

Performance of Test Supermartingale Confidence Intervals for the Success Probability of Bernoulli Trials

Peter Wills¹, Emanuel Knill^{2,3}, Kevin Coakley² and Yanbao Zhang^{4,5}

¹ Department of Applied Mathematics, University of Colorado Boulder, Boulder, CO 80309, USA

² National Institute of Standards and Technology, Boulder, CO 80305, USA

³ Center for Theory of Quantum Matter, University of Colorado Boulder, Boulder, CO 80309, USA

⁴ Institute for Quantum Computing and Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada

⁵ NTT Basic Research Laboratories, NTT Corporation, 3-1 Morinosato-Wakamiya, Atsugi, Kanagawa 243-0198, Japan

peter.wills@colorado.edu
emanuel.knill@nist.gov
kevin.coakley@nist.gov
yanbaoz@gmail.com

Given a composite null hypothesis \mathcal{H}_0 , test supermartingales are non-negative supermartingales with respect to \mathcal{H}_0 with an initial value of 1. Large values of test supermartingales provide evidence against \mathcal{H}_0 . As a result, test supermartingales are an effective tool for rejecting \mathcal{H}_0 , particularly when the p -values obtained are very small and serve as certificates against the null hypothesis. Examples include the rejection of local realism as an explanation of Bell test experiments in the foundations of physics and the certification of entanglement in quantum information science. Test supermartingales have the advantage of being adaptable during an experiment and allowing for arbitrary stopping rules. By inversion of acceptance regions, they can also be used to determine confidence sets. We used an example to compare the performance of test supermartingales for computing p -values and confidence intervals to Chernoff-Hoeffding bounds and the “exact” p -value. The example is the problem of inferring the probability of success in a sequence of Bernoulli trials. There is a cost in using a technique that has no restriction on stopping rules, and, for a particular test supermartingale, our study quantifies this cost.

Key words: asymptotics; Bernoulli trials; Chernoff-Hoeffding bounds; confidence intervals; hypothesis tests; large deviations; p -values; test supermartingales.

Accepted: September 17, 2019

Published: February 5, 2020

<https://doi.org/10.6028/jres.125.003>

Contents

1	Introduction	2
2	Basic Concepts	4
3	Bernoulli Hypothesis Tests	6
4	Comparison of p-Values	9
5	Comparison of Confidence Intervals	12
6	Conclusion	21
7	Appendix	21
	7.1 Preliminaries	21
	7.2 Closed-Form Expression for P_{PBR}	22
	7.3 Log- p -Value Approximations	22
	7.4 Asymptotic Normality of the $\log(p)$ -Values and Their Differences	26
	7.5 Confidence Interval Endpoints	27
8	References	34

1. Introduction

Experiments in physics require very high confidence to justify claims of discovery or to unambiguously exclude alternative explanations [1]. Particularly striking examples in the foundations of physics are experiments to demonstrate that theories based on local hidden variables, called local realist (LR) theories, cannot explain the statistics observed in quantum experiments called Bell tests. See Ref. [2] for a review and Refs. [3–6] for the most definitive experiments to date. Successful Bell tests imply the presence of some randomness in the observed statistics. As a result, one of the most notable applications of Bell tests is to randomness generation [7]. In this application, it is necessary to certify the randomness generated, and these certificates are equivalent to extremely small significance levels in an appropriately formulated hypothesis test. In general, such extreme significance levels are frequently required in protocols for communication or computation to ensure performance.

Bell tests consist of a sequence of “trials”, each of which gives a result M_i . LR models restrict the statistics of the M_i and therefore constitute a composite null hypothesis to be rejected. Traditionally, data has been analyzed by estimating the value of a Bell function \hat{B} and its standard error $\hat{\sigma}$ from the collective result statistics (see [8, 9]). Under the null hypothesis, \hat{B} is expected to be negative, so a large value of \hat{B} compared to $\hat{\sigma}$ is considered to be strong evidence against the null hypothesis. This method suffers from several problems, including the failure of the Gaussian approximation in the extreme tails and the fact that the trials are observably not independent and identically distributed (i.i.d.) [9].

For the earliest discussion of martingales for analysis of Bell test experiments, see Refs. [10, 11]. In Ref. [9] a method was introduced that can give rigorous p -value bounds against LR. These p -value bounds are memory-robust, that is, without any assumptions on dependence of trial statistics on earlier trials. The method can be seen as an application of test supermartingales as defined in Ref. [12]. Test supermartingales

were first considered, and many of their basic properties were proved, by Ville [13] in the same work that introduced the notion of martingales. The method involves constructing a non-negative stochastic process V_i determined by $(M_j)_{j \leq i}$ such that the initial value is $V_0 = 1$ and, under LR models, the expectations conditional on all past events are non-increasing. As explained further below, the final value of V_i in a sequence of n trials has expectation bounded by 1, so its inverse $p = 1/V_n$ is a p -value bound according to Markov's inequality. A large observed value of such a test supermartingale thus provides evidence against LR models. Refs. [9, 14] give methods to construct V_i that achieve asymptotically optimal gain rate $\mathbb{E}(-\log(p)/n)$ for i.i.d. trials, where $\mathbb{E}(\dots)$ is the expectation functional. This is typically an improvement over other valid memory-robust Bell tests. Additional benefits are that V_i can be constructed adaptively based on the observed statistics, and the p -value bounds remain valid even if the experiment is stopped based on the current value of V_i . These techniques were successfully applied to experimental data from a Bell test with photons where other methods fail [15].

Although the terminology is apparently relatively recent, test supermartingales have traditionally played a major theoretical role. Carefully constructed test supermartingales contribute to the asymptotic analysis of distributions and proofs of large deviation bounds. They can be constructed for any convex-closed null hypothesis viewed as a set of distributions, so they can be used for memory- and stopping-robust adaptive hypothesis tests in some generality. The application to Bell tests shows that at least in a regime where high significance results are required, test supermartingales can perform well or better than other methods. Here we compare the performance of test supermartingales directly to (1) the standard large deviation bounds based on the Chernoff-Hoeffding inequality [16, 17], and (2) "exact" p -value calculations. Our comparison is for a case where all calculations can be performed efficiently, namely for testing the success probability in Bernoulli trials. The three p -value bounds thus obtained have asymptotically optimal gain rates. Not surprisingly, for any given experiment, test supermartingales yield systematically worse p -value bounds, but the difference is much smaller than the experiment-to-experiment variation. This effect can be viewed as the cost of robustness against arbitrary stopping rules. For ease of calculation, we do not use an optimal test supermartingale construction, but we expect similar results no matter which test supermartingale is used.

Any hypothesis test parametrized by ϕ can be used to construct confidence regions for ϕ by acceptance region inversion (see Ref. [18], Sec. 7.1.2). Motivated by this observation, we consider the use of test supermartingales for determining confidence regions. We expect that they perform well in the high-confidence regime, with an increase in region size associated with robustness against stopping rules. We therefore compared the methods mentioned above for determining confidence intervals for the success probability in Bernoulli trials. After normalizing the difference between the interval endpoints and the success probability by the standard deviation, which is $O(1/\sqrt{n})$, we find that while large deviation bounds and exact regions differ by a constant at fixed confidence levels, the test supermartingale's normalized endpoint deviation is $\Omega(\sqrt{\log(n)})$ instead of $O(1)$. This effect was noted in Ref. [12] and partially reflects a suboptimal choice of supermartingale. To maintain robustness against stopping rule, one expects $\Omega(\sqrt{\log \log(n)})$ according to the law of the iterated logarithm. However, we note that if the number of trials n is fixed in advance, the normalized endpoint deviation can be reduced to $O(1)$ with an adaptive test supermartingale. So although the increase in confidence region necessitated by stopping rule robustness is not so large for reasonably sized n , when n is known ahead of time it can, in principle, be avoided without losing the ability to adapt the test supermartingale on the fly during the experiment in non-i.i.d. situations.

The remainder of the paper is structured as follows. We establish the notation to be used and define the basic concepts in Sec. 2. Here we also explain how adaptivity can help reject hypotheses for stochastic processes. We introduce the three methods to be applied to Bernoulli trials in Sec. 3. Here we also establish the basic monotonicity properties and relationships of the three p -value bounds obtained. In Sec. 4 we determine the behavior of the p -value bounds in detail, including their asymptotic behavior. In Sec. 5 we

introduce the confidence intervals obtained by acceptance region inversion. We focus on one-sided intervals determined by lower bounds but note that the results apply to two-sided intervals. The observations in Secs. 4 and 5 are based on theorems whose proofs can be found in the Appendix. While many of the observations in these sections can ignore asymptotically small terms, the results in the Appendix uncompromisingly determine interval bounds for all relevant expressions, with explicit constants. Concluding remarks can be found in Sec. 6.

2. Basic Concepts

We use the usual conventions for random variables (RVs) and their values. RVs are denoted by capital letters such as X, Y, \dots and their values by the corresponding lower case letters x, y, \dots . All our RVs are finite valued. Probabilities and expectations are denoted by $\mathbb{P}(\dots)$ and $\mathbb{E}(\dots)$, respectively. For a formula ϕ , the expression $\{\phi\}$ refers to the event where the formula is true. The notation $\mu(X)$ refers to the distribution of X induced on its space of values. We use the usual conventions for conditional probabilities and expectations. Also, $\mu(X|\phi)$ denotes the probability distribution induced by X conditional on the event $\{\phi\}$.

We consider stochastic sequences of RVs such as $\mathbf{X} = (X_i)_{i=1}^n$ and $\mathbf{X}_{\leq k} = (X_i)_{i=1}^k$. We think of the X_i as the outcomes from a sequence of trials. For our study, we consider $\mathbf{B} = (B_i)_{i=1}^n$, where the B_i are $\{0, 1\}$ -valued RVs. The standard $\{0, 1\}$ -valued RV with parameter θ is the Bernoulli RV B satisfying $\mathbb{E}(B) = \theta$. The parameter θ is also referred to as the success probability. We denote the distribution of B by ν_θ . We define $S_k = \sum_{i=1}^k B_i$ and $\hat{\Theta}_k = S_k/k$. We extend the RV conventions to the Greek letter $\hat{\Theta}_k$. That is, $\hat{\theta}_k = s_k/k = \sum_{i=1}^k b_i/k$ is the value of the RV $\hat{\Theta}_k$ determined by the values b_i of B_i . We may omit subscripts on statistics such as S_n and $\hat{\Theta}_n$ when they are based on the full set of n samples. Some expressions involving $\hat{\Theta}_n$ require that $n\hat{\Theta}_n$ is an integer, which is assured by the definition.

A null hypothesis for X is equivalent to a set \mathcal{H}_0 of distributions of X , which we refer to as the “null”. For our study of Bernoulli RVs, we consider the nulls

$$\mathcal{B}_\varphi = \{\nu_\theta | \theta \leq \varphi\} \tag{1}$$

parametrized by $0 \leq \varphi \leq 1$. This is the set of distributions of Bernoulli RVs with $\mathbb{P}(B = 1) \leq \varphi$. One can test the null hypothesis based on special statistics called p -value bounds. A statistic $P = P(X) \geq 0$ is a p -value bound for \mathcal{H}_0 if for all $\mu \in \mathcal{H}_0$ and $p \geq 0$, $\mathbb{P}_\mu(P \leq p) \leq p$. Here, the subscript μ on $\mathbb{P}_\mu(\dots)$ indicates the distribution with respect to which the probabilities are to be calculated. We usually just write “ p -value” instead of “ p -value bound”, even when the bounds are not achieved by a member of \mathcal{H}_0 . Small p -values are strong evidence against the null. Since we are interested in very small p -values, we preferentially use their negative logarithm $-\log(P)$ and call this the $\log(p)$ -value. In this work, logarithms are base e by default.

A general method for constructing p -values is to start with an arbitrary real-valued RV Q jointly distributed with X . Usually Q is a function of X . Define the worst-case tail probability of Q as $P(q) = \sup_{\mu \in \mathcal{H}_0} \mathbb{P}_\mu(Q \geq q)$. Then $P(Q)$ is a p -value for \mathcal{H}_0 . The argument is standard. Define $F_\mu(q) = \mathbb{P}_\mu(Q \geq q)$. The function F_μ is non-increasing. We need to show that for all $\mu \in \mathcal{H}_0$, $\mathbb{P}_\mu(P(Q) \leq p) \leq p$. Since $F_\mu(q) \leq P(q)$, we have $\mathbb{P}_\mu(P(Q) \leq p) \leq \mathbb{P}_\mu(F_\mu(Q) \leq p)$. The set $\{q : F_\mu(q) \leq p\}$ is either of the form $[q_{\min}, \infty)$ or (q_{\min}, ∞) for some q_{\min} . In the first case, $\mathbb{P}_\mu(F_\mu(Q) \leq p) = \mathbb{P}_\mu(Q \geq q_{\min}) = F_\mu(q_{\min}) \leq p$. In the second, $\mathbb{P}_\mu(F_\mu(Q) \leq p) = \mathbb{P}_\mu(\bigcup_n \{q : q \geq q_{\min} + 1/n\}) = \lim_n \mathbb{P}_\mu(\{q : q \geq q_{\min} + 1/n\}) = \lim_n \mathbb{P}_\mu(F_\mu(Q) \leq q_{\min} + 1/n) \leq p$, with σ -additivity applied to the countable monotone union.

When referring to \mathcal{H}_0 as a null for \mathbf{X} , we mean that \mathcal{H}_0 consists of the distributions where the X_i are i.i.d., with X_i distributed according to μ for some fixed μ independent of i . To go beyond i.i.d., we extend \mathcal{H}_0 to the set of distributions of \mathbf{X} that have the property that for all $\mathbf{x}_{\leq i-1}$, $\mu(X_i | \mathbf{X}_{\leq i-1} = \mathbf{x}_{\leq i-1}) = \mu_i$ for

some $\mu_i \in \mathcal{H}_0$, where μ_i depends on i and $\mathbf{x}_{\leq i-1}$. We denote the extended null by $\overline{\mathcal{H}_0}$. In particular,

$$\overline{\mathcal{B}_\varphi} = \{\mu : \text{for all } i \text{ and } \mathbf{b}_{\leq i-1}, \mu(B_i | \mathbf{B}_{\leq i-1} = \mathbf{b}_{\leq i-1}) = v_\theta \text{ for some } \theta \leq \varphi\}. \quad (2)$$

The LR models mentioned in the introduction constitute a particular null \mathcal{H}_{LR} for sequences of trials called Bell tests. In Ref. [9], a technique called the probability-based ratio (PBR) method was introduced to construct p -values P_n that achieve asymptotically optimal gain rates defined as $\mathbb{E}(\log(1/P_n))/n$. The method is best understood as a way of constructing a test supermartingale for \mathcal{H}_{LR} . A test supermartingale of \mathbf{X} for \mathcal{H}_0 is a stochastic sequence $\mathbf{T} = (T_i)_{i=0}^n$ where T_i is a function of $\mathbf{X}_{\leq i}$, $T_0 = 1$, $T_i \geq 0$ and for all $\mu \in \mathcal{H}_0$, $\mathbb{E}_\mu(T_{i+1} | \mathbf{X}_{\leq i}) \leq T_i$. In this work, to avoid unwanted boundary cases, we further require T_i to be positive. The definition of test supermartingale used here is not the most general one because we consider only discrete time and avoid the customary increasing sequence of σ -algebra by making it dependent on an explicit stochastic sequence \mathbf{X} . Every test supermartingale defines a p -value by $P_n = 1/T_n$. This follows from $\mathbb{E}(T_n) \leq T_0 = 1$ (one of the characteristic properties of supermartingales) and Markov's inequality for non-negative statistics, according to which $\mathbb{P}(T_n \geq \kappa) \leq \mathbb{E}(T_n)/\kappa \leq 1/\kappa$. From martingale theory, the stopped process T_τ for any stopping rule τ with respect to \mathbf{X} also defines a p -value by $P = 1/T_\tau$. Further, $P_n^* = 1/\max_{i=1}^n T_i$ also defines a p -value. See Ref. [12] for a discussion and examples.

A test supermartingale \mathbf{T} can be viewed as the running product of the $F_i = T_i/T_{i-1}$, which we call the test factors of \mathbf{T} . The defining properties of \mathbf{T} are equivalent to having $F_i > 0$ and $\mathbb{E}(F_i | \mathbf{X}_{\leq i-1}) \leq 1$ for all distributions in the null, for all i . The PBR method adaptively constructs F_i as a function of the next trial outcome X_i from the earlier trial outcomes $\mathbf{X}_{\leq i-1}$. The method is designed for testing $\overline{\mathcal{H}_0}$ for a closed convex null \mathcal{H}_0 , where asymptotically optimally gain rates are achieved when the trials are i.i.d. with a trial distribution ν not in \mathcal{H}_0 . If ν were known, the optimal test factor would be given by $x \mapsto \nu(x)/\mu(x)$, where $\mu \in \mathcal{H}_0$ is the distribution in \mathcal{H}_0 closest to ν in Kullback-Leibler (KL) divergence $\text{KL}(\nu | \mu) = \sum_x \nu(x) \log(\nu(x)/\mu(x))$ [19]. Since ν is not known, the PBR method obtains an empirical estimate $\hat{\nu}$ of ν from $\mathbf{x}_{\leq i-1}$ and other information available before the i 'th trial. It then determines the KL-closest $\mu \in \mathcal{H}_0$ to $\hat{\nu}$. The test factor F_i is then given by $F_i(x) = \hat{\nu}(x)/\mu(x)$. The test factors satisfy $\mathbb{E}_{\mu'}(F_i) \leq 1$ for all $\mu' \in \mathcal{H}_0$, see Ref. [9] for a proof and applications to the problem of testing LR.

The ability to choose test factors adaptively helps reject extended nulls when the distributions vary as the experiment progresses, both when the distributions are still independent (so only the parameters vary) and when the parameters depend on past outcomes. Suppose that the distributions are sufficiently stable so that the empirical frequencies over the past k trials are statistically close to the next trial's probability distribution. Then we can adaptively compute the test factor to be used for the next trial from the past k trials' empirical frequencies, for example by following the strategy outlined in the previous paragraph. The procedure now has an opportunity to reject an extended null provided only that there is a sufficiently long period where the original null does not hold. For example, consider the extended null $\overline{\mathcal{B}_\varphi}$. The true success probabilities θ_i at the i 'th trial may vary, maybe as a result of changes in experimental parameters that need to be calibrated. Suppose that the goal is to calibrate for $\theta_i > \varphi$. If we use adaptive test factors and find at some point that we cannot reject $\overline{\mathcal{B}_\varphi}$ according to the running product of the test factors, we can recalibrate during the experiment. If the recalibration succeeds at pushing θ_i above φ for the remaining trials, we may still reject the extended null by the end of the experiment. In many cases, the analysis is performed after the experiment, or it may not be possible to stop the experiment for recalibration. For this situation, if the frequencies for a run of k trials clearly show that $\theta_i < \varphi$, the adaptive test factors chosen would tend to be trivial (equal to 1), in which case the next trials do not contribute to the final test factor product. This is in contrast to a hypothesis test based on the final sum of the outcomes for which all trials contribute equally.

Let φ be a parameter of distributions of X . Here, φ need not determine the distributions. There is a close relationship between methods for determining confidence sets for φ and hypothesis tests. Let \mathcal{H}_φ be a null

such that for all distributions μ with parameter φ , $\mu \in \mathcal{H}_\varphi$. Given a family of hypothesis tests with p -values P_φ for \mathcal{H}_φ , we can construct confidence sets for φ by inverting the acceptance region of P_φ , see Ref. [18], Sec. 7.1.2. According to this construction, the confidence set C_a at level a is given by $\{\varphi | P_\varphi(X) \geq a\}$ and is a random quantity. The defining property of a level a confidence set is that its coverage probability satisfies $\mathbb{P}_\mu(\varphi \in C_a) \geq 1 - a$ for all distributions $\mu \in \mathcal{H}_\varphi$. When we use this construction for sequences \mathbf{B} of i.i.d. Bernoulli RVs with the null \mathcal{B}_φ , we obtain one-sided confidence intervals of the form $[\varphi_0, 1]$ for $\theta = \mathbb{E}(B_i)$. When the confidence set is a one-sided interval of this type, we refer to φ_0 as the confidence lower bound or endpoint. If \mathbf{B} has a distribution μ that is not necessarily i.i.d., we can define $\Theta_{\max} = \max_{i \leq n} \mathbb{E}_\mu(B_i | \mathbf{B}_{\leq i-1})$. If we use acceptance region inversion with the extended null $\overline{\mathcal{B}_\varphi}$, we obtain a confidence region for Θ_{\max} . Note that Θ_{\max} is an RV whose value is covered by the confidence set with probability at least $1 - a$. The confidence set need not be an interval in general, but including everything between its infimum and its supremum increases the coverage probability, so the set can be converted into an interval if desired.

We focus on construction of one-sided confidence intervals. Given this machinery, we then construct two-sided intervals. For our example, we can obtain confidence upper bounds at level a by symmetry, for example by relabeling the Bernoulli outcomes $0 \mapsto 1$ and $1 \mapsto 0$. To obtain a two-sided interval at level a , we compute lower and upper bounds at level $a/2$. The two-sided interval is the interval between the bounds. The coverage probability of the two-sided interval is valid according to the union bound applied to maximum non-coverage probabilities of the two one-sided intervals.

3. Bernoulli Hypothesis Tests

We compare three hypothesis tests for the nulls \mathcal{B}_φ or the extended nulls $\overline{\mathcal{B}_\varphi}$: The “exact” test with p -value P_X , the Chernoff-Hoeffding test with p -value P_{CH} and a PBR test with p -value P_{PBR} . In discussing properties of these tests with respect to the hypothesis parameter φ , the true success probability θ and the empirical success probability $\hat{\Theta}$, we generally assume that these parameters are in the interior of their range. In particular, $0 < \varphi < 1$, $0 < \theta < 1$, and $0 < \hat{\Theta} < 1$. When discussing purely functional properties with respect to values $\hat{\theta}$ of $\hat{\Theta}$, we use the variable t instead of $\hat{\theta}$. By default nt is a positive integer.

The p -value for the exact test is obtained from the tail for i.i.d. Bernoulli RVs:

$$P_{X,n}(\hat{\Theta}|\varphi) = \sum_{k \geq \hat{\Theta}n} \binom{n}{k} \varphi^k (1 - \varphi)^{n-k}, \quad (3)$$

where $\hat{\Theta} = S_n/n = \sum_{i=1}^n B_i/n$ as defined in Sec. 2. Note that unlike the other p -values we consider, $P_{X,n}$ is not just a p -value bound. It is achieved by a member of the null. The quantity $P_{X,n}(t|\varphi)$ is decreasing as a function of t , given $0 < \varphi < 1$. It is smooth and monotonically increasing as a function of φ , given $t > 0$. To

see this, compute

$$\begin{aligned}
 \frac{d}{d\varphi} P_{X,n}(t|\varphi) &= \sum_{i=nt}^n \varphi^i (1-\varphi)^{n-i} \binom{n}{i} \left(\frac{i}{\varphi} - \frac{n-i}{1-\varphi} \right) \\
 &= n \sum_{i=nt}^n \varphi^{i-1} (1-\varphi)^{n-i} \binom{n-1}{i-1} - n \sum_{i=nt}^{n-1} \varphi^i (1-\varphi)^{n-i-1} \binom{n-1}{i} \\
 &= n \left(\sum_{i=nt-1}^{n-1} \varphi^i (1-\varphi)^{n-1-i} \binom{n-1}{i} - \sum_{i=nt}^{n-1} \varphi^i (1-\varphi)^{n-1-i} \binom{n-1}{i} \right) \\
 &= n \varphi^{nt-1} (1-\varphi)^{n(1-t)} \binom{n-1}{nt-1}.
 \end{aligned} \tag{4}$$

This is positive for $\varphi \in (0, 1)$. The probability that $S_n \geq tn$, given that all B_i are distributed as v_θ with $\theta \leq \varphi$, is bounded by $P_{X,n}(t|\theta) \leq P_{X,n}(t|\varphi)$. That P_X is a p -value for the case where the null is restricted to i.i.d. distributions now follows from the standard construction of p -values from worst-case (over the null) tails of statistics (here S_n) as explained in the previous section. That P_X is a p -value for the extended null $\overline{\mathcal{B}_\varphi}$ follows from the observations that the tail probabilities of S_n are linear functions of the distribution parameters $\theta_1, \theta_2, \dots, \theta_n$ where $\theta_i \leq \varphi, i = 1, 2, \dots, n$, the extremal distributions in $\overline{\mathcal{B}_\varphi}$ have B_i independent with $\mathbb{P}(B_i = 1) = \theta_i \leq \varphi$, and the tail probabilities of S_n are monotonically increasing in $\mathbb{P}(B_i = 1)$ for each i separately. See section C of Appendix of Ref. [20].

Define $\hat{\Theta}_{\max} = \max(\hat{\Theta}, \varphi)$. The p -value for the Chernoff-Hoeffding test is the optimal Chernoff-Hoeffding bound [16, 17] for a binary random variable given by

$$\begin{aligned}
 P_{\text{CH},n}(\hat{\Theta}|\varphi) &= \left(\frac{\varphi}{\hat{\Theta}_{\max}} \right)^{n\hat{\Theta}_{\max}} \left(\frac{1-\varphi}{1-\hat{\Theta}_{\max}} \right)^{n(1-\hat{\Theta}_{\max})} \\
 &= \begin{cases} \left(\frac{\varphi}{\hat{\Theta}} \right)^{n\hat{\Theta}} \left(\frac{1-\varphi}{1-\hat{\Theta}} \right)^{n(1-\hat{\Theta})} & \text{if } \hat{\Theta} \geq \varphi, \\ 1 & \text{otherwise.} \end{cases}
 \end{aligned} \tag{5}$$

This is a p -value for our setting because $P_{\text{CH},n}(t|\varphi) \geq P_{X,n}(t|\varphi)$, see Ref. [17]. For $\varphi \leq t$, we have $-\log(P_{\text{CH},n}(t|\varphi)) = n\text{KL}(v_t|v_\varphi)$. We abbreviate $\text{KL}(v_t|v_\varphi)$ by $\text{KL}(t|\varphi)$. For $\varphi \leq t < 1$, $P_{\text{CH},n}(t|\varphi)$ is monotonically increasing in φ , and decreasing in t . For $0 \leq t \leq \varphi$, it is constant.

The p -value for the PBR test that we use for comparison is constructed from a p -value for the point null $\{v_\varphi\}$ defined as

$$P_{\text{PBR},n}^0(\hat{\Theta}|\varphi) = \varphi^{n\hat{\Theta}} (1-\varphi)^{n(1-\hat{\Theta})} (n+1) \binom{n}{n\hat{\Theta}}. \tag{6}$$

The PBR test's p -value for $\overline{\mathcal{B}_\varphi}$ is

$$P_{\text{PBR},n}(\hat{\Theta}|\varphi) = \max_{0 \leq \varphi' \leq \varphi} P_{\text{PBR},n}^0(\hat{\Theta}|\varphi'). \tag{7}$$

That P_{PBR} is a p -value for $\overline{\mathcal{B}_\varphi}$ is shown below. As a function of φ , $P_{\text{PBR},n}^0(t|\varphi)$ has an isolated maximum at $\varphi = t$. This can be seen by differentiating $\log(\varphi^t (1-\varphi)^{1-t}) = t \log(\varphi) + (1-t) \log(1-\varphi)$. Thus in Eq. 7

when $\varphi \geq \hat{\Theta}$, the maximum is achieved by $\varphi' = \hat{\Theta}$. We can therefore write Eq. 7 as

$$P_{\text{PBR},n}(\hat{\Theta}|\varphi) = \begin{cases} P_{\text{PBR},n}^0(\hat{\Theta}|\varphi) & \text{if } \hat{\Theta} \geq \varphi, \\ P_{\text{PBR},n}^0(\hat{\Theta}|\hat{\Theta}) & \text{otherwise.} \end{cases} \quad (8)$$

By definition, $P_{\text{PBR},n}(t|\varphi)$ is non-decreasing in φ and strictly increasing for $\varphi \leq t$. As a function of t , it is strictly decreasing for $t \geq \varphi$ (integer-valued nt). To see this, consider $k = nt \geq n\varphi$ and compute the ratio of successive values as follows:

$$\begin{aligned} P_{\text{PBR},n}^0((k+1)/n|\varphi)/P_{\text{PBR},n}^0(k/n|\varphi) &= \frac{\varphi}{1-\varphi} \frac{n-k}{k+1} \\ &= \frac{\varphi}{1-\varphi} \frac{1-t}{t+1/n} \\ &\leq \frac{\varphi}{1-\varphi} \frac{1-t}{t} \\ &\leq 1. \end{aligned} \quad (9)$$

The expression for $P_{\text{PBR},n}^0$ is the final value of a test supermartingale obtained by constructing test factors F_{k+1} from S_k . Define

$$\tilde{\Theta}_k = \frac{1}{k+2} (S_k + 1). \quad (10)$$

Thus, $\tilde{\Theta}_k$ would be an empirical estimate of θ if there were two initial trials B_{-1} and B_0 with values 0 and 1, respectively. The test factors are given by

$$F_{k+1}(B_{k+1}) = \left(\frac{\tilde{\Theta}_k}{\varphi}\right)^{B_{k+1}} \left(\frac{1-\tilde{\Theta}_k}{1-\varphi}\right)^{1-B_{k+1}}. \quad (11)$$

One can verify that $\mathbb{E}_{v_\theta}(F_{k+1}) = 1$ for $\theta = \varphi$. More generally, set $\delta = \theta - \varphi$ and compute

$$\begin{aligned} \mathbb{E}_{v_\theta}(F_{k+1}|\tilde{\Theta}_k = t) &= \theta \frac{t}{\varphi} + (1-\theta) \frac{1-t}{1-\varphi} \\ &= 1 + \delta \left(\frac{t}{\varphi} - \frac{1-t}{1-\varphi}\right) \\ &= 1 + \delta \frac{t-\varphi}{\varphi(1-\varphi)}. \end{aligned} \quad (12)$$

As designed, $T_n = \prod_{k=1}^n F_k$ is a test supermartingale for the point null $\{v_\varphi\}$. Theorem 5 in section 7.2 of the Appendix, establishes that $T_n = 1/P_{\text{PBR},n}^0(\hat{\Theta}|\varphi)$. The definition of $P_{\text{PBR},n}(\hat{\Theta}|\varphi)$ as a maximum of p -values for $v_{\varphi'}$ with $\varphi' \leq \varphi$ in Eq. 7 ensures that $P_{\text{PBR},n}(\hat{\Theta}|\varphi)$ is a p -value for \mathcal{B}_φ .

To show that P_{PBR} is a p -value for \mathcal{B}_φ , we establish that for all t (integer-valued nt), $P_{\text{PBR},n}(t|\varphi) \geq P_{\text{CH},n}(t|\varphi)$. By direct calculation for both $\varphi \leq t$ and $t \leq \varphi$, we have

$$P_{\text{PBR},n}(t|\varphi)/P_{\text{CH},n}(t|\varphi) = t^{nt}(1-t)^{n(1-t)}(n+1) \binom{n}{nt}. \quad (13)$$

The expression $t^k(1-t)^k \binom{n}{k}$ is maximized at $k = nt$ as can be seen by considering ratios for successive

values of k and the calculation in Eq. 9, now applied also for $k < nt$. Therefore,

$$t^{nt}(1-t)^{n(1-t)}(n+1)\binom{n}{nt} = \sum_{k=0}^n t^{nt}(1-t)^{n(1-t)}\binom{n}{nt} \geq \sum_{k=0}^n t^k(1-t)^k\binom{n}{k} = 1. \quad (14)$$

A better choice for test factors to construct a test supermartingale to test $\overline{\mathcal{B}_\varphi}$ would be

$$T'_{k+1} = \begin{cases} T_{k+1} & \text{if } \tilde{\Theta}_k \geq \varphi, \\ 1 & \text{otherwise.} \end{cases} \quad (15)$$

This choice ensures that $\mathbb{E}_{\nu_\theta}(F_{k+1}|\mathbf{B}_{\leq k}) \leq 1$ for all $\theta \leq \varphi$ but the final value of the test supermartingale obtained by multiplying these test factors is not determined by S_n , which would complicate our study.

We summarize the observations about the three tests in the following theorem.

Theorem 1. *We have*

$$P_X \leq P_{CH} \leq P_{PBR}. \quad (16)$$

The three tests satisfy the following monotonicity properties for $0 < \varphi < 1$ and $0 < t < 1$ with nt integer-valued:

$P_X(t|\varphi)$ is strictly increasing in φ and strictly decreasing as a function of t .

$P_{CH}(t|\varphi)$ is strictly increasing in φ for $\varphi \leq t$, constant in φ for $\varphi \geq t$, strictly decreasing in t for $t \geq \varphi$ and constant in t for $t \leq \varphi$.

$P_{PBR}(t|\varphi)$ is strictly increasing in φ for $\varphi \leq t$, constant in φ for $\varphi \geq t$ and strictly decreasing in t for $t \geq \varphi$.

4. Comparison of p -Values

We begin by determining the relationships between P_X , P_{CH} and P_{PBR} more precisely. Since we are interested in small p -values, it is convenient to focus on the $\log(p)$ -values instead and determine their differences to $O(1/\sqrt{n})$. Because of the identity $-\log(P_{CH,n}(t, \varphi)) = n\text{KL}(t|\varphi)$, we reference all $\log(p)$ -values to $-\log(P_{CH,n})$. Here we examine the differences for $t \geq \varphi$ determined by the following theorem:

Theorem 2. *For $0 < \varphi < t < 1$,*

$$-\log(P_{PBR,n}(t|\varphi)) = -\log(P_{CH,n}(t|\varphi)) - \frac{1}{2}\log(n+1) + \frac{1}{2}\log(2\pi t(1-t)) + O\left(\frac{1}{n}\right), \quad (17)$$

$$-\log(P_{X,n}(t|\varphi)) = -\log(P_{CH,n}(t|\varphi)) + \frac{1}{2}\log(n) - \log\left(\sqrt{\frac{t}{2\pi(1-t)}}\frac{1-\varphi}{t-\varphi}\right) + O\left(\frac{1}{n}\right). \quad (18)$$

The theorem follows from Thms. 6, 7 and Cor. 8 proven in the Appendix, where explicit interval expressions are obtained for these $\log(p)$ -value differences. The order notation assumes fixed $t > \varphi$. The bounds are not uniform, see the expressions in the Appendix for details.

The most notable observation is that there are systematic gaps of $\log(n)/2 + O(1)$ between the $\log(p)$ -values. As we already knew, there is no question that the exact test is the best of the three for this simple application. While these gaps may seem large on an absolute scale, representing factors close to \sqrt{n} , they are in fact much smaller than the experiment-to-experiment variation of the p -values. To determine this

variation, we consider the asymptotic distributions. We can readily determine that the $\log(p)$ -values are asymptotically normal with standard deviations proportional to \sqrt{n} , which is transferred from the variance of $\hat{\Theta}$. Compared to these standard deviations the gaps are negligible. The next theorem determines the specific way in which asymptotic normality holds. Let $N(\mu, \sigma^2)$ denote the normal distribution with mean μ and variance σ^2 . The notation $X_n \xrightarrow{D} N(\mu, \sigma^2)$ means that X_n converges in distribution to the normal distribution with mean μ and variance σ^2 .

Theorem 3. *Assume $0 < \varphi < \theta < 1$. For $P = P_{CH,n}$, $P = P_{PBR,n}$ or $P = P_{X,n}$, the $\log(p)$ -value $-\log(P)$ converges in distribution according to*

$$\sqrt{n}(-\log(P)/n - \text{KL}(\theta|\varphi)) \xrightarrow{D} N(0, \sigma_G^2), \quad (19)$$

with

$$\sigma_G^2 = \theta(1 - \theta) \left(\log \left(\frac{\theta}{1 - \theta} \frac{1 - \varphi}{\varphi} \right) \right)^2.$$

The theorem is proven in the Appendix, see Thm. 10. For the rest of the paper, we write P or P_n for the p -values of any one of the tests when it does not matter which one.

We display the behavior described in the above theorems for $n = 100$ and $\theta = 0.5$ in Fig. 1. We conclude that the phenomena discussed above are already apparent for small numbers of trials. For Fig. 1, we computed the quantiles of the $\log(p)$ -values numerically using the formulas provided in the previous section, substituting for t the corresponding quantile of $\hat{\Theta}$ given that $\mathbb{P}(B = 1) = \theta$. To be explicit, let $t_{r,n}(\theta)$ be the r -quantile of $\hat{\Theta}$ defined as the minimum value $\hat{\theta}$ of $\hat{\Theta}$ satisfying $\mathbb{P}(\hat{\Theta} \leq \hat{\theta}) \geq r$. (For simplicity we do not place the quantile in the middle of the relevant gap in the distribution.) For example, $t_{0.5,n}(\theta)$ is the median. Then, by the monotonicity properties of the tests, the r -quantile of $-\log(P_n(\hat{\Theta}|\varphi))$ is given by $-\log(P_n(t_{r,n}(\theta)|\varphi))$.

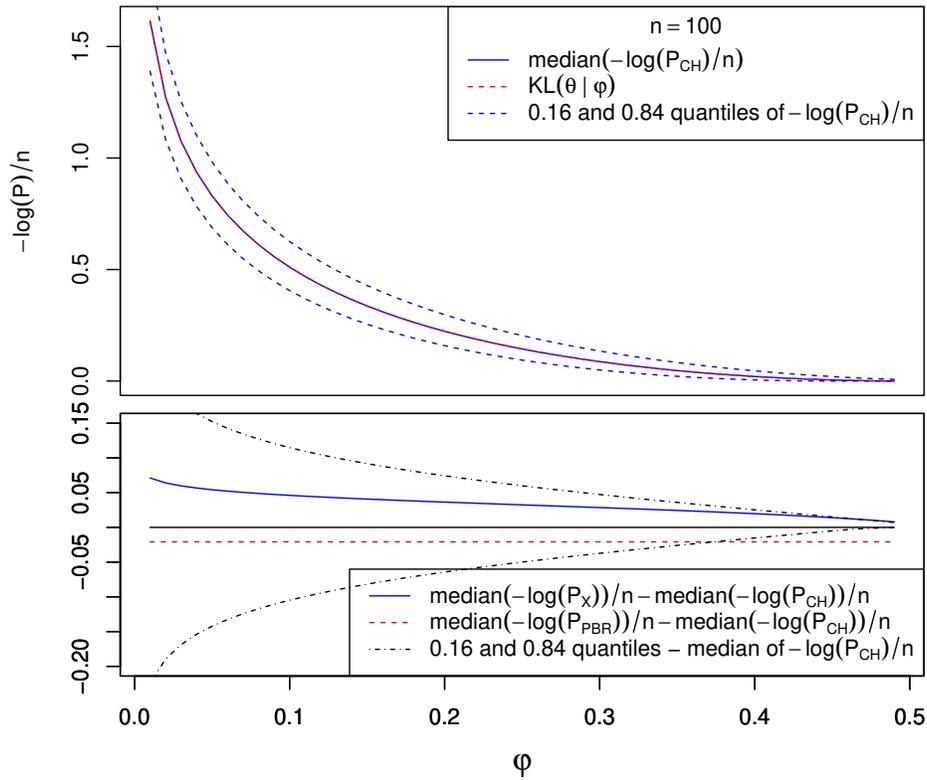


Fig. 1. Comparison of $\log(p)$ -values at $n = 100$ and $\theta = 0.5$. The top half of the figure shows the median, and the 0.16 and 0.84 quantiles of $-\log(P_{CH,n}(\hat{\Theta}|\varphi))/n$. For $\theta = 0.5$, the median agrees with $KL(\theta|\varphi)$ by symmetry. The lower half shows the median differences $-\log(P(\hat{\Theta}|\varphi))/n + \log(P_{CH,n}(\hat{\Theta}|\varphi))/n$ for $P = P_{PBR,n}$ and $P = P_{X,n}$. The difference between the 0.16 and 0.84 quantiles and the median for $-\log(P_{CH,n}(\hat{\Theta}|\varphi))/n$ are also shown where they are within the range of the plot; even for n as small as 100, they dominate the median differences, except where φ approaches $\theta = 0.5$, where the absolute p -values are no longer extremely small.

As noted above, the gaps between the $\log(p)$ -values are of the form $\log(n)/2 + O(1)$. In fact, it is possible to determine the asymptotic behavior of these gaps. After accounting for the explicitly given $O(1)$ terms in Thm. 2, they are asymptotically normal with variances of order $O(1/n)$. The standard deviations of the gaps are therefore small compared to their size. The precise statement of their asymptotic normality is Thm. 11 in the Appendix. Figure 2 shows how these gaps depend on the value $\hat{\theta}$ of $\hat{\Theta}$ given φ . The gaps are scaled by $\log(n)$ so that they can be compared to $\log(n)/2$ visually for different values of n . The deviation from $\log(n)/2$ is most notable near the boundaries, where convergence is also slower, particularly for P_X . This behavior is consistent with the divergences as t approaches φ in the explicit interval bounds in Thm. 7 and Cor. 8.

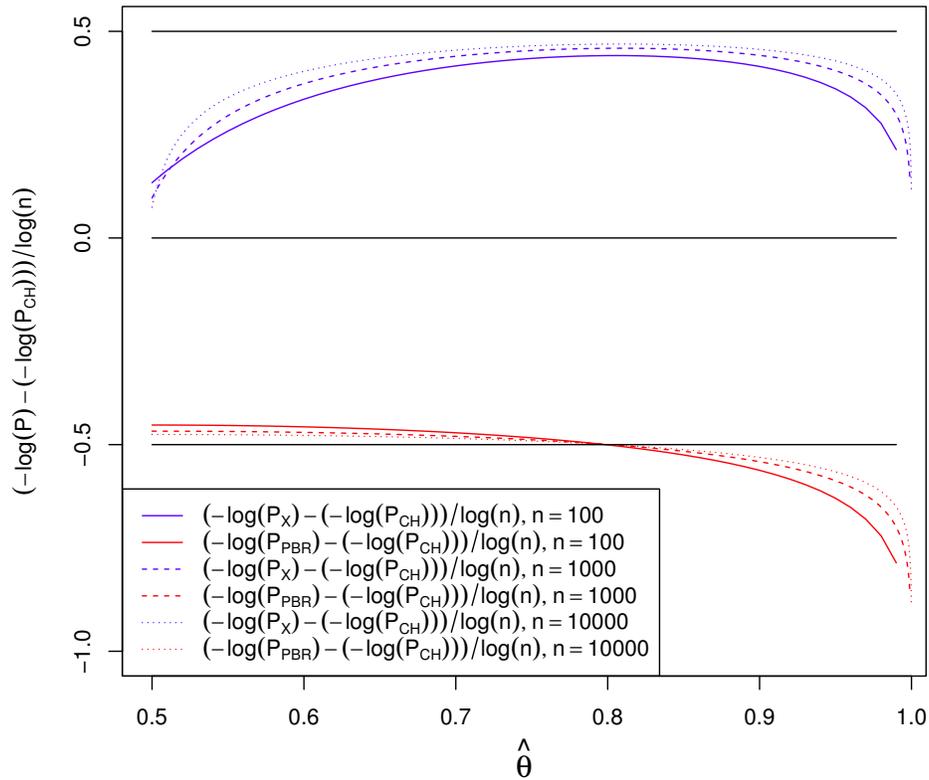


Fig. 2. Gaps between the $\log(p)$ -values depending on $\hat{\theta}$ at $\varphi = 0.5$. We show the normalized differences $(-\log(P_n(\hat{\theta}|\varphi)) + \log(P_{CH,n}(\hat{\theta}|\varphi)))/\log(n)$ for $P = P_{CH}$ and $P = P_X$ at $n = 100, 1000$, and 10000 . For large n , at constant $\hat{\theta}$ with $0.5 < \hat{\theta} < 1$, the PBR test's normalized difference converges to -0.5 , and the exact test's normalized difference converges to 0.5 . The horizontal lines at ± 0.5 indicate this limit. The lowest order normalized asymptotic differences from ± 0.5 are $O(1/\log(n))$ and diverge at $\hat{\theta} = 0.5$ and $\hat{\theta} = 1$.

5. Comparison of Confidence Intervals

Before presentation of our technical results, we remark that there are many excellent publications on construction of one-sided and two-sided confidence intervals for the success probability of binomial trials for the case of i.i.d. observations including [21] and [22].

Let P be one of $P_{CH,n}$, $P_{PBR,n}$ or $P_{X,n}$. Given a value $\hat{\theta}$ of $\hat{\Theta}$, the level- a confidence set determined by the test with p -value P is $I = \{\varphi | P(\hat{\theta}|\varphi) \geq a\}$. By the monotonicity properties of P , the closure of I is an interval $[\varphi_a(\hat{\theta}; P), 1]$. We can compute the endpoint φ_a by numerically inverting the exact expressions for P . An example is shown in Fig. 3, where we show the endpoints according to each test for $a = 0.01$ and $\hat{\theta} = 0.5$ as a function of n . All tests' endpoints converge to 0.5 as the number of trials grows. Notably, the relative separation between the endpoints is not large at level $a = 0.01$.

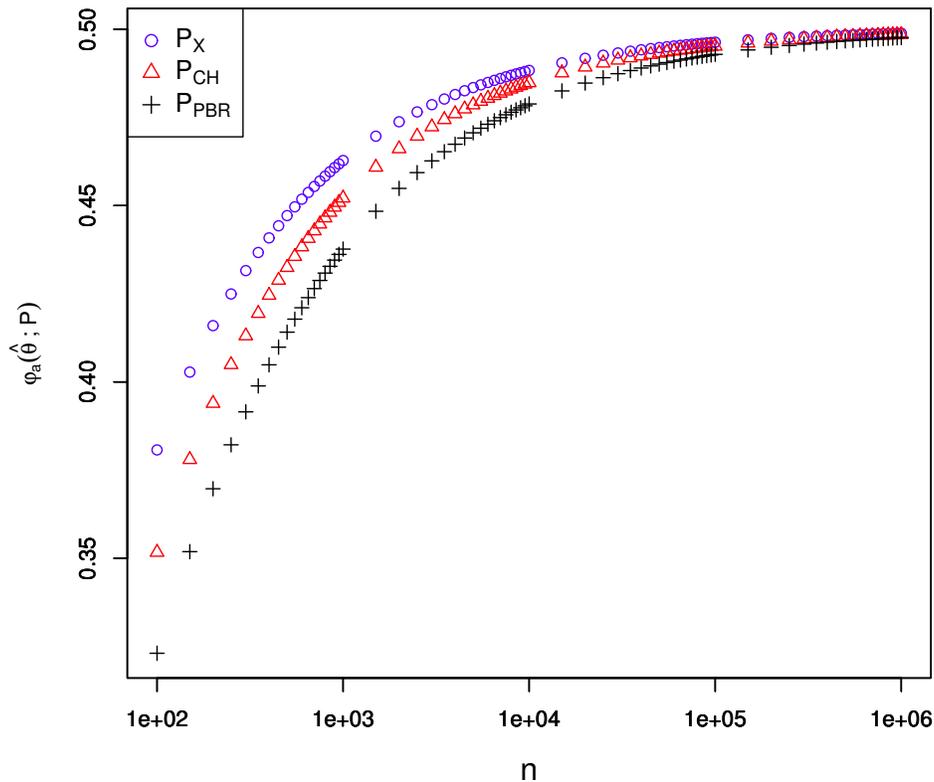


Fig. 3. Lower endpoints for the level 0.01 confidence set of the three tests as a function of n , where $\hat{\theta} = 0.5$.

To quantify the behavior of the endpoints for the different tests, we normalize by the empirical standard deviation $\hat{\sigma} = \sqrt{\hat{\theta}(1 - \hat{\theta})/n}$. The empirical endpoint deviation is then defined as

$$\gamma_a(\hat{\theta}; P) = \frac{\hat{\theta} - \varphi_a(\hat{\theta}; P)}{\hat{\sigma}}. \quad (20)$$

For the exact test and for large n , we expect this quantity to be determined by the tail probabilities of a standard normal distribution. That is, if the significance a is the probability that a normal RV of variance 1 exceeds κ , we expect $\gamma_a(\hat{\theta}; P_X) \approx \kappa$.

We take the point of view that the performance of a test is characterized by the size of the endpoint deviation. If the relative size of the deviations for two tests is close to 1 then they perform similarly for the purpose of characterizing the parameter θ . Another way of comparing the intervals obtained is to consider their coverage probabilities. For our situation, the coverage probability for test P at a can be approximated by determining a' such that $\gamma_{a'}(\theta; P_X) = \gamma_a(\theta; P)$. From Thm. 4 below, one can infer that the coverage probability is then approximately $1 - a' \geq 1 - a$. The coverage probabilities can be very conservative (larger than $1 - a$), particularly for small a and $P = P_{PBR}$.

We determined interval bounds for the empirical endpoint deviation for all three tests. The details are in section 7.5 of the Appendix. The next theorem summarizes the results asymptotically.

Theorem 4. Let $q(x) = -\log(\mathbb{P}_{N(0,1)}(N \geq x))$ be the negative logarithm of the tail of the standard normal. Fix $\hat{\theta} \in (0, 1)$. Write $\alpha = |\log(a)|$. There is a constant c (depending on $\hat{\theta}$) such that for $\alpha \in (1, c\alpha]$, γ_a

satisfies

$$\gamma_a(\hat{\theta}; P_{\text{CH}}) = \sqrt{2\alpha} + O(\alpha/\sqrt{n}), \quad (21)$$

$$\gamma_a(\hat{\theta}; P_{\text{PBR}}) = \sqrt{2\alpha + \log(n)/2 - \log(2\pi\hat{\theta}(1-\hat{\theta}))/2} + O(\alpha/\sqrt{n}), \quad (22)$$

$$\gamma_a(\hat{\theta}; P_{\text{X}}) = q^{-1}(\alpha) + O(\alpha/\sqrt{n}). \quad (23)$$

The last expression has the following approximation relevant for sufficiently large α :

$$\gamma_a(\hat{\theta}; P_{\text{X}}) = \sqrt{2\alpha - \log(2\pi) - \log(2\alpha - \log(2\pi))} + O(\log(\alpha)/\alpha^{3/2}) + O(\alpha/\sqrt{n}). \quad (24)$$

For $\alpha = o(\sqrt{n})$, the relative error of the approximation in the first two identities goes to zero as n grows. This is not the case for the last identity, where the relative error for large n is dominated by the term $O(\log(\alpha)/\alpha^{3/2})$, and large α is required for a small relative error.

Proof. The expression for $\gamma_a(\hat{\theta}; P_{\text{CH}})$ is obtained from Thm. 12 in the Appendix by changing the relative approximation errors into absolute errors.

To obtain the expression for $\gamma_a(\hat{\theta}; P_{\text{PBR}})$, note that the term Δ in Thm. 13 satisfies $\Delta = \log(n)/2 - \log(2\pi\hat{\theta}(1-\hat{\theta}))/2 + O(1/n)$, see Thm. 6. The $O(1/n)$ under the square root pulls out to an $O(1/(\sqrt{\max(\alpha, \log(n))n}))$ term that is dominated by $O(\alpha/\sqrt{n})$ because $\alpha \geq 1$ by assumption.

To obtain the expressions for $\gamma_a(\hat{\theta}; P_{\text{X}})$, we refer to Thm. 14, where the lower bound on α implies $\alpha \geq 1 > \log(2)$. The intervals in Thm. 14 give relative errors that need to be converted to absolute quantities. By positivity and monotonicity of q^{-1} , for sufficiently large n and for some positive constants u and v , we have

$$\gamma_a(\hat{\theta}; P_{\text{X}}) \in [q^{-1}(\alpha(1-u\sqrt{\alpha}/\sqrt{n}))(1-v\sqrt{\alpha}/\sqrt{n}), q^{-1}(\alpha(1+u\sqrt{\alpha}/\sqrt{n}))(1+v\sqrt{\alpha}/\sqrt{n})]. \quad (25)$$

Explicit values for u and v can be obtained from Thm. 14. We simplified the argument of q^{-1} by absorbing the additive terms in the theorem into the term $u\alpha\sqrt{\alpha}/\sqrt{n}$ with the constant u chosen to be sufficiently large. Consider Eq. 94 with $\delta = u\sqrt{\alpha}/\sqrt{n}$. For sufficiently large n , the expression in the denominator of the approximation error on the right-hand side exceeds a constant multiple of α . From this, with some new constant u' ,

$$\gamma_a(\hat{\theta}; P_{\text{X}}) \in [q^{-1}(\alpha)(1-u'\sqrt{\alpha}/\sqrt{n})(1-v\sqrt{\alpha}/\sqrt{n}), q^{-1}(\alpha)(1+u'\sqrt{\alpha}/\sqrt{n})(1+v\sqrt{\alpha}/\sqrt{n})], \quad (26)$$

which, with order notation simplifies further to

$$\gamma_a(\hat{\theta}; P_{\text{X}}) = q^{-1}(\alpha)(1 + O(\sqrt{\alpha}/\sqrt{n})). \quad (27)$$

It now suffices to apply $q^{-1}(\alpha) = O(\sqrt{\alpha})$ (see the proof of Eq. 24 below) and Eq. 23 is obtained.

For Eq. 24, we bound $x = q^{-1}(\alpha)$, which we can do via bounds for $\alpha = q(x)$. From the expression $q(x) = x^2/2 + \log(2\pi)/2 - \log(Y(x)) = x^2/2 + \log(2\pi)/2 + \log(x) - \log(xY(x))$ in the statement of Thm. 14 and the bounds in Eq. 58, we have the two inequalities

$$q(x) = x^2/2 + \log(2\pi)/2 + \log(x) - \log(xY(x)) \geq x^2/2 + \log(2\pi)/2 + \log(x), \quad (28)$$

$$q(x) \leq x^2/2 + \log(2\pi)/2 + \log(x) + 1/x^2. \quad (29)$$

Let $l(x) = x^2/2 + \log(2\pi)/2 + \log(x)$, which is monotonically increasing, as is q . The first inequality implies that $q^{-1} \leq l^{-1}$. We need a bound of the form $q(x) \leq dx^2$, from which we can conclude that $x^2 \geq \alpha/d$. A

bound of this type can be obtained from Eq. 91 in the Appendix. For definiteness, we restrict to $\alpha \geq 6$ and show that the bound holds with $d = 1$. By Eq. 29, it suffices to establish that for $x \geq \sqrt{6}$, $l(x) + 1/x^2 \leq x^2$. Since $\log(2\pi)/2 \leq 1$, we have $\log(2\pi)/2 + \log(x) + 1/x^2 \leq 1 + \log(1 + (x - 1)) + 1/x^2 \leq x + 1/x^2$. For $x \geq 9/4$, $x + 1/x^2 \leq x^2/2$. To finish the argument, apply the inequality $\sqrt{6} \geq 9/4$.

Given the bound $x^2 \geq \alpha$, Eq. 29 becomes $q(x) = \alpha \leq l(x) + 1/\alpha$. With Eq. 28 we get $\alpha = q(x) \in l(x) + [0, 1]/\alpha$. Equivalently,

$$l(x) \in \alpha + \frac{1}{\alpha}[-1, 0]. \quad (30)$$

Applying the monotone l^{-1} on both sides gives

$$x = l^{-1}(l(x)) \in l^{-1}\left(\alpha + \frac{1}{\alpha}[-1, 0]\right). \quad (31)$$

Let α' satisfy $x = l^{-1}(\alpha')$ with $\alpha' = \alpha + \delta$ and $\delta \in [-1, 0]/\alpha$. Write $z = x^2$ and $c = \log(2\pi)$. We have $l(z^{1/2}) = z/2 + c/2 + \log(z)/2 = \alpha'$, which we can write as a fixed point equation $z = f(z)$ for z with $f(z) = 2\alpha' - c - \log(z)$. We can accomplish our goal by determining lower and upper bounds on the fixed point. Since $\frac{d}{dy}f(y) = -1/y < 0$ for $y > 0$, the iteration $z_0 = 2\alpha' - c$ and $z_k = f(z_{k-1})$ is alternating around the fixed point z , provided $z_k > 0$ for all k . Provided $z_0 > 1$, $z_1 = f(z_0) < z_0$, from which we conclude that $z_1 \leq z < z_0$. Since we are assuming that $\alpha \geq 6$ and from above $z_0 \geq 2(\alpha - 1/\alpha) - c$, the condition $z_0 > 1$ is satisfied. If $z_1 \geq 1$, then $0 > \frac{d}{dy}f(y) > -1$ between z_1 and z_0 , which implies that z_0 and z_1 are in the region where the iteration converges to z . For our bounds, we only require $z_1 > 0$, so that we can bound z according to $z_1 \leq z \leq z_2$. That $z_1 > 0$ follows from $\log(y) < y$ for $y > 0$. We have

$$\begin{aligned} z_2 - z_1 &= z_0 - \log(z_1) - (z_0 - \log(z_0)) \\ &= \log(z_0/z_1) \\ &= \log(z_0/(z_0 - \log(z_0))) \\ &= -\log(1 - \log(z_0)/z_0) \\ &= O(\log(z_0)/z_0) = O(\log(\alpha')/\alpha') = O(\log(\alpha)/\alpha), \end{aligned} \quad (32)$$

where $z_0 = 2\alpha' - c \in 2\alpha - c + 2[-1, 0]/\alpha$, and so $-\log(z_0) = -\log(2\alpha - c) + O(1/\alpha^2)$. For z_1 we get $z_1 = z_0 - \log(z_0) = 2\alpha - c - \log(2\alpha - c) + O(1/\alpha)$. Applying Eq. 32 and from the definitions,

$$q^{-1}(\alpha) = x = \sqrt{2\alpha - c - \log(2\alpha - c) + O(\log(\alpha)/\alpha)}. \quad (33)$$

The approximation error in Eq. 24 is obtained by expanding the square root. We could have used Newton's method starting from z_0 to obtain better approximations in one step, but the resulting expression is more involved. □

The expression for $\gamma_a(\hat{\theta}; P_X)$ confirms our expectation that it approaches the expected value for a standard normal distribution and may be compared to the Berry-Esseen theorem [23]. The empirical endpoint deviation of the CH test approaches that of the exact test for small a (large α). Their squares differ by a term of order $\log(\alpha) = \log|\log(a)|$. Notably, the ratio of the PBR and CH tests' empirical endpoint deviation grows as $\Theta(\sqrt{\log(n)/\alpha})$. The relationships are visualized in Figs. 4, 5 and 6 for different values of a . The figures show that the relative sizes of the empirical endpoint deviations tend toward 1 with smaller a . The $\Theta(\sqrt{\log(n)/\alpha})$ relative growth of the PBR test's endpoint deviations leads to less than a doubling of the deviations relative to the exact test's at $a = 0.01$ and $a = 0.001$ even for $n = 10^6$. So while the test's

coverage probabilities are much closer to 1 than the nominal value of $1 - a$, we believe that it does not lead to unreasonably conservative results in many applications.

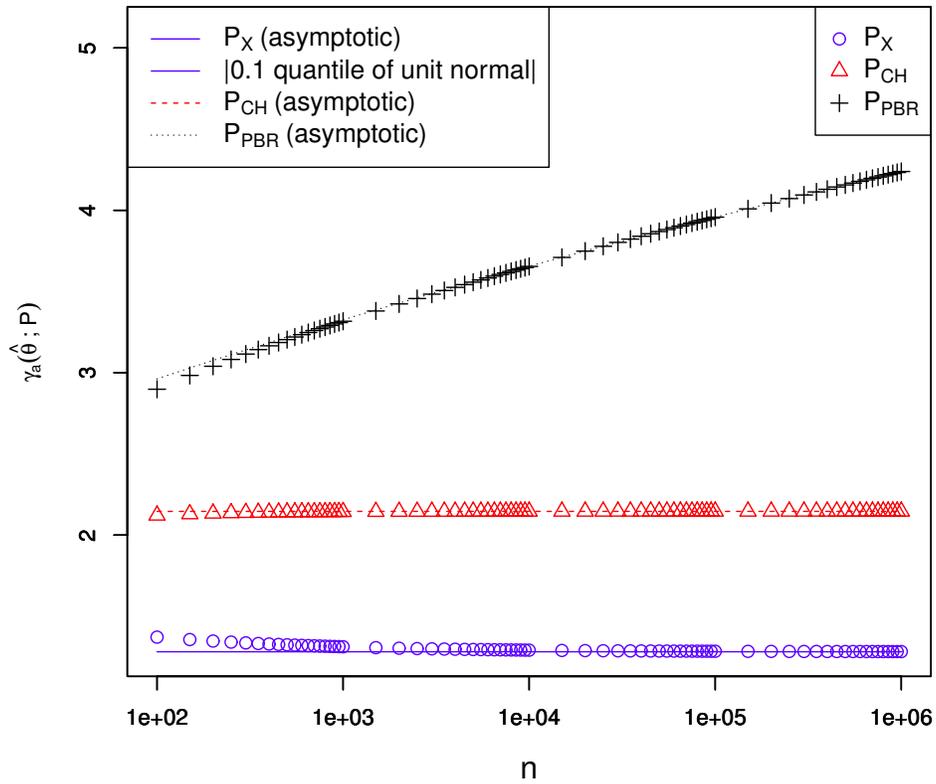


Fig. 4. Empirical confidence set endpoint deviations at level $a = 0.1$ for $\hat{\theta} = 0.5$ as a function of n . The continuous lines show the expressions obtained after dropping the $O(1/\sqrt{n})$ terms. For the exact test, these expressions are the same as the normal approximation and therefore match the absolute value of the 0.1 quantile of a unit normal.

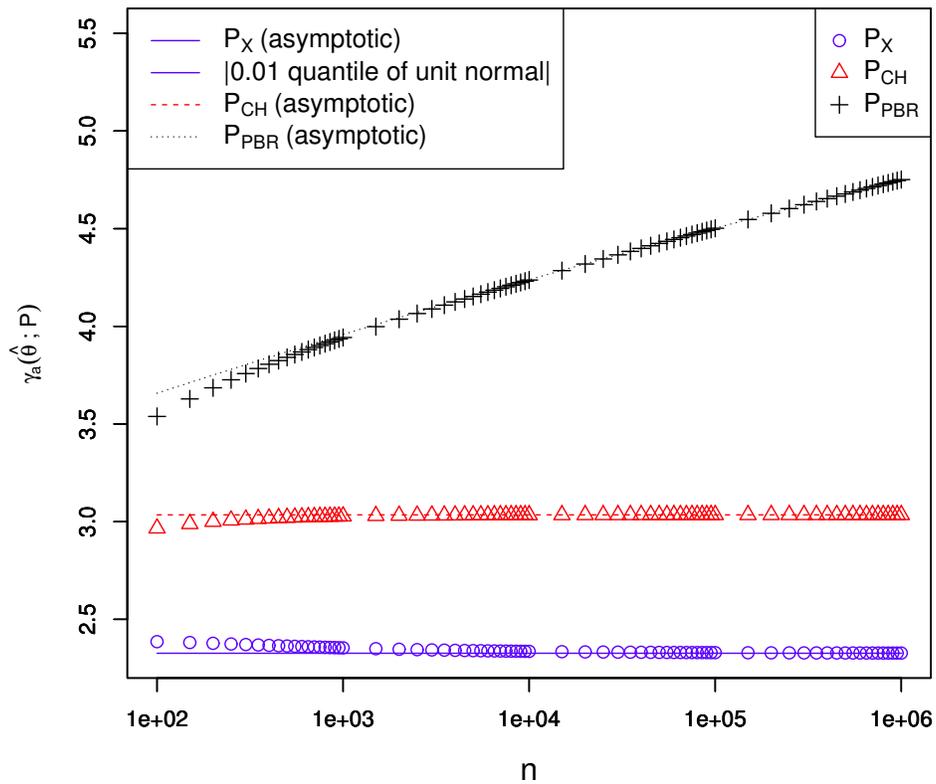


Fig. 5. Empirical confidence set endpoint deviations at level $a = 0.01$ for $\hat{\theta} = 0.5$ as a function of n . See the caption of Fig. 4.

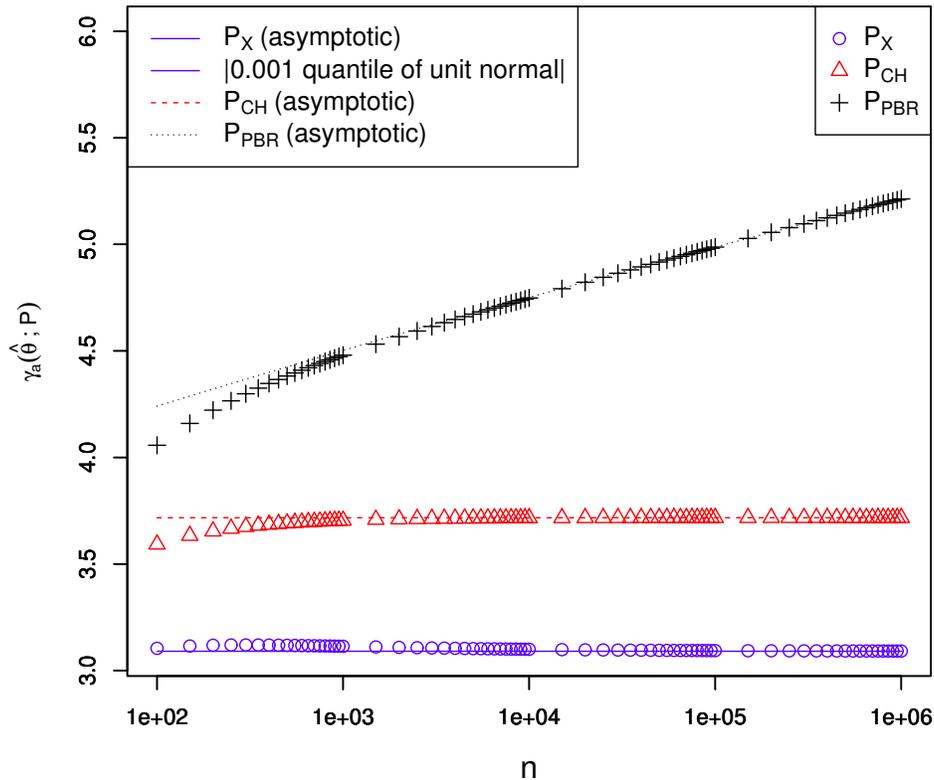


Fig. 6. Empirical confidence set endpoint deviations at level $a = 0.001$ for $\hat{\theta} = 0.5$ as a function of n . See the caption of Fig. 4.

Next we consider the behavior of the true endpoint deviations given by the normalized difference of the true success probability θ and the endpoint obtained from one of the tests. Let $\sigma = \sqrt{\theta(1-\theta)/n}$ be the true standard deviation and define the true endpoint deviation determined by test P as

$$\begin{aligned} \tilde{\gamma}_a(\hat{\Theta}|P) &= (\theta - \varphi_a(\hat{\Theta}|P))/\sigma \\ &= (\theta - \hat{\Theta})/\sigma + \gamma_a(\hat{\Theta}|P)\hat{\sigma}/\sigma. \end{aligned} \quad (34)$$

The true endpoint deviations show how the inferred endpoint compares to θ and therefore directly exhibits the statistical fluctuations of $\hat{\Theta}$. In contrast, the empirical endpoint deviations are to lowest order independent of $\hat{\theta} - \theta$.

We take the view that two tests' endpoints perform similarly if their true endpoint deviations differ by an amount that is small compared to the width of the distribution of the true endpoint deviations. To compare the three tests on this basis, we consider the quantiles for $\hat{\Theta}$ corresponding to $\pm\kappa$ Gaussian standard deviations from θ with κ constant. The quantiles satisfy $\theta_{\pm\kappa} = \theta \pm \kappa\sigma(1 + O(1/\sqrt{n}))$, by the Berry-Esseen theorem or from Thm. 14. Since $\hat{\sigma} = \sigma(1 + O(1/\sqrt{n}))$, we can also see that $\gamma_a(\theta_{\pm\kappa}|P) = \gamma_a(\hat{\theta}|P) + O(1/\sqrt{n})$, and so by substituting into the definition,

$$\tilde{\gamma}_a(\theta_{\pm\kappa}|P) = \gamma_a(\theta|P) \pm \kappa + O(1/\sqrt{n}), \quad (35)$$

where the implicit constants depend on κ . For large α , the CH and exact tests' endpoints are close and are

dominated by κ , so they perform similarly. But this does not hold for the comparison of the CH or the exact test's endpoints to those of the PBR test, since the latter's endpoint deviation grows as $\sqrt{\log(n)}/2$.

The PBR test's robustness to stopping rules requires that endpoint deviations must grow. Qualitatively, we expect a growth of at least $\Omega(\sqrt{\log \log(n)})$ due to the law of the iterated logarithm. This growth is slower than the $\sqrt{\log(n)}/2$ growth found above, suggesting that improvements are possible, as observed in Ref. [12]. In many applications, the number of trials to be acquired can be determined ahead of time, so full robustness to stopping rules is not necessary. However, the ability to adapt to changing experimental conditions may still be helpful, as the example in Sec. 2 shows. If we know the number of trials ahead of time, we can retain the ability to adapt while avoiding the asymptotic growth of the endpoint deviations of the PBR test.

A strategy for avoiding the asymptotic growth of the PBR test's endpoint deviations is to set aside the first $m = \lambda n$ of the trials for training to infer the probability of success, and then use this to determine the test factor to be used on the remaining $(1 - \lambda)n$ of the trials. With this strategy, the endpoint deviations are bounded on average and typically. We formalize the training strategy as follows: Modify Eq. 11 by setting $F_{k=1} = 1$ for $k < m$ and for $k \geq m$,

$$F_{k+1}(B_{k+1}) = F(B_{k+1}) = \left(\frac{\hat{\Theta}_m}{\varphi}\right)^{B_{k+1}} \left(\frac{1 - \hat{\Theta}_m}{1 - \varphi}\right)^{1 - B_{k+1}}. \quad (36)$$

Let $G = F$ if $\varphi \leq \hat{\Theta}_m$ and $G = 1$ otherwise. The G_{k+1} are valid test factors for the null \mathcal{B}_φ . A p -value for testing \mathcal{B}_φ is given by

$$P_\lambda(\mathbf{B}|\varphi) = G(1)^{-(n-m)\hat{\Theta}'_m} G(0)^{-(n-m)(1-\hat{\Theta}'_m)} \quad (37)$$

where $\hat{\Theta}'_m$ is defined by $(n - m)\hat{\Theta}'_m + m\hat{\Theta}_m = n\hat{\Theta}_n$. We call this the P_λ test.

Define

$$Q_\lambda(\mathbf{B}|\varphi) = \left(\frac{\varphi}{\hat{\Theta}_m}\right)^{(n-m)\hat{\Theta}'_m} \left(\frac{1 - \varphi}{1 - \hat{\Theta}_m}\right)^{(n-m)(1-\hat{\Theta}'_m)}. \quad (38)$$

Then for $\varphi \leq \hat{\Theta}_m$, $Q_\lambda(\mathbf{B}|\varphi) = P_\lambda(\mathbf{B}|\varphi)$. To investigate the behavior of these quantities, we consider values \mathbf{b} , $\hat{\theta}$, $\hat{\theta}_m$ and $\hat{\theta}'_m$ of the corresponding RVs. As a function of φ , $Q_\lambda(\mathbf{b}|\varphi)$ is maximized at $\varphi = \hat{\theta}'_m$ and monotone on either side of $\hat{\theta}'_m$. If $\hat{\theta}_m \leq \varphi \leq \hat{\theta}'_m$, then $Q_\lambda(\mathbf{b}|\varphi) \geq 1 = P_\lambda(\mathbf{b}|\varphi)$. So for $\varphi \leq \max(\hat{\theta}_m, \hat{\theta}'_m)$, we can use Q_λ instead of P_λ without changing endpoint calculations.

For determining the endpoint of a level- α one-sided confidence interval from P_λ , we seek the maximum φ such that for all $\varphi' \leq \varphi$, $P_\lambda(\mathbf{b}|\varphi') \leq \alpha$. This maximum value of φ satisfies that $\varphi \leq \min(\hat{\theta}'_m, \hat{\theta}_m)$: For $\hat{\theta}_m \leq \hat{\theta}'_m$, this follows from $P_\lambda(\mathbf{b}|\hat{\theta}_m) = 1$. For $\hat{\theta}_m \geq \hat{\theta}'_m$, the location of the maximum of Q_λ implies that $P_\lambda(\mathbf{b}|\hat{\theta}'_m) \geq P_\lambda(\mathbf{b}|\hat{\theta}_m) = 1$.

We show that endpoint deviations from the P_λ test are typically a constant factor larger than those of the CH test. For large α , the factor approaches $1/\sqrt{1 - \lambda}$, approximating the endpoint deviations for a CH test with $(1 - \lambda)n$ trials. We begin by comparing P_λ to $P_{\text{CH},(1-\lambda)n}$ with the latter applied to the last $(1 - \lambda)n$ trials and $\varphi \leq \hat{\theta}'_m$, where we can use Q_λ in place of P_λ .

$$Q_\lambda(\mathbf{b}|\varphi)/P_{\text{CH},(1-\lambda)n}(\hat{\theta}'_m|\varphi) = \left(\frac{\hat{\theta}'_m}{\hat{\theta}_m}\right)^{(1-\lambda)n\hat{\theta}'_m} \left(\frac{1 - \hat{\theta}'_m}{1 - \hat{\theta}_m}\right)^{(1-\lambda)n(1-\hat{\theta}'_m)}. \quad (39)$$

Or, for the $\log(p)$ -value difference l_p ,

$$l_p = -\log(Q_\lambda(\mathbf{b}|\varphi)) + \log(P_{\text{CH},(1-\lambda)n}(\hat{\theta}'_m|\varphi)) = -(1-\lambda)n\text{KL}(\hat{\theta}'_m|\hat{\theta}_m), \quad (40)$$

which is non-positive. By expanding to second order,

$$\begin{aligned} \text{KL}(t+x|t+y) &= (t+x)(\log(1+x/t) - \log(1+y/t)) \\ &\quad + (1-t-x)(\log(1-x/(1-t)) - \log(1-y/(1-t))) \\ &= \frac{(x-y)^2}{2t(1-t)} + O(\max(|x|, |y|)^3). \end{aligned} \quad (41)$$

Let $\Delta = \hat{\theta}_m - \theta$ and $\Delta' = \hat{\theta}'_m - \theta$. From the above expansion with $t = \theta$, $x = \delta'$ and $y = \delta$ (where δ and δ' are values of Δ and Δ')

$$l_p = -(1-\lambda)n \left(\frac{(\delta - \delta')^2}{2\theta(1-\theta)} + O(\max(|\delta|, |\delta'|^3)) \right). \quad (42)$$

The RVs Δ and Δ' are independent with means 0 and variances σ^2/λ and $\sigma^2/(1-\lambda)$. Furthermore, $\sqrt{n}\Delta$ and $\sqrt{n}\Delta'$ are asymptotically normal with variances $\theta(1-\theta)/\lambda$ and $\theta(1-\theta)/(1-\lambda)$. Consequently, the RV $\sqrt{n}(\Delta - \Delta')$ is asymptotically normal with variance $v = \theta(1-\theta)/(\lambda(1-\lambda))$. Accordingly, the probability that $n(\Delta - \Delta')^2 \geq \kappa^2\theta(1-\theta)/(\lambda(1-\lambda))$ is asymptotically given by the two-sided tail for κ standard deviations of the standard normal. For determining typical behavior, we consider $(\delta - \delta')^2 = \kappa^2\theta(1-\theta)/(n\lambda(1-\lambda))$ with $\kappa \geq 0$ constant for asymptotic purposes. Observe that $n\Delta^3$ and $n\Delta'^3$ are $\tilde{O}(1/\sqrt{n})$ with probability 1, where the “soft- O ” notation \tilde{O} subsumes the polylogarithmic factor from the law of the iterated logarithm. We can now write

$$l_p = -\frac{\kappa^2}{2\lambda} + \tilde{O}(1/\sqrt{n}). \quad (43)$$

Fix the level a and thereby also $\alpha = |\log(a)|$. Define $\hat{\sigma}' = \sqrt{\hat{\theta}'_m(1-\hat{\theta}'_m)/(1-\lambda)n}$, and let $\varphi' = \hat{\theta}'_m - \gamma'\hat{\sigma}'$ be the smallest solution of $-\log(Q_\lambda(\mathbf{b}|\varphi')) = \alpha$. Because

$$-\log(Q_\lambda(\mathbf{b}|\varphi')) = -\log(P_{\text{CH},(1-\lambda)n}(\hat{\theta}'_m|\varphi')) + l_p, \quad (44)$$

we can estimate γ' as $\gamma' = \gamma'_{a',(1-\lambda)n}(\hat{\theta}'_m; P_{\text{CH}}) = \sqrt{2(\alpha - l_p)} + O(\alpha/\sqrt{n})$ with $a' = e^{-(\alpha - l_p)}$. Here, the subscript $(1-\lambda)n$ of $\gamma'_{a'}$ makes the previously implicit number of trials explicit.

To finish our comparison, we express the endpoint φ' relative to $\hat{\theta}$. For this, we write

$$\begin{aligned} \varphi' &= \hat{\theta}'_m - \gamma'\hat{\sigma}' \\ &= \hat{\theta} + (\hat{\theta}'_m - \hat{\theta}) - \gamma'\hat{\sigma}'\sqrt{\frac{\hat{\theta}'_m(1-\hat{\theta}'_m)}{(1-\lambda)\hat{\theta}(1-\hat{\theta})}} \\ &= \hat{\theta} + (\hat{\theta}'_m - \hat{\theta}) - \frac{\gamma'}{\sqrt{1-\lambda}}\hat{\sigma}'(1 + O(|\hat{\theta} - \hat{\theta}'_m|)). \end{aligned} \quad (45)$$

We have $\hat{\theta}'_m - \hat{\theta} = \lambda(\hat{\theta}'_m - \hat{\theta}_m) = \lambda(\delta' - \delta)$, and we are considering the case

$\lambda|\delta' - \delta| = \kappa\sqrt{\lambda\theta(1-\theta)/(n(1-\lambda))}$, so

$$\varphi' = \hat{\theta} - \frac{\gamma'}{\sqrt{1-\lambda}} \hat{\sigma} (1 + O(1/\sqrt{n})). \quad (46)$$

We can therefore identify

$$\begin{aligned} \gamma_a(\hat{\theta}|P_\lambda) &= \frac{\gamma'}{\sqrt{1-\lambda}} (1 + O(1/\sqrt{n})) \\ &= \frac{\sqrt{2(\alpha + \kappa^2/(2\lambda)) + \tilde{O}(1/\sqrt{n})} + O(\alpha/\sqrt{n})}{\sqrt{1-\lambda}} (1 + O(1/\sqrt{n})) \\ &= \frac{\sqrt{2(\alpha + \kappa^2/(2\lambda))}}{\sqrt{1-\lambda}} + \tilde{O}(\alpha/\sqrt{n}), \end{aligned} \quad (47)$$

which compares as promised to $\gamma_a(\hat{\theta}; P_{CH}) = \sqrt{2\alpha} + O(\alpha/\sqrt{n})$.

6. Conclusion

It is clear that for the specific problem of one-sided hypothesis testing and confidence intervals for Bernoulli RVs, it is always preferable to use the exact test in the ideal case, where the trials are i.i.d. For general nulls, exact tests are typically not available, so approximations are used. The approximations often do not take into account failure of underlying distributional assumptions. The approximation errors can be large at high significance. Thus trustworthy alternatives such as those based on large deviation bounds or test supermartingales are desirable. Our goal here is not to suggest that these alternatives are better for the example of Bernoulli RVs, but to determine the gap between them and an exact test, in a case where an exact test is known and all tests are readily calculable. The suggestion is that for high significance applications, the gaps are relatively small on the relevant logarithmic scale. For p -values, they are within what is expected from experiment-to-experiment variation, even for moderate significances. For confidence intervals, the increase in size is bounded by a constant if the number of trials is known ahead of time, but there is a slowly growing cost with number of trials if we allow for arbitrary stopping-rules.

7. Appendix

7.1 Preliminaries

Notation and definitions are as introduced in the text. The p -value bounds obtained by the three tests investigated are denoted by P_X for the exact, P_{CH} for the Chernoff-Hoeffding, and P_{PBR} for the PBR test. They depend on n , φ and $\hat{\Theta}$. For reference, here are the definitions again.

$$\begin{aligned} P_X(\hat{\Theta}|\varphi, n) &= \sum_{i=n\hat{\Theta}}^n \varphi^i (1-\varphi)^{n-i} \binom{n}{i}, \\ P_{CH}(\hat{\Theta}|\varphi) &= \begin{cases} \left(\frac{\varphi}{\hat{\Theta}}\right)^{n\hat{\Theta}} \left(\frac{1-\varphi}{1-\hat{\Theta}}\right)^{n(1-\hat{\Theta})} & \text{if } \hat{\Theta} \geq \varphi, \\ 1 & \text{otherwise.} \end{cases} \\ P_{PBR}(\hat{\Theta}|\varphi) &= \begin{cases} \varphi^{n\hat{\Theta}} (1-\varphi)^{n(1-\hat{\Theta})} (n+1) \binom{n}{n\hat{\Theta}} & \text{if } \hat{\Theta} \geq \varphi, \\ \hat{\Theta}^{n\hat{\Theta}} (1-\hat{\Theta})^{n(1-\hat{\Theta})} (n+1) \binom{n}{n\hat{\Theta}} & \text{otherwise.} \end{cases} \end{aligned} \quad (48)$$

The gain per trial for a p -value bound P_n is $G_n(P_n) = -\log(P_n)/n$. The values of φ , $\hat{\theta}$ and θ are usually constrained. Unless otherwise stated, we assume that $0 < \varphi, \hat{\theta}, \theta < 1$ and $n \geq 1$.

Most of this appendix is dedicated to obtaining upper and lower bounds on $\log(p)$ -values and lower bounds on endpoints of confidence intervals. We make sure that the upper and lower bounds differ by quantities that converge to zero as n grows. Their differences are $O(1/n)$ for $\log(p)$ -values and $O(1/\sqrt{n})$ for confidence lower bounds. We generally aim for simplicity when expressing these bounds, so we do not obtain tight constants.

7.2 Closed-Form Expression for P_{PBR}

Theorem 5. *Define*

$$\begin{aligned}\tilde{\Theta}_k &= \frac{1}{k+2} (S_k + 1), \\ F_{k+1} &= \left(\frac{\tilde{\Theta}_k}{\varphi}\right)^{B_{k+1}} \left(\frac{1-\tilde{\Theta}_k}{1-\varphi}\right)^{1-B_{k+1}}.\end{aligned}\quad (49)$$

Then

$$\frac{1}{\prod_{k=1}^n F_k} = \varphi^{n\hat{\Theta}} (1-\varphi)^{n(1-\hat{\Theta})} (n+1) \binom{n}{n\hat{\Theta}}.\quad (50)$$

Proof. The proof proceeds by induction. Write P_k for the right-hand side of Eq. 50. For $n = 0$, $P_0 = 1$, and the left-hand side of Eq. 50 evaluates to 1 as required, with the usual convention that the empty product evaluates to 1.

Now suppose that Eq. 50 holds at trial $n = k$. For $n = k + 1$ we can use $(k+1)\hat{\Theta}_{k+1} = S_{k+1} = S_k + B_{k+1}$. We expand the binomial expression to rewrite the right-hand side as

$$\begin{aligned}P_{k+1} &= \varphi^{k\hat{\Theta}_k + B_{k+1}} (1-\varphi)^{k(1-\hat{\Theta}_k) + (1-B_{k+1})} (k+2) \binom{k+1}{k\hat{\Theta}_k + B_{k+1}} \\ &= \varphi^{k\hat{\Theta}_k} (1-\varphi)^{k(1-\hat{\Theta}_k)} (k+1) \binom{k}{k\hat{\Theta}_k} \\ &\quad \cdot \varphi^{B_{k+1}} (1-\varphi)^{1-B_{k+1}} (k+2) (k - k\hat{\Theta}_k + 1)^{-(1-B_{k+1})} (k\hat{\Theta}_k + 1)^{-B_{k+1}}.\end{aligned}\quad (51)$$

Since $\tilde{\Theta}_k = (S_k + 1)/(k+2) = (k\hat{\Theta}_k + 1)/(k+2)$ and $1 - \tilde{\Theta}_k = (k - S_k + 1)/(k+2) = (k - k\hat{\Theta}_k + 1)/(k+2)$, the identity simplifies to

$$P_{k+1} = P_k \cdot \frac{1}{F_{k+1}},\quad (52)$$

thus establishing the induction step. \square

The expression in Eq. 50 can be seen as the inverse of a positive martingale for $\mathcal{H}_0 = \{v_\varphi\}$ determined by S_n . The complete family of such martingales was obtained by Ville [13], Chapter 5, Sec. 3, Eq. 21. Ours is obtained from Ville's with $dF(t) = dt$ as the probability measure.

7.3 Log- p -Value Approximations

We use $-\log(P_{\text{CH},n}(t|\varphi)) = n\text{KL}(t|\varphi)$ as our reference value. According to Theorem 1, the $\log(p)$ -values are ordered according to $-\log(P_{\text{PBR}}) \leq -\log(P_{\text{CH}}) \leq -\log(P_{\text{X}})$. To express the asymptotic

differences between the $\log(p)$ -values, we use auxiliary functions. The first is

$$\begin{aligned} H_n(t) &= -\log\left(t^{nt}(1-t)^{n(1-t)}\binom{n}{nt}\sqrt{n+1}\right) \\ &= -nt\log(t) - n(1-t)\log(1-t) - \log\binom{n}{nt} - \frac{1}{2}\log(n+1). \end{aligned} \quad (53)$$

The first two terms of this expression can be recognized as the Shannon entropy of n independent random bits, each with probability t for bit value 1. For $t \in [1/n, 1 - 1/n]$ and with Stirling's approximation $\sqrt{2\pi n}(n/e)^n e^{1/(12n+1)} \leq n! \leq \sqrt{2\pi n}(n/e)^n e^{1/(12n)}$ applied to the binomial coefficient, we get

$$\begin{aligned} \log\binom{n}{nt} &= \log\left(\frac{n!}{(tn)!(1-t)n!}\right) \\ &\in \log\left(\frac{\sqrt{2\pi n}}{\sqrt{2\pi tn}\sqrt{2\pi(1-t)n}}\right) + \log\left(\frac{(n/e)^n}{(tn/e)^{tn}((1-t)n/e)^{(1-t)n}}\right) \\ &\quad + \left[\frac{1}{12n+1}, \frac{1}{12n}\right] + \left[-\frac{1}{12tn} - \frac{1}{12(1-t)n}, -\frac{1}{12tn+1} - \frac{1}{12(1-t)n+1}\right] \\ &= -\frac{1}{2}\log(2\pi t(1-t)) - \frac{1}{2}\log(n) - tn\log(t) - (1-t)n\log(1-t) \\ &\quad + \left[\frac{1}{12n+1} - \frac{1}{12t(1-t)n}, \frac{1}{12n} - \frac{12n+2}{(12tn+1)(12(1-t)n+1)}\right]. \end{aligned} \quad (54)$$

We can increase the interval to simplify the bounds while preserving convergence for large n . For the lower bound, we use $-1/(12t(1-t)n)$. For the upper bound, note that $(12tn+1)(12(1-t)n+1)$ is maximized at $t = 1/2$. We can therefore increase the upper bound according to

$$\frac{1}{12n} - \frac{12n+2}{(12tn+1)(12(1-t)n+1)} \leq \frac{1}{12n} - \frac{2}{6n+1} \leq 0 \quad (55)$$

for $n \geq 1$. From this we obtain the interval expression

$$H_n(t) \in \frac{1}{2}\log(2\pi t(1-t)) - \frac{1}{2}\log(1+1/n) + \left[0, \frac{1}{12nt(1-t)}\right], \quad (56)$$

valid for $t \in [1/n, 1 - 1/n]$. The boundary values of H_n at $t = 0$ and $t = 1$ are $-\log(n+1)/2$.

The next auxiliary function is

$$Y(t) = \frac{1}{e^{-t^2/2}} \int_t^\infty e^{-s^2/2} ds \in \left(\frac{t}{1+t^2}, \frac{1}{t}\right) \text{ for } t > 0, \quad (57)$$

where the bounds are from Ref. [24]. See this reference for a summary of all properties of Y mentioned here, or Ref. [25] for more details. The function Y is related to the tail of the standard normal distribution, the Q -function, by $Q(t) = e^{-t^2/2}Y(t)/\sqrt{2\pi}$. The function Y is monotonically decreasing, convex, $Y(0) = \sqrt{\pi/2}$, and it satisfies the differential equation $\frac{d}{dt}Y(t) = tY(t) - 1$. We make use of the following bounds involving Y :

$$-\log tY(t) \in \left[0, \frac{1}{t^2}\right]. \quad (58)$$

The lower bound comes from the upper bound $1/t$ for $Y(t)$. The upper bound is from the lower bound $t/(1+t^2)$ for $Y(t)$. Specifically, we compute

$$-\log(Y(t)) \leq -\log(t/(1+t^2)) = \log(t) + \log(1+1/t^2) \leq \log(t) + 1/t^2.$$

With these definitions, we can express the $\log(p)$ -values in terms of their difference from $-\log(P_{CH})$.

Theorem 6. For $0 < \varphi \leq t < 1$,

$$-\log(P_{PBR,n}(t|\varphi)) = -\log(P_{CH,n}(t|\varphi)) - \frac{1}{2} \log(n+1) + H_n(t) \quad (59)$$

$$\begin{aligned} &\in -\log(P_{CH,n}(t|\varphi)) - \frac{1}{2} \log(n+1) + \frac{1}{2} \log(2\pi t(1-t)) - \frac{1}{2} \log(1+1/n) \\ &\quad + \left[0, \frac{1}{12nt(1-t)} \right] \end{aligned} \quad (60)$$

Proof. The theorem is obtained by substituting definitions and then applying the bounds of Eq. 56 on $H_n(t)$. Here are the details.

$$\begin{aligned} -\log(P_{PBR,n}(t|\varphi)) &= -\log\left(\varphi^t(1-\varphi)^{n(1-t)}(n+1)\binom{n}{nt}\right) \\ &= -\log\left(\left(\frac{\varphi}{t}\right)^{nt}\left(\frac{1-\varphi}{1-t}\right)^{n(1-t)}\right) \\ &\quad -\log\left(t^{nt}(1-t)^{n(1-t)}(n+1)\binom{n}{nt}\right) \\ &= -\log(P_{CH,n}(t|\varphi)) - \frac{1}{2} \log(n+1) \\ &\quad -\log\left(t^{nt}(1-t)^{n(1-t)}\sqrt{n+1}\binom{n}{nt}\right) \\ &= -\log(P_{CH,n}(t|\varphi)) - \frac{1}{2} \log(n+1) + H_n(t). \end{aligned} \quad (61)$$

It remains to substitute the interval expression for $H_n(t)$. □

Theorem 7. Define

$$\text{IE}_n(t|\varphi) = \min\left((t-\varphi)\sqrt{\frac{\pi n}{8\varphi(1-\varphi)}}, 1\right). \quad (62)$$

Then for $0 < \varphi < t < 1$,

$$\begin{aligned} -\log(P_{X,n}(t|\varphi)) &\in -\log(P_{PBR,n}(t|\varphi)) + \log(n+1) - \log\left(t\sqrt{\frac{(1-\varphi)}{\varphi}}\right) \\ &\quad -\log\left(\sqrt{n}Y\left(\sqrt{\frac{n}{\varphi(1-\varphi)}}(t-\varphi)\right)\right) + \left[-\frac{\text{IE}_n(t|\varphi)}{n(t-\varphi)}, 0\right], \end{aligned} \quad (63)$$

$$\begin{aligned} -\log(P_{X,n}(t|\varphi)) &\in -\log(P_{CH,n}(t|\varphi)) + \frac{1}{2} \log(n) - \log\left(\sqrt{\frac{t(1-\varphi)}{2\pi(1-t)\varphi}}\right) \\ &\quad -\log\left(\sqrt{n}Y\left(\sqrt{\frac{n}{\varphi(1-\varphi)}}(t-\varphi)\right)\right) + \left[-\frac{\text{IE}_n(t|\varphi)}{n(t-\varphi)}, \frac{1}{12nt(1-t)}\right]. \end{aligned} \quad (64)$$

Observe that $\text{IE}_n(t|\varphi)$ is $O(1)$ with respect to n for $t > \varphi$ constant. The first term in the defining minimum is smaller than 1 only for φ within less than one standard deviation (which is $O(1/\sqrt{n})$) of t . It is defined so that the primary dependence on the parameters is visible in the interval bounds.

Proof. For approximating P_X , we apply Theorem 2 of Ref. [24] with the following sequence of substitutions, the first four of which expand the definitions in the reference:

$$\begin{aligned}
 B(k; n, p) &\leftarrow \sum_{j=k}^n b(j; n, p), \\
 b(k-1; n-1, p) &\leftarrow \binom{n-1}{k-1} p^{k-1} (1-p)^{n-k}, \\
 x &\leftarrow (k-pn)/\sigma, \\
 \sigma &\leftarrow \sqrt{np(1-p)}, \\
 p &\leftarrow \varphi, \\
 k &\leftarrow nt.
 \end{aligned} \tag{65}$$

With the given substitutions and $Y(t)$ as defined by Eq. 57, we obtain for $t \geq \varphi$,

$$\begin{aligned}
 -\log(P_X) &\in -\log\left(\sqrt{np(1-\varphi)}\varphi^{nt-1}(1-\varphi)^{n(1-t)}\binom{n-1}{nt-1}\right) \\
 &\quad -\log\left(Y\left(\frac{\sqrt{n}(t-\varphi)}{\sqrt{\varphi(1-\varphi)}}\right)\right) + \left[-\frac{\mathbb{I}E_n(t|\varphi)}{n(t-\varphi)}, 0\right] \\
 &= -\log\left(\varphi^{nt}(1-\varphi)^{n(1-t)}(n+1)\binom{n}{nt}\right) - \log\left(\frac{nt\sqrt{np(1-\varphi)}}{\varphi n(n+1)}\right) \\
 &\quad -\log\left(Y\left(\frac{\sqrt{n}(t-\varphi)}{\sqrt{\varphi(1-\varphi)}}\right)\right) + \left[-\frac{\mathbb{I}E_n(t|\varphi)}{n(t-\varphi)}, 0\right] \\
 &= -\log(P_{\text{PBR}}) + \log(n+1) - \log\left(t\sqrt{\frac{(1-\varphi)}{\varphi}}\right) \\
 &\quad -\log\left(\sqrt{n}Y\left(\frac{\sqrt{n}(t-\varphi)}{\sqrt{\varphi(1-\varphi)}}\right)\right) + \left[-\frac{\mathbb{I}E_n(t|\varphi)}{n(t-\varphi)}, 0\right].
 \end{aligned} \tag{66}$$

The second identity of the theorem follows by substituting the expression from Theorem 6. □

We can eliminate the function Y from the expressions by applying the bounds from Eq. 58.

Corollary 8. *With the assumptions of Theorem 7,*

$$\begin{aligned}
 -\log(P_{X,n}(t|\varphi)) &\in -\log(P_{\text{CH},n}(t|\varphi)) + \frac{1}{2}\log(n) - \log\left(\frac{1-\varphi}{t-\varphi}\sqrt{\frac{t}{2\pi(1-t)}}\right) \\
 &\quad + \left[-\frac{\mathbb{I}E_n(t|\varphi)}{n(t-\varphi)}, \frac{\varphi(1-\varphi)}{(t-\varphi)^2n} + \frac{1}{12nt(1-t)}\right].
 \end{aligned} \tag{67}$$

Proof. Define $c = (t - \varphi) / \sqrt{\varphi(1 - \varphi)}$. In view of Eq. 58, we have

$$\begin{aligned} -\log\left(\sqrt{n}Y\left(\sqrt{\frac{n}{\varphi(1-\varphi)}}(t-\varphi)\right)\right) &= -\log(\sqrt{n}Y(c\sqrt{n})) \\ &= \log(c) - \log(c\sqrt{n}Y(c\sqrt{n})) \\ &\in \log(c) + \left[0, \frac{1}{c^2n}\right]. \end{aligned} \quad (68)$$

Substituting in Eq. 64 and simplifying the expression gives the desired result. \square

7.4 Asymptotic Normality of the $\log(p)$ -Values and Their Differences

The main tool for establishing the asymptotic distribution of the $\log(p)$ -values is the “delta method”. A version sufficient for our purposes is Theorem 1.12 and Corollary 1.1 of Ref. [18]. The notation $X_n \xrightarrow{D} N(\mu, \sigma^2)$ means that X_n converges in distribution to the normal distribution with mean μ and variance σ^2 . By the central limit theorem, $\hat{\Theta}_n = S_n/n$ satisfies $\sqrt{n}(\hat{\Theta}_n - \theta) \xrightarrow{D} N(0, \theta(1 - \theta))$. An application of the delta method therefore yields the next lemma.

Lemma 9. *Let $F : \mathbb{R} \rightarrow \mathbb{R}$ be differentiable at θ , with $F'(\theta) \neq 0$. Then*

$$\sqrt{n}(F(\hat{\Theta}_n) - F(\theta)) \xrightarrow{D} N(0, F'(\theta)^2 \theta(1 - \theta))$$

Theorem 10. *For $P = P_{\text{CH}}$, $P = P_{\text{PBR}}$ or $P = P_X$, and $0 < \varphi < \theta < 1$ constant, the gain per trial $G_n(P)$ converges in distribution according to*

$$\sqrt{n}(G_n(P) - \text{KL}(\theta|\varphi)) \xrightarrow{D} N(0, \sigma_G^2), \quad (69)$$

with

$$\sigma_G^2 = \theta(1 - \theta) \left(\log\left(\frac{\theta}{1 - \theta} \frac{1 - \varphi}{\varphi}\right) \right)^2.$$

Proof. Consider $P = P_{\text{CH}}$ first. In Lemma 9, define

$F(x) = \text{KL}(x|\varphi) = x \log(x/\varphi) + (1 - x) \log((1 - x)/(1 - \varphi))$ so that $F(\hat{\Theta}_n) = G_n(P_{\text{CH}})$. For the derivative of F at $x = \theta$, we get

$$F'(\theta) = \log\left(\frac{\theta}{1 - \theta} \frac{1 - \varphi}{\varphi}\right). \quad (70)$$

The theorem now follows for P_{CH} by applying Lemma 9.

Theorem 6 and the law of large numbers imply that $(-\log(P_{\text{PBR}})/\sqrt{n}) - (-\log(P_{\text{CH}})/\sqrt{n})$ converges in probability to 0. Corollary 8 implies the same for P_X , namely that $(-\log(P_X)/\sqrt{n}) - (-\log(P_{\text{CH}})/\sqrt{n})$ converges in probability to 0. In general, if $X_n - Y_n$ converges in probability to 0 and $Y_n \xrightarrow{D} \mu$, then $X_n \xrightarrow{D} \mu$, see Ref. [26], Proposition 6.3.3. The statement of the theorem to be proven now follows for $P = P_{\text{PBR}}$ and $P = P_X$ by comparison of $\sqrt{n}G_n(P_{\text{PBR}})$ and $\sqrt{n}G_n(P_X)$ to $\sqrt{n}G_n(P_{\text{CH}})$. \square

The differences of the $\log(p)$ -values have much tighter distributions. They are also asymptotically normal with scaling and variances given in the next theorem. The differences are $\Omega(\log(n))$ with standard deviations $O(1/\sqrt{n})$.

Theorem 11. Let $0 < \varphi < \theta < 1$ be constant. If $\theta \neq 1/2$, then $P_{\text{PBR}}/(\sqrt{n}P_{\text{CH}})$ satisfies

$$-\sqrt{n} \log \left(\frac{\sqrt{2\pi\theta(1-\theta)}P_{\text{PBR}}}{\sqrt{n}P_{\text{CH}}} \right) \xrightarrow{D} N \left(0, \frac{(1-2\theta)^2}{4\theta(1-\theta)} \right). \quad (71)$$

If $\varphi \neq \theta(2\theta - 1)$, then $\sqrt{n}P_{\text{X}}/P_{\text{CH}}$ satisfies

$$-\sqrt{n} \log \left(\frac{\theta - \varphi}{1 - \varphi} \sqrt{\frac{2\pi(1-\theta)}{\theta}} \frac{\sqrt{n}P_{\text{X}}}{P_{\text{CH}}} \right) \xrightarrow{D} N \left(0, \frac{(\theta(1-2\theta) + \varphi)^2}{4(\theta - \varphi)^2\theta(1-\theta)} \right), \quad (72)$$

Proof. From Theorem 6, Eq. 60 and the law of large numbers, we see that

$$\sqrt{n} \left(-\log \left(\frac{P_{\text{PBR}}}{\sqrt{n}P_{\text{CH}}} \right) - \log \left(\sqrt{2\pi\hat{\theta}(1-\hat{\theta})} \right) \right) \quad (73)$$

converges in probability to zero. From Lemma 9 and

$$\frac{d}{dx} \log(x(1-x))/2 = \frac{1}{2x} - \frac{1}{2(1-x)} = \frac{1-2x}{2x(1-x)}, \quad (74)$$

we conclude

$$\sqrt{n} \left(\log \left(\sqrt{2\pi\hat{\theta}(1-\hat{\theta})} \right) - \log \left(\sqrt{2\pi\theta(1-\theta)} \right) \right) \xrightarrow{D} N \left(0, \left(\frac{1-2\theta}{2\theta(1-\theta)} \right)^2 \theta(1-\theta) \right). \quad (75)$$

Combining the above observations gives Eq. 71.

Similarly, from Corollary 8 and taking note of the definition of $\text{IE}_n(t|\varphi)$,

$$\sqrt{n} \left(-\log \left(\frac{\sqrt{n}P_{\text{X}}}{P_{\text{CH}}} \right) - \log \left(\frac{\hat{\theta} - \varphi}{1 - \varphi} \sqrt{\frac{2\pi(1-\hat{\theta})}{\hat{\theta}}} \right) \right) \quad (76)$$

converges in probability to zero. The relevant derivative is

$$\frac{d}{dx} (\log(x - \varphi) + \log((1-x)/x))/2 = \frac{1}{x - \varphi} - \frac{1}{2(1-x)} - \frac{1}{2x} = \frac{x(1-2x) + \varphi}{2(x - \varphi)x(1-x)}, \quad (77)$$

from which

$$\begin{aligned} \sqrt{n} \left(\log \left(\frac{\hat{\theta} - \varphi}{1 - \varphi} \sqrt{\frac{2\pi(1-\hat{\theta})}{\hat{\theta}}} \right) - \log \left(\frac{\theta - \varphi}{1 - \varphi} \sqrt{\frac{2\pi(1-\hat{\theta})}{\hat{\theta}}} \right) \right) \\ \xrightarrow{D} N \left(0, \left(\frac{\theta(1-2\theta) + \varphi}{2(\theta - \varphi)\theta(1-\theta)} \right)^2 \theta(1-\theta) \right), \end{aligned} \quad (78)$$

and combining the two observations gives Eq. 72. □

7.5 Confidence Interval Endpoints

For the one-sided confidence intervals, we need to determine the lower boundaries of acceptance regions, that is the confidence lower bounds. By monotonicity of the p -values in φ , it suffices to solve

equations of the form $-\log(P(\hat{\theta}, \varphi)) = \alpha$, where $a = e^{-\alpha}$ is the desired significance level. Here we obtain lower and upper bounds on the solutions φ .

To illuminate the asymptotic behavior of solutions φ of $-\log(P(\hat{\theta}, \varphi)) = \alpha$, we reparametrize the log- p -values so that our scale is set by an empirical standard deviation, namely $\hat{\sigma} = \sqrt{\hat{\theta}(1-\hat{\theta})/n}$. Thus we express the solution as

$$\varphi(\gamma, \hat{\theta}) = \hat{\theta} - \hat{\sigma}\gamma, \quad (79)$$

in terms of a scaled deviation down from $\hat{\theta}$. Inverting for γ we get

$$\gamma = \gamma(\varphi, \hat{\theta}) = \frac{\hat{\theta} - \varphi}{\hat{\sigma}}. \quad (80)$$

Theorem 12. *Let $0 < \hat{\theta} < 1$ and $\alpha > 0$. Suppose that $\alpha \leq n\hat{\theta}^2(1-\hat{\theta})^2/8$. Then there is a solution $\gamma_\alpha > 0$ of the identity $-\log(P_{\text{CH}}(\hat{\theta}, \varphi(\gamma_\alpha, \hat{\theta}))) = \alpha$ satisfying*

$$\gamma_\alpha \in \sqrt{2\alpha} \left(1 + \frac{5}{2} \frac{\sqrt{\alpha}}{\sqrt{n\hat{\theta}(1-\hat{\theta})}} [-1, 1] \right)^{-1/2}. \quad (81)$$

The constants in this theorem and elsewhere are chosen for convenience, not for optimality; better constants can be extracted from the proofs. Note that the upper bound on α ensures that the reciprocal square root is bounded away from zero. However, for the relative error to go to zero as n grows requires $\alpha = o(n)$.

Proof. Consider the parametrized bound $\alpha \leq 2n\hat{\theta}^2(1-\hat{\theta})^2(1-a_1)^2$, where later we set $a_1 = 3/4$ to match the theorem statement. Let $F(\gamma) = -\log(P_{\text{CH}}(\hat{\theta}, \varphi(\gamma, \hat{\theta})))$. F is continuous and monotone increasing. A standard simplification of the Chernoff-Hoeffding bound noted in Ref. [17] is

$$P_{\text{CH}} \leq e^{-2n(\hat{\theta}-\varphi)^2} = e^{-2\hat{\theta}(1-\hat{\theta})\gamma^2}. \quad (82)$$

For $\varphi = \varphi(\gamma_\alpha, \hat{\theta})$ solving the desired equation, we have $(\hat{\theta} - \varphi) \leq \sqrt{\alpha/(2n)}$ (by monotonicity), which in turn is bounded above according to $\sqrt{\alpha/2n} \leq \hat{\theta}(1-\hat{\theta})(1-a_1) \leq \hat{\theta}(1-a_1)$, according to our assumed bound. We conclude that $\varphi \geq a_1\hat{\theta}$. For the solution γ_α , we get

$$\gamma_\alpha \leq \sqrt{\alpha/(2\hat{\theta}(1-\hat{\theta}))} \leq \sqrt{n\hat{\theta}(1-\hat{\theta})(1-a_1)}.$$

We now Taylor expand $\text{KL}(\hat{\theta}|\varphi)$ with remainder at third order around $\varphi = \hat{\theta}$. Write $f(x) = \text{KL}(\hat{\theta}|\hat{\theta} - x)$, where we can restrict x according to $\hat{\theta} \geq \hat{\theta} - x = \varphi \geq a_1\hat{\theta}$. The derivatives of f can be written explicitly as follows:

$$f^{(k)}(x) = (k-1)! \frac{\hat{\theta}}{(\hat{\theta}-x)^k} - (-1)^{k-1} (k-1)! \frac{1-\hat{\theta}}{(1-\hat{\theta}+x)^k}. \quad (83)$$

We have

$$\begin{aligned}
 f^{(1)}(0) &= 0, \\
 f^{(2)}(0) &= \frac{1}{\hat{\theta}} + \frac{1}{1-\hat{\theta}} = \frac{1}{\hat{\theta}(1-\hat{\theta})}, \\
 f^{(3)}(x) &= 2\frac{\hat{\theta}}{(\hat{\theta}-x)^3} - 2\frac{1-\hat{\theta}}{(1-\hat{\theta}+x)^3}, \\
 f^{(3)}(x) &\leq 2\frac{\hat{\theta}}{(\hat{\theta}-x)^3} \leq 2\frac{\hat{\theta}}{a_1^3\hat{\theta}^3} = 2\frac{1}{a_1^3\hat{\theta}^2}, \\
 f^{(3)}(x) &\geq -2\frac{1-\hat{\theta}}{(1-\hat{\theta}+x)^3} \geq -2\frac{1-\hat{\theta}}{(1-\hat{\theta})^3} = -2\frac{1}{(1-\hat{\theta})^2},
 \end{aligned} \tag{84}$$

since $0 < a_1 < 1$. We use the bounds on $f^{(3)}(x)$ to bound the remainder in the Taylor expansion, where, to get cleaner expressions, we can decrease $\hat{\theta}$ and $1-\hat{\theta}$ to $\hat{\theta}(1-\hat{\theta})$ in the denominators.

$$\text{KL}(\hat{\theta}|\hat{\theta}-x) \in \frac{x^2}{2\hat{\theta}(1-\hat{\theta})} + \frac{x^3}{3(\hat{\theta}(1-\hat{\theta}))^2} \left[-1, \frac{1}{a_1^3} \right]. \tag{85}$$

Substituting $x = \gamma_\alpha \sqrt{\hat{\theta}(1-\hat{\theta})/n}$ gives

$$\alpha = -\log(P_{\text{CH}}(\hat{\theta}, \varphi(\gamma_\alpha, \hat{\theta}))) = n\text{KL}(\hat{\theta}|\hat{\theta}-x) \in \frac{\gamma_\alpha^2}{2} \left(1 + \frac{2\gamma_\alpha}{3\sqrt{n\hat{\theta}(1-\hat{\theta})}} \left[-1, \frac{1}{a_1^3} \right] \right). \tag{86}$$

For $\hat{\theta} \leq 1/2$, $f^{(4)}(x)$ and $f^{(3)}(0)$ are non-negative, so we could have taken the lower bound in the interval to be zero for $\theta \leq 1/2$. For the theorem, we prefer not to separate the cases.

We substitute the bound $\gamma \leq \sqrt{n\hat{\theta}(1-\hat{\theta})}(1-a_1)$ for the γ multiplying the interval in Eq. 86 and use the lower bound in the interval for the inequality

$$\alpha \geq \frac{\gamma^2}{2} \left(1 - \frac{2(1-a_1)}{3} \right). \tag{87}$$

For the theorem, we have $a_1 = 3/4$, so $1 - 2(1-a_1)/3 = 5/6$. Inverting the inequality for γ gives $\gamma \leq 2\sqrt{3/5}\sqrt{\alpha}$. Now substituting this bound on γ for the γ multiplying the interval in Eq. 86 gives

$$\alpha \in \frac{\gamma^2}{2} \left(1 + \frac{4}{\sqrt{15}} \frac{\sqrt{\alpha}}{\sqrt{n\hat{\theta}(1-\hat{\theta})}} \left[-1, \frac{4^3}{3^3} \right] \right). \tag{88}$$

By monotonicity of the appropriate operations,

$$\gamma \in \sqrt{2\alpha} \left(1 + \frac{4}{\sqrt{15}} \frac{\sqrt{\alpha}}{\sqrt{n\hat{\theta}(1-\hat{\theta})}} \left[-1, \frac{4^3}{3^3} \right] \right)^{-1/2}. \tag{89}$$

For the theorem statement, we simplify the bounds with $1 \leq 4^3/3^3$ and $4^4/(3^3\sqrt{15}) \leq 5/2$. □

Theorem 13. Let $0 < \hat{\theta} < 1$ and $\alpha > 0$. Define $\Delta = \log(n+1)/2 - H_n(\hat{\theta})$. Suppose that $\alpha + \Delta \leq n\hat{\theta}^2(1 - \hat{\theta})^2/8$. Then there is a solution $\gamma_\alpha > 0$ of the identity $-\log(P_{\text{PBR}}(\hat{\theta}, \varphi(\gamma_\alpha, \hat{\theta}))) = \alpha$ satisfying

$$\gamma_\alpha \in \sqrt{2(\alpha + \Delta)} \left(1 + \frac{5}{2} \frac{\sqrt{\alpha + \Delta}}{\sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1] \right)^{-1/2}.$$

Proof. By Theorem 6, $-\log(P_{\text{CH}}) - (-\log(P_{\text{PBR}})) = \Delta$. If we define $\tilde{\alpha} = \alpha + \Delta$, then solving $-\log(P_{\text{PBR}}) = \alpha$ is equivalent to solving $-\log(P_{\text{CH}}) = \tilde{\alpha}$. Since Δ depends only on n and $\hat{\theta}$, $\tilde{\alpha}$ does not depend on γ . We can therefore apply Theorem 12 to get the desired bounds. \square

Theorem 14. For $x \geq 0$, let $q(x) = -\log(e^{-x^2/2}Y(x)/\sqrt{2\pi}) = x^2/2 + \log(2\pi)/2 - \log(Y(x))$. Suppose that $0 < \hat{\theta} < 1$, and $\log(2) < \alpha \leq n\hat{\theta}^2(1 - \hat{\theta})^2/8$. Then there is a solution γ_α of the identity $-\log(P_X(\hat{\theta}, \varphi(\gamma_\alpha, \hat{\theta}))) = \alpha$ satisfying

$$\begin{aligned} \gamma_\alpha \in \max & \left(0, q^{-1} \left(\alpha \left(1 + \frac{64\sqrt{\alpha}/(15\sqrt{15})}{\sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1] \right) + \frac{\sqrt{\pi/6} + 8\sqrt{\alpha}/\sqrt{15}}{\sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1] \right) \right) \\ & \times \left(1 + \frac{2\sqrt{\alpha}/\sqrt{5}}{\sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1] \right), \end{aligned} \quad (90)$$

where we extend q^{-1} to negative values by $q^{-1}(y) = -\infty$ for $y \leq 0$ (if necessary) when evaluating this interval expression.

The function $q(x)$ is the negative logarithm of the Q -function, which is the tail of the standard normal distribution. The lower bound on α in Theorem 14 ensures that there is a solution with $\gamma_\alpha > 0$, because $q(0) = \log(2)$. For reference, the constants multiplying the interval expressions are $64/(15\sqrt{15}) \approx 1.102$, $8/\sqrt{15} \approx 2.066$, $\sqrt{\pi/6} \approx 0.724$, $2/\sqrt{5} \approx 0.894$. Note that in the large n limit, where the $O(1/\sqrt{n})$ terms are negligible, the value of γ_α in Theorem 14 corresponds to the $(1 - e^{-\alpha})$ -quantile of the standard normal.

By monotonicity of q^{-1} , the explicit bounds in Eq. 90 are obtained by combining the lower or the upper bounds in intervals in the expression. We remark that q^{-1} behaves well with respect to relative error for α large enough because of the inequalities

$$\begin{aligned} q^{-1}(y)/(1 + q^{-1}(y)^2) &\leq \frac{d}{dy} q^{-1}(y) \leq 1/q^{-1}(y), \\ q^{-1}(y)^2 &\geq y - q(1) + 1, && \text{for } y \geq q(1) \approx 1.841, \\ q^{-1}(y)^2 &\leq 2(y - \log(2)), && \text{for } y \geq q(0) = \log(2), \end{aligned} \quad (91)$$

which we now establish. By implicit differentiation and from the properties of Y noted after Eq. 57, $\frac{d}{dy} q^{-1}(y)|_{y=q(x)} = Y(x) \in [x/(1+x^2), 1/x]$. Therefore $q^{-1}(y)/(1 + q^{-1}(y)^2) \leq \frac{d}{dy} q^{-1}(y) \leq 1/q^{-1}(y)$. For $y \geq \log(2)$, we can integrate $\frac{d}{dz} q^{-1}(z)^2 = 2q^{-1}(z) \frac{d}{dz} q^{-1}(z) \leq 2$ from $z = \log(2)$ to y to show that $q^{-1}(y)^2 = q^{-1}(y)^2 - q^{-1}(\log(2))^2 \leq 2(y - \log(2))$, making use of the identity $q^{-1}(\log(2)) = 0$. Consider $y, z \geq q(1)$. Since $q^{-1}(z)$ and $0 \leq x \mapsto x^2/(1+x^2)$ are monotone increasing, $q^{-1}(z)^2/(1 + q^{-1}(z)^2) \geq q^{-1}(q(1))^2/(1 + q^{-1}(q(1))^2) = 1/2$, so the integral of $\frac{d}{dz} q^{-1}(z)^2$ from $z = q(1)$ to y with the lower bound on $\frac{d}{dz} q^{-1}(z)$ gives $q^{-1}(y)^2 - q^{-1}(q(1))^2 = q^{-1}(y)^2 - 1 \geq y - q(1)$.

From the inequality $\frac{d}{dy}q^{-1}(y) \leq 1/q^{-1}(y)$ in Eq. 91, integration and monotonicity, for $0 \leq z \leq \delta$,

$$\begin{aligned} q^{-1}(\alpha - z) &\geq q^{-1}(\alpha) - \frac{z}{q^{-1}(\alpha - \delta)} \leq q^{-1}(\alpha) \left(1 - \frac{z}{q^{-1}(\alpha - \delta)^2}\right), \\ q^{-1}(\alpha + z) &\leq q^{-1}(\alpha) + \frac{z}{q^{-1}(\alpha - \delta)} \geq q^{-1}(\alpha) \left(1 + \frac{z}{q^{-1}(\alpha - \delta)^2}\right). \end{aligned} \quad (92)$$

To determine the relative error, write $\delta' = \delta/\alpha$ to obtain the interval inclusion

$$q^{-1}(\alpha(1 + \delta'[-1, 1])) \subseteq q^{-1}(\alpha) \left(1 + \frac{\alpha\delta'}{q^{-1}(\alpha(1 - \delta'))^2}[-1, 1]\right). \quad (93)$$

For $\alpha(1 - \delta') > q(1)$, the interval relationship can be weakened to

$$q^{-1}(\alpha(1 + \delta'[-1, 1])) \subseteq q^{-1}(\alpha) \left(1 + \frac{\alpha\delta'}{\alpha(1 - \delta') - q(1) + 1}[-1, 1]\right). \quad (94)$$

The relative error on the right-hand side is given by the term multiplying the interval, and can be written as $\alpha\delta'/(\alpha - (\alpha\delta' + q(1) - 1))$. If $\alpha\delta' + q(1) - 1 \leq \alpha/2$, then the relative error is bounded by $2\delta'$ which is twice the relative error of α . Of course, for the interval bounds to converge, we need $\alpha = o(n)$.

Proof. As in the proof of Theorem 12, consider the parametrized bound $\alpha \leq 2n\hat{\theta}^2(1 - \hat{\theta})^2(1 - a_1)^2$, where later we set $a_1 = 3/4$ to match the statement of Theorem 14. From the Chernoff-Hoeffding bound, we get

$$\varphi \geq a_1\hat{\theta} \text{ and } \gamma_\alpha \leq \sqrt{\alpha/(2\hat{\theta}(1 - \hat{\theta}))} \leq \sqrt{n\hat{\theta}(1 - \hat{\theta})(1 - a_1)}.$$

Define $\tilde{\gamma} = (\hat{\theta} - \varphi)/\sqrt{\varphi(1 - \varphi)/n}$. We start from Eq. 64, rewritten as follows:

$$\begin{aligned} -\log(P_X) &\in n\text{KL}(\hat{\theta}|\varphi) + \frac{1}{2}\log(2\pi) - \log Y(\tilde{\gamma}) - \frac{1}{2}\log\left(\frac{\hat{\theta}(1 - \varphi)}{(1 - \hat{\theta})\varphi}\right) \\ &\quad + \left[-\frac{\text{IE}_n(\hat{\theta}|\varphi)}{n(\hat{\theta} - \varphi)}, \frac{1}{12n\hat{\theta}(1 - \hat{\theta})} \right]. \end{aligned} \quad (95)$$

If $\tilde{\gamma} \geq \sqrt{8/\pi} \approx 1.6$, $\text{IE}_n(\hat{\theta}|\varphi) = 1$. For better bounds at small values of $\tilde{\gamma}$, we use the other alternative in the definition of IE_n , according to which the lower bound in the last interval of Eq. 95 is

$$-\frac{\text{IE}_n(\hat{\theta}|\varphi)}{n(\hat{\theta} - \varphi)} \geq -\frac{\sqrt{\pi/8}}{\sqrt{n\varphi(1 - \varphi)}} \geq -\frac{\sqrt{\pi/8}}{\sqrt{na_1\hat{\theta}(1 - \varphi)}} \geq -\frac{\sqrt{\pi/8}}{\sqrt{na_1\hat{\theta}(1 - \hat{\theta})}}. \quad (96)$$

Next we approximate $n\text{KL}(\hat{\theta}|\varphi)$ in terms of $\tilde{\gamma}$ instead of γ . We still write the interval bounds in terms of

γ . Let $f(x) = \text{KL}(\varphi + x|\varphi)$. We are concerned with the range $0 \leq x \leq \hat{\theta} - \varphi$, with $\varphi \geq a_1 \hat{\theta}$. We have

$$\begin{aligned} f^{(1)}(x) &= \log((\varphi + x)/\varphi) - \log((1 - \varphi - x)/(1 - \varphi)) \\ f^{(2)}(x) &= \frac{1}{\varphi + x} + \frac{1}{1 - \varphi - x} \\ &= \frac{1}{(\varphi + x)(1 - \varphi - x)} \\ f^{(3)}(x) &= -\frac{1}{(\varphi + x)^2} + \frac{1}{(1 - \varphi - x)^2} \\ &= -\frac{1 - 2(\varphi + x)}{(\varphi + x)^2(1 - \varphi - x)^2} \\ |f^{(3)}(x)| &\leq \frac{1}{a_1^2 \hat{\theta}^2 (1 - \hat{\theta})^2}, \end{aligned} \tag{97}$$

yielding

$$\text{KL}(\varphi + x|\varphi) \in \frac{x^2}{2\varphi(1 - \varphi)} + \frac{x^3}{6a_1^2 \hat{\theta}^2 (1 - \hat{\theta})^2} [-1, 1], \tag{98}$$

and with $x = \tilde{\gamma} \sqrt{\varphi(1 - \varphi)/n} = \gamma \sqrt{\hat{\theta}(1 - \hat{\theta})/n}$,

$$n\text{KL}(\hat{\theta}|\varphi) \in \frac{\tilde{\gamma}^2}{2} + \frac{\gamma^3}{6a_1^2 \sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1]. \tag{99}$$

For the fourth term on the right-hand side of Eq. 95,

$$\frac{d}{dx} \log \left(\frac{\hat{\theta}(1 - \hat{\theta} + x)}{(1 - \hat{\theta})(\hat{\theta} - x)} \right) = \frac{1}{1 - \hat{\theta} + x} + \frac{1}{\hat{\theta} - x} = \frac{1}{(1 - \hat{\theta} + x)(\hat{\theta} - x)}, \tag{100}$$

whose absolute value is bounded by $1/(a_1 \hat{\theta}(1 - \hat{\theta}))$ for x in the given range. Thus

$$\log \left(\frac{\hat{\theta}(1 - \varphi)}{(1 - \hat{\theta})\varphi} \right) \in \frac{\gamma}{a_1 \sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1]. \tag{101}$$

Since $P_X \leq P_{\text{CH}}$, we can also use the bound $\gamma \leq 2\sqrt{3/5}\sqrt{\alpha}$ obtained in the proof of Theorem 12. Substituting $a_1 = 3/4$ as needed, the equation to solve is now

$$\begin{aligned} \alpha \in & \frac{\tilde{\gamma}^2}{2} + \frac{1}{2} \log(2\pi) - \log Y(\tilde{\gamma}) \\ & + \frac{8}{\sqrt{15}} \frac{\sqrt{\alpha}}{\sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1] + \frac{64}{15\sqrt{15}} \frac{\sqrt{\alpha}^3}{\sqrt{n\hat{\theta}(1 - \hat{\theta})}} [-1, 1] \\ & + \left[-\frac{\sqrt{\pi/6}}{\sqrt{n\hat{\theta}(1 - \hat{\theta})}}, \frac{1}{12n\hat{\theta}(1 - \hat{\theta})} \right]. \end{aligned} \tag{102}$$

The sum of the first three terms evaluates to $q(\tilde{\gamma})$. The remaining terms are now independent of γ and are of

order $1/\sqrt{n}$. They can be merged by means of common bounds using $2n\hat{\theta}(1-\hat{\theta}) \geq \sqrt{n\hat{\theta}(1-\hat{\theta})}$, since $n\hat{\theta}(1-\hat{\theta}) \geq 1/2$ for our standing assumptions that $n \geq 1$ and $\hat{\theta}n$ is an integer different from 0 and n . Consequently, $12n\hat{\theta}(1-\hat{\theta}) \geq 6\sqrt{n\hat{\theta}(1-\hat{\theta})} \geq \sqrt{6/\pi}\sqrt{n\hat{\theta}(1-\hat{\theta})}$. The interval bounds then combine conservatively to

$$\frac{\sqrt{\pi/6} + 8\sqrt{\alpha}/\sqrt{15} + 64\sqrt{\alpha^3}/(15\sqrt{15})}{\sqrt{n\hat{\theta}(1-\hat{\theta})}}. \quad (103)$$

We can now write

$$\alpha \in q(\tilde{\gamma}) + \frac{\sqrt{\pi/6} + 8\sqrt{\alpha}/\sqrt{15} + 64\sqrt{\alpha^3}/(15\sqrt{15})}{\sqrt{n\hat{\theta}(1-\hat{\theta})}}[-1, 1], \quad (104)$$

which holds iff

$$q(\tilde{\gamma}) \in \alpha \left(1 + \frac{64\sqrt{\alpha}/(15\sqrt{15})}{\sqrt{n\hat{\theta}(1-\hat{\theta})}}[-1, 1] \right) + \frac{\sqrt{\pi/6} + 8\sqrt{\alpha}/\sqrt{15}}{\sqrt{n\hat{\theta}(1-\hat{\theta})}}[-1, 1]. \quad (105)$$

By monotonicity of q and extending q^{-1} to negative arguments as mentioned in the statement of Theorem 14 if necessary, the constraint is equivalent to

$$\tilde{\gamma} \in q^{-1} \left(\alpha \left(1 + \frac{64\sqrt{\alpha}/(15\sqrt{15})}{\sqrt{n\hat{\theta}(1-\hat{\theta})}}[-1, 1] \right) + \frac{\sqrt{\pi/6} + 8\sqrt{\alpha}/\sqrt{15}}{\sqrt{n\hat{\theta}(1-\hat{\theta})}}[-1, 1] \right). \quad (106)$$

For $\alpha > \log(2)$, we know that $\tilde{\gamma} > 0$, so we can add $\max(0, \dots)$ as in the theorem statement.

To determine the interval equation for γ , we have $\gamma = \tilde{\gamma}\sqrt{\varphi(1-\varphi)/(\hat{\theta}(1-\hat{\theta}))}$. We use the first-order remainder to bound the factor on the right-hand side. For this consider the numerator, and write

$g(x) = \sqrt{(\hat{\theta}-x)(1-\hat{\theta}+x)}$ with $0 \leq x \leq \hat{\theta} - \varphi$. We have

$$g^{(1)}(x) = \frac{2(\hat{\theta}-x)-1}{2\sqrt{(\hat{\theta}-x)(1-\hat{\theta}+x)}}, \quad (107)$$

$$\begin{aligned} |g^{(1)}(x)| &\leq \frac{1}{2\sqrt{a_1\hat{\theta}(1-\hat{\theta})}} \\ &= \frac{1}{\sqrt{3\hat{\theta}(1-\hat{\theta})}}, \end{aligned} \quad (108)$$

$$g(x) \in \sqrt{\hat{\theta}(1-\hat{\theta})} + \frac{x}{\sqrt{3\hat{\theta}(1-\hat{\theta})}}[-1, 1]. \quad (109)$$

With $x = \gamma\sqrt{\hat{\theta}(1-\hat{\theta})}/n$ and the bound of $\gamma \leq 2\sqrt{3/5}\sqrt{\alpha}$, we get

$$\gamma \in \tilde{\gamma} \left(1 + \frac{2\sqrt{\alpha}/\sqrt{5}}{\sqrt{n\hat{\theta}(1-\hat{\theta})}}[-1, 1] \right). \quad (110)$$

The theorem follows by composing this constraint with Eq. 106. □

Acknowledgments

Y. Z. acknowledges support from the Ontario Research Fund (ORF), the Natural Sciences and Engineering Research Council of Canada (NSERC), and Industry Canada.

8. References

- [1] Dorigo T (2015) Extraordinary claims: the 0.000029% solution. *EPJ Web of Conferences* 95:02003. <https://doi.org/10.1051/epjconf/20159502003>
- [2] Genovese M (2005) Research on hidden variable theories: A review of recent progresses. *Physics Reports* 413(6):319–396. <https://doi.org/10.1016/j.physrep.2005.03.003>
- [3] Giustina M, Versteegh MAM, Wengerowsky S, Handsteiner J, Hochrainer A, Phelan K, Steinlechner F, Kofler J, Larsson J-A, Abellán C, Amaya W, Pruneri V, Mitchell MW, Beyer J, Gerrits T, Lita AE, Shalm LK, Nam SW, Scheidl T, Ursin R, Wittmann B, Zeilinger A (2015) Significant-loophole-free test of Bell's theorem with entangled photons. *Physical Review Letters* 115:250401. <https://doi.org/10.1103/PhysRevLett.115.250401>
- [4] Hensen B (2015) Loophole-free Bell inequality violation using electron spins separated by 1.3 km. *Nature* 526:682. <https://doi.org/10.1038/nature15759>
- [5] Rosenfeld W, Burchardt D, Garthoff R, Redeker K, Ortel N, Rau M, Weinfurter H (2017) Event-ready Bell-test using entangled atoms simultaneously closing detection and locality loopholes. *Physical Review Letters* 119:010402. <https://doi.org/10.1103/PhysRevLett.119.010402>
- [6] Shalm LK, Meyer-Scott E, Christensen BG, Bierhorst P, Wayne MA, Stevens MJ, Gerrits T, Glancy S, Hamel DR, Allman MS, Coakley KJ, Dyer SD, Hodge C, Lita AE, Verma VB, Lambrocco C, Tortorici E, Migdall AL, Zhang Y, Kumor DR, Farr WH, Marsili F, Shaw MD, Stern JA, Abellán C, Amaya W, Pruneri V, Jennewein T, Mitchell MW, Kwiat PG, Bienfang JC, Mirin RP, Knill E, Nam SW (2015) Strong loophole-free test of local realism *Physical Review Letters* 115:250402 <https://doi.org/10.1103/PhysRevLett.115.250402>
- [7] Acín A, Masanes L (2016) Certified randomness in quantum physics. *Nature* 540:213–219. <https://doi.org/10.1038/nature20119>
- [8] Larsson JA (2014) Loopholes in Bell inequality tests of local realism. *Journal of Physica A* 47:424003. <https://doi.org/10.1088/1751-8113/47/42/424003>
- [9] Zhang Y, Glancy S, Knill E (2011) Asymptotically optimal data analysis for rejecting local realism. *Physical Review A* 84:062118. <https://doi.org/10.1103/PhysRevA.84.062118>
- [10] Gill RD (2003) Accardi contra Bell (cum mundi): The impossible coupling. *Lecture Notes—Monograph Series* 42:133–154. Available at <https://www.jstor.org/stable/4356235>
- [11] Gill RD (2003) Time, finite statistics, and Bell's fifth position. *Foundations of Probability and Physics 2* (Växjö, Sweden, Växjö University Press), p 179. Available at <https://arxiv.org/abs/quant-ph/0301059>
- [12] Shafer G, Shen A, Vereshchagin N, Vovk V (2011) Test martingales, Bayes factors and p -values. *Statistical Science* 26:84–101. <https://doi.org/10.1214/10-ST347>
- [13] Ville J (1939) *Etude Critique de la Notion de Collectif* (Gauthier-Villars, Paris).
- [14] Zhang Y, Glancy S, Knill E (2013) Efficient quantification of experimental evidence against local realism. *Physical Review A* 88:052119. <https://doi.org/10.1103/PhysRevA.88.052119>
- [15] Christensen BG, Hill A, Kwiat PG, Knill D, Nam SW, Coakley K, Glancy S, Shalm LK, Zhang Y (2015) Analysis of coincidence-time loopholes in experimental Bell tests. *Physical Review A*, 92:032130. <https://doi.org/10.1103/PhysRevA.92.032130>
- [16] Chernoff H (1952) A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics* 23:493–507. <https://doi.org/10.1214/aoms/1177729330>
- [17] Hoeffding W (1963) Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58:13. <https://doi.org/10.2307/2282952>
- [18] Shao J (2003) *Mathematical Statistics* (Springer, New York), 2nd Ed. <https://doi.org/10.1007/b97553>
- [19] Kullback S, Leibler RA (1951) On information and sufficiency. *Annals of Mathematical Statistics* 22:79. <https://doi.org/10.1214/aoms/1177729694>
- [20] Bierhorst P (2015) A robust mathematical model for a loophole-free Clauser-Horne experiment. *Journal of Physica A* 48:195302. <https://doi.org/10.1088/1751-8113/48/19/195302>
- [21] Cai T (2005) One-sided confidence intervals in discrete distributions. *Journal of Statistical Planning and Inference* 131:63–68. <https://doi.org/10.1016/j.jspi.2004.01.005>

- [22] Brown LD, Cai T, Dasgupta A (2001) Interval estimation for a binomial proportion. *Statistical Science* 16(2):101–117. <https://doi.org/10.1214/ss/1009213286>
- [23] Nagaev SV, Chebotarev VI (2011) On the bound of proximity of the binomial distribution to the normal one. *Doklady Mathematics* 436(1):26–18 <https://doi.org/10.1137/S0040585X97985364>
- [24] McKay BD (1989) On Littlewood’s estimate for the binomial distribution. *Advances in Applied Probability* 21(2):475–478. <https://doi.org/10.2307/1427172>
- [25] Patel JK, Read CB (1996) *Handbook of the Normal Distribution* (CRC Press, Boca Raton, FL), Vol. 150.
- [26] Brockwell PJ, Davis RA (1991) *Time Series: Theory and Methods* (Springer Verlag, New York), 2nd Ed. <https://doi.org/10.1007/978-1-4419-0320-4>

About the authors: Peter Wills is an applied mathematician. Kevin Coakley is a mathematical statistician in the Statistical Engineering Division at NIST. Emanuel Knill is a mathematician in the Applied and Computational Mathematics Division at NIST. Yanbao Zhang is a postdoctoral researcher working on quantum information theory at NTT Basic Research Laboratories. The National Institute of Standards and Technology is an agency of the U.S. Department of Commerce.