

Practical valid inferences for the two-sample binomial problem

Michael P. Fay and Sally A. Hunsberger

*Biostatistics Research Branch,
National Institute of Allergy and Infectious Diseases,
Bethesda, MD, USA*
e-mail: mfay@niaid.nih.gov; Sally.Hunsberger@nih.gov

Abstract: Our interest is whether two binomial parameters differ, which parameter is larger, and by how much. This apparently simple problem was addressed by Fisher in the 1930's, and has been the subject of many review papers since then. Yet there continues to be new work on this issue and no consensus solution. Previous reviews have focused primarily on testing and the properties of validity and power, or primarily on confidence intervals, their coverage, and expected length. Here we evaluate both. For example, we consider whether a p-value and its matching confidence interval are compatible, meaning that the p-value rejects at level α if and only if the $1 - \alpha$ confidence interval excludes all null parameter values. For focus, we only examine non-asymptotic inferences, so that most of the p-values and confidence intervals are valid (i.e., exact) by construction. Within this focus, we review different methods emphasizing many of the properties and interpretational aspects we desire from applied frequentist inference: validity, accuracy, good power, equivariance, compatibility, coherence, and parameterization and direction of effect. We show that no one method can meet all the desirable properties and give recommendations based on which properties are given more importance.

MSC2020 subject classifications: Primary 62F03; secondary 62F25.

Keywords and phrases: 2 by 2 table, Barnard's test, Fisher's exact test, unconditional exact test.

Received June 2020.

Contents

1	Introduction	73
2	Overview: Failure of normal intuition	74
	2.1 Frequentist inferences	74
	2.2 Background and notation	75
	2.3 Standard frequentist inference: Normal intuition	76
	2.4 Two-sample binomial: Failure of normal intuition	77
3	Choosing the effect measure	80
4	Properties of frequentist triples	84
	4.1 Defining a matched triple	84
	4.2 Implications of compatible inferences	84

4.3	Directional inferences	85
5	Methods for creating one-sided exact unconditional testing procedures	86
5.1	Basic procedure for defining p-values	86
5.2	Improving power by breaking ties: Refinement of ordering functions	88
5.3	Ordering functions for ratio and odds ratio	88
5.4	Other improvements: E+M and Berger-Boos	89
5.5	Ordering functions that depend on significance level	90
5.6	Ordering functions that depend on hypothesis space boundaries .	90
6	One-sided conditional exact tests	91
7	Melded confidence intervals	92
8	Non-central confidence intervals and associated tests	95
9	Mid-p methods: Improving accuracy by sacrificing validity	98
10	Computational issues	99
11	Power and efficiency comparisons	99
12	Recommendations	102
A	Proof of Theorem 4.1	105
B	Barnard's CSM ordering	105
	Acknowledgments	106
	References	106

1. Introduction

Suppose we observe two independent binomial variates with parameters (n_1, θ_1) and (n_2, θ_2) . Two questions are: are θ_1 and θ_2 equal? and how much larger is one θ parameter than the other? To answer these two questions, the frequentist typically presents an estimate of an effect, a confidence interval (CI) on that effect, and a p-value to test that there is no effect. Surprisingly, there is no consensus method for testing and creating confidence intervals for this problem. New methods continue to be developed for this problem [see e.g., 38, 57, 58, 21, 27]. Many review papers focus on testing alone [see 40, 48], or confidence intervals alone [see 46, 53, 17]. Here we focus on both.

We limit the scope of this paper by considering only frequentist approaches (so Bayesian methods are not covered), and by not considering asymptotic methods or other approximations. Many review papers or books [see e.g., 46, 40, 17, 47] cover and compare many of those approximations. Sometimes those approximations are closed-form expressions and can be useful for deriving simple sample size formulas or when the test is applied many times such as in genomics. But often the approximations are unnecessary with modern computers. Non-asymptotic methods are often called exact, but in this paper we reserve the term exact for non-asymptotic methods that are valid, meaning tests that control the type I error rate, and confidence intervals that cover the parameter with at least the nominal value. See Section 2.2 for further discussion of the term exact. A class of important non-asymptotic tests that are not valid are mid-p methods (Section 9), which are sometimes called quasi-exact [29] and are included in our review because, for confidence intervals, sometimes we want av-

erage coverage close to the nominal value instead of guaranteed coverage that on average is conservative.

Here is an outline of our paper. Section 2 begins by contrasting the two-sample binomial problem with the two-sample difference in normal distributions with the same variance, in which there is an accepted solution: the two-sample t-test. This allows us to define inferential properties of interest as well as highlight why there is no single accepted solution to the two-sample binomial problem. Newcombe [47] takes a similar approach. Section 3 discusses the choice of effect measure (e.g., difference in binomial parameters, ratio of parameters, or odds ratio of parameters). Section 4 defines a frequentist triple as a parameter estimator, an associated confidence interval procedure, and a p-value function. We then formally discuss some properties of triples, such as whether the confidence interval and p-value match and are compatible, and whether directional inferences may be made from the triple. The idea of matched triples is discussed in Hirji [28, p. 77] in a less formal way as a “unified report”. Our review says very little about parameter estimators, and mostly focuses on properties of p-values and confidence intervals and the compatibility of p-values with confidence intervals. Our discussion of directional inferences is motivated from the three decision rule of Neyman [see e.g., 26]. We describe methods for defining valid one-sided decision rules in Sections 5 (unconditional methods) and 6 (conditional methods), including the associated p-values and confidence intervals. Much of Sections 5 and 6 was thoroughly reviewed in [48] but is included in this paper for completeness; however, Section 5.3 presents some new ideas on informativeness of ordering functions. Section 7 reviews the melded confidence intervals of Fay, Proschan and Brittain [21] which are compatible with the one-sided conditional method (i.e., Fisher’s exact test) p-values. Section 8 discusses non-central confidence intervals and associated tests, with a new focus on the relationship of these intervals to directional inferences. Section 9 discusses mid-p methods, which are non-asymptotic methods that relax the validity assumption in order to achieve better accuracy. Section 10 discusses the computational aspects of various methods. Section 11 discusses power and efficiency of methods, including some new calculations. Section 12 presents recommendations. Briefly, we recommend using exact central confidence intervals (those with equal error bounds on both sides) because it is better for directional inferences. For fast calculations, use exact conditional tests with compatible confidence intervals, but for more power consider exact unconditional tests using the version that orders by the one-sided mid-p Fisher’s exact p-values. If validity is not vital, use mid p-values on the exact conditional test, which are often a good approximation to the exact unconditional tests.

2. Overview: Failure of normal intuition

2.1. Frequentist inferences

We define a frequentist triple (or just a triple) as an estimator of a parameter of interest, a confidence interval, and a p-value function. This approach allows us

to compare different triples by examining not just properties of each component (i.e., comparing powers of different p-value functions or expected lengths of different confidence intervals), but also to examine properties of the triples as a whole. For example, within a triple, we examine inferential agreement between the p-value function and confidence interval procedure. Additionally, we examine what directional inferential statements we can make from the triple, such as whether θ_2 is significantly larger than θ_1 , and at what significance level.

Although in some different statistical settings (e.g., two-sample normal problem) the standard triple will automatically give inferential agreement between p-values and confidence intervals as well as automatically give directional inferential statements, in the two-sample binomial problem those inferential properties are not automatic. Thus, before discussing the binomial problem, we review the two-sample problem with normally distributed responses with the same variance. We consider the latter problem first, because there is some consensus that one triple (the difference in means, and the confidence interval and p-value associated with the t-test) is appropriate for this problem. In the normal case, this t-test triple meets some regularity properties that lead to inferences that are intuitive and easy to understand. Because these properties form the basis for a certain statistical intuition about how frequentist inferences ought to be, and because the example uses normal distributional assumptions, we call these properties the “normal intuition”. We show later how the normal intuition breaks down for the two-sample binomial problem, although many of the properties may approximately hold for large samples.

2.2. Background and notation

Consider a general frequentist problem, where we observe data, \mathbf{x} , and denote its random variable as \mathbf{X} . Assume some probability model for \mathbf{X} that depends on a parameter vector θ , but we are interested in a function of θ that returns a scalar, $b(\theta) = \beta$. We partition the possible values of θ into two sets, the null hypothesis space, Θ_0 , and the alternative hypothesis space, Θ_1 .

In this paper, we consider only three classes of partitions, where the null and alternative space is defined by β , and separated by a value β_0 on the boundary between the null and alternative hypothesis spaces. The first of these three classes are two-sided hypotheses,

$$\begin{aligned} H_0 : & \quad \beta = \beta_0 \\ H_1 : & \quad \beta \neq \beta_0 \end{aligned}$$

which can be equivalently written as

$$\begin{aligned} H_0 : & \quad \theta \in \Theta_0 \text{ where } \Theta_0 = \{\theta : b(\theta) = \beta_0\} \\ H_1 : & \quad \theta \in \Theta_1 \text{ where } \Theta_1 = \{\theta : b(\theta) \neq \beta_0\}. \end{aligned}$$

The other two classes are the one-sided hypotheses,

<u>Alternative is Less</u>	<u>Alternative is Greater</u>
$H_0 : \beta \geq \beta_0$	$H_0 : \beta \leq \beta_0$
$H_1 : \beta < \beta_0$	$H_1 : \beta > \beta_0$

Let $p(\mathbf{x}, \Theta_0)$ be a p-value associated with the null hypothesis space, Θ_0 . Typically, we assume a class of hypotheses and write (with a slight abuse of notation) $p(\mathbf{x}, \beta_0)$ as a p-value associated with the null hypothesis indexed by β_0 . We reject the null hypothesis at level α if $p(\mathbf{x}, \beta_0) \leq \alpha$. Following Berger and Boos [5], we define a p-value procedure as *valid* if

$$P_\theta [p(\mathbf{X}, \beta_0) \leq \alpha] \leq \alpha,$$

for all $\alpha \in (0, 1)$ and all $\theta \in \Theta_0$. (Ripamonti, *et al* [48] call a valid p-value procedure a *guaranteed* p-value.) The term *exact* is often used to describe tests that give valid p-values, but be aware that the term ‘exact’ is used in at least 4 different ways in the literature: (i) methods not based on asymptotic or other approximations [see 28, p.450], (ii) valid methods [see 29, 40, 21], (iii) methods where the size is equal to the significance level (only possible with randomized tests for discrete data) [16], or (iv) methods where the p-values are the smallest p-values among a class of valid p-values [48, equation 2.5], specifically, p-value procedures such that

$$\sup_{\theta \in \Theta_0} P_\theta [p(\mathbf{X}, \beta_0) \leq p(\mathbf{x}, \beta_0)] = p(\mathbf{x}, \beta_0) \text{ for all } \mathbf{x}, \quad (2.1)$$

In this review, we use the term exact only in the sense of (ii).

Following Röhmel [49], we define a p-value procedure as *coherent* if for every \mathbf{x} , $p(\mathbf{x}, \Theta_0^*) \leq p(\mathbf{x}, \Theta_0)$ if $\Theta_0^* \subseteq \Theta_0$.

For the classes of hypotheses above, we can invert the p-value function to get its associated $100(1 - \alpha)\%$ confidence region,

$$C(\mathbf{x}, 1 - \alpha) = \{\beta : p(\mathbf{x}, \beta) > \alpha\}. \quad (2.2)$$

We define a confidence region as *valid* if it is guaranteed to have at least nominal coverage for every θ (and hence every $b(\theta) = \beta$); in other words,

$$P_\theta [\beta \in C(\mathbf{X}, 1 - \alpha)] \geq 1 - \alpha.$$

This paper considers non-asymptotic methods, and all are valid except the mid-p methods described in Section 9.

2.3. Standard frequentist inference: Normal intuition

Consider the two-sample problem, where the a th group has n_a independent and normally distributed responses, with mean μ_a and variance σ^2 , for $a = 1, 2$. Let $\theta = [\mu_1, \mu_2, \sigma]$, and suppose we are interested in $\beta = b(\theta) = \mu_2 - \mu_1$. The t-test is valid for testing the null that $\beta = \beta_0$ and it is the uniformly most powerful

(UMP) unbiased test [37, p. 160]. UMP unbiasedness means that among the class of unbiased tests (i.e., tests for which power for each parameter value in the alternative space is at least as large as the power for every parameter value in the null space), the t-test is the most powerful test for each $\theta \in \Theta_1$.

We study this case first to define “normal intuition” about frequentist inferences. This normal intuition is a series of properties, that if they are not met, conflict with many statisticians’ intuitive feeling of how p-values and confidence regions ought to work. Here are those properties met by the triple: difference in sample means, $\hat{\beta}$; the two-sided p-value from the t-test, p ; and the $100(1 - \alpha)\%$ confidence interval on β associated with that p-value, (L, U) .

Reproducibility: Application of the method by two independent statisticians to the same data always gives the same results (as opposed to randomized tests).

Confidence region is an interval: The confidence region created from p through equation 2.2 is an interval, meaning it can be written as (L, U) .

Compatible Inferences: $p \leq \alpha$ if and only if the $(1 - \alpha)$ confidence interval does not contain β_0 .

Accuracy (of coverage): Taken over repeated applications, the probability that the $100(1 - \alpha)\%$ confidence interval procedure includes β is equal to $(1 - \alpha)$ for all values of θ such that $b(\theta) = \beta$.

Centrality (of CI): The $100(1 - \alpha)\%$ CI is a central one, meaning $P[L > \beta] \leq \alpha/2$ and $P[U < \beta] \leq \alpha/2$.

One-sided p-value from Two-sided p-value: Half of the two-sided p-value can be interpreted as a one-sided p-value in the apparent direction of the effect. For example, if $\hat{\beta} > \beta_0$ then we can reject $H_0 : \beta \leq \beta_0$ at level $p/2$.

Directional Coherence (of p-value): Call a two-sided p-value function *directionally coherent* if the p-values are decreasing as β_0 gets farther from $\hat{\beta}$. In other words, directionally coherent two-sided p-values have $p(\mathbf{x}, \beta_0^*) \leq p(\mathbf{x}, \beta_0)$ when either $\beta_0^* < \beta_0 < \hat{\beta}$ or $\hat{\beta} < \beta_0 < \beta_0^*$. A two-sided p-value with this property can be interpreted as a coherent one-sided p-value in the appropriate direction. For example, if $\hat{\beta} > \beta_0$ then we can reject $H_0 : \beta \leq \beta_0$ at level p . (And for the t-test p-value, we can also reject at a level of $p/2$.)

Monotonicity (of power): Under the alternative hypothesis, power increases as the sample size increases.

Nestedness (of CIs): If $(1 - \alpha^*) > (1 - \alpha)$, then the $100(1 - \alpha^*)\%$ confidence interval, (L^*, U^*) , would contain the $100(1 - \alpha)\%$ one, (L, U) ; in other words, $L^* \leq L < U \leq U^*$.

2.4. Two-sample binomial: Failure of normal intuition

Now we turn to the two-sample binomial problem, where $X_1 \sim \text{Binomial}(n_1, \theta_1)$ and independently $X_2 \sim \text{Binomial}(n_2, \theta_2)$. Here the parameter of interest is typically one of three functions of $\theta = [\theta_1, \theta_2]$: the difference ($\beta_d = \theta_2 - \theta_1$), the ratio ($\beta_r = \theta_2/\theta_1$), or the odds ratio ($\beta_{or} = \{\theta_2(1 - \theta_1)\} / \{\theta_1(1 - \theta_2)\}$). In this

problem, the inferential methods do not necessarily follow the properties that we would expect from normal intuition. We list several examples using several different valid tests, valid confidence intervals, or triples.

Failure of Reproducibility: The uniformly most powerful unbiased (UMPU) test of $H_0 : \theta_1 \geq \theta_2$ versus $H_1 : \theta_1 < \theta_2$ is a randomized version of a one-sided Fisher’s exact test [see e.g., 37, 23]. Testing this hypothesis at the one-sided $\alpha = 0.025$ level for the data $x_1/n_1 = 1/6$ and $x_2/n_2 = 7/9$, the UMPU test rejects 70.3% of the time. So, provided they are not using the same pseudo-random number generator, there is a 41.7% chance that two researchers applying the UMPU test to those data will have different accept/reject decisions.

Associated confidence region not an interval: There are two versions of the two-sided Fisher’s exact test and the most common is the Fisher-Irwin test (default in current versions of SAS [version 9.4] and R [version 4.0.4], see Section 8 for definition). The test was designed to test $H_0 : \beta_{or} = 1$, but it can be generalized to test other null hypotheses. Consider the data $x_1/n_1 = 7/262$ and $x_2/n_2 = 30/494$ [see 19, Supplement, Section 3.1]. The two-sided p-value for testing $H_0 : \beta_{or} = 1$ is $p = 0.04996$, which rejects the null hypothesis at the $\alpha = 0.05$ level. If we slightly change the null and test $H_0 : \beta_{or} = 0.99$, we get $p = 0.05005$, and we fail to reject. But counter-intuitively, if we change the null the other way and test $H_0 : \beta_{or} = 1.01$, we also fail to reject, $p = 0.05006$. So if we create the 95% confidence region by inverting the p-value procedure, this region is not contiguous,

$$C(\mathbf{x}, 0.95) = \{\beta : \beta \in (0.177, 0.993) \text{ or } \beta \in (1.006, 1.014)\}.$$

and includes values of β_{or} both larger and smaller than 1. The cause of this behaviour is the lack of unimodality of the p-value function; see Figure 1.

Incompatible inferences: If the confidence region is not an interval, we can create a valid CI by using the interval that covers the whole confidence region. But this will not give compatible inferences with the p-value function. Returning to the Fisher’s exact test confidence region example, we can create a 95% confidence interval by “filling in the hole” as $(0.177, 1.014)$ to create the *matching* confidence interval [see Section 4.1 or Ref. 6]. In this case, the two-sided p-value rejects the null that $\beta_{or} = 1$ at the 0.05 level, but the matching 95% confidence interval includes $\beta_{or} = 1$. This issue is different from the incompatible inferences that often occur by using different methods to calculate p-values and confidence intervals, which can be quite prevalent in this application. For example, the default for R (`fisher.test` in base R, version 4.0.4) and SAS (`exact` option in Proc Freq, version 9.4) uses the Fisher-Irwin two-sided p-value, but calculates the two-sided confidence interval on β_{or} by inverting two one-sided Fisher exact p-values [see e.g., 19, 20].

Imperfect Accuracy of Coverage: Because of discreteness, the valid confidence interval must have coverage larger than the nominal level for some values of θ , in order to ensure validity for all values of θ . Remember, the

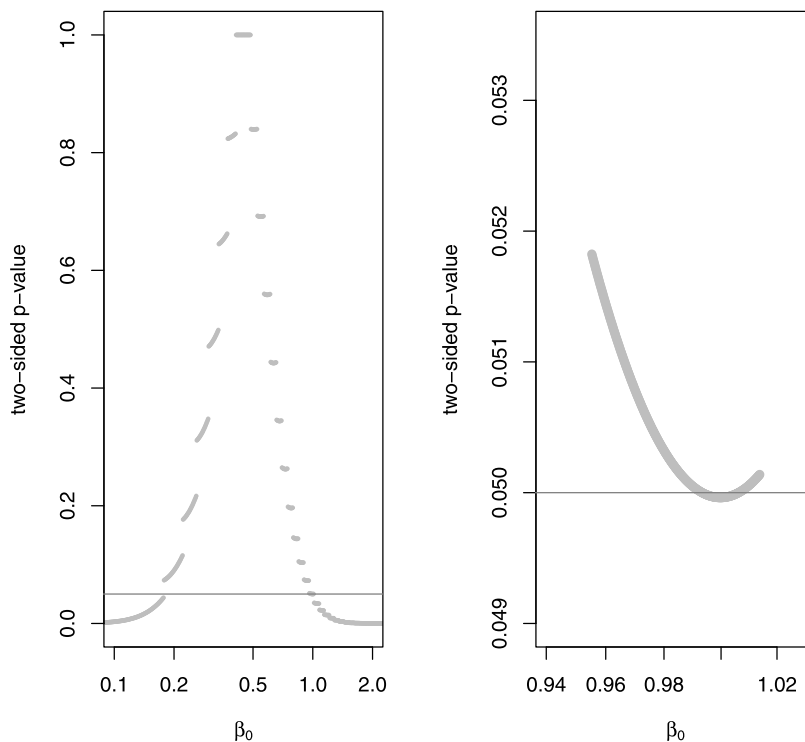


FIG 1. Two-sided Fisher's exact test (Fisher-Irwin version) p-values by β_0 for $x_1/n_1 = 7/262$ and $x_2/n_2 = 30/494$. Right panel is an enlargement of part of the left panel. Reference line is 0.05.

term “exact” is often used to mean valid (see Section 2.2), so an “exact” confidence interval may have coverage greater than the nominal level and not, as the term might imply, have coverage exactly equal to the nominal level. Section 9 discusses relaxing the requirement of validity in order to have coverage closer to the nominal level “on average”, slightly greater than nominal for some parameter values and slightly less for others.

Non-Centrality of Confidence Interval: Although central $(1 - \alpha)$ CIs for the binomial problem are important, much has been written on non-central intervals. Agresti and Min [1] showed that inverting certain two-sided tests, produces shorter confidence intervals than central ones. For the difference in proportions, this strategy often uses an unconditional exact (i.e., valid) version of a two-sided score test [see 17]. For $x_1/n_1 = 5/9$ and $x_2/n_2 = 7/7$, the difference in proportions is $\hat{\beta}_d = 0.444$ with 95% confidence interval $(0.005, 0.749)$ and the associated two-sided exact p-value for testing $\beta_d = 0$ is $p = 0.0496$. Because the 95% confidence interval is based on inverting a two-sided test, we cannot use $p/2 = 0.0248$ as a one-sided p-value to show that $\beta_d > 0$ at the 0.025 level. In fact, to ensure validity, we can only use

the two-sided p-value as an upper bound on that one-sided p-value.

Non-monotonicity of power: Continuing with the previous example ($x_1/n_1 = 5/9$ and $x_2/n_2 = 7/7$ using the unconditional exact two-sided score test), if we add one more observation to group 2 the two-sided p-value increases regardless of whether the extra observation is a failure (giving $x_2/n_2 = 7/8$ and $p = 0.172$), or success (giving $x_2/n_2 = 8/8$ and $p = 0.0510$) [this example comes from 56]. Thus, it is not surprising that the power to reject at the two-sided 0.05 level when $\theta_1 = .4$ and $\theta_2 = .9$ is higher for $n_1 = 9, n_2 = 7$ (power= 61.9%) than for $n_1 = 9, n_2 = 8$ (power=53.7%). Power non-monotonicity can also exist for common one-sided tests. Using a one-sided Fisher’s exact test (to reject $H_0 : \theta_1 \leq \theta_2$) at the 0.025 level, when $n_1 = n_2 = 4$ the only way to reject is the most extreme case, $x_1/n_1 = 4/4$ and $x_2/n_2 = 0/4$, and the power to reject is $(1 - \theta_1)^4 \theta_2^4$. When $n_1 = 4$ and $n_2 = 5$ then still the only way to reject at the 0.025 level is the most extreme case, $x_1/n_1 = 4/4$ and $x_2/n_2 = 0/5$, and the power is $(1 - \theta_1)^4 \theta_2^5$. Then, $(1 - \theta_1)^4 \theta_2^4 > (1 - \theta_1)^4 \theta_2^5$ for all $0 < \theta_1 < 1$ and $0 < \theta_2 < 1$, and the power with the larger sample size is smaller. A similar decrease in power occurs by instead adding one to the other group: $n_1 = 5$ and $n_2 = 4$.

Non-nested Confidence Intervals: Wang [57] proposed a method for constructing the smallest one-sided confidence interval for the difference of two proportions. Consider $x_1/n_1 = 2/7$ and $x_2 = 2/5$. The lower one-sided 95% interval on the difference, β_d , is $(-0.467, 1)$, but the 96% interval by the same method is $(-0.442, 1)$. See Figure 2 and Section 5.5.

Non-Coherence: For testing for non-inferiority on a difference in proportions, Chan and Zhang [10] recommend the exact unconditional test based on the score test. Röhmel [49] gives the following illustrative example: the proportion of failures on control is $x_1/n_1 = 130/248$ and on new treatment is $x_2/n_2 = 76/170$, with the failure rate slightly lower on new treatment, $\hat{\beta}_d = -0.077$. If we want to show that $H_1 : \beta_d < 0.025$ the p-value is $p = 0.0226$, but if we want to show an even less stringent margin, $H_1 : \beta_d < 0.026$ the p-value non-intuitively increases to $p = 0.0240$ (see Figure 3 and Section 5.6).

For the two-sample binomial problem, many attempts to increase power or get the smallest expected length CI result in violations of some of these “normal intuition” properties.

3. Choosing the effect measure

Choosing the effect measure is dependent on the application, so we examine a real application to discuss the issues. Coulibaly et al. [14] studied a parasite called *Mansonella perstans* that infects people in parts of Africa. The usual drugs that kill other similar parasites had not been working on killing *M. perstans*. Coulibaly et al. [14] realized that in this case there was a symbiotic bacteria, *Wolbachia*, that helped the *M. perstans* live. They suspected that if




FIG 2. Thick gray lines are lower limits for the smallest one-sided $100(1 - \alpha)\%$ confidence limits for β_d from Wang [57] for $x_1/n_1 = 2/7$ and $x_2/n_2 = 2/5$. Solid black lines show one-sided 95% limit of -0.467 , while dotted black lines show one-sided 96% limit of -0.442 .

they gave a common antibiotic, doxycycline, to kill the bacteria, it may in fact help cure the patient of *M. perstans*. Patients were randomized to the treatment group (doxycycline) or the control group (no treatment). There are issues of missing data that we ignore for simplicity. At 12 months, $x_2 = 67$ out of $n_2 = 69$ subjects who received doxycycline and $x_1 = 10$ out of $n_1 = 63$ had cleared the *M. perstans* from their blood. There are several reasonable choices for how to measure the effect: the difference in clearance rates, the ratio of clearance rates, the ratio of failure probabilities, and the odds ratio of clearance rates. Although the choice is often dominated by what is most natural to the intended audience, there are some statistical issues related to this choice.

Without loss of generality, consider effect measures that measure how much larger θ_2 is than θ_1 . The opposite effect can be measured by switching group labels. But we could also simultaneously switch group labels *and* switch the response and failure labels. If the effect remains the same after this double switching, we say that the measure has symmetry equivariance. The measures β_d and β_{or} have symmetry equivariance; however, β_r does not have it, as we demonstrate with the example. Let $\hat{\theta}_2 = 67/69 \approx 0.97$ and $\hat{\theta}_1 = 10/63 \approx 0.16$. An estimate of the rate ratio for success (cleared parasites at 12 months) is




FIG 3. One-sided exact unconditional p -value using the score statistic ordering, $p_U(\mathbf{x}, \beta_0)$. Virtual data example from Röhmel [49]: $x_1/n_1 = 130/248$ and $x_2/n_2 = 76/170$. Thick black line is the p -value function. Thin black lines depict the test of the null $H_0 : \beta \geq 0.026$ and dotted lines depict the test of the null $H_0 : \beta \geq 0.025$.

$\hat{\theta}_2/\hat{\theta}_1 \approx 6.12$. The rate ratio is often called the relative risk, but in this case the “risk” is the risk of getting cured. A different expression of the same data measures the ratio of the rates of failures (those still having detectable parasites at 12 months). Let $\hat{\theta}_{F2} = 2/69 \approx 0.03$ and $\hat{\theta}_{F1} = 53/63 \approx 0.84$. Then an estimate of the relative risk of failure is $\hat{\theta}_{F1}/\hat{\theta}_{F2} \approx 29.0$. In this latter case the control group looks about 29 times worse than the treatment group, while if we look at the rate ratios for success, the treatment group looks only about 6 times better than the control group. So how many times better treatment is than control depends on which way we measure risk. This is a violation of symmetry equivariance. Despite this, the rate ratio is often used because it is easy to understand [see e.g., 14], or because it has become the parameter of choice within a field, so that its use facilitates comparisons between studies.

The difference has symmetry equivariance. If we measured the difference in rates of disease rather than the difference in rates of cure we get exactly the negative difference as we might expect. Similar to the relative risk, the difference is often used because it is easy to understand. Additionally, the sample difference in rates is always defined, unlike the ratio which is undefined when $\hat{\theta}_1 = \hat{\theta}_2 = 0$.




FIG 4. Four simple ordering functions. Dark blue means $\hat{\theta}_2$ is much larger than $\hat{\theta}_1$ and dark red is the opposite. White means both treatments appear the same. The functions are based on $n_1 = n_2 = 8$ and using functions of the sample proportions, $\hat{\theta}_1 = x_1/n_1$ and $\hat{\theta}_2 = x_2/n_2$. The sample space is depicted by a 9×9 grid of responses, ranked by the ordering functions: difference in success proportions (upper left), odds ratio (upper right), ratio of success proportions (lower left), and ratio of failure proportions (lower right). Colors rank the functions from the highest values (dark blue) indicating larger θ_2 , to middle values (white) indicating $\theta_1 = \theta_2$, to lowest values (dark red), with black indicating no information.

Figure 4 plots the three statistics using $\hat{\theta}_2$ and $\hat{\theta}_1$ with $n_1 = n_2 = 8$. The plots go from dark blue ($\hat{\theta}_2$ is larger) to white ($\hat{\theta}_1 = \hat{\theta}_2$) to dark red ($\hat{\theta}_1$ is larger), with black denoting indeterminate. Because of the indeterminate black areas, the ordering of the sample space for the ratio and odds ratio is not straightforward (see Section 5.3). The ordering of the measures on the parameters themselves would give a continuous version of Figure 4, and the black regions would reduce to points at $(\theta_1, \theta_2) = (0, 0)$ or $(1, 1)$. The bottom panels show the lack of symmetry equivariance for β_r . Comparing the panel for β_{or} with the two different ratio panels, we see that the lower left hand corner of the β_{or} panel is similar to the lower left hand corner of $\hat{\beta}_r = \hat{\theta}_2/\hat{\theta}_1$. For small θ , $\hat{\beta}_{or}$ is a good approxima-

tion to $\hat{\beta}_r$. Similarly for both θ values close to 1, $\hat{\beta}_{or}$ is a good approximation of $(1 - \theta_1)/(1 - \theta_2)$ (right bottom panel).

The odds ratio is the most complicated of the three measures, but it has some nice properties. It is very important for the case-control design used to study rare diseases, because the odds ratio of disease given exposure is equal to the odds ratio of exposure given disease [see 8]. Also for performing regression on binary observations, logistic regression allows linear predictors to be used to model the log odds, and effects of binary covariates can be expressed as odds ratios. An advantage of the odds ratio for the two-sample binomial case is that by conditioning on the total number of successes in both groups, the probability distribution reduces to a noncentral hypergeometric distribution which is a function of β_{or} . This is discussed more in Section 6.

4. Properties of frequentist triples

4.1. Defining a matched triple

Once we choose an effect measure, we choose an appropriate *triple* (an estimator, confidence interval, and p-value function) for inferences. We will not specify the estimator except to require that it is within the confidence interval. We focus mostly on choosing the CI and p-value function. Except in Section 9, we only consider triples that are valid (i.e., the CI and p-value are both valid) and reproducible. Because we require reproducibility, the triple based on the UMP unbiased (and randomized) test is not allowed. We focus on triples where the p-value function and the confidence interval are derived from the same procedure. We call this a *matched triple*.

Here is a precise definition of a matched triple. If we start with $p(\mathbf{x}, \beta_0)$, an associated confidence region is given by equation 2.2, and the matching CI is smallest interval that contains that confidence region. In other words, if the confidence region has holes in it, then those holes are “filled in”. On the other hand, if we start with $(L, U) = C_I(\mathbf{x}, 1 - \alpha)$, then the matching p-value function is the smallest α such that β_0 is outside $C_I(\mathbf{x}, 1 - a)$ for all $a \geq \alpha$.

4.2. Implications of compatible inferences

Theorem 4.1. *Consider a valid, reproducible, and matched triple. The triple has compatible inferences*

1. *if and only if the CI is equal to the confidence region associated with the p-value, and*
2. *only if the CI is nested, and*
3. *only if the the p-value function is coherent (for one-sided p-values), or directionally coherent (for two-sided p-values).*

The formal proof of the theorem is in Appendix A. The theorem says we must have nested CIs and coherent p-values in order to have compatible inferences.

These ideas are best understood graphically. Figure 1 shows lack of directional coherence; for every β_0 there is only one p-value, and the two-sided p-value function is not unimodal (i.e., as β_0 increases, the p-value function does not increase to the global maximum, then decrease after that; see the right panel). Similarly, Figure 3 shows lack of coherence. Figure 2 shows non-nestedness; for every α there is only one lower limit, and the lower limit is not a monotonic function of the level.

4.3. Directional inferences

Typically, a researcher who finds a significant difference from the two-sided p-value suggesting that $\beta \neq \beta_0$ is almost always interested in interpreting the result in terms of whether $\beta > \beta_0$ or $\beta < \beta_0$. In other words, the two-sided hypothesis test is often treated as a three-decision rule: (1) fail to reject $\beta = \beta_0$, (2) reject $\beta = \beta_0$ and conclude $\beta > \beta_0$, or (3) reject $\beta = \beta_0$ and conclude $\beta < \beta_0$. If the two-sided p-value has directional coherence, then if we reject $H_0 : \beta = \beta_0$ at level α , we can additionally reject at level α either $H_0 : \beta \leq \beta_0$ (if $\beta_0 < \hat{\beta}$) or $H_0 : \beta \geq \beta_0$ (if $\beta_0 > \hat{\beta}$).

Consider comparing two triples that both have compatible inferences, one with a central CI, and one with a non-central CI. For the non-central triple (i.e., the one with the non-central CI) the associated two-sided hypothesis test may be slightly more powerful, but if the non-central triple is applied also to a subsequent one-sided hypothesis (as in the three decision rule), it can be quite a bit less powerful than the central one. To see this, start with a nested central CI, say (L, U) , and pair it with its matching two-sided p-value, say p_C . By Theorem 4.1, this means that whenever the $100(1 - \alpha)\%$ CI excludes β_0 then $p_C \leq \alpha$, and we can reject $H_0 : \beta = \beta_0$ at level α . After rejecting the two-sided hypothesis at level α , we can reject one of the one-sided hypotheses at level $\alpha/2$; if $\beta_0 < L$ we reject $H_0 : \beta \leq \beta_0$, while if $\beta_0 > U$ we reject $H_0 : \beta \geq \beta_0$. A non-central CI does not allow one-sided rejections at the $\alpha/2$ level. Freedman [26] discusses this issue in terms of clinical trials, and, using these arguments as well as some Bayesian motivation, [26] recommends performing two one-sided tests at the $\alpha/2$ level, which is another way of describing the use of central CI methods for three decision rules.

In summary, if we desire directional inferences, and we want to compare the power to detect a one-sided effect in a fair way (i.e., both methods bound the one-sided type I error rates of the three decision rule at the same level), then we need to compare a method with a two-sided p-value and its matching $100(1 - \alpha)\%$ non-central CI, with a pair of one-sided p-values and its matching $100(1 - 2\alpha)\%$ central CI. This means that when comparing expected lengths of CIs, if directionality of effect is important, we should compare the expected length of a $100(1 - \alpha)\%$ non-central CI with the expected length of a $100(1 - 2\alpha)\%$ central CI. Because directionality is usually important, our default recommendation is to use central confidence intervals and perform three-sided inferences as described above.

5. Methods for creating one-sided exact unconditional testing procedures

5.1. Basic procedure for defining p-values

Suppose larger values of θ are better. We want to know if treatment 2 is better than treatment 1 ($\theta_2 > \theta_1$), and by how much. Let $T(\mathbf{x})$ be a function of the data, where larger values of $T(\mathbf{x})$ indicate that treatment 2 is better than treatment 1, and $T(\mathbf{X})$ is defined for all possible values of \mathbf{X} . For example, a simple $T(\mathbf{x})$ is the difference in observed proportions (see Figure 4 upper left). For this section and the next (Section 5.2), we require that T is a function of \mathbf{x} only. Later in Section 5.5, T may depend on α , and in Section 5.6, T may depend on β_0 . Barnard [4] outlined convexity conditions which ensure that larger values of T suggest treatment 2 is better. Barnard's convexity (BC) conditions are:

$$\begin{aligned} \text{if } x_2^* > x_2 \quad \text{then } T([x_1, x_2^*]) &\geq T([x_1, x_2]) \\ &\text{and} \\ \text{if } x_1^* < x_1 \quad \text{then } T([x_1^*, x_2]) &\geq T([x_1, x_2]). \end{aligned} \tag{5.1}$$

Many choices for T satisfy the BC conditions. For example, $T(\mathbf{x}) = \hat{\theta}_2 - \hat{\theta}_1$ meets the BC conditions.

Once we have decided on the ordering function, T , we can create valid unconditional one-sided p-values: p_U for testing the null H_{U0} (defined as $H_0 : \beta \geq \beta_0$) and p_L for testing H_{L0} ($H_0 : \beta \leq \beta_0$) using

$$\begin{aligned} p_U(\mathbf{x}, \beta_0) &= \sup_{\theta: b(\theta) \geq \beta_0} P_\theta [T(\mathbf{X}) \leq T(\mathbf{x})] \\ &\text{and} \\ p_L(\mathbf{x}, \beta_0) &= \sup_{\theta: b(\theta) \leq \beta_0} P_\theta [T(\mathbf{X}) \geq T(\mathbf{x})]. \end{aligned} \tag{5.2}$$

These p-values are valid since

$$\sup_{\theta \in \Theta_0} P_\theta [p(\mathbf{X}, \beta_0) \leq p(\mathbf{x}, \beta_0)] \leq p(\mathbf{x}, \beta_0)$$

where $\Theta_0 = \{\theta : b(\theta) \geq \beta_0\}$ for p_U and $\Theta_0 = \{\theta : b(\theta) \leq \beta_0\}$ for p_L . The p-values are also 'exact' by the terminology of [48] (see equation 2.1, or Theorem 1 of [38]). Thus, any other valid p-values with ordering based on T are inadmissible (that is, they have values that are never less than the valid unconditional p-values and are greater for at least one \mathbf{x}) [38, Theorem 2].

These valid one-sided p-values can be inverted to create two $100(1 - \alpha/2)$ one-sided confidence limits using

$$U(\mathbf{x}) = \begin{cases} \sup \{\beta_0 : p_U(\mathbf{x}, \beta_0) > \alpha/2\}, & \text{if } \exists \text{ a } \beta_0 \\ & \text{with } p_U(\mathbf{x}, \beta_0) > \alpha/2 \\ \beta_{max} & \text{otherwise} \end{cases}$$

$$L(\mathbf{x}) = \begin{cases} \inf \{ \beta_0 : p_L(\mathbf{x}, \beta_0) > \alpha/2 \}, & \text{if } \exists \text{ a } \beta_0 \\ & \text{with } p_L(\mathbf{x}, \beta_0) > \alpha/2 \\ \beta_{min} & \text{otherwise} \end{cases} \quad \text{and} \quad (5.3)$$

where $(\beta_{min}, \beta_{max}) = (-1, 1)$ for β_d and $(0, \infty)$ for β_r or β_{or} . A central $100(1 - \alpha)$ confidence interval is the union of the one-sided ones, $(L(\mathbf{x}), U(\mathbf{x}))$, and a central p-value is $p_C(\mathbf{x}, \beta_0) = \min(1, 2p_L, 2p_U)$. These confidence limits are called exact unconditional [see e.g., 44] or Buehler confidence limits [see 39]. Lloyd and Kabaila [39] and Wang [57] show two results about these one-sided intervals. First, the lower and upper one-sided confidence limits retain a logical ordering analogous to Barnard's convexity conditions. Specifically, $(L, U) \in \mathcal{O}_T$, where \mathcal{O}_T is the class of valid central confidence intervals such that if $T(\mathbf{x}_1) < T(\mathbf{x}_2)$ then $L(\mathbf{x}_1) \leq L(\mathbf{x}_2)$ and $U(\mathbf{x}_1) \leq U(\mathbf{x}_2)$. Second, (L, U) calculated in this manner is the smallest confidence interval within \mathcal{O}_T . In other words, any other valid central confidence interval (L^*, U^*) in \mathcal{O}_T must have $L^*(\mathbf{x}) \leq L(\mathbf{x})$ and $U(\mathbf{x}) \leq U^*(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$.

Barnard [4] proposed the two-sided CSM test (we discuss the name later). We define the CSM test more generally using an ordering method which may be used for one or two-sided tests and confidence intervals, and we begin with the one-sided versions. Briefly, Barnard's CSM one-sided ordering starts from the most extreme point and incrementally adds more points to the order such that (1) the new point(s) and all previous points meet the BC conditions, and (2) the new point(s) have the lowest one-sided p-value among the possible new points that meet the BC conditions. Details are in Appendix B. Once the appropriate one-sided CSM ordering function, $T(\mathbf{x})$, is defined, we use the above definitions for the p-values (equation 5.2) and confidence intervals (equation 5.3). The CSM stands for convexity, symmetry, and maximum. *Convexity* refers to the BC condition that each new point must meet, and *maximum* refers to the maximization of the null hypothesis space in the definition of the p-value (see sup expression in equations (5.2)). The *symmetry* condition only applies to the two-sided version of the CSM test, but nevertheless we use "CSM" to describe all versions. The symmetry condition states that whenever a point $[x_1^*, x_2^*]$ is added to the order, one must simultaneously add $[n_1 - x_1^*, n_2 - x_2^*]$ and give it the same T value in the ordering (see Appendix B and Section 8 for more discussion of two-sided tests).

In a different paper, Barnard [3] outlined the general exact unconditional test, and those tests are sometimes referred to as "Barnard's test" [see e.g., 54, 15], but we do not use that terminology to avoid confusion with Barnard's CSM test. Röhmel and Kieser [50] discussed one-sided exact unconditional tests using Barnard's CSM p-value ordering, except with breaking more ties to get higher power, an idea discussed in the next section.

Martín Andrés, Sánchez Quevedo and Silva Mato [41] proposed a good all-purpose ordering, which is to base the ordering on the one-sided mid p-value from Fisher's exact test (see equation 9.1). We explore the power properties of this ordering in Section 11. Alternatively, the ordering can be tailored to a specific application. For example, Gabriel et al. [27] proposed an ordering to

optimize power for certain types of animal experiments where θ_1 , the parameter for the control group, is expected to be nearly 1.

5.2. Improving power by breaking ties: Refinement of ordering functions

One important way to improve the power of some unconditional exact tests based on a function T is to break any ties that exist in the ordering function. If T is an ordering function with ties, and T^* is an ordering function that gives the same ordering of T at all the untied values and additionally breaks some ties, then we say T^* is a *refinement* of T . Then the unconditional exact p-values formed with T^* are always less than or equal to those formed with T [see 51, p. 158]. Similarly, one-sided exact unconditional lower confidence limits formed using T^* are always at least as large as the ones formed using T [35, 57].

We describe one specific refinement or tie breaking algorithm for the difference in proportions next, which as far as we are aware, has not been specifically described in the literature and has not been available in software (although there are some closely related methods). We can order within each set of tied values using Wald statistics for $\hat{\beta}_d$, i.e., ordering by

$$Z(\mathbf{x}) = \frac{\hat{\beta}_d}{\sqrt{\widehat{\text{var}}_0(\hat{\beta}_d)}} = \frac{\hat{\theta}_2 - \hat{\theta}_1}{\sqrt{\hat{\theta}(1 - \hat{\theta})(1/n_1 + 1/n_2)}}$$

where $\hat{\theta} = (x_1 + x_2)/(n_1 + n_2)$. This leaves the ties for $\hat{\beta}_d = 0$, but otherwise defines points with more precision as more extreme, where extreme is further away from zero. Not all the values with $\hat{\beta}_d \neq 0$ break all the ties. For example, consider the ties at $\hat{\beta}_d = 5/8$ that happen at the \mathbf{x} values $[0, 5]$, $[1, 6]$, $[2, 7]$, and $[3, 8]$, for $n_1 = n_2 = 8$. This method still leaves tied the two pairs of points, $\{[0, 5], [3, 8]\}$ and $\{[1, 6], [2, 7]\}$. These remaining ties we argue should remain tied in order for the ordering to retain symmetry equivariance. Note that this suggested ordering is similar, but not equivalent to just ordering the entire sample space by $Z(\mathbf{x})$ [as was studied in 44].

If we break the ties in this way, then the BC conditions are still met, because only at the boundaries (where the ties are broken according to the BC conditions) do the ties occur at two points \mathbf{x}_a and \mathbf{x}_b with $x_{a1} = x_{b1}$ or $x_{a2} = x_{b2}$. All of the other ties will not have any $x_{a1} = x_{b1}$ or $x_{a2} = x_{b2}$ so they can be broken in any manner and the overall ordering function, T^* , will meet the BC conditions. This is important for computation (see Section 10). Further, the proposed T^* (tie-breaking on difference in proportions) does not depend on α or β_0 like some score test based methods (see Sections 5.5 and 5.6) so avoids problems with nesting and coherence.

5.3. Ordering functions for ratio and odds ratio

Performing exact unconditional tests on β_r or β_{or} is not straightforward. We consider β_r first since it is simpler. One problem is that $\mathbf{x} = [0, 0]$ could occur

with high probability if the true ratio was 100 or if it was 1/100 as long as both θ_1 and θ_2 were very small. So if $T(\mathbf{x})$ is designed so that larger values suggest $\theta_2 > \theta_1$, it is not clear how to define $T([0, 0])$ if our interest is in β_r .

Since $\mathbf{x} = [0, 0]$ gives no information about β_r , we must deal with $[0, 0]$ in a special way; set the p-value at $\mathbf{x} = [0, 0]$ to 1 for tests of β_r regardless of the null hypothesis. This means that $\mathbf{x} = [0, 0]$ is placed “deepest” within the null. Following equations 5.2, this implies $T([0, 0])$ can be thought of as the largest value when calculating $p_U(\mathbf{x}, \beta_0)$ and the smallest value when calculating $p_L(\mathbf{x}, \beta_0)$. A similar issue applies to the odds ratio, except in that case, the point $\mathbf{x} = [n_1, n_2]$ also has no information about β_{or} .

For clarity, we rewrite equations 5.2 applied to all three parameters. Let \mathcal{X}_I denote the set of \mathbf{X} values with information about β . Then if $\mathbf{x} \notin \mathcal{X}_I$ set $p_U(\mathbf{x}, \beta_0)$ and $p_L(\mathbf{x}, \beta_0)$ to 1, otherwise let $p_U(\mathbf{x}, \beta_0)$ be

$$\sup_{\theta: b(\theta) \geq \beta_0} P_\theta [T(\mathbf{X}) \leq T(\mathbf{x}) | \mathbf{X} \in \mathcal{X}_I] P_\theta [\mathbf{X} \in \mathcal{X}_I]$$

and analogously, let $p_L(\mathbf{x}, \beta_0)$ be

$$\sup_{\theta: b(\theta) \leq \beta_0} P_\theta [T(\mathbf{X}) \geq T(\mathbf{x}) | \mathbf{X} \in \mathcal{X}_I] P_\theta [\mathbf{X} \in \mathcal{X}_I].$$

Since we never reject when $\mathbf{x} \notin \mathcal{X}_I$, these definitions give valid p-values, and additionally when $\mathbf{x} \notin \mathcal{X}_I$ we do not need to define $T(\mathbf{x})$.

The simple ordering function by the estimate of β_r or β_{or} (even when using a tie breaking ordering similar to what was done for β_d) is not very powerful (see Section 11), and is not recommended. Typically, we order using a score function (see Section 5.6) since it gives more reasonable power.

5.4. Other improvements: *E+M* and *Berger-Boos*

Another method to apparently improve the ordering statistic for any efficacy parameter (difference, ratio, or odds ratio) is the estimated and maximized (*E+M*) p-value [38]. In this method, we replace an ordering statistic, T , with T^* , where T^* is an estimated p-value when testing H_{L0} (or the negative estimated p-value when testing H_{U0}). We estimate the p-value by plugging in $\hat{\theta}_0$ instead of taking the supremum of θ under the null, where $\hat{\theta}_0$ is the maximum likelihood estimator of $\theta \in \Theta_0$. For example, the approximation for p_L in expression 5.2 uses $\hat{p}_L(\mathbf{x}, \beta_0) = P_{\hat{\theta}_0} [T(\mathbf{X}) \leq T(\mathbf{x})]$. Then we “maximize” using $T^*(\mathbf{x}) = \hat{p}_L(\mathbf{x}, \beta_0)$ instead of T as the ordering function. That is, we calculate the exact conditional p-value using expression 5.2 by taking the supremum. Lloyd [38] studied this method and observed that when T^* (the approximate p-value) is used as the ordering statistic, the resulting exact unconditional p-value is generally smaller than the exact unconditional p-value on T . The process can be repeated (replace T^* by its approximate p-value), but the additional reduction appears to be minimal.

Berger and Boos [5] introduced a popular adjustment that tends to reduce exact unconditional p-values. Instead of taking the supremum over the entire null hypotheses parameter space, take the supremum only over C_γ , a $100(1 - \gamma)\%$ confidence set of θ restricted to be in the null space, then add γ to ensure validity. This is usually done by reexpressing the parameter space (θ_1, θ_2) as (β, ψ) , where ψ is a nuisance parameter, then defining C_γ as the intersection of $\theta \in \Theta_0$ and the set of θ values with ψ in its $100(1 - \gamma)\%$ confidence interval. A Berger-Boos version of p_U of expression 5.2, uses

$$p_{U\gamma}(\mathbf{x}, \beta_0) = \gamma + \sup_{\theta \in C_\gamma} P_\theta [T(\mathbf{X}) \geq T(\mathbf{x})].$$

This is not optimal, since we may be able to improve it by using $p_{U\gamma}(\mathbf{x}, \beta_0)$ as an ordering function. Nevertheless, it usually provides some reduction in p-values [see e.g., 38].

5.5. Ordering functions that depend on significance level

Kabaila and Lloyd [34] showed that for one-sided $100(1 - \alpha/2)\%$ exact unconditional upper confidence limits, the ordering function, T , that maximizes the asymptotic efficiency is an approximate $100(1 - \alpha/2)\%$ one-sided upper confidence limit itself. A different ordering function is used for the upper and lower limit, and for different confidence levels.

Wang [57] and Wang and Shan [58] also proposed an ordering function to give the smallest CI, and the calculation of the ordering function itself is iterative and quite involved, similar to the CSM test of Barnard [4]. The precise definition of the ordering is notationally cumbersome, but the idea is roughly as follows. Consider the lower $100(1 - \alpha/2)\%$ one-sided limit. Start from the most extreme point $\mathbf{x} = [0, n_2]$. Then add points one at a time, picking the point, \mathbf{x}_a , that gives the largest $L(\mathbf{x}_a, 1 - \alpha/2)$ and belongs to the set of closest neighboring points with the already included points, where closest neighbor is defined in terms of the BC conditions. The algorithm ensures that the lower limit function meets the BC conditions. Because each added $L(\mathbf{x})$ value is as large as possible, if the resulting ordering function T gives the finest partition (there are no ties), then any valid $100(1 - \alpha/2)\%$ one-sided lower limit that meets the BC conditions and uses T for ordering, say L^* , has $L^*(\mathbf{x}) \leq L(\mathbf{x})$ for all \mathbf{x} [see 57, 58].

The price for this optimality property is that the ordering function depends on α . Different ordering functions arise for different α , which can lead to non-nestedness (see Figure 2).

5.6. Ordering functions that depend on hypothesis space boundaries

Basing the ordering statistic on a score test can increase power over using simple Wald-type Z statistics [see 9]. Although this increased power has been shown in several simulation studies, it is not clear whether the increase is due to fewer ties for the score test, or from some other difference between the ordering statistics. A problem with the score statistic is that the induced ordering may change based on the β_0 , since score statistics use β_0 in their calculation, whereas most

other test statistics do not include β_0 in the calculation. This can produce non-coherence as was shown in Section 2.4 and Figure 3.

Although the exact unconditional p-values and confidence intervals of this section can be powerful, they are more difficult to calculate than the exact conditional ones described in the next two sections: Section 6 for p-values, and Section 7 for compatible confidence intervals.

6. One-sided conditional exact tests

Yates [61] argues that conditioning on the total number of failures is the proper strategy for this problem, and most of the discussants of the paper agreed with this (including Barnard, who first suggested the unconditional approach). One of the main reasons that others had recommended the unconditional approach is an overemphasis on the fixed significance level and the resulting power, which when used leads to more power for unconditional tests because the sample space has more values and hence is less discrete. Yates [61] argues (in his Section 9) that over reliance on the nominal significance level is not a good reason to prefer the unconditional test, and that p-values should be reported instead of accept/reject decisions. Yates [61] also argues for conditioning on the total number of events ($X_1 + X_2$), because that statistic is approximately ancillary to the effects of interest. Chernoff [11] quantifies the approximate ancillarity, showing that the absolute amount of information “is quite small unless $[\theta_1$ and $\theta_2]$ are very far apart”, and the proportion of information in the margins decreases with the sample size. Recent reviews [e.g., 40] have emphasized power arguments, and we review the choice of test from that perspective in Section 11. Historically, conditional tests have been important because of their much smaller computational burden compared to unconditional tests. The computational burden for unconditional tests has become less important, although for some applications it may be a non-trivial concern (e.g., big data applications with small sample sizes but very many covariates being tested).

For the unconditional one-sided exact method, to calculate p-values we need to take the supremum of the probability that $T(\mathbf{X})$ is at least as extreme than the observed $T(\mathbf{x})$ over the parameter space Θ_0 (see e.g., equation 5.2). This is a difficult calculation (see Section 10). An alternative method conditions on the sum $s = x_1 + x_2$, and calculates the conditional probability. The resulting conditional distribution is the extended hypergeometric distribution [32] also called Fisher’s noncentral hypergeometric distribution [25], which depends only on β_{or} . Additionally, because s is fixed, we can write the ordering function in terms of X_2 only. In fact, the only unique ordering function that makes sense and meets the BC conditions is X_2 itself (ordering on $n_1 - X_1$ will be equivalent). So this simplifies the calculations if the effect measure is β_{or} . For example, for testing $H_0 : \beta_{or} \geq \beta_0$ use

$$\begin{aligned} p_{Uc}(\mathbf{x}, \beta_0) &= \sup_{\theta \in \Theta_0} P_\theta [T(\mathbf{X}) \geq T(\mathbf{x}) | S] = \sup_{\beta_{or} : \beta_{or} \geq \beta_0} P_{\beta_{or}} [X_2 \geq x_2 | S] \\ &= P_{\beta_0} [X_2 \geq x_2 | S], \end{aligned} \quad (6.1)$$

where the last step follows because the conditional distribution is monotone in β_{or} [45]. The other conditional one-sided p-value, p_{Lc} is calculated similarly except by reversing the inequality. These conditional p-values for testing $H_0 : \beta_{or} = 1$ (or equivalently $H_0 : \theta_1 = \theta_2$) are Fisher's exact one-sided p-values. We calculate the central confidence intervals on β_{or} using equation 5.3 except using the conditional exact one-sided intervals instead of the unconditional ones.

Now consider the other measures, β_d and β_r . At the boundary of equality, the one-sided hypotheses are equivalent. For example, the following three null hypotheses give equivalent Θ_0 : (odds ratio) $H_{0U} : \beta_{or} \geq 1$, (ratio) $H_{0U} : \beta_r \geq 1$, and (difference) $H_{0U} : \beta_d \geq 0$. Analogously for the other one-sided p-value. But for boundaries not representing equality, Θ_0 changes depending on the effect measure. The simplification of the p-value calculation only works for the odds ratio. For example, for the difference in proportions (i.e., $\beta = \beta_d$) there is no simplification analogous to equation 6.1. Figure 5 shows that the exact one-sided conditional confidence limit on β_d is not efficient, because the conditional distribution depends on β_{or} . The upper $100(1 - \alpha/2)\%$ limit for β_d , say U_d , based on the upper limit for β_{or} , say U_{or} , is [see 52, Section 2]

$$U_d = \begin{cases} 0 & \text{if } U_{or} \leq 1 \\ \frac{\sqrt{U_{or}-1}}{\sqrt{U_{or}+1}} & \text{if } U_{or} > 1 \end{cases}$$

There are better ways to get confidence intervals on β_d and β_r that provide compatible inferences with the one-sided p-values with β_0 representing $\theta_1 = \theta_2$. We show these in the next section.

7. Melded confidence intervals

Fay, Proschan and Brittain [21] developed melded confidence intervals, a general method for creating confidence intervals for the two-sample case, that is closely related to the confidence distribution (CD) approach [59]. Confidence distributions are a frequentist analog to the Bayesian posterior with a non-informative prior. These melded confidence intervals give compatible inferences with the central conditional tests.

Before discussing the binomial case, we consider the normal case because it is more straightforward. Consider the difference in means between two normal samples with different variances. Let μ_a , \bar{y}_a , n_a , and s_a^2 be, respectively, the mean, the sample mean, sample size, and unbiased sample variance estimate for group a . The two $100(1 - \alpha/2)\%$ one-sided confidence intervals for the mean in group a , are

$$(L_{\mu_a}(1 - \alpha/2), \quad \infty) \quad \text{and} \\ (-\infty, \quad U_{\mu_a}(1 - \alpha/2))$$

with

$$L_{\mu_a}(1 - \alpha/2) = \bar{y}_a - F_{n_a-1}^{-1}(1 - \alpha/2) \frac{s_a}{\sqrt{n_a}},$$




FIG 5. 97.5% Confidence region based on one-sided conditional test of odds ratio (gray shaded area). Data is $x_1/n_1 = 4/12$ and $x_2/n_2 = 8/15$. Upper 97.5% exact conditional limit on β_{or} is $U = 2.664$ (dotted line) and on β_d is $U = 0.240$ (solid line). The confidence region based on the upper limit for β_d is the gray region plus the white space between the dotted and solid line. We see that because the conditional probability depends on β_{or} alone, that white space represents the lack of efficiency of basing the confidence region on β_d instead of β_{or} .

and (7.1)

$$U_{\mu_a}(1 - \alpha/2) = \bar{y}_a + F_{n_a-1}^{-1}(1 - \alpha/2) \frac{s_a}{\sqrt{n_a}},$$

where $F_{df}^{-1}(q)$ is the q th quantile of the t-distribution with df degrees of freedom. The central $100(1 - \alpha)\%$ confidence interval is the intersection of the two one-sided intervals,

$$(L_{\mu_a}(1 - \alpha/2), U_{\mu_a}(1 - \alpha/2)).$$

The confidence distribution approach is a way to re-express the confidence interval. Let A and B be two independent uniform $(0, 1)$ random variables. Let $M_{L_{\mu_a}} \sim L_{\mu_a}(A)$ and $M_{U_{\mu_a}} \sim U_{\mu_a}(B)$ be the lower and upper confidence distribution random variables for μ_a , where the randomness comes from A and B , while \bar{y}_a and $s_a/\sqrt{n_a}$ are treated as constants. From (7.1) and the probability

integral transformation, we re-express those random variables as

$$\begin{aligned} M_{L_{\mu_a}} &= \bar{y}_a - T_A \frac{s_a}{\sqrt{n_a}}, \\ \text{and} \\ M_{U_{\mu_a}} &= \bar{y}_a + T_B \frac{s_a}{\sqrt{n_a}}, \end{aligned}$$

where T_A and T_B are independent and distributed t with $n_a - 1$ degrees of freedom. Because of the symmetry of the t distribution about 0, $M_{L_{\mu_a}}$ and $M_{U_{\mu_a}}$ have the same distribution, so in this case the lower and upper confidence distributions are equivalent, and we let $M_{\mu_a} \equiv M_{L_{\mu_a}} \equiv M_{U_{\mu_a}}$ be the confidence distribution random variable associated with μ_a . In terms of the CD-RV, the $100(1 - \alpha)\%$ confidence interval for μ_a is

$$\{q(\alpha/2, M_{\mu_a}), q(1 - \alpha/2, M_{\mu_a})\}$$

where $q(a, M)$ is the a th quantile of the random variable M .

The confidence distribution approach appears to be a confusing and round-about way to express the confidence interval. The advantage comes when we want a confidence interval for $\mu_2 - \mu_1$, based on a two-sample problem with independent samples. Then we can write the $100(1 - \alpha)\%$ confidence interval for $\mu_2 - \mu_1$ as

$$(q(\alpha/2, M_{\mu_2} - M_{\mu_1}), q(1 - \alpha/2, M_{\mu_2} - M_{\mu_1})), \quad (7.2)$$

which can be estimated with Monte Carlo simulation. Expression 7.2 is equivalent to the Behrens-Fisher confidence interval, and the confidence distribution approach gives a simple way to conceptualize it [21]. The traditional approach calculates the Behrens-Fisher statistic,

$$T_{BF} = \frac{\bar{y}_2 - \bar{y}_1}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$$

and calculates its distribution, which depends on n_1, n_2 and s_1/s_2 [36].

For the binomial problem the lower and upper confidence distributions are not equal. Let $X_a \sim \text{Binomial}(n_a, \theta_a)$, for $a = 1, 2$. Let the $100(1 - \alpha)\%$ exact central confidence interval for θ_a (i.e., the Clopper-Pearson interval [12]) be

$$\{L_{\theta_a}(1 - \alpha/2), U_{\theta_a}(1 - \alpha/2)\},$$

where $L_{\theta_a}(1 - \alpha/2)$ and $U_{\theta_a}(1 - \alpha/2)$ are exact $100(1 - \alpha/2)\%$ one-sided confidence limits, for θ_a for $a = 1, 2$. The lower and upper CD random variables for group a are $W_{L_a} = L_{\theta_a}(A_{a1})$ and $W_{U_a} = U_{\theta_a}(A_{a2})$, where A_{ai} are independent uniform random variables. This gives, $W_{L_a} \sim \text{Beta}(x_a, n_a - x_a + 1)$ with expectation $x_a/(n_a + 1)$, and $W_{U_a} \sim \text{Beta}(x_a + 1, n_a - x_a)$ with expectation $(x_a + 1)/(n_a + 1)$, and using limits of parameters going to zero we define $\text{Beta}(0, n + 1)$ as a point mass at 0 and $\text{Beta}(n + 1, 0)$ as a point mass at 1. The

lower CD-RV is stochastically smaller than the upper CD-RV. In CD form, the $100(1 - \alpha)\%$ Clopper-Pearson interval is

$$\{q(\alpha/2, W_{La}), q(1 - \alpha/2, W_{Ua})\}.$$

The $100(1 - \alpha)\%$ melded confidence interval for $\theta_2 - \theta_1$ is

$$\{q(\alpha/2, W_{L2} - W_{U1}), q(1 - \alpha/2, W_{U2} - W_{L1})\},$$

where in order to be conservative for the lower limit, we use the lower CD-RVs for θ_2 but the upper CD-RV for θ_1 , and vice versa for the upper limit. We can generalize this to other functions of $\theta = [\theta_1, \theta_2]$. Let $b(\theta)$ be a monotonic function of the parameters, such that $\beta = b(\theta)$ is increasing in θ_2 and decreasing in θ_1 , within the allowable range of the parameters. For the binomial problem all three parameters (β_d , β_r and β_{or}) meet the monotonicity requirements, while for the normal two-sample problem the ratio of means (and odds ratios of means) does not meet those requirements. In general form, the $100(1 - \alpha)\%$ (two-sided) melded confidence interval is given by

$$(q\{\alpha/2, b([W_{U1}, W_{L2}])\}, q\{1 - \alpha/2, b([W_{L1}, W_{U2}])\}).$$

Fay, Proschan and Brittain [21] conjectured that if the one-sample confidence interval procedures are valid, central, and nested, and $\beta = b(\theta)$ is increasing in θ_2 for fixed θ_1 and decreasing in θ_1 for fixed θ_2 (such as β_d , β_r , and β_{or}), then the melded confidence interval is valid, nested and central. Some mathematical results, simulations in several situations, and extensive numeric calculations in the binomial case supported this conjecture. A rigorous proof of the conjecture is still needed.

Let $p_{Um}(\mathbf{x}, \beta_0)$ and $p_{Lm}(\mathbf{x}, \beta_0)$ be the one-sided melded p-values, the p-values that match with the one-sided melded confidence limits. Then for the binomial case, Fay, Proschan and Brittain [21] showed that the one-sided melded p-values equal the exact one-sided conditional p-values when testing the null with margin β_0 which implies $\theta_1 = \theta_2$. For example, for testing $H_0 : \beta_d \geq 0$, we have $p_{Um}(\mathbf{x}, 0) = p_{Uc}(\mathbf{x}, 0)$, and for testing $H_0 : \beta_r \geq 1$, we have $p_{Um}(\mathbf{x}, 1) = p_{Uc}(\mathbf{x}, 1)$. Because the melded confidence intervals are nested, by Theorem 4.1 the melded confidence intervals are compatible with the p-values from the one-sided Fisher's exact test.

The melded CIs for β_{or} are very close to the exact conditional ones, but the melded CIs for β_d are more efficient (lower are larger, and upper are smaller) than the exact conditional ones (see Figure 6).

8. Non-central confidence intervals and associated tests

Let $T_{ts}(\mathbf{x}) \equiv T_{ts}(\mathbf{x}, \alpha, \beta_0)$ be an ordering function for testing the two-sided null $H_0 : \beta = \beta_0$, with smaller values suggesting β further away from the null. We can create exact unconditional two-sided p-values using

$$p_{ts}(\mathbf{x}, \beta_0) = \sup_{\theta \in \Theta(\beta_0)} P_\theta [T_{ts}(\mathbf{X}) \leq T_{ts}(\mathbf{x})]$$

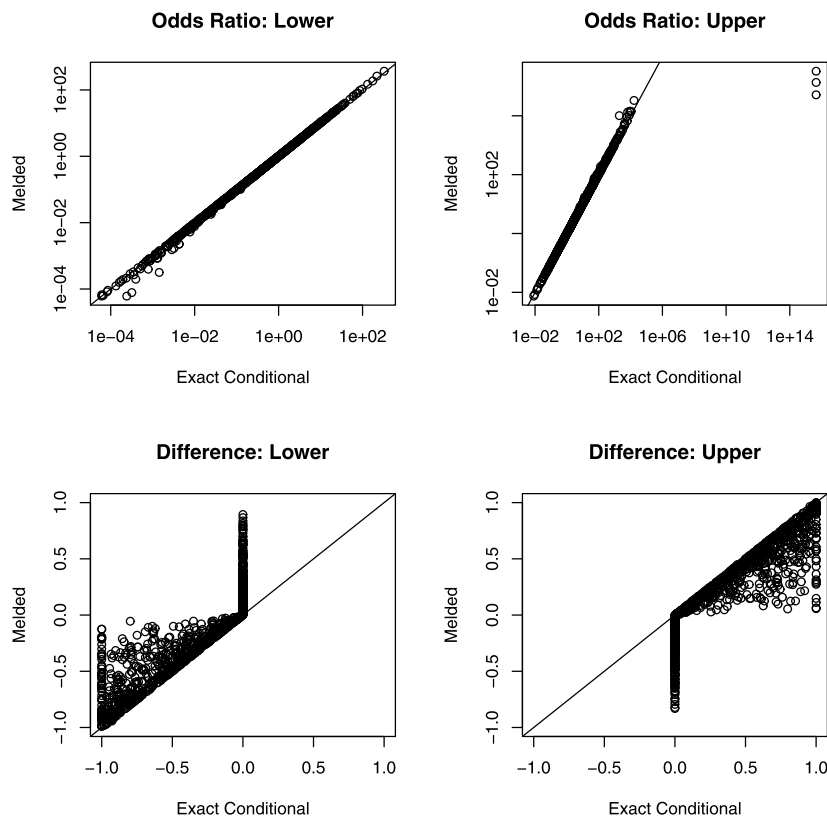


FIG 6. Lower and Upper limits associated with 95% central confidence intervals by exact conditional method and melding method. Simulated data where n_a is simulated from uniform on 1 to 100, and x_a is uniform on 0 to n_a , 1000 replications. Calculation used the `exact2x2` R package for melded confidence limits and `fisher.test` from the `stats` package for the exact conditional limits. The limits for β_{or} agree well; the identity line for the upper limits does not reach the top right corner of the graph because of some extreme data (e.g., $x_1/n_1 = 1/68$ and $x_2/n_2 = 57/61$) perhaps caused by numeric issues in the computation. In contrast, the limits for β_d show that the melded are shorter intervals (lower is larger, upper is smaller).

and exact conditional two-sided p-values using

$$p_{ts}(\mathbf{x}, \beta_0) = \sup_{\theta \in \Theta(\beta_0)} P_{\theta} [T_{ts}(\mathbf{X}) \leq T_{ts}(\mathbf{x}) | S = s].$$

which simplifies to

$$p_{ts}(\mathbf{x}, \beta_0) = P_{\beta_0} [T_{ts}(\mathbf{X}) \leq T_{ts}(\mathbf{x}) | S = s], \quad (8.1)$$

if $\beta = \beta_{or}$.

For example, consider $T_{ts}(\mathbf{x}, \beta_0) = f(\mathbf{x}, \beta_0)$, where f is the probability mass function for the extended hypergeometric distribution with parameter $\beta_{or} = \beta_0$.

The associated exact conditional p-value when $\beta_0 = 1$ is the usual Fisher's exact test, called the *Fisher-Irwin test* since it was proposed by Irwin [31] and to distinguish it from the central Fisher's exact test created by doubling the minimum of the one-sided Fisher's exact p-values. Using Fisher's exact p-values (either Fisher-Irwin or central version) as an ordering function in an unconditional exact test gives a version of Boschloo's test. Boschloo [7] showed that using the Fisher-Irwin p-values in this way is uniformly more powerful than the Fisher-Irwin test. In an analogous way, Using either one-sided or central Fisher's exact p-values as ordering functions in an unconditional test, is also uniformly more powerful than the original Fisher's exact versions of those tests [40].

Blaker [6] studied non-central confidence sets that always are subsets of the central confidence sets in one parameter distributions. To translate into this problem, we consider only the conditional distribution based on $S = s$ and $\beta = \beta_{or}$. Start with $T(\mathbf{x}) = x_2$, a one-sided ordering function for the conditional problem (see Section 6). Define

$$\gamma(\mathbf{x}, \beta) = \min \{P_\beta[X_2 \leq x_2|S = s], P_\beta[X_2 \geq x_2|S = s]\}.$$

Let the two-sided ordering function for Blaker's test be

$$T_B(\mathbf{x}, \beta) = P_\beta[\gamma(\mathbf{X}, \beta) \leq \gamma(\mathbf{x}, \beta)|S = s].$$

Blaker's test two-sided p-value is $p_B(\mathbf{x}, \beta_0)$ from equation 8.1 using $T_{ts} = T_B$, and the associated $100(1 - \alpha)\%$ confidence region is

$$C_B(\mathbf{x}, 1 - \alpha) = \{\beta : p_B(\mathbf{x}, \beta) > \alpha\}.$$

Blaker [6] showed that this gives smaller confidence sets than the central CIs. Specifically, $C_B(\mathbf{x}, 1 - \alpha) \subset C_c(\mathbf{x}, 1 - \alpha)$, where C_c is the exact conditional central CI using the one-sided ordering function $T(\mathbf{x}) = x_2$. Let the $100(1 - \alpha)\%$ matching confidence interval to p_B be the smallest interval that contains C_B .

Consider the conditional two-sided tests for $\beta_0 = 1$ when $x_1/n_1 = 8/14$ and $x_2/n_2 = 1/7$. Conditionally on $X_1 + X_2 = 9$, the support of X_2 is $\{0, 1, \dots, 7\}$. In Table 1 we give the values of $f(\mathbf{x}, 1)$, $\gamma(\mathbf{x}, 1)$, and $T_B(\mathbf{x}, 1)$. Note, $f([8, 1], 1) = f([4, 5], 1)$. Suppressing the $\beta_0 = 1$ term in the functions, the p-value for the Fisher-Irwin test is,

$$p_{FI}([8, 1]) = f([9, 0]) + f([8, 1]) + f([4, 5]) + f([3, 6]) + f([2, 7]) = 0.159,$$

for the Blaker test is

$$p_B([8, 1]) = f([9, 0]) + f([8, 1]) + f([3, 6]) + f([2, 7]) = 0.087,$$

and for the central Fisher's exact test is

$$p_c([8, 1]) = 2 * \{f([9, 0]) + f([8, 1])\} = 0.157.$$

This example was chosen to clarify the differences between the tests, but often the Fisher-Irwin and Blaker tests give the same p-values. The calculation of

the matching 95% confidence intervals involves calculating a series of p-value functions for changing β_0 , which may not be unimodal for Blaker's test or the Fisher-Irwin test (see e.g., Figure 1), so the algorithm is not simple [19]. The 95% confidence intervals are: $C_{FI}(\mathbf{x}, 0.95) = (0.005, 1.53)$, $C_B(\mathbf{x}, 0.95) = (0.005, 1.53)$, and $C_c(\mathbf{x}, 0.95) = (0.002, 1.62)$. The original (i.e., two-sided) CSM test of Barnard [4] is more difficult to calculate (see Appendix B), it gives a two-sided p-value of 0.089, and we know of no software to calculate the matching confidence interval.

TABLE 1
The hypergeometric probability mass function (f), $\gamma(\mathbf{x}, 1)$ function, and the two-sided ordering function for Blaker's exact test ($T_B(\mathbf{x}, 1)$), for $x_1/n_1 = 8/14$ and $x_2/n_2 = 1/7$.

	$x_2 = 0$	$x_2 = 1$	$x_2 = 2$	$x_2 = 3$	$x_2 = 4$	$x_2 = 5$	$x_2 = 6$	$x_2 = 7$
$f(\mathbf{x}, 1)$	0.007	0.072	0.245	0.358	0.238	0.072	0.009	0.000
$\gamma(\mathbf{x}, 1)$	0.007	0.078	0.324	0.676	0.319	0.080	0.009	0.000
$T_B(\mathbf{x}, 1)$	0.007	0.087	0.642	1.000	0.397	0.159	0.016	0.000

Agresti and Min [1] showed that to create two-sided CIs with shorter expected length, it is generally better to invert p-values from two-sided hypothesis tests that are not central. This makes sense because centrality is a restriction, and two-sided tests without that restriction will leave room for improving expected CI length. For the two-sample binomial problem, basing $T_{ts}(\mathbf{x}, \beta_0)$ on score tests gives good expected CI length; see Chan and Zhang [10] for β_d and Agresti and Min [2] for β_{or} . Despite this apparent improvement, if directional inferences are needed, then central confidence intervals are recommended (see Section 4.3).

9. Mid-p methods: Improving accuracy by sacrificing validity

The mid p-value is a modification of a p-value for discrete data. Instead of calculating the probability of observing equal or more extreme responses, the mid p-value is 0.5 times the probability of equality plus the probability of more extreme. For example, the conditional exact p-value of equation 6.1 becomes

$$p_{UC-mid}(\mathbf{x}, \beta_0) = P_{\beta_0} [X_2 > x_2 | S] + \frac{1}{2} P_{\beta_0} [X_2 = x_2 | S]. \quad (9.1)$$

Hwang and Yang [30] gave some optimality criteria for the mid-p approach applied to one parameter situations, which applies to the conditional test using β_{or} since the conditional probability is completely described by only the β_{or} parameter. They show that for one-sided or two-sided hypothesis tests, the loss based on squared error between an indicator that $\beta \in \{b(\theta) : \theta \in \Theta_0\}$ and the p-value function, and shows that for all $\beta \in \{b(\theta) : \theta \in \Theta_1\}$ (and $\beta = \beta_0$) the expected loss is less than or equal to (strictly less than) the expected loss from any randomized exact p-value function (Theorem 3.3 and 4.3 with Yang, Lee and Hwang [60]). Fellows [22] showed minimaxity under squared error and linear loss, and also showed that of all non-randomized ordered decision rules, the mid-p version is the only one that has expectation 1/2 under a point null hypothesis.

10. Computational issues

Overall, conditional p-values are much easier to calculate than unconditional ones, since they do not require taking the supremum over the null space. The melded confidence intervals allow matching CIs to conditional tests of $\theta_1 = \theta_2$, and are very quick to calculate, since they use numeric integration. There may be some precision issues in the numeric integration for extreme data sets.

The main computational speed issues apply to unconditional tests, since they require computing the supremum. Röhmel and Mansmann [51, p. 161] showed that for ordering statistics, T , that meet the BC conditions, the supremum in the p-value calculation is on the boundary between hypotheses. For example,

$$\sup_{\theta \in \Theta_0} P_\theta [T(\mathbf{X}) \geq T(\mathbf{x})] = \sup_{\theta: b(\theta) = \beta_0} P_\theta [T(\mathbf{X}) \geq T(\mathbf{x})].$$

For example, the score statistic on β_d [18] has been shown to follow the BC conditions for fixed β_0 [49]. Further, if T meets the BC conditions and does not depend on β_0 , then Theorem 3.1 of Kabaila [33] shows that the exact unconditional one-sided p-values based on T are either nonincreasing (for $p_U(\mathbf{x}, \beta_0)$) or nondecreasing (for $p_L(\mathbf{x}, \beta_0)$) in β_0 for fixed \mathbf{x} . This property means that for these p-values, the associated $100(1 - \alpha/2)$ one-sided confidence intervals can be easily calculated by finding the value β_0 where the p-value equals $\alpha/2$.

Calculation using Barnard's CSM p-value ordering can be very slow, because determining the ordering itself requires p-value calculation. Röhmel and Kieser [50] discussed one-sided exact unconditional tests using Barnard's CSM p-value ordering, except with breaking ties in a manner that does not worry about symmetry equivariance. They also do not worry about the exact ordering for very small p-values. This can speed up the calculations substantially.

Table 2 reviews different methods, their properties of centrality and compatible inferences, and approximate ranking of computational speed and power. The last column gives some software availability for the methods; it is not a comprehensive list, and only considers SAS 9.4, R (with packages), and StatXact 11.

11. Power and efficiency comparisons

A comprehensive simulation or calculation comparing different methods with respect to power or efficiency is beyond the scope of this review. Here we review a few of the best of those types of papers and add an example and some graphical calculation results to supplement the previous literature on the topic. In essence this section gives some detailed justification for the rough power/efficiency classifications listed in Table 2.

In general conditional tests (e.g., Fisher's exact tests) are less powerful than the best of the unconditional tests, because the latter tests are less discrete [40]. Martín Andrés and Silva Mato [43] provide a very comprehensive power comparison of several valid unconditional tests (including tests based on either

TABLE 2. *Valid (and Mid-p adjusted) Methods for Two-Sample Binomial Problem, and some Properties, References, and Software*

Method	Central	Compat. Infer.	Comput. Speed*	Power/Efficiency**	References	Sect.	Software***
Smallest CI	yes	no	3	1	Wang [57] (for β_d) Wang and Shan [58] (for β_r, β_{or})	5.5	Rpkg:ExactCIdiff (for β_d CI only)
Barnard's CSM	both	?	3	1	Barnard [4]	5.1	Rpkg: Exact(p-value only)
Boschloo Test	both	both	2	2	Boschloo [7]	8	Rpkg: Exact (p-value only), exact2x2
Uncond Exact Score Stat (square T)	no	no	2	2	Chan and Zhang [10] (for β_d) Agresti and Min [1] (for β_d, β_r) Agresti and Min [2] (for β_{or})	8	StatXact-11 (only β_d, β_r), SAS 9.4 (only β_d, β_r), Rpkg: exact2x2 (tsmethod="square")
Uncond Exact β Estimates with tie break	yes	yes	2	$2(\beta_d)$ $5(\beta_r)$ $5(\beta_{or})$		5.2 5.3	Rpkg: exact2x2 (tsmethod="central")
Uncond Exact Wald Stat (T^2)	no	no	2	2	Mehrotra, Chan and Berger [44]	8	StatXact-11 (only β_d) Rpkg: exact2x2 (tsmethod="square")
Uncond Exact β Estimates	yes	yes	2	$3(\beta_d)$ $5(\beta_r, \beta_{or})$	Barnard [3] Mehrotra, Chan and Berger [44]	5.1 5.3	Rpkg: exact2x2 (tsmethod="central")
Cond Exact with Fisher-Irwin Exact Test	no	no	1	3	Fisher [24] (for p-value) Fay [19] (for CI)	6	Rpkg:exact2x2
Cond Exact with Blaker Method	no	no	1	3	Blaker [6] Fay [19]	8	Rpkg: exact2x2
Cond Exact with Melded CIs	yes	yes	1	4	Fisher [24] (for p-value) Fay, Proschan and Brittain [21] (for CI)	7	Rpkg: exact2x2
Cond exact with tail approach CI (only for β_{or})	yes	yes	1	4	Agresti and Min [1] Fay [19]	6	SAS 9.4 (use double one-sided Fisher's exact p-values) StatXact-11, Rpkg: exact2x2
Adjustment	Notes					Sect.	Software
Berger-Boos	Adjustment by Berger and Boos [5] applies to unconditional exact tests and generally increases power					5.4	StatXact-11, Rpkg: exact2x2 Rpkg: Exact(p-values only)
E+M	Adjustment by Lloyd [38] applies to unconditional exact tests and generally increases power					5.4	Rpkg: exact2x2
Mid-p	Applies to any method, increases power at the cost of validity					9	Rpkg: exact2x2 SAS 9.4 (not all tests)

* Approximate computation speed: 1=fast, 2=moderate, 3=slow. ** Approximate power/efficiency: 1=higher power/shorter CI, . . . , 5=lower power/larger CI.

*** Software (not comprehensive, only considered R, SAS and StatXact): R packages available at <https://cran.r-project.org/>.

For SAS the methods are available in PROC FREQ using exact option. The value "both" denotes there could be versions with and without the property, and "?" denotes that it is not clear if the matching confidence intervals are compatible with the p-values because confidence intervals have not been studied with that test (although it is likely the method will not be compatible because it is similar to the smallest CI method).

an ordering function of the difference in sample proportions, or on some test-based ordering functions related to Fisher's exact p-value, the unpooled Z test, or Barnard's CSM test). They only considered ordering functions that do not depend on α or β_0 (since they only consider power to show $\theta_2 > \theta_1$ [i.e., with $\beta_0 = 0$ for the difference or $\beta_0 = 1$ for the ratio or odds ratio] the ordering functions automatically do not depend on β_0). Martín Andrés and Silva Mato [43] based power comparisons on expected power assuming bivariate uniformly distributed (θ_1, θ_2) . They found that Barnard's CSM test was the most powerful on average, and that ordering by either the unpooled Z statistics for the difference in means or Fisher's exact p-values (i.e., a Boschloo-type test) gave the next best power. Martín Andrés and Silva Mato [43] did not include a pooled Z test, but Mehrotra, Chan and Berger [44] did, and they showed that the pooled Z test can have much better power with unequal sample sizes. So in general we can recommend ordering by the pooled Z instead of the unpooled Z. Since Barnard's CSM test is difficult to calculate, Martín Andrés, Sánchez Quevedo and Silva Mato [42] compared many approximations to that value. They concluded that the mid-p Fisher's p-value was the best approximation to the CSM test, although it could be conservative for very small samples. Hirji, Tan and Elashoff [29] did extensive calculations finding the type I error rate for the exact conditional mid-p one-sided and two-sided (Fisher-Irwin-type) tests. They found that out of 3125 sample size and parameter situations (all with $\theta_1 = \theta_2$), typically 90-95% of both types of the mid-p p-value when used to test at a 5% significance level, had type I error rates less than or equal to 5%. Further, Lydersen, Fagerland and Laake [40] stated that the mid-p version of the Fisher-Irwin test approximates the Fisher-Boschloo test well, and the latter test (or the exact unconditional test on Pearson's chi-squared test) was their recommendation.

For confidence intervals, we focus on two papers. Chan and Zhang [10] compared unconditional confidence intervals based on estimates or tests on the difference: the difference in proportions, the unpooled Z statistic, the score statistic (which they called the δ -Projected Z statistic), and the likelihood ratio statistic. They tried all with and without the Berger and Boos [5] adjustment. They showed the score statistic with no adjustment generally gave shorter expected confidence interval length. Santner et al. [53] did a very comprehensive set of calculations for β_d confidence intervals, calculating expected coverage and confidence interval length for a 100×100 grid of values of (θ_1, θ_2) . They compared three valid methods and two approximate methods, including the unconditional method based on a two-sided score test, the unconditional method based on two one-sided score tests, and an approximate method of Coe and Tamhane [13]. The results show that of the valid methods, the unconditional method based on the two-sided score test statistic had the lowest expected length, while the central unconditional method based on two one-sided score tests had larger expected length. However, if directional inferences are important, then the proper comparison should be the former method using $100(1 - \alpha)\%$ intervals compared to the latter method using $100(1 - 2\alpha)\%$ intervals (see Section 4.3). Further, score tests may lack coherence (see Figure 3). Santner et al. [53] recommended the approximate method of Coe and Tamhane [13], which had shorter expected

length confidence intervals and gave coverage above the nominal except in less than 0.6% of the cases. Fagerland, Lydersen and Laake [17] also recommends for small samples the exact unconditional confidence intervals with the ordering function the two-sided score test statistic. Fagerland, Lydersen and Laake [17] mentions using one-sided tests if direction is important.

We now compare score tests to other tests not included in the previous simulations. Between unconditional tests applied to β_r and β_{or} , the ordering based on score tests or the ordering based on one-sided mid-p Fisher's exact p-values [41] perform much better than ordering by estimates with tie breaks as in Section 5.3. For example, with $n_1 = n_2 = 20$, $\theta_1 = 0.4$, $\theta_2 = 0.8$, and a one-sided 0.025 significance level, power is 73% for score-based or mid-p Fisher-based tests of both β_r and β_{or} but is very small for the test that orders by estimates with tie breaks (power ≈ 0 for β_r and power $\approx 1\%$ for β_{or}). Power increases slightly for the latter tests with a Berger and Boos adjustment and $\gamma = 10^{-6}$ (power is 11% for β_r and 16% for β_{or}). In contrast, for β_d in that example all three methods of ordering with or without the Berger-Boos adjustment give 73% power.

Figure 7 compares powers on the two-sided 0.05 level central tests that $\beta_d = 0$. Powers are calculated on a 99×99 grid of values of (θ_1, θ_2) . We plot the difference in powers between all pairs of three tests: two unconditional exact tests (one based on the score test for the difference in proportions, and one based on the difference in proportions with a tie break) and the conditional test (the central Fisher's exact test). We find, as expected, that unconditional tests do better, and that the simple method with a tie break does well when the sample sizes are not equal [see e.g., 44, for a different set of simulations showing a similar result for the two-sided test].

Figure 8 compares unconditional exact tests ordered by score statistics (on either $\beta_d = 0$, $\beta_{or} = 1$, or $\beta_r = 1$) compared to unconditional exact tests based on the mid p-values from the one-sided Fisher's exact test. We find that the latter tests are generally more powerful.

12. Recommendations

There are many ways to perform frequentist inferences on the two-sample binomial problem. Our extensive review focused on valid inferences and highlighted practical properties of tests. We give a few recommendations.

1. Use central confidence intervals with either a central p-value, or the minimum of the one-sided p-values. Using non-central two-sided CIs can slightly decrease expected CI length, but at a cost in terms of allowable one-sided inferences. Since we usually care about the direction of effect, non-central CIs are not routinely recommended.
2. Avoid maximizing power or minimizing the expected length of the confidence interval, because it increases computational burden and can lead to incoherent p-values and non-nested CIs.
3. For fast calculations use one-sided conditional exact tests and melded confidence intervals.




FIG 7. Comparison of powers for testing $\theta_1 = \theta_2$ using central tests at the two-sided 0.05 level. The three tests compared are “score”= unconditional exact test based on the score test of the difference in proportions, “simple TB”= unconditional exact test based on the difference in proportions using a simple tie-break (see Section 5.2), and “Fisher”= tests based on central Fisher’s exact test. For columns labeled Test 1 vs Test 2, the result is power of Test 1 minus Power of Test 2, so that positive values (pink and gray) indicate that Test 1 is more powerful. White indicates that powers are within 0.025 of each other. Colors are smoothly changing so that light blue (-0.2 to -0.3) are next to yellow and gray (0.2 to 0.30) is next to pink.

4. For more power use unconditional one-sided valid p-values and associated central CIs. For inferences on β_d , order based on the difference in sample proportions, except break ties while maintaining the BC conditions, and do not let the ordering function depend on β_0 or α . This will ensure monotonicity of p-values as a function of β_0 , allowing for relatively fast

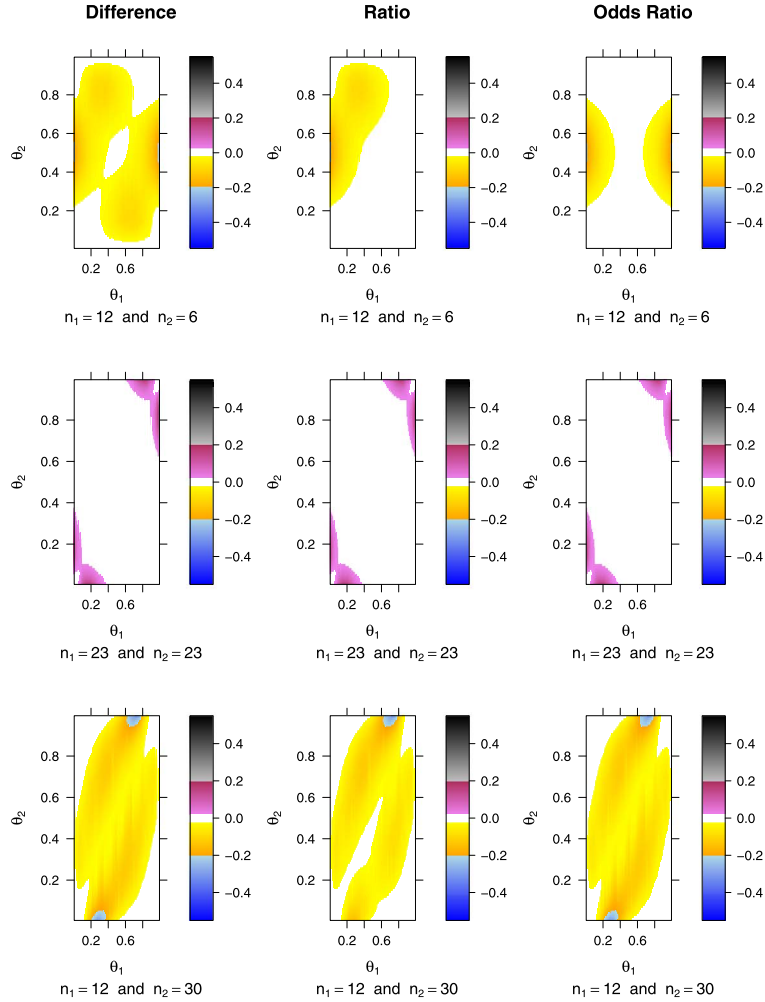


FIG 8. Power of unconditional exact score test minus power of unconditional exact test based on ordering by Fisher's exact test one-sided mid p -value. Negative values (yellow and blue) denote parameter values in which the latter test is more powerful. The unconditional exact score tests are defined based on testing either $H_0 : \beta_d = 0$ (first column), $H_0 : \beta_{or} = 1$ (second column), or $H_0 : \beta_r = 1$ (third column). White indicates that the two powers are within 0.025 of each other. Colors are smoothly changing so that light blue (-0.2 to -0.3) are next to yellow and gray (0.2 to 0.30) is next to pink. Additional calculations with $n_1 = 12, n_2 = 12$ showed nearly equal powers (all white) for all three columns and are not plotted.

calculations, while preserving coherence and nestedness. For inferences on β_r and β_{or} , using the simple function with a tie breaking ordering has much smaller power than the score method or ordering based on one-sided mid- p Fisher's exact p -values. The score method causes incoherence or non-nestedness, while the mid- p Fisher p -value ordering does not. Because the

latter method only uses the mid p-values for ordering within the exact unconditional test framework, the resulting p-values are valid. Further, for inferences on β_d , the mid-p ordering meets the BC conditions and is relatively fast to calculate.

5. If validity is not vital, then the mid-p conditional tests are a good approximation to the more powerful of the unconditional exact ones. Additionally, with a large proportion of situations with $\theta_1 = \theta_2$, the mid-p conditional tests still have type I error rates less than the nominal value.

Appendix A: Proof of Theorem 4.1

Proof of statement 1 :

(Compatible Inferences) \Rightarrow ($C_I = C$): If the confidence region associated with a p-value is not an interval, then there must be an α and β_0 such that $p(\mathbf{x}, \beta_0) \leq \alpha$ and $\beta_0 \in C_I(\mathbf{x}, 1 - \alpha)$, which contradicts the compatible inferences, therefore $C_I(\mathbf{x}, 1 - \alpha) = C(\mathbf{x}, 1 - \alpha)$.

($C_I = C$) \Rightarrow (Compatible Inferences): If the confidence region associated with the p-value is the matching confidence interval, then the inferences are compatible by definition (equation 2.2).

Proof of statement 2, (Compatible Inferences) \Rightarrow (Nested CI): We show the contrapositive. If a method has non-nested CIs, then there exists some $\alpha_1 < \alpha_2$ and some β_0 such that $\beta_0 \notin C_I(\mathbf{x}, 1 - \alpha_1)$ and $\beta_0 \in C_I(\mathbf{x}, 1 - \alpha_2)$. If the method had compatible inferences, then $p(\mathbf{x}, \beta_0) \equiv p \leq \alpha_1$ and $p > \alpha_2$. This leads to the contradiction, $p \leq \alpha_1 < \alpha_2 < p$, so the method must not have compatible inferences, and we have proven the result.

Proof of statement 3, (Compatible Inferences) \Rightarrow (Coherence): From statement 2, compatible inferences imply nested CIs. For one-sided p-values, compatible inferences with nested CIs imply that the p-values are non-decreasing as the null space expands (e.g., β_0 gets larger when $H_0 : \beta \leq \beta_0$), and hence are coherent by definition. For two-sided p-values, because of compatible inferences and nested CIs, the p-values are increasing (i.e., non-decreasing) as $1 - \alpha$ decreases. This is directional coherence by definition.

Appendix B: Barnard's CSM ordering

Because Barnard [4] defined his CSM ordering as a two-sided ordering, there may be more than one way to generalize the idea to a one-sided ordering. We present two one-sided orderings here [see 55, for alternative algorithmic details]. Consider first the bottom-up CSM ordering, where we start with the point with the lowest value of $T(\mathbf{x})$, which is $[x_1, x_2] = [n_1, 0]$, and make that the first rejection region, say R_1 , and let $T([n_1, 0]) = 1$. Then repeat the following algorithm to create the j th rejection region (for $j = 2, \dots$) until all points have been ordered:

1. Let Q_j be the set of points such that when each individual point is added to R_{j-1} , the resulting set meets the BC condition. Let the elements of Q_j be $\{\mathbf{x}_{j1}, \dots, \mathbf{x}_{jm_j}\}$. For example, when $j = 1$, then $Q_2 = \{[n_1, 1], [n_1 - 1, 0]\}$, with $\mathbf{x}_{21} = [n_1, 1]$ and $\mathbf{x}_{22} = [n_1 - 1, 0]$.
2. Calculate $p_U(\mathbf{x}_{jk})$ for each member of Q_j , where

$$p_U(\mathbf{x}_{jk}) = \sup_{\theta: \theta_2 = \theta_1} P_\theta [T(\mathbf{X}) \leq T(\mathbf{x}_{jk})]$$

and $T(\mathbf{x}) = \infty$ for all points not yet added to the rejection region (i.e., not in $\cup_{h=1}^{j-1} R_h$), and $T(\mathbf{x}_{jk'}) = j$ if $k' = k$ and $T(\mathbf{x}_{jk'}) = \infty$ if $k' \neq k$. Note, because of the BC conditions (see Section 10), this is equivalent to p_U as defined in (5.2) when $H_0 : \theta_2 \geq \theta_1$ (e.g., $b(\theta) = \theta_2 - \theta_1$ and $\beta_0 = 0$).

3. Define R_j as R_{j-1} combined with the point with the lowest value of $p_U(\mathbf{x}_{jk})$, and if there are ties include all tied points, and define the associated T function for all included points as j .

The top-down CSM ordering is analogous, starts from the highest value of $T(\mathbf{x})$ (i.e., $[x_1, x_2] = [0, n_2]$), and uses the other one-sided p-value function, p_L . It is not obvious whether the bottom-up and top-down CSM orderings are equivalent or not.

Barnard's original two-sided CSM ordering is similar, except whenever a point $[x_1^*, x_2^*]$ is included in R_j , its symmetric point $[n_1 - x_1^*, n_2 - x_2^*]$ is also included.

Acknowledgments

The authors thank Erica Brittain and anonymous reviewers for comments that improved the paper.

References

- [1] AGRESTI, A. and MIN, Y. (2001). On small-sample confidence intervals for parameters in discrete distributions. *Biometrics* **57** 963–971. [MR1863460](#)
- [2] AGRESTI, A. and MIN, Y. (2002). Unconditional small-sample confidence intervals for the odds ratio. *Biostatistics* **3** 379–386.
- [3] BARNARD, G. (1945). A new test for 2×2 tables. *Nature* **156** 177. [MR0013274](#)
- [4] BARNARD, G. (1947). Significance tests for 2×2 tables. *Biometrika* **34** 123–138. [MR0019285](#)
- [5] BERGER, R. L. and BOOS, D. D. (1994). P values maximized over a confidence set for the nuisance parameter. *Journal of the American Statistical Association* **89** 1012–1016. [MR1294746](#)
- [6] BLAKER, H. (2000). Confidence curves and improved exact confidence intervals for discrete distributions. *Canadian Journal of Statistics* **28** 783–798. [MR1821434](#)

- [7] BOSCHLOO, R. (1970). Raised conditional level of significance for the 2×2 -table when testing the equality of two probabilities. *Statistica Neerlandica* **24** 1–9. [MR0264827](#)
- [8] BRESLOW, N. E. (1996). Statistics in epidemiology: the case-control study. *Journal of the American Statistical Association* **91** 14–28. [MR1394064](#)
- [9] CHAN, I. S. (2003). Proving non-inferiority or equivalence of two treatments with dichotomous endpoints using exact methods. *Statistical methods in medical research* **12** 37–58. [MR1977234](#)
- [10] CHAN, I. S. and ZHANG, Z. (1999). Test-based exact confidence intervals for the difference of two binomial proportions. *Biometrics* **55** 1202–1209.
- [11] CHERNOFF, H. (2004). Information, for testing the equality of two probabilities, from the margins of the 2×2 table. *Journal of statistical planning and inference* **121** 209–214. [MR2038817](#)
- [12] CLOPPER, C. and PEARSON, E. S. (1934). The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika* **26** 404–413.
- [13] COE, P. R. and TAMHANE, A. C. (1993). Small sample confidence intervals for the difference, ratio and odds ratio of two success probabilities. *Communications in Statistics-Simulation and Computation* **22** 925–938.
- [14] COULIBALY, Y. I., DEMBELE, B., DIALLO, A. A., LIPNER, E. M., DOUMBIA, S. S., COULIBALY, S. Y., KONATE, S., DIALLO, D. A., YALCOUYE, D., KUBOFCHIK, J. et al. (2009). A randomized trial of doxycycline for *Mansonella perstans* infection. *New England Journal of Medicine* **361** 1448–1458.
- [15] CYTEL (2010). StatXact 9.
- [16] DUDEWICZ, E. J., MA, Y., MAI, E. S. and SU, H. (2007). Exact solutions to the Behrens–Fisher Problem: asymptotically optimal and finite sample efficient choice among. *Journal of statistical planning and inference* **137** 1584–1605. [MR2339261](#)
- [17] FAGERLAND, M. W., LYDERSEN, S. and LAAKE, P. (2015). Recommended confidence intervals for two independent binomial proportions. *Statistical methods in medical research* **24** 224–254. [MR3336292](#)
- [18] FARRINGTON, C. P. and MANNING, G. (1990). Test statistics and sample size formulae for comparative binomial trials with null hypothesis of non-zero risk difference or non-unity relative risk. *Statistics in medicine* **9** 1447–1454.
- [19] FAY, M. P. (2010a). Confidence intervals that match Fisher’s exact or Blaker’s exact tests. *Biostatistics* **11** 373–374.
- [20] FAY, M. P. (2010b). Two-sided exact tests and matching confidence intervals for discrete data. *R journal* **2** 53–58.
- [21] FAY, M. P., PROSCHAN, M. A. and BRITAIN, E. (2015). Combining one-sample confidence procedures for inference in the two-sample case. *Biometrics* **71** 146–156. [MR3335359](#)
- [22] FELLOWS, I. (2010). The minimaxity of the mid p-value under linear and squared loss functions. *Communications in Statistics-Theory and Methods* **40** 244–254. [MR2765853](#)
- [23] FINNER, H. and STRASSBURGER, K. (2001). UMP (U)-Tests for a Binomial

- Parameter: A Paradox. *Biometrical journal* **43** 667–675. [MR1862604](#)
- [24] FISHER, R. A. (1934). *Statistical Methods for Research Workers, 5th edition*. Oliver and Boyd, Edinburgh.
- [25] FOG, A. (2008). Sampling methods for Wallenius' and Fisher's noncentral hypergeometric distributions. *Communications in Statistics—Simulation and Computation*® **37** 241–257. [MR2422884](#)
- [26] FREEDMAN, L. S. (2008). An analysis of the controversy over classical one-sided tests. *Clinical Trials* **5** 635–640.
- [27] GABRIEL, E. E., NASON, M., FAY, M. P. and FOLLMANN, D. A. (2018). A boundary-optimized rejection region test for the two-sample binomial problem. *Statistics in medicine* **37** 1047–1058. [MR3777958](#)
- [28] HIRJI, K. F. (2006). *Exact Analysis of Discrete Data*. Chapman and Hall/CRC, New York. [MR2193238](#)
- [29] HIRJI, K. F., TAN, S.-J. and ELASHOFF, R. M. (1991). A quasi-exact test for comparing two binomial proportions. *Statistics in Medicine* **10** 1137–1153.
- [30] HWANG, J. G. and YANG, M.-C. (2001). An optimality theory for mid p-values in 2×2 contingency tables. *Statistica Sinica* 807–826. [MR1863164](#)
- [31] IRWIN, J. (1935). Tests of significance for differences between percentages based on small numbers. *Metron* **12** 84–94.
- [32] JOHNSON, N. L., KEMP, A. W. and KOTZ, S. (2005). *Univariate Discrete Distributions, 3rd edition*. John Wiley & Sons, New York. [MR2163227](#)
- [33] KABAILA, P. (2005). Computation of exact confidence limits from discrete data. *Computational Statistics* **20** 401–414. [MR2236616](#)
- [34] KABAILA, P. and LLOYD, C. J. (2003). The efficiency of Buehler confidence limits. *Statistics & probability letters* **65** 21–28. [MR2012621](#)
- [35] KABAILA, P. and LLOYD, C. J. (2006). Improved Buehler limits based on refined designated statistics. *Journal of statistical planning and inference* **136** 3145–3155. [MR2256221](#)
- [36] KIM, S.-H. and COHEN, A. S. (1996). Table of percentage points of the Behrens-Fisher distribution. *Journal of Statistical Computation and Simulation* **55** 181–187.
- [37] LEHMANN, E. L. and ROMANO, J. P. (2005). *Testing Statistical Hypotheses, third edition*. Springer. [MR2135927](#)
- [38] LLOYD, C. J. (2008). Exact p-values for discrete models obtained by estimation and maximization. *Australian & New Zealand Journal of Statistics* **50** 329–345. [MR2474195](#)
- [39] LLOYD, C. J. and KABAILA, P. (2003). On the optimality and limitations of Buehler bounds. *Australian & New Zealand Journal of Statistics* **45** 167–174. [MR1983349](#)
- [40] LYDERSEN, S., FAGERLAND, M. W. and LAAKE, P. (2009). Recommended tests for association in 2×2 tables. *Statistics in medicine* **28** 1159–1175. [MR2662203](#)
- [41] MARTÍN ANDRÉS, A., SÁNCHEZ QUEVEDO, M. and SILVA MATO, A. (1998). Fisher's mid p-value arrangement in 2×2 comparative trials. *Computational statistics & data analysis* **29** 107–115.

- [42] MARTÍN ANDRÉS, A., SÁNCHEZ QUEVEDO, M. and SILVA MATO, A. (2002). Asymptotical tests in 2×2 comparative trials:(unconditional approach). *Computational statistics & data analysis* **40** 339–354. [MR1924014](#)
- [43] MARTÍN ANDRÉS, A. and SILVA MATO, A. (1994). Choosing the optimal unconditioned test for comparing two independent proportions. *Computational statistics & data analysis* **17** 555–574.
- [44] MEHROTRA, D. V., CHAN, I. S. and BERGER, R. L. (2003). A cautionary note on exact unconditional inference for a difference between two independent binomial proportions. *Biometrics* **59** 441–450. [MR1982586](#)
- [45] MEHTA, C. R., PATEL, N. R. and GRAY, R. (1985). Computing an exact confidence interval for the common odds ratio in several 2×2 contingency tables. *Journal of the American Statistical Association* **80** 969–973. [MR0819602](#)
- [46] NEWCOMBE, R. G. (1998). Interval estimation for the difference between independent proportions: comparison of eleven methods. *Statistics in medicine* **17** 873–890.
- [47] NEWCOMBE, R. G. (2013). *Confidence intervals for proportions and related measures of effect size*. CRC Press, Boca Raton, FL. [MR2986377](#)
- [48] RIPAMONTI, E., LLOYD, C. and QUATTO, P. (2017). Contemporary frequentist views of the 2×2 binomial trial. *Statistical Science* **32** 600–615. [MR3730524](#)
- [49] RÖHMEL, J. (2005). Problems with Existing Procedures to Calculate Exact Unconditional P-Values for Non-Inferiority/Superiority and Confidence Intervals for Two Binomials and How to Resolve Them. *Biometrical Journal* **47** 37–47. [MR2135888](#)
- [50] RÖHMEL, J. and KIESER, M. (2013). Investigations on non-inferiority—the Food and Drug Administration draft guidance on treatments for nosocomial pneumonia as a case for exact tests for binomial proportions. *Statistics in medicine* **32** 2335–2348. [MR3067387](#)
- [51] RÖHMEL, J. and MANSMANN, U. (1999). Unconditional non-asymptotic one-sided tests for independent binomial proportions when the interest lies in showing non-inferiority and/or superiority. *Biometrical Journal* **41** 149–170. [MR1693980](#)
- [52] SANTNER, T. J. and SNELL, M. K. (1980). Small-sample confidence intervals for p $1-p$ 2 and p $1/p$ 2 in 2×2 contingency tables. *Journal of the American Statistical Association* **75** 386–394. [MR0577370](#)
- [53] SANTNER, T. J., PRADHAN, V., SENCHAUDHURI, P., MEHTA, C. R. and TAMHANE, A. (2007). Small-sample comparisons of confidence intervals for the difference of two independent binomial proportions. *Computational Statistics & Data Analysis* **51** 5791–5799. [MR2407677](#)
- [54] SAS (2012). SAS/STAT 12.1 User’s Guide, Proc Freq.
- [55] SILVIA MATO, A. and MARTÍN ANDRÉS, A. (1997). Simplifying the calculation of the P-value for Barnard’s test and its derivatives. *Statistics and Computing* **7** 137–143.
- [56] VOS, P. W. and HUDSON, S. (2008). Problems with Binomial Two-sided Tests and the Associated Confidence Intervals. *Australian & New Zealand*

- Journal of Statistics* **50** 81–89. [MR2414657](#)
- [57] WANG, W. (2010). On construction of the smallest one-sided confidence interval for the difference of two proportions. *The Annals of Statistics* **38** 1227–1243. [MR2604711](#)
- [58] WANG, W. and SHAN, G. (2015). Exact confidence intervals for the relative risk and the odds ratio. *Biometrics* **71** 985–995. [MR3436724](#)
- [59] XIE, M.-G. and SINGH, K. (2013). Confidence Distribution, the Frequentist Distribution Estimator of a Parameter: A Review (with discussion). *International Statistical Review* **81** 3–77. [MR3047496](#)
- [60] YANG, M.-C., LEE, D.-W. and HWANG, J. G. (2004). The equivalence of the mid p-value and the expected p-value for testing equality of two balanced Binomial proportions. *Journal of statistical planning and inference* **126** 273–280. [MR2090697](#)
- [61] YATES, F. (1984). Tests of significance for 2×2 contingency tables. *Journal of the Royal Statistical Society. Series A (General)* **147** 426–463. [MR0769998](#)