mathbf{x}, 1-\alpha / 2)$, then the central interval is

$$
C_{c}(\mathbf{x}, 1-\alpha)=\{L(\mathbf{x}, 1-\alpha / 2), U(\mathbf{x}, 1-\alpha / 2)\}
$$

For the mid-p confidence regions, we replace the p-values with the mid-p values.

## 2 Mid-p Modifications with Binomial Melding

For a single binomial response, the mid p-value and associated central confidence interval can be represented using confidence distribution random variables. Suppose that the exact central 100(1- $\alpha$ ) percent binomial confidence interval for a single binomial random variable (i.e., the default in binom.test) is $(L(1-\alpha / 2), U(1-\alpha / 2))$. Then the lower and upper confidence distribution random variables are respectively, $W_{L}=L\left(A_{1}\right)$ and $W_{U}=U\left(A_{2}\right)$, where $A_{1}$ and $A_{2}$ and independent uniform random variables. Let $B$ be an independent Bernoulli random variable with parameter $1 / 2$. Then the 95 percent central mid-p confidence interval for the binomial parameter is the middle 95 percent of the distribution of $W=B * W_{L}+(1-B) * W_{U}$. This is shown in the appendix of Fay and Brittain (2016).

The way the midp=TRUE option is done in binomMeld.test is to replace the upper and lower confidence distribution random variables in the usual melding equations, with the "midp" confidence distribution random variable (CD-RV) analogous to $W$ for each group. For example if the lower and upper CD-RVs for group 1 are $W_{1 L}$ and $W_{1 U}$, then the mid-p CD-RV is $W_{1}=B_{1} * W_{1 L}+\left(1-B_{1}\right) * W_{1 U}$, where $B_{1}$ is a Bernoulli random variable with parameter $1 / 2$. The mid-p CD-RV $W_{2}$ is defined analogously. It is fairly simple to program a Monte Carlo estimate of the "mid" p-value and associated confidence interval. Let $g\left(\theta_{1}, \theta_{2}\right)$ be the parameter of interest (e.g., $g\left(\theta_{1}, \theta_{2}\right)=\theta_{2}-\theta_{1}$ for parmtype="difference"). The one-sided pvalues are the proportion of times that $g\left(W_{1}, W_{2}\right)$ is $\leq$ nullparm (for alternative="greater") or $\geq$ nullparm (for alternative="less"). The confidence intervals just use the appropriate quantiles of the Monte Carlo values of $g\left(W 1, W_{2}\right)$.

When $n m c=0$, we estimate the one-sided $p$-values with numeric integration. Conceptually, the usual melded $p$-value might be, for example when alternative $=$ "greater" and nullparm $=\beta_{0}$ :

$$
\operatorname{Pr}\left[g\left(W_{1 U}, W_{2 L}\right) \leq \beta_{0}\right]=\int \operatorname{Pr}\left[g\left(W_{1}, w_{2}\right) \leq \beta_{0} \mid W_{2}=w_{2}\right] \operatorname{Pr}\left[W_{2}=w_{2}\right]
$$

where $W_{1 U}$ is the upper confidence distribution random variable (CD-RV) for group 1 and $W_{2 L}$ is the lower CD-RV for group 2. These CD-RVs are beta distributions (see Fay, Proschan, and Brittain, 2015). For the mid-p version, we use

$$
\begin{aligned}
\operatorname{Pr}\left[g\left(W_{1}, W_{2}\right) \leq \beta_{0}\right]= & \frac{1}{4} \int \operatorname{Pr}\left[g\left(W_{1 L}, w_{2}\right) \leq \beta_{0} \mid W_{2 L}=w\right] \operatorname{Pr}\left[W_{2 L}=w\right]+ \\
& \frac{1}{4} \int \operatorname{Pr}\left[g\left(W_{1 L}, w_{2}\right) \leq \beta_{0} \mid W_{2 U}=w\right] \operatorname{Pr}\left[W_{2 U}=w\right]+ \\
& \frac{1}{4} \int \operatorname{Pr}\left[g\left(W_{1 U}, w_{2}\right) \leq \beta_{0} \mid W_{2 L}=w\right] \operatorname{Pr}\left[W_{2 L}=w\right]+ \\
& \frac{1}{4} \int \operatorname{Pr}\left[g\left(W_{1 U}, w_{2}\right) \leq \beta_{0} \mid W_{2 U}=w\right] \operatorname{Pr}\left[W_{2 U}=w\right] .
\end{aligned}
$$

The integration simplifies for special cases (e.g., when $x_{1}=0$ ), and in other case we just use the integrate function. For the confidence intervals we solve for the $\beta_{0}$ values such that the p-values equal either $\alpha$ (for one-sided alternatives) or $\alpha / 2$ (for two-sided alternatives), where alpha $=1$-conf.level. If there is no $\beta_{0}$ value that solves that, we set the confidence limit to the appropriate extreme.

It is known that the p-values that match the melded confidence intervals for two independent binomial observations exactly equal the one-sided Fisher's exact p-values (see Fay, et al, 2015). For example,

```
> x1<-6
> n1<-12
> x2<-15
> n2<- 17
> exact2x2(matrix(c(x2,n2-x2,x1,n1-x1),2,2), tsmethod="central", midp=FALSE)
Central Fisher's Exact Test
```

```
data: matrix(c(x2, n2 - x2, x1, n1 - x1), 2, 2)
p-value = 0.06506
alternative hypothesis: true odds ratio is not equal to 1
95 percent confidence interval:
    0.9119249 89.4167455
sample estimates:
odds ratio
    6.924704
> binomMeld.test(x1,n1,x2,n2, parmtype="oddsratio", midp=FALSE)
    melded binomial test for oddsratio
data: sample 1:(6/12), sample 2:(15/17)
proportion 1 = 0.5, proportion 2 = 0.88235, p-value = 0.06506
alternative hypothesis: true oddsratio is not equal to 1
95 percent confidence interval:
    0.909023 106.265540
sample estimates:
odds ratio {p2(1-p1)}/{p1(1-p2)}
7.5
```

Note, the confidence intervals for the two methods are not equal.
This does not necessarily mean that the midp versions give equivalent p-values:

```
> x1<-6
> n1<-12
> x2<-15
> n2<- 17
> exact2x2(matrix(c(x2,n2-x2,x1,n1-x1),2,2), tsmethod="central", midp=TRUE)
    Central Fisher's Exact Test (mid-p version)
data: matrix(c(x2, n2 - x2, x1, n1 - x1), 2, 2)
p-value = 0.03578
alternative hypothesis: true odds ratio is not equal to 1
95 percent confidence interval:
    1.12685 62.05021
sample estimates:
odds ratio
    6.924704
> binomMeld.test(x1,n1,x2,n2, parmtype="oddsratio", midp=TRUE)
melded binomial test for oddsratio, mid-p version
```

```
data: sample 1:(6/12), sample 2:(15/17)
proportion 1 = 0.5, proportion 2 = 0.88235, p-value = 0.02899
alternative hypothesis: true oddsratio is not equal to 1
95 percent confidence interval:
    1.214721 66.148301
sample estimates:
odds ratio {p2(1-p1)}/{p1(1-p2)}
                                    7.5
```


## References

Fay, MP, and Brittain, EH (2016). "Finite sample pointwise confidence intervals for a survival distribution with right-censored data." Statistics in Medicine 35: 2726-2740.

Fay, MP, Proschan, MA, and Brittain, E (2015). "Combining One-Sample Confidence Procedures for Inference in the Two-Sample Case." Biometrics 71: 146-156.
hirji, K. F. (2006). Exact Analysis of Discrete Data. Chapman and Hal/CRC, New York.
Lancaster, H.O. (1961). Significance Tests in Discrete Distributions. Journal of the American Statistical Association. 56: 223-234.

