

Testing with p^* -values: Between p-values, mid p-values, and e-values

RUODU WANG^{1,a}

¹*Department of Statistics and Actuarial Science, University of Waterloo,* ^awang@uwaterloo.ca

We introduce the notion of p^* -values (p^* -variables), which generalizes p-values (p-variables) in several senses. The new notion has four natural interpretations: operational, probabilistic, Bayesian, and frequentist. A main example of a p^* -value is a mid p-value, which arises in the presence of discrete test statistics. A unified stochastic representation for p-values, mid p-values, and p^* -values is obtained to illustrate the relationship between the three objects. We study several ways of merging arbitrarily dependent or independent p^* -values into one p-value or p^* -value. Admissible calibrators of p^* -values to and from p-values and e-values are obtained with nice mathematical forms, revealing the role of p^* -values as a bridge between p-values and e-values. The notion of p^* -values becomes useful in many situations even if one is only interested in p-values, mid p-values, or e-values. In particular, deterministic tests based on p^* -values can be applied to improve some classic methods for p-values and e-values.

MSC2020 subject classifications: Primary 62G10, 62F03; secondary 62C15

Keywords: Arbitrary dependence; average of p-values; mid p-values; posterior predictive p-values; test martingale

1. Introduction

Hypothesis testing is usually conducted with the classic notion of p-values. We introduce the abstract notion of p^* -variables, with p^* -values as their realizations, defined via a simple inequality in stochastic order, in a way similar to p-variables and p-values. As generalized p-variables, p^* -variables are motivated by mid p-values, and they admit four natural interpretations: operational, probabilistic, Bayesian, and randomized testing, arising in various statistical contexts.

The most important and practical example of p^* -values is the class of mid p-values (Lancaster [14]), arising from discrete test statistics. A mid p-value can be computed from $F(T-)/2 + F(T)/2$, where T is a test statistic (with smaller values indicating evidence against the null hypothesis) and F its cumulative distribution function. Mid p-variables are not p-variables, but they are p^* -variables. Discrete test statistics appear in many applications, especially when data represent frequencies or counts; see e.g., Döhler et al. [6] in the context of false discovery rate control. Another example of discrete p-values is the conformal p-values (Vovk et al. [33]); see e.g., the recent study of Bates et al. [1] on detecting outliers using conformal p-values. Using mid p-values is one way to address discrete test statistics; another way is using randomized p-values. We refer to Habiger [12] for mid and randomized p-values in multiple testing, and Rubin-Delanchy et al. [23] for probability bounds on combining independent mid p-values based on convex order.

In addition to mid p-values, p^* -values are also naturally connected to e-values (Vovk and Wang [36], Grünwald et al. [11]), and they have several advantages in contrast to p-values, especially via their connections to Bayes factors and test martingales (Shafer et al. [28]), betting scores (Shafer [27]), universal inference (Wasserman et al. [40]), anytime-valid tests (Grünwald et al. [11]), conformal tests (Vovk [32]), and false discovery rate under dependence (Wang and Ramdas [39], Vovk and Wang [37]). A recent review on this topic is Ramdas et al. [22].




Figure 1: Calibration among p-values, p*-values and e-values, where $f : [0, 1] \rightarrow [0, \infty]$ is left-continuous and decreasing with $f(0) = \infty$ and $\int_0^1 f(t) dt = 1$.

In discussions where the probabilistic specification as random variables is not emphasized, we will loosely use the term “p/p*/e-values” for both p/p*/e-variables and their realizations (as commonly done in the literature), and this should be clear from the context.

The relationship between p*-values and mid p-values is studied in Section 3. We obtain a new stochastic representation for mid p-values (Theorem 3.1), which unifies the classes of p-, mid p-, and p*-values. The set of p*-values is closed under several types of operations, and this closure property is not shared by that of p-values or mid p-values (Proposition 3.3). Based on these results, we find that p*-values serve as an abstract and generalized version of mid p-values which is mathematically more convenient to work with.

There are several equivalent definitions of p*-variables: by stochastic order (Definition 2.2), by averaging p-variables (Theorem 4.1), by conditional probability (Proposition 4.3), and by randomized tests (Proposition 4.5); each of them represents a natural statistical path to a generalization of p-values, and these paths lead to the same mathematical object of p*-values. Moreover, a p*-value is a posterior predictive p-value of Meng [19] in the Bayesian context. The p*-test in Section 4.4 is a randomized version of the traditional p-test; the randomization is needed because p*-values are weaker than p-values.

Merging methods are useful in multiple hypothesis testing for both p-values and e-values. Merging several p-values or e-values are used, either directly or implicitly, in false discovery control procedures in Genovese and Wasserman [9] and Goeman and Solari [10] and generalized Bonferroni-Holm procedures (see [35] for p-values and [36] for e-values). We study merging functions for p*-values in Section 5, which turn out to have convenient structures. In particular, we find that a (randomly) weighted geometric average of arbitrary p*-variables multiplied by $e \approx 2.718$ is a p-variable (Theorem 5.2), allowing for a simple combination of p*-values under unknown dependence; a similar merging function for p-values is obtained by [35]. In the setting of merging independent p*-values, inequalities obtained by [23] on mid p-values can be directly applied to p*-values.

We explore in Sections 6 the connections among p-values, p*-values, and e-values by establishing results for admissible calibrators. Figure 1 summarizes these calibrators where the ones between p-values and e-values are obtained in [36]. Notably, for an e-value e , $(2e)^{-1}$ is a calibrated p*-value, which has an extra factor of $1/2$ compared to the standard calibrated p-value e^{-1} . A composition of the e-to-p* calibration $p^* = (2e)^{-1} \wedge 1$ and the p*-to-p calibration $p = (2p^*) \wedge 1$ leads to the unique admissible e-to-p calibration $p = e^{-1} \wedge 1$, thus showing that p*-values serve as a bridge between e-values and p-values.

In classic statistical settings where precise (uniform on $[0, 1]$) p-values are available, p*-values may not be directly useful as their properties are weaker than p-values. Nevertheless, applying a p*-test in

situations where precise p-values are unavailable leads to several improvements on classic methods for p-values and e-values. An application on testing with e-values is discussed and numerically illustrated in Section 7. Based on calibration to p^* -values, an extra factor of 2 can be obtained for the product of independent e-values when one of them satisfies a decreasing density condition. An application on testing with discrete test statistics and mid p-values is presented in Section 8, where our p^* -merging methods become useful when the test statistics are dependent. (Another application is presented in Appendix B.) From these examples, we find that the tool of p^* -values is useful even when one is primarily interested in p-values, mid p-values, or e-values.

The paper is written such that p-values, p^* -values and e-values are treated as abstract measure-theoretical objects, following the setting of [35,36]. Our null hypothesis is a generic and unspecified one, and it can be simple or composite; nevertheless, for the discussions of our results, it would be harmless to keep a simple hypothesis in mind as a primary example. All proofs are put in Appendix A and the randomized p^* -test is discussed in Appendix B.

2. P-values, p^* -values, and e-values

Following the setting of [36], we directly work with a fixed atomless probability space $(\Omega, \mathcal{A}, \mathbb{P})$, where our null hypothesis is set to be the singleton $\{\mathbb{P}\}$. As explained in Appendix D of [36], no generality is lost as all mathematical results (of the kind in this paper) are valid also in the contexts of general composite hypotheses and multiple testing of a global null. We assume that $(\Omega, \mathcal{A}, \mathbb{P})$ is rich enough so that we can find a uniform random variable independent of a given random vector as we wish. We first define stochastic orders, which will be used to formulate the main objects in the paper. All terms like “increasing” and “decreasing” are in the non-strict sense.

Definition 2.1. *Let X and Y be two random variables.*

1. X is first-order stochastically smaller than Y , written as $X \leq_1 Y$, if $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$ for all increasing real functions f such that the expectations exist.
2. X is second-order stochastically smaller than Y , written as $X \leq_2 Y$, if $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$ for all increasing concave real functions f such that the expectations exist.

Recall that the defining property of a p-value, realized by a random variable P , is that $\mathbb{P}(P \leq \alpha) \leq \alpha$ for all $\alpha \in (0, 1)$, meaning that the type-I error of rejecting P based on $P \leq \alpha$ is at most α (see e.g., [36]). The above is equivalent to the statement that P is first-order stochastically larger than a uniform random variable on $[0, 1]$ (e.g., Section 1.A of [29]). Motivated by this simple observation, we define p-variables and e-variables, and our new concept, called p^* -variables, via stochastic orders.

Definition 2.2. *Let U be a uniform random variable on $[0, 1]$.*

1. A random variable P is a p-variable if $U \leq_1 P$.
2. A random variable P is a p^* -variable if $U \leq_2 P$.
3. A random variable E is an e-variable if $0 \leq_2 E \leq_2 1$.

We allow both p-variables and p^* -variables to take values above one, although such values are uninteresting, and one may safely truncate them at 1; moreover, we allow an e-variable E to take the value ∞ (but with probability 0 under the null), which corresponds to a p-variable taking the value 0 (also with probability 0).

Since \leq_1 is stronger than \leq_2 , a p-variable is also a p^* -variable. Due to the close proximity between p-variables and p^* -variables, we often use P for both of them; this should not create any confusion.

We refer to *p-values* as realizations of p-variables, *p*-values* as those of p*-variables, and *e-values* as those of e-variables. By definition, both a p-variable and a p*-variable have a mean at least 1/2.

The classic definition of an e-variable E is via $\mathbb{E}[E] \leq 1$ and $E \geq 0$ a.s. ([36]). This is equivalent to our Definition 2.2 because

$$0 \leq_2 E \iff 0 \leq E \text{ a.s.}; \quad E \leq_2 1 \iff \mathbb{E}[E] \leq 1.$$

We choose to express our definition via stochastic orders to make an analogy among the three concepts, and stochastic orders will be a main technical tool for results in this paper.

Our main focus is the notion of p*-values, which will be motivated from five perspectives in Sections 3 and 4.

Remark 2.3. There are many equivalent conditions for the stochastic order $U \leq_2 P$. One of the most convenient conditions, which will be used repeatedly in this paper, is (see e.g., Theorem 4.A.3 of [29])

$$U \leq_2 P \iff \int_0^v G_P(u) du \geq \frac{v^2}{2} \text{ for all } v \in (0, 1), \quad (1)$$

where G_P is the left-quantile function of P , defined as

$$G_P(u) = \inf\{x \in \mathbb{R} : \mathbb{P}(P > x) \leq u\}, \quad u \in (0, 1].$$

3. Mid p-values and discrete test statistics

An important motivation for p*-values is the use of mid p-values and discrete test statistics. We first recall the usual practice to obtain p-values. Let T be a test statistic which is a function of the observed data. Here, a smaller value of T represents stronger evidence against the null hypothesis. The p-variable P is usually computed from the conditional probability

$$P = \mathbb{P}(T' \leq T | T) = F(T), \quad (2)$$

where F is the distribution of T , and T' is an independent copy of T ; here and below, a copy of T is a random variable identically distributed as T .

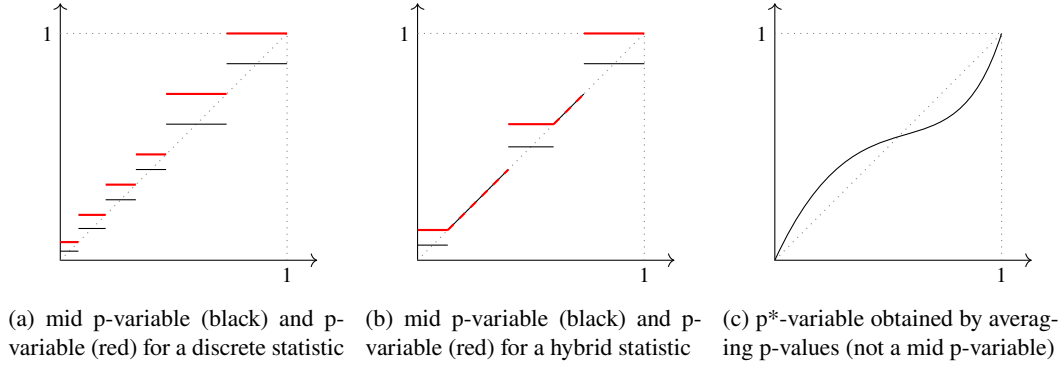
If T has a continuous distribution, then the p-variable P defined by (2) has a standard uniform distribution. If the test statistic T is discrete, e.g., when testing a binomial model $\text{Binomial}(n, \pi)$, P is strictly first-order stochastically larger than a uniform random variable on $[0, 1]$.

The discreteness of T leads to a conservative p-value; in particular, $\mathbb{E}[P] > 1/2$. One way to address this is to randomize the p-value to make it uniform on $[0, 1]$; however, randomization is generally undesirable in testing. As the most natural alternative, *mid p-values* (Lancaster [14]) arise in the presence of discrete test statistics. For the test statistic T , the mid p-value is given by

$$P_T = \frac{1}{2} \mathbb{P}(T' \leq T | T) + \frac{1}{2} \mathbb{P}(T' < T | T) = \frac{1}{2} F(T-) + \frac{1}{2} F(T), \quad (3)$$

where $F(t-) = \lim_{s \uparrow t} F(s)$ for $t \in \mathbb{R}$. Clearly, $P_T \leq F(T)$ and $\mathbb{E}[P_T] = 1/2$. If T is continuously distributed, then $P_T = F(T)$ is uniform on $[0, 1]$. In case T is discrete, P_T is not a p-variable. In Figure 2 we present some examples of quantile functions of p-, mid p- and p*-variables.

Similarly to the case of p-variables in Definition 2.2, we formally define a *mid p-variable* as a random variable $P \geq_1 P_T$ for some test statistic T . Often we have equality (i.e., $P =_1 P_T$), and a strict inequality may appear due to, e.g., composite hypotheses, similarly to the case of p-variables or e-variables.

Figure 2: Examples of quantile functions of p-, mid p- and p^* -variables

It is straightforward to verify that mid p-variables are p^* -variables; see [23] where convex order is used in replace of our second-order stochastic dominance. The following theorem establishes a new unified stochastic representation for p-, mid p-, and p^* -variables using the relation

$$P \geq_1 \mathbb{E}[U|f(U)], \text{ where } U \text{ is a standard uniform random variable.} \quad (4)$$

Theorem 3.1. For a random variable P ,

- (i) P is a p-variable if and only if (4) holds for some strictly increasing f ;
- (ii) P is a mid p-variable if and only if (4) holds for some increasing f ;
- (iii) P is a p^* -variable if and only if (4) holds for some measurable f .

Remark 3.2. In Theorem 3.1 above, conditional expectations and conditional probabilities are defined in the a.s. sense, as usual in probability theory. Whereas the first statement in Theorem 3.1 follows directly by definition and the second statement can be checked by construction, the third statement follows from a highly non-trivial recent result of Nutz et al. [18] on the existence of backward Monge martingale transport. Recall that Strassen's theorem [31] yields the form $P \geq_1 \mathbb{E}[U|V]$ for some random variable V , but changing V to $f(U)$ requires a much more sophisticated analysis in [18].

From Theorem 3.1, it is clear that mid p-variables are special cases of p^* -variables but the converse is not true. As a direct consequence, all results later on p^* -variables can be directly applied to mid p-values. The stochastic representations in Theorem 3.1 may not be directly useful in statistical inference; nevertheless they reveal a deep connection between mid p-values and p^* -values, allowing us to analyze possible improvements of methods designed for p^* -variables when applied to mid p-values, and vice versa.

In the next result, we summarize some closure property of the sets of p-, mid p and p^* -variables.

- Proposition 3.3.**
- (i) The set of p^* -variables is closed under convex combinations and under distribution mixtures.
 - (ii) The set of p-variables is closed under distribution mixtures but not under convex combinations.
 - (iii) The set of mid p-variables is neither closed under convex combinations nor under distribution mixtures.
 - (iv) The three sets above are all closed under convergence in distribution.

Proposition 3.3 suggests that the set of p^* -variables has the nicest closure properties among the three. Moreover, we will see in Theorem 4.1 below that the set of p^* -variables is precisely the convex hull of the set of p -variables, and hence, it is also the convex hull of the set of mid p -variables which contains all p -variables.

4. Four formulations of p^* -values

In this section, we will present four further equivalent definitions of p^* -values, each arising from a different statistical context, and providing several interpretations of p^* -values.

4.1. Averages of p -values

Our first interpretation of p^* -variables is operational: we will see that a p^* -variable is precisely the arithmetic average of some p -values which are obtained from possibly different sources and arbitrarily dependent. This characterization relies on a recent technical result of Mao et al. [17] on the sum of standard uniform random variables.

Theorem 4.1. *A random variable is a p^* -variable if and only if it is the convex combination of some p -variables. Moreover, any p^* -variable can be expressed as the arithmetic average of three p -variables.*

Remark 4.2. As implied by Theorem 5 of [17], a p^* -variable can always be written as the arithmetic average of n dependent p -variables for any $n \geq 3$, but the statement is not true for $n = 2$ ([17, Proposition 1]).

As a consequence of Theorem 4.1, the set of p^* -variables is the convex hull of the set of p -variables, as briefly mentioned in Section 3. Testing with the arithmetic average of dependent p -values has been studied by [35]. We further discuss in Appendix B an application of p^* -values which improves some tests based on arithmetic averages of p -values.

4.2. Conditional probability of exceedance

Our second interpretation is probabilistic: we will interpret both p -variables and p^* -variables as conditional probabilities. Let T be a test statistic which is a function of the observed data, represented by a vector X . Recall that a p -variable P in (2) has the form

$$P = F(T) = \mathbb{P}(T' \leq T|T) = \mathbb{P}(T' \leq T|X).$$

It turns out that p^* -variables have a similar representation, where the only difference is that T defining a p^* -variable may not be a function of X ; instead, T can include some unobservable randomness or additional randomness in the statistical experiment (but the p^* -variable itself is deterministic from the data).

Proposition 4.3. *For a $\sigma(X)$ -measurable random variable P ,*

- (i) *P is a p -variable if and only if there exists a $\sigma(X)$ -measurable T such that $P \geq \mathbb{P}(T' \leq T|X)$ where T' is a copy of T independent of X ;*

- (ii) P is a p*-variable if and only if there exists a random variable T such that $P \geq \mathbb{P}(T' \leq T|X)$ where T' is a copy of T independent of (T, X) .

In (ii) above, $\mathbb{P}(T' \leq T|X)$ can be safely replaced by $\mathbb{P}(T' \leq T|X)/2 + \mathbb{P}(T' < T|X)/2$.

Proposition 4.3 suggests that p*-variables are very similar to p-variables when interpreted as conditional probabilities; the only difference is whether extra randomness is allowed in T .

Remark 4.4. Both Proposition 4.3 and Theorem 3.1 give stochastic representations for p- and p*-variables. They are similar and with a few differences. First, one is stated via stochastic order whereas the other via inequalities between random variables. Second, one involves a uniform random variable U whereas the other does not as T may be discrete. Third, measurability conditions are different as Proposition 4.3 specifies $\sigma(X)$.

4.3. Posterior predictive p-values

In the Bayesian context, the posterior predictive p-value of Meng [19] is a p*-value. Let X be the data vector in Section 4.2. The null hypothesis H_0 is given by $\{\psi \in \Psi_0\}$ where Ψ_0 is a subset of the parameter space Ψ on which a prior distribution is specified. The posterior predictive p-value is defined as the realization of the random variable

$$P_B := \mathbb{P}(D(X', \psi) \geq D(X, \psi)|X),$$

where D is a function (taking a similar role as test statistics), X' and X are iid conditional on ψ , and the probability is computed under the joint posterior distribution of (X', ψ) . Note that P_B can be rewritten as

$$P_B = \int \mathbb{P}(D(X', y) \geq D(X, y)|X, y) d\Pi(y|X)$$

where Π is the posterior distribution of ψ given the data X . One can check that P_B is a p*-variable by using Jensen's inequality; see Theorem 1 of [19] where $D(X, \psi)$ is assumed to be continuously distributed conditional on ψ .

In this formulation, p*-variables are obtained by integrating p-variables over the posterior distribution of some unobservable parameter. Since p*-variables are treated as measure-theoretic objects in this paper, we omit a detailed discussion of the Bayesian interpretation; nevertheless, it is reassuring that p*-values have a natural appearance in the Bayesian context as put forward by Meng [19]. One of our later results is related to an observation of [19] that two times a p*-variable is a p-variable (see Proposition 5.1).

4.4. Randomized tests with p*-values

Recall that the defining property of a p-variable P is that the standard p-test

$$\text{rejecting the null hypothesis} \iff P \leq \alpha \tag{5}$$

has size (i.e., probability of type-I error) at most α for each $\alpha \in (0, 1)$. Since p*-values are a weaker version of p-values, one cannot guarantee that the test (5) for a p*-variable P has size at most α . Nevertheless, a randomized version of the test (5) turns out to be valid. Moreover, this randomized test yields a defining property for p*-variables, just like p-variables are defined by the deterministic p-test (5).

Proposition 4.5. *Let $V_\alpha \sim U[0, 2\alpha]$ and a random variable P be independent of V_α , $\alpha \in (0, 1/2]$. Then $\mathbb{P}(P \leq V_\alpha) \leq \alpha$ for all $\alpha \in (0, 1/2]$ if and only if P is a p^* -variable.*

Proposition 4.5 implies that p^* -variables are precisely test statistics which can pass the *randomized p^* -test* (rejection via $P \leq V_\alpha$) with the specified level α , thus a further equivalent definition of p^* -variables. The drawback of the randomized p^* -test is obvious: an extra layer of randomization is needed. This undesirable feature is the price one has to pay when a p -variable is weakened to a p^* -variable. More details and applications of the randomized p^* -test are put in Appendix B. Since randomization is undesirable, we omit the detailed discussions here.

Remark 4.6. The random threshold V_α in Proposition 4.5 can be replaced by any random variable V with mean α and a decreasing density function on $(0, 1)$. The uniform distribution on $[0, 2\alpha]$ turns out to be the one with the smallest variance among all valid choices of the random threshold (see Proposition B.4).

5. Merging p^* -values and mid p -values

Merging p -values and e -values is extensively studied in the literature of multiple hypothesis testing; see the recent studies [16,36,34] and the references therein. We will be interested in merging p^* -values (including mid p -values) into both a p^* -value and a p -value. The following proposition gives a convenient conversion rule between p^* -values and p -values. The fact that two times a p^* -variable is a p -variable is already observed by [19].

Proposition 5.1. *A p -variable is a p^* -variable, and the sum of two p^* -variables is a p -variable.*

Proposition 5.1 implies that, in order to obtain a valid p -value from several p^* -values, a naive method is to multiply each p^* -value by 2 and then apply a valid method for merging p -values (under the corresponding assumptions). We will see in the next few results that we can often obtain stronger results than this.

As argued by Efron [8, p.50-51], dependence assumptions are difficult to verify in multiple hypothesis testing. We will first focus on the case of arbitrarily dependent p^* -values, that is, without making any assumptions on the dependence structure of the p^* -variables, and then turn to the case of independent or positively dependent p^* -variables.

5.1. Arbitrarily dependent p^* -values

We first provide a new method which merges several p^* -values into a p -value based on geometric averaging. Vovk and Wang [35] showed that the geometric average of p -variables may fail to be a p -variable, but it yields a p -variable when multiplied by $e := \exp(1)$. The constant $e \approx 2.718$ is practically the best-possible (smallest) multiplier (see [35, Table 2]) that provides validity against all dependence structures. In the next result, we show that a similar but stronger result holds for p^* -values: the geometric average of p^* -variables multiplied by e is not only a p^* -variable, but also a p -variable, and this holds also for randomly weighted geometric averages. For using weighted p -values in multiple testing, see e.g., Benjamini and Hochberg [2].

In what follows, the (randomly) weighted geometric average \tilde{P} of P_1, \dots, P_K for random weights w_1, \dots, w_K is given by

$$\tilde{P} = \prod_{k=1}^K P_k^{w_k} = \exp\left(\sum_{k=1}^K w_k \log P_k\right),$$

where w_1, \dots, w_K satisfy

$$w_1, \dots, w_K \geq 0, \text{ independent of } (P_1, \dots, P_K), \text{ and } \sum_{k=1}^K w_k = 1. \quad (6)$$

If $w_1 = \dots = w_K = 1/K$, then \tilde{P} is the unweighted geometric average of P_1, \dots, P_K .

Theorem 5.2. *Let \tilde{P} be a weighted geometric average of p^* -variables. Then $e\tilde{P}$ is a p -variable. That is, for arbitrary p^* -variables P_1, \dots, P_K and weights w_1, \dots, w_K satisfying (6),*

$$\mathbb{P}\left(\prod_{k=1}^K P_k^{w_k} \leq \alpha\right) \leq e\alpha \text{ for all } \alpha \in (0, 1). \quad (7)$$

In Theorem 5.2, the random weights are allowed to be arbitrarily dependent. These random weights may come from preliminary experiments. One way to obtain such weights is to use scores such as e -values from preliminary data. Using e -values to compute weights is quite natural as a main motivation of e -values is an accumulation of evidence between consecutive experiments; see [11], [36] and [39].

Theorem 5.2 generalizes the result of [35, Proposition 4] which considered unweighted geometric average of p -values. When dependence is unspecified, testing with (randomly) weighted geometric averages of p^* -values has the same critical values α/e as those with unweighted p -values.

Next, we will study two methods which merge p^* -values into a p^* -value. Since two times p^* -variable is a p -variable, probability guarantee can also be obtained from these merging functions. A p^* -merging function in dimension K is an increasing Borel function M on $[0, \infty)^K$ such that $M(P_1, \dots, P_K)$ is a p^* -variable for all p^* -variables P_1, \dots, P_K ; p -merging and e -merging functions are defined analogously; see [36]. A p^* -merging function M is *admissible* if it is not strictly dominated by another p^* -merging function.

Proposition 5.3. *The arithmetic average M_K is an admissible p^* -merging function in any dimension K .*

Proposition 5.3 illustrates that p^* -values are very easy to combine using an arithmetic average; recall that M_K is not a valid p -merging function since the average of p -values is not necessarily a p -value (instead, $2M_K$ is a p -merging function). On the other hand, M_K is an admissible e -merging function which essentially dominates all other symmetric admissible e -merging functions ([35, Proposition 3.1]).

Another benchmark merging function is the Bonferroni merging function

$$M_B : (p_1, \dots, p_K) \mapsto \left(K \bigwedge_{k=1}^K p_k\right) \wedge 1.$$

The next result shows that M_B is an admissible p^* -merging function. The Bonferroni merging function M_B is known to be an admissible p -merging function ([36, Proposition 6.1]), whereas its transformed form (via $e = 1/p$) is an e -merging function but not admissible; see [36, Section 6] for these claims.

Proposition 5.4. *The Bonferroni merging function M_B is an admissible p^* -merging function in any dimension K .*

Combining Propositions 5.3 and 5.4, p^* -merging is admissible via both the arithmetic average (admissible for e-merging, invalid for p-merging) and the Bonferroni correction (admissible for p-merging, inadmissible for e-merging).

5.2. Independent p^* -variables

We next turn to the problem of merging independent p^* -variables. Merging independent mid p-values is studied by Rubin-Delanchy et al. [23] based on arguments of convex order. Since our p^* -variables are defined via the order \geq_2 which is closely related to convex order, the bounds in [23] can be directly applied to the case of p^* -variables. More precisely, for any p^* -variable P , using Strassen's theorem in the form of [29, Theorems 4.A.5 and 4.A.6], there exists a random variable Z such that $Z \leq P$ and Z satisfies the convex order relation used in [23]. In particular, for the arithmetic average \bar{P} of independent p^* -variables, using [23, Theorem 1] leads to the probability bound

$$\mathbb{P}(\bar{P} \leq 1/2 - t) \leq \exp(-6Kt^2) \quad \text{for all } t \in [0, 1/2]. \quad (8)$$

Three probability bounds for the geometric average of independent p^* -variables is obtained by [23, Theorem 2]; recall that Fisher's combination method uses the geometric average of independent p-values. More precisely, for the geometric average \tilde{P} of independent p^* -variables, [23, Theorem 2] gives, for $y \geq 1$,

$$\mathbb{P}(\tilde{P} \leq \exp(-y)) \leq \min \left\{ H_{2K}(2K(y - \log 2)), \frac{1}{1 + (y - 1)^2}, \exp(K(1 - y + \log y)) \right\}, \quad (9)$$

where H_{2K} is the survival function of a chi-square distribution with $2K$ degrees of freedom.

It is well-known that statistical validity of Fisher's method or other methods based on concentration inequalities can be fragile when independence does not hold; see also our simulation results in Section 8. Since independence is difficult to verify in multiple hypothesis testing (see e.g., [8]), these independence-based methods (for either p-values or p^* -values) need to be applied with caution.

There are, nevertheless, some methods which work well for independent p-values and are relatively robust to dependence assumptions. In addition to the Bonferroni correction which is valid for all dependence structures, the most famous such method is perhaps that of Simes [30]. Define the function

$$S_K : [0, \infty)^K \rightarrow [0, \infty), \quad S_K(p_1, \dots, p_k) = \bigwedge_{k=1}^K \frac{K p_{(k)}}{k},$$

where $p_{(k)}$ is the k -th smallest order statistic of p_1, \dots, p_K . A celebrated result of [30] is that if P_1, \dots, P_K are independent p-variables, then the Simes inequality holds

$$\mathbb{P}(S_K(P_1, \dots, P_K) \leq \alpha) \leq \alpha \quad \text{for all } \alpha \in (0, 1). \quad (10)$$

Further, if P_1, \dots, P_K are iid uniform on $[0, 1]$, then $S_K(P_1, \dots, P_K)$ is again uniform on $[0, 1]$. The Simes inequality (10) holds also under some notion of positive dependence, in particular, positive regression dependence (PRD); see Benjamini and Yekutieli [3] and Ramdas et al. [21].

One may wonder whether $S_K(P_1, \dots, P_K)$ yields a p^* -variable or p-variable for independent or PRD p^* -variables P_1, \dots, P_K . It turns out that this is not the case, as illustrated by the following example,

where $S_2(P_1, P_2)$ fails to be a p^* -variable or a p -variable, even in case that P_1 and P_2 are iid p^* -variables.

Example 5.5. Let P_1 be a random variable satisfying $\mathbb{P}(P_1 = 0.2) = 0.4$ and $\mathbb{P}(P_1 = 0.7) = 0.6$. It is straightforward to verify that P_1 is a p^* -variable (indeed, it is a mid p -variable by Theorem 3.1). Let P_2 be an independent copy of P_1 and $P := S_2(P_1, P_2)$. We can check that $\mathbb{P}(P = 0.2) = 0.16$, $\mathbb{P}(P = 0.4) = 0.48$ and $\mathbb{P}(P = 0.7) = 0.36$. It follows that $\mathbb{E}[P] = 0.476 < 1/2$, and hence P is not a p^* -variable.

Since twice a p^* -variable is a p -variable (Proposition 5.1), it is safe (and conservative) to use $2S_K(P_1, \dots, P_K)$ which is a p -variable under independence or PRD (note that PRD is preserved under linear transformations).

Other methods on p -values that are relatively robust to dependence include the harmonic mean p -value of Wilson [41] and the Cauchy combination method of Liu and Xie [16]. As shown by Chen et al. [4], the three methods of Simes, harmonic mean, and Cauchy combinations are closely related and similar in several senses.

Obviously, more robustness to dependence leads to a more conservative method. Indeed, all p -merging methods designed for arbitrary dependence are quite conservative in some situations; see the comparative study in [4]. Thus, there is a trade-off between power and robustness to dependence. For p^* -merging methods, the bound (8) is the most stringent on the independence assumption. Using $2S_K$ is valid for independent or PRD p^* -variables. Finally, all methods in Section 5.1 work for any dependence structure among the p^* -variables.

Remark 5.6. Any function M which merges iid standard uniform random variables U_1, \dots, U_K into a standard uniform one, such as the functions in the methods of Simes, Fisher's and the Cauchy combination, satisfies

$$\mathbb{P}(M(P_1, \dots, P_K) \leq \alpha) \leq \alpha \text{ for all } \alpha \in (0, 1) \quad (11)$$

for any independent p -variables P_1, \dots, P_K . However, they generally cannot satisfy (11) for all independent p^* -variables (or mid p -variables) P_1, \dots, P_K , since $M(P_1, \dots, P_K) \geq_1 M(U_1, \dots, U_K)$ does not hold for some choices of P_1, \dots, P_K . Therefore, some form of penalty always needs to be paid when relaxing p -values to p^* -values or mid p -values for these methods.

6. Calibration between p -values, p^* -values, and e -values

In this section, we discuss calibration between p -, p^* -, and e -values. Calibration between p -values and e -values is one of the main topics of [36].

6.1. Calibration between p -values and p^* -values

Calibration between p -values and p^* -values is relatively simple. A p -to- p^* calibrator is an increasing function $f : [0, 1] \rightarrow [0, \infty)$ that transforms p -variables to p^* -variables, and a p^* -to- p calibrator is an increasing function $g : [0, 1] \rightarrow [0, \infty)$ which transforms in the reverse direction. Clearly, the values of p -values larger than 1 are irrelevant, and hence we restrict the domain of all calibrators in this section to be $[0, 1]$; in other words, input p -variables and p^* -variables larger than 1 will be treated as 1. A calibrator is said to be *admissible* if it is not strictly dominated by another calibrator of the same kind (for calibration to p -values and p^* -values, f dominates g means $f \leq g$, and for calibration to e -values in Section 6.2 it is the opposite inequality).

Theorem 6.1. (i) *The p^* -to- p calibrator $u \mapsto (2u) \wedge 1$ dominates all other p^* -to- p calibrators.*
(ii) *An increasing function f on $[0, 1]$ is an admissible p -to- p^* calibrator if and only if f is left-continuous, $f(0) = 0$, $\int_0^v f(u) du \geq v^2/2$ for all $v \in (0, 1)$, and $\int_0^1 f(u) du = 1/2$.*

Theorem 6.1 (i) states that a multiplier of 2 is the best calibrator that works for all p^* -values. This observation justifies the deterministic threshold $\alpha/2$ in the test (26) for p^* -values, as mentioned in Section 4.4. Although Theorem 6.1 (ii) implies that there are many admissible p -to- p^* calibrators, it seems that there is no obvious reason to use anything other than the identity in Proposition 5.1 when calibrating from p -values to p^* -values. Finally, we note that the conditions in Theorem 6.1 (ii) imply that the range of f is contained in $[0, 1]$, an obvious requirement for an admissible p -to- p^* calibrator.

For a single p^* -value computed from the mid p -value $P = F(T^-)/2 + F(T)/2$ as in (3), it is never beneficial to use the p -variable $2P$ over the p -variable $F(T)$, since $F(T) \leq 2P$ always holds true.

6.2. Calibration between p^* -values and e -values

Next, we discuss calibration between e -values and p^* -values, which has a richer structure. A p^* -to- e calibrator is a decreasing function $f : [0, 1] \rightarrow [0, \infty]$ that transforms p^* -variables to e -variables, and an e -to- p^* calibrator $g : [0, \infty] \rightarrow [0, 1]$ is a decreasing function which transforms in the reverse direction. We include $e = \infty$ in the calibrators, which corresponds to $p = 0$.

First, since a p -variable is a p^* -variable, any p^* -to- e calibrator is also a p -to- e calibrator. Hence, the set of p^* -to- e calibrators is contained in the set of p -to- e calibrators. By Proposition 2.1 of [36], any admissible p -to- e calibrator $f : [0, 1] \rightarrow [0, \infty]$ is a decreasing function such that $f(0) = \infty$, f is left-continuous, and $\int_0^1 f(t) dt = 1$. We will see below that some of these admissible p -to- e calibrators are also p^* -to- e calibrators.

Regarding the other direction of e -to- p^* calibrators, we first recall that there is a unique admissible e -to- p calibrator, given by $e \mapsto e^{-1} \wedge 1$, as shown by [36]. Since the set of p^* -values is larger than that of p -values, the above e -to- p calibrator is also an e -to- p^* calibrator. The interesting questions are whether there is any e -to- p^* calibrator stronger than $e \mapsto e^{-1} \wedge 1$, and whether an admissible e -to- p^* calibrator is also unique. The constant map $e \mapsto 1/2$ is an e -to- p^* calibrator since $1/2$ is a constant p^* -variable. If there exists an e -to- p^* calibrator f which dominates all other e -to- p^* calibrators, then it is necessary that $f(e) \leq 1/2$ for all $e \geq 0$; however this would imply $f = 1/2$ since any p^* -variable has mean at least $1/2$. Since $e \mapsto 1/2$ does not dominate $e \mapsto e^{-1} \wedge 1$, we conclude that there is no e -to- p^* calibrator which dominates all others, in contrast to the case of e -to- p calibrators.

Nevertheless, some refined form of domination can be helpful. We say that an e -to- p^* calibrator f *essentially dominates* another e -to- p^* calibrator f' if $f(e) \leq f'(e)$ whenever $f'(e) < 1/2$. That is, we only require dominance when the calibrated p^* -value is useful (relatively small); this consideration is similar to the essential domination of e -merging functions in [36]. It turns out that the e -to- p calibrator $e \mapsto e^{-1} \wedge 1$ can be improved by a factor of $1/2$, which essentially dominates all other e -to- p^* calibrators.

The following theorem summarizes the validity and admissibility results on both directions of calibration.

Theorem 6.2. (i) *A convex p -to- e calibrator is a p^* -to- e calibrator; an admissible convex p -to- e calibrator is an admissible p^* -to- e calibrator.*
(ii) *An admissible p -to- e calibrator is a p^* -to- e calibrator if and only if it is convex.*
(iii) *The e -to- p^* calibrator $e \mapsto (2e)^{-1} \wedge 1$ essentially dominates all other e -to- p^* calibrators.*

All practical examples of p-to-e calibrators are convex and admissible; see [36, Section 2 and Appendix B] for a few classes (which are all convex). By Theorem 6.2, all of these calibrators are admissible p^* -to-e calibrators. A popular class of p-to-e calibrators is given by, for $\kappa \in (0, 1)$,

$$p \mapsto \kappa p^{\kappa-1}, \quad p \in [0, 1]. \quad (12)$$

Another simple choice, proposed by Shafer [27], is

$$p \mapsto p^{-1/2} - 1, \quad p \in [0, 1]. \quad (13)$$

Clearly, the p-to-e calibrators in (12) and (13) are convex and thus they are p^* -to-e calibrators.

The result in Theorem 6.2 (iii) shows that the unique admissible e-to-p calibrator $e \mapsto e^{-1} \wedge 1$ can actually be achieved by a two-step calibration: first use $p^* = (2e)^{-1} \wedge 1$ to get a p^* -value, and then use $p = (2p^*) \wedge 1$ to get a p-value.

On the other hand, all p-to-e calibrators f in [36] are convex, and they can be seen as a composition of the calibrations $p^* = p$ and $e = f(p^*)$. Therefore, p^* -values serve as an intermediate step in both directions of calibration between p-values and e-values, although one of the directions is less interesting since the p-to- p^* calibrator is an identity. Figure 1 in the Introduction illustrates our recommended calibrators among p-values, p^* -values and e-values based on Theorems 6.1 and 6.2, and they are all admissible.

Example 6.3. Suppose that U is uniformly distributed on $[0, 1]$. Using the calibrator (12), for $\kappa \in (0, 1)$, $E := \kappa U^{\kappa-1}$ is an e-variable. By Theorem 6.2 (iii), we know that $P := (2E)^{-1}$ is a p^* -variable. Below we check this directly. The left-quantile function G_P of P satisfies

$$G_P(u) = \frac{u^{1-\kappa}}{2\kappa}, \quad u \in (0, 1).$$

Using $\kappa(2 - \kappa) \leq 1$ for all $\kappa \in (0, 1)$, we have

$$\int_0^v G_P(u) du = \frac{v^{2-\kappa}}{2\kappa(2-\kappa)} \geq \frac{v^{2-\kappa}}{2} \geq \frac{v^2}{2}, \quad v \in (0, 1).$$

Hence, P is a p^* -variable by verifying (1). Moreover, for $\kappa \in (0, 1/2]$, P is even a p-variable, since $G_P(u) \geq u$ for $u \in (0, 1)$.

In the next result, we show that a p^* -variable obtained from the calibrator in Theorem 6.2 (iii) is a p-variable under a further condition (DE):

(DE) $E \leq_1 E'$ for some e-variable E' which has a decreasing density on $(0, \infty)$.

In particular, condition (DE) is satisfied if E itself has a decreasing density on $(0, \infty)$. Examples of e-variables satisfying (DE) are those obtained from applying a non-constant convex p-to-e calibrator f with $f(1) = 0$ to any p-variable, e.g., the p-to-e calibrator (13) but not (12); this is because convexity of the calibrator yields a decreasing density when applied to a uniform p-variable.

Proposition 6.4. For any e-variable E , $P := (2E)^{-1} \wedge 1$ is a p^* -variable, and if E satisfies (DE), then P is a p-variable.

Remark 6.5. In a spirit similar to Proposition 6.4, smoothing techniques leading to an extra factor of 2 in the Markov inequality have been studied by Huber [13].

7. Testing with e-values and martingales

In this section we discuss applications of p*-values to tests with e-values and martingales. E-values and test martingales are usually used for purposes more than rejecting a null hypothesis while controlling type-I error; in particular, they offer anytime validity and different interpretations of statistical evidence (e.g., [11]). We compare the power of several methods here for a better understanding of their performance, while keeping in mind that single-run detection power (which is maximized by p-values if they are available) is not the only purpose of e-values.

Suppose that E is an e-variable, usually obtained from likelihood ratios or stopped test supermartingales (e.g., [28], [27]). A traditional e-test is

$$\text{rejecting the null hypothesis} \iff E \geq \frac{1}{\alpha}. \quad (14)$$

Using the fact that $(2E)^{-1}$ is a p*-variable in Theorem 6.2 (iii), we can design the randomized test

$$\text{rejecting the null hypothesis} \iff 2E \geq \frac{1}{V}, \quad (15)$$

where $V \sim U[0, 2\alpha]$ is independent of E (Proposition 4.5). The test (15) has 3/4 chance of being more powerful than the traditional choice of testing E against $1/\alpha$ in (14). Randomization is undesirable, but (15) inspires us to look for alternative deterministic methods.

Suppose that one has two independent e-variables E_1 and E_2 for a null hypothesis. As shown by [36], it is optimal in a weak sense to use the combined e-variable $E_1 E_2$ for testing the null. Assume further that one of E_1 and E_2 satisfies condition (DE).

Using (15) with the random threshold αE_2 and Proposition B.5 in Appendix B, we get $\mathbb{P}((2E_1)^{-1} \leq \alpha E_2) \leq \alpha$ (note that the positions of E_1 and E_2 are symmetric here). Hence, the test

$$\text{rejecting the null hypothesis} \iff 2E_1 E_2 \geq \frac{1}{\alpha}, \quad (16)$$

has size at most α . The threshold of the test (16) is half the one obtained by directly applying (14) to the e-variable $E_1 E_2$. Thus, the test statistic is boosted by a factor of 2 via condition (DE) on either E_1 or E_2 . No assumption is needed for the other e-variable. In particular, by setting $E_2 = 1$, we get a p-variable $(2E_1)^{-1}$ if E_1 satisfies (DE), as we see in Proposition 6.4.

E-values calibrated from p-values are useful in the context of testing randomness online (see [32]) and designing test martingales (see [7]). More specifically, for a possibly infinite sequence of independent p-variables $(P_t)_{t \in \mathbb{N}}$ and a sequence of p-to-e calibrators $(f_t)_{t \in \mathbb{N}}$, the stochastic process

$$X_t = \prod_{k=1}^t f_k(P_k), \quad t = 0, 1, \dots$$

is a supermartingale (with respect to the filtration of $(P_t)_{t \in \mathbb{N}}$) with initial value $X_0 = 0$ (it is a martingale if P_t , $t \in \mathbb{N}$ are standard uniform and f_t , $t \in \mathbb{N}$ are admissible). As a supermartingale, $(X_t)_{t \in \mathbb{N}}$ satisfies anytime validity, i.e., X_τ is an e-variable for any stopping time τ ; moreover, Ville's inequality gives

$$\mathbb{P}\left(\sup_{t \in \mathbb{N}} X_t \geq \frac{1}{\alpha}\right) \leq \alpha \text{ for any } \alpha > 0. \quad (17)$$

The process $(X_t)_{t \in \mathbb{N}}$ is called an *e-process* by [39]. Anytime validity is crucial in the design of online testing where evidence arrives sequentially in time, and scientific discovery is reported at a stopping time considered with sufficient evidence.

Notably, the most popular choice of p-to-e calibrators are those in (12) and (13) (see e.g., [32]), which are convex. Theorem 6.2 implies that if the inputs are not p-values but p^* -values, we can still obtain e-processes using convex calibrators such as (12) and (13), without calibrating these p^* -values to p-values. This observation becomes useful when each observed P_t is only a p^* -variable, e.g., a mid p-value or an average of several p-values from parallel experiments.

Moreover, for a fixed $t \in \mathbb{N}$, if there is a convex f_s for some $s \in \{1, \dots, t\}$ with $f_s(1) = 0$, and P_s is a p-variable (the others can be p^* -variables with any p^* -to-e calibrators), then (DE) is satisfied by $f_s(P_s)$, and we have $\mathbb{P}(X_t \geq 1/\alpha) \leq \alpha/2$ by using the test (16); see our numerical experiments below.

Finally, for *sequential e-variables* E_1, \dots, E_t , that is, $\mathbb{E}[E_k | E_1, \dots, E_{k-1}] \leq 1$ for $k = 1, \dots, t$ (this is weaker than independence), if the conditional distribution of E_t given each realization of (E_1, \dots, E_{t-1}) satisfies (DE), then, by writing $X_s = \prod_{k=1}^s E_k$ for $s = 1, \dots, t$, we have, for $\alpha \in (0, 1)$,

$$\mathbb{P}\left(X_t \geq \frac{1}{2\alpha}\right) = \mathbb{E}\left[\mathbb{P}\left(2E_t \geq \frac{1}{\alpha X_{t-1}} \mid E_1, \dots, E_{t-1}\right)\right] \leq \mathbb{E}[\alpha X_{t-1}] \leq \alpha,$$

where we use the fact in (16) that $\mathbb{P}(2E_t \geq 1/x) \leq x$ for $x > 0$ if E_t satisfies (DE).

Simulation experiments

In the simulation results below, we generate test martingales following [36]. Similarly to Section B.2, the null hypothesis H_0 is $N(0, 1)$ and the alternative is $N(\delta, 1)$ for some $\delta > 0$. We generate iid X_1, \dots, X_n from $N(\delta, 1)$. Define the e-variables from the likelihood ratios of the alternative to the null density,

$$E_t := \frac{\exp(-(X_t - \delta)^2/2)}{\exp(-X_t^2/2)} = \exp(\delta X_t - \delta^2/2), \quad t = 1, \dots, n. \quad (18)$$

The e-process $S = (S_t)_{t=1, \dots, n}$ is defined as $S_t = \prod_{s=1}^t E_s$. Such an e-process S is *growth optimal* in the sense of Shafer [27], as it maximizes the expected log growth among all test martingales built on the data (X_1, \dots, X_n) ; indeed, S is Kelly's strategy under the betting interpretation. Here, we constructed the e-process S assuming that we know δ ; otherwise we can use universal test martingales (e.g., [5]) by taking a mixture of S over δ under some probability measure.

Note that each E_t is log-normally distributed and it does not satisfy (DE). Hence, (16) cannot be applied to S_n . Nevertheless, we can replace E_1 by another e-variable E'_1 which satisfies (DE). We choose E'_1 by applying the p-to-e calibrator (13) to the p-variable $P_1 = 1 - \Phi(X_1)$, namely, $E'_1 = (P_1)^{-1/2} - 1$.

Replacing E_1 by E'_1 , we obtain the new e-process $S' = (S'_t)_{t=1, \dots, n}$ by $S'_t = E'_1 \prod_{s=2}^t E_s$. The e-process S' is not growth optimal, but as E'_1 satisfies (DE), we can test via the rejection condition $2S'_n \geq 1/\alpha$, thus boosting the terminal value by a factor of 2. Let $V \sim U[0, 2\alpha]$ be independent of the test statistics. We compare five different tests, all with size at most α :

- (a) applying (14) to S_n : reject H_0 if $S_n \geq 1/\alpha$ (benchmark case);
- (b) applying (15) to S_n : reject H_0 if $2S_n \geq 1/V$;
- (c) applying (16) to S'_n : reject H_0 if $2S'_n \geq 1/\alpha$;
- (d) applying a combination of (15) and (16) to S'_n : reject H_0 if $2S'_n \geq 1/V$;
- (e) applying (17) to the maximum of S : reject H_0 if $\max_{1 \leq t \leq n} S_t \geq 1/V$.

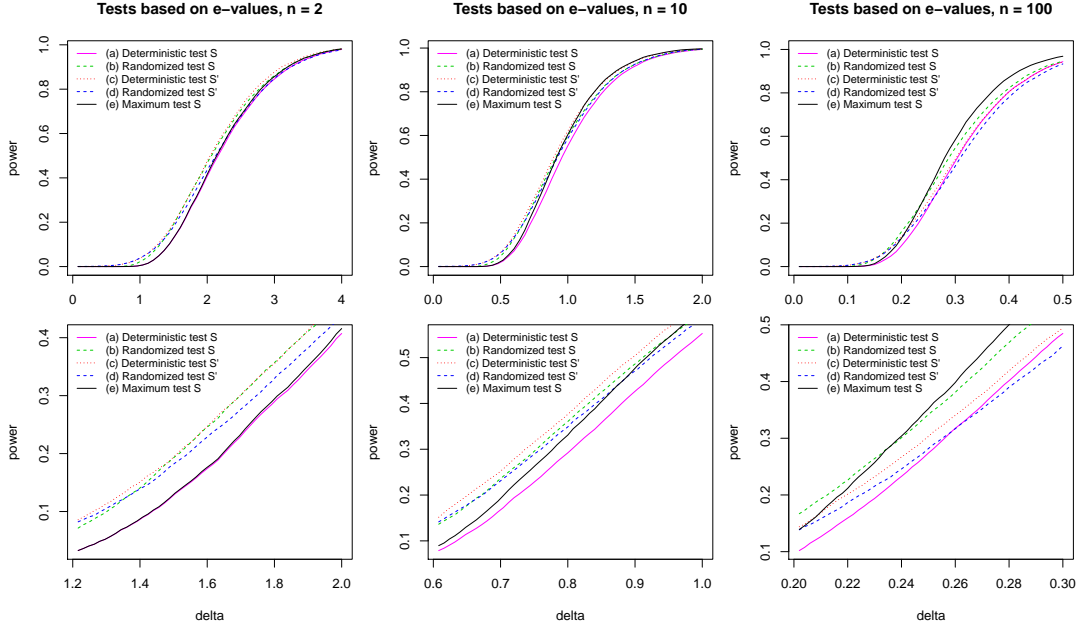


Figure 3: Tests based on e-values; the second row is zoomed in from the first row

Since test (a) is strictly dominated by test (e), we do not need to use (a) in practice; nevertheless we treat it as a benchmark for comparison on tests based on e-values as it is built on the fundamental connection between e-values and p-values: the e-to-p calibrator $e \mapsto e^{-1} \wedge 1$.

The significance level α is set to be 0.01. The power of the five tests is computed from the average of 10,000 replications for varying signal strength δ and for $n \in \{2, 10, 100\}$. Results are reported in Figure 3. For most values of δ and n , either the deterministic test (c) for S' or the maximum test (e) has the best performance. The deterministic test (c) performs very well in the cases $n = 2$ and $n = 10$, especially for weak signals; this may be explained by the factor of 2 being substantial when the signal is weak. If n is large and the signal is not too weak, the effect of using the maximum of S in (e) is dominating; this is not surprising. Although the randomized test (b) usually improves the performance from the benchmark case (a), the advantages seem to be quite limited, especially in view of the extra randomization, often undesirable. The randomized test (d) does not have better performance than the deterministic test (c) because both of them have the same factor of 2 in front of S'_n .

8. Testing with combined mid p-values

We compare by simulation the performance of a few tests via merging mid p-values. The (global) null hypothesis H_0 is that the test statistic T follows a binomial distribution $\text{Binomial}(n, \pi)$, and K tests are conducted. We set $n = 40$ and $\pi = 0.3$, so that the obtained p-values are considerably discrete. We denote by P_1, \dots, P_K the obtained p-values via (2), and P_1^*, \dots, P_K^* the obtained mid p-values via (3). Let \bar{P}^* and \tilde{P}^* be the arithmetic average and the geometric average of P_1^*, \dots, P_K^* , respectively.

The true data probability generating the test statistics is a binomial distribution $\text{Binomial}(n, (1 - \theta)\pi)$, where $\theta \in [0, 1]$. The case $\theta = 0$ means that the null hypothesis H_0 is true, and a larger θ indicates a stronger signal.

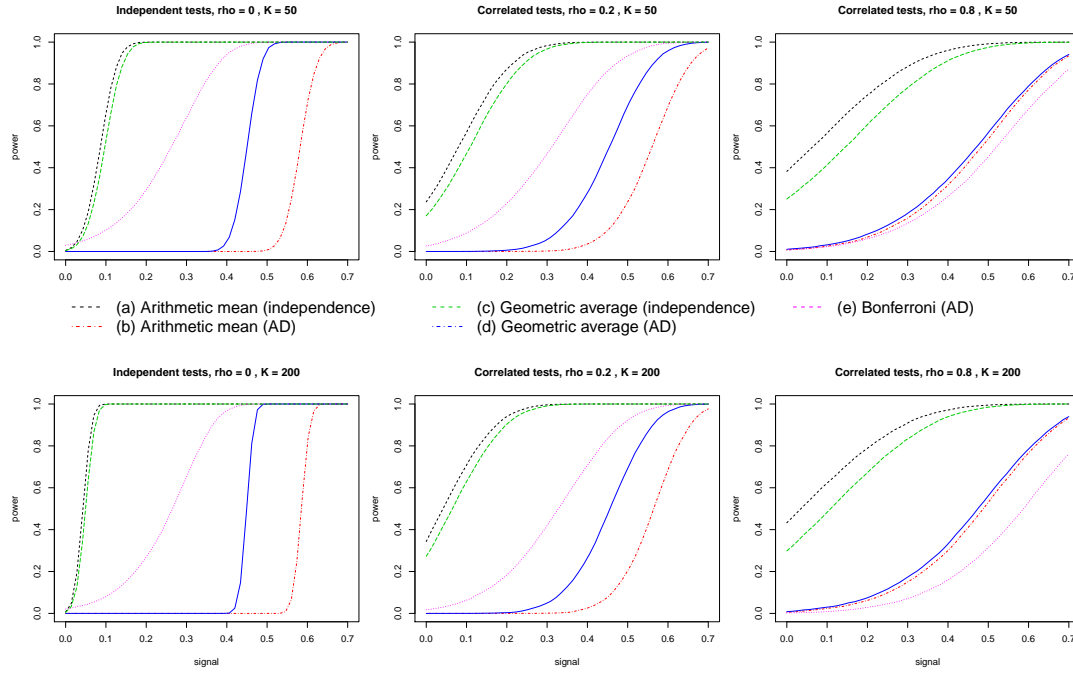


Figure 4: Tests based on combining p-values or mid p-values

We allow the test statistics to be correlated, and this is achieved by simulating from a Gaussian copula with common pairwise correlation parameter ρ (more precisely, we first simulate from a Gaussian copula, and then obtain the observable discrete test statistics by a quantile transform). We consider the following tests (there are other tests possible for this setting, and we only compare these five to illustrate a few relevant points):

- (a) the bound (8) on the arithmetic mean in [23]: reject H_0 if $\bar{P}^* \leq 1/2 - (-\log \alpha / (6K))^{1/2}$;
- (b) the arithmetic mean times 2 using Proposition 5.3: reject H_0 if $\bar{P}^* \leq \alpha/2$;
- (c) the bound (9) on the geometric mean in [23]: reject H_0 if $\tilde{P}^* \leq \exp(-y)$ where $y \geq 1$ is such that the right-hand side of (9) is α .
- (d) the geometric average of p^* -values using Theorem 5.2: reject H_0 if $\tilde{P}^* \leq \alpha/e$;
- (e) the Bonferroni correction: reject H_0 if $\min(P_1, \dots, P_K) \leq \alpha/K$.

Note that tests (a), (b), (c) and (d) use mid p-values based on methods for p^* -values, and (e) uses p-values. All of (b), (d) and (e) are valid tests under arbitrary dependence (AD) whereas the validity of (a) and (c) requires independence. Therefore, we expect (a) and (c) to perform very well in case independence holds. All other methods are valid but conservative, as there is a big price to pay to gain robustness against all dependence structures.

The significance level α is set to be 0.05 for good visibility. The power of the four tests is reported in Figure 4 which is computed from the average of 10,000 replications for varying signal strength $\theta \in [0, 1]$ and for $\rho = 0$ (independence), $\rho = 0.2$ (mild dependence) and $\rho = 0.8$ (strong dependence). The situation of $\rho = 0.8$ is the most relevant to us as averaging methods are designed mostly for situations where the presence of strong or complicated dependence is suspected.

As we can see from Figure 4, the test (a) relying on independence has the strongest power as expected. Its size becomes largely inflated as soon as mild dependence is present, and hence it can only

be applied in situations where independence among obtained mid p-values can be justified. Indeed, the size of test (a) can be ≈ 0.4 in case $\rho = 0.2$, which is clearly not useful. The test (c) performs similarly to (a) with slightly less power. Among the three methods that are valid for arbitrary dependence, the geometric average (d) has stronger power for large ρ and the Bonferroni correction (e) has stronger power for small ρ . The arithmetic average (b) performs poorly unless p-values are very strongly correlated. These observations on merging mid p-values are consistent with those in [35] on merging p-values. In conclusion, the geometric average (d) can be useful when the dependence among mid p-values is speculated to be strong or complicated, although unknown to the decision maker.

9. Conclusion

In this paper we introduced p*-values (p*-variables) as an abstract measure-theoretic object. The notion of p*-values is a multifold generalization of p-values, and it enjoys many attractive theoretical properties in contrast to p-values. In particular, mid p-values, which arise in the presence of discrete test statistics, form an important subset of p*-values. Merging methods for p*-values are studied. In particular, a weighted geometric average of arbitrarily dependent p*-values multiplied by $e \approx 2.718$ yields a valid p-value, which can be useful when multiple mid p-values are possibly strongly correlated.

Results on calibration between p*-values and e-values reveal that p*-values serve as an intermediate step in both the standard e-to-p and p-to-e calibrations. Although a direct test with p*-values may involve randomization, we find that p*-values are useful in the design of deterministic tests with averages of p-values, mid p-values, and e-values. From the results in this paper, the concept of p*-values serves a useful technical tool which enhances the extensive and growing applications of p-values, mid p-values, and e-values.

Appendix A: Proofs of all results

In this appendix, we collect proofs of all theorems and propositions.

A.1. Proofs of results in Section 3

Proof of Theorem 3.1. (i) For f being a strictly increasing function, we have $\mathbb{E}[U|f(U)] = U$, and the equivalence statement follows directly from the definition of p-variables.

(ii) We first show the “if” statement. Write $V = f(U)$ for an increasing function f . Denote by F and G the distribution function and the left-quantile function of V , respectively, and let Z be a standard uniform random variable independent of V . Moreover, let

$$U_V = F(V-) + Z(F(V) - F(V-)),$$

which is uniformly distributed on $[0, 1]$ and satisfies $G(U_V) = V$ a.s. (e.g., [25, Proposition 1.3]). Since $G(U_V) = V = f(U)$ a.s., and both G and f are increasing, we know that the functions G and f differ on a set of Lebesgue measure 0. Therefore, $\mathbb{E}[U|V] = \mathbb{E}[U|G(U)]$, which is identically distributed as $\mathbb{E}[U_V|G(U_V)]$. Moreover,

$$\mathbb{E}[U_V|G(U_V)] = \mathbb{E}[U_V|V] = \frac{1}{2}F(V-) + \frac{1}{2}F(V). \quad (19)$$

Therefore, $P \geq_1 \mathbb{E}[U|V]$ implies $P \geq_1 (F(V-) + F(V))/2$, and thus P is a mid p-variable.

The “only if” statement follows from $P \geq_1 (F(T-) + F(T))/2$ and (19) by choosing $U = U_V$ and $f = G$.

- (iii) We first show the “if” statement. Note that $P \geq_1 \mathbb{E}[U|f(U)] \geq_2 U$ where the second inequality is guaranteed by Jensen’s inequality. Since \geq_1 is stronger than \geq_2 and \geq_2 is transitive, we get $P \geq_2 U$ and hence P is a p^* -variable.

Next, we show the “only if” statement. By using [29, Theorems 4.A.5], there exist a standard uniform random variable U and a random variable V such that $P \geq_1 \mathbb{E}[U|V]$. By [18, Proposition 3.4], there exists a measurable function f such that $\mathbb{E}[U|f(U)]$ and $\mathbb{E}[U|V]$ are identically distributed. \square

Proof of Proposition 3.3. By Theorem 4.1, the set of p^* -variables is the convex hull of the set of p -variables, and thus convex. This also implies that none of the set of p -variables or that of mid p -variables is convex.

To show that the set of p^* -variables is closed under distribution mixtures, it suffices to note that the stochastic orders \leq_1 and \leq_2 (indeed, any order induced by inequalities via integrals) is closed under distribution mixture.

To see that the set of mid p -variables is not closed under distribution mixtures, we note from (3) that any mid p -variable with mean $1/2$ and a point-mass at $1/2$ must not have any density in a neighbourhood of $1/2$. Hence, the mixture of uniform distribution on $[0, 1]$ and a point-mass at $1/2$ is not the distribution of a mid p -variable.

Closure under convergence for \leq_1 is justified by Theorem 1.A.3 of [29], and closure under convergence for \leq_2 is justified by Theorem 1.5.9 of [20]. Closure under convergence for the set of mid p -values follows by noting that the set of distributions of P_T in (3) is closed under convergence in distribution, which can be checked by definition. \square

A.2. Proofs of results in Section 4

Proof of Theorem 4.1. We first show that a convex combination of p -variables is a p^* -variable. Let U be a uniform random variable on $[0, 1]$, P_1, \dots, P_K be K p -variables, $(\lambda_1, \dots, \lambda_K)$ be an element of the standard K -simplex, and f be an increasing concave function. By monotonicity and concavity of f , we have

$$\mathbb{E} \left[f \left(\sum_{k=1}^K \lambda_k P_k \right) \right] \geq \mathbb{E} \left[\sum_{k=1}^K \lambda_k f(P_k) \right] \geq \mathbb{E} \left[\sum_{k=1}^K \lambda_k f(U) \right] = \mathbb{E}[f(U)].$$

Therefore, $\sum_{k=1}^K \lambda_k P_k \geq_2 U$ and thus $\sum_{k=1}^K \lambda_k P_k$ is a p^* -variable.

Next, we show the second statement that any p^* -variable can be written as the average of three p -variables, which also justifies the “only if” direction of the first statement.

Let P be a p^* -variable satisfying $\mathbb{E}[P] = 1/2$. Note that $P \geq_2 U$ and $\mathbb{E}[P] = \mathbb{E}[U]$ together implies $P \geq_{cv} U$ (see e.g., [29, Theorem 4.A.35]), where \leq_{cv} is the concave order, meaning that $\mathbb{E}[f(P)] \geq \mathbb{E}[f(U)]$ for all concave f . Theorem 5 of [17] says that any $P \geq_{cv} U$, there exist three standard uniform random variables P_1, P_2, P_3 such that $3P = P_1 + P_2 + P_3$ (this statement is highly non-trivial). This implies that P can be written as the arithmetic average of three p -variables P_1, P_2, P_3 .

Finally, assume that the p^* -variable P satisfies $\mathbb{E}[P] > 1/2$. In this case, using Strassen’s Theorem in the form of [29, Theorems 4.A.5 and 4.A.6], there exists a random variable Z such that $U \leq_{cv} Z \leq P$. As we explained above, there exist p -variables P_1, P_2, P_3 such that $3Z = P_1 + P_2 + P_3$. For $i = 1, 2, 3$, let $\tilde{P}_i := P_i + (P - Z) \geq_1 P_i$. Note that $\tilde{P}_1, \tilde{P}_2, \tilde{P}_3$ are p -variables and $3P = \tilde{P}_1 + \tilde{P}_2 + \tilde{P}_3$. Hence, P can be written as the arithmetic average of three p -variables. \square

Proof of Proposition 4.3. For the “only-if” statement in (i), since P is a p-variable, we know that its distribution F satisfies $F(t) \leq t$ for $t \in (0, 1)$. Therefore, by setting $T = P$,

$$P \geq F(P) = F(T) = \mathbb{P}(T' \leq T|T) = \mathbb{P}(T' \leq T|X),$$

where the last equality holds since T' is independent of X . To check the “if” direction of (i), we have $\mathbb{P}(T' \leq T|X) = \mathbb{P}(T' \leq T|T) = F(T)$ where F is the distribution of T . Note that $F(T)$ is stochastically larger than or equal to a uniform random variable on $[0, 1]$, and hence $\mathbb{P}(F(T) \leq t) \leq t$.

Next, we show (ii). First, suppose that $P \geq \mathbb{P}(T' \leq T|X)$. Let U be a uniform random variable on $[0, 1]$. By Jensen’s inequality, we have $\mathbb{E}[F(T)|X] \geq_2 F(T)$. Hence, $P \geq_2 \mathbb{P}(T' \leq T|X) = \mathbb{E}[F(T)|X] \geq_2 F(T) \geq_2 U$, and thus P is a p*-variable.

For the converse direction, suppose that P is a p*-variable. By Strassen’s Theorem in the form of [29, Theorem 3.A.4], there exists a uniform random variable U' on $[0, 1]$ and a random variable P' identically distributed as P such that $P' \geq \mathbb{E}[U'|P']$. Let $G(\cdot|p)$ be the left-quantile function of a regular conditional distribution of U' given $P' = p \in [0, 1]$. Further, let V be a uniform random variable on $[0, 1]$ independent of (P, X) , and $U := G(V|P)$. It is clear that (U, P) has the same law as (U', P') . Therefore, $P \geq \mathbb{E}[U|P]$. Moreover, $\mathbb{E}[U|X] = \mathbb{E}[G(V|P)|X] = \mathbb{E}[G(V|P)|P]$ since V is independent of X . Hence, $P \geq \mathbb{E}[U|P] = \mathbb{E}[U|X]$. Let V' be another uniform random variable on $[0, 1]$ independent of (U, X) . We have

$$\mathbb{P}(V' \leq U|X) = \mathbb{E}[\mathbf{1}_{\{V' \leq U\}}|X] = \mathbb{E}[U|X].$$

Hence, the representation $P \geq \mathbb{P}(T' \leq T|X)$ holds with $T' = V'$ and $T = U$.

Finally, we show the last statement on replacing $\mathbb{P}(T' \leq T|X)$ by $\mathbb{P}(T' \leq T|X)/2 + \mathbb{P}(T' < T|X)/2$. The “only-if” direction follows from the argument for (ii) by noting that U constructed above has a continuous distribution. The “if” direction follows from

$$\begin{aligned} \frac{1}{2}(\mathbb{P}(T' \leq T|X) + \mathbb{P}(T' < T|X)) &= \frac{1}{2}(\mathbb{E}[F(T)|X] + \mathbb{E}[F(T-)|X]) \\ &\geq_2 \frac{1}{2}(F(T) + F(T-)) \geq_2 U, \end{aligned}$$

where the second-last inequality is Jensen’s, and the last inequality is implied by Theorem 3.1. \square

Proof of Proposition 4.5. The “if” statement is implied by Proposition B.1. To show the “only if” statement, denote by F_P the distribution function of P and F_U be the distribution function of a uniform random variable U on $[0, 1]$. We have

$$\alpha \geq \mathbb{P}(P \leq V_\alpha) = \int_0^{2\alpha} \frac{F_P(u)}{2\alpha} du.$$

Therefore, for $v \in (0, 1]$, we have

$$\int_0^v F_P(u) du \leq \frac{v^2}{2} = \int_0^v u du = \int_0^v F_U(u) du.$$

By Theorem 4.A.2 of [29], the above inequality implies $U \leq_2 P$. Hence, P is a p*-variable. \square

A.3. Proofs of results in Section 5

Proof of Proposition 5.1. Let U be a uniform random variable on $[0, 1]$. The first statement is trivial by definition. For the second statement, let P_1, P_2 be two p^* -variables. For any $\epsilon \in (0, 1)$,

$$\begin{aligned} \mathbb{P}(P_1 + P_2 \leq \epsilon) &= \mathbb{E}[\mathbb{1}_{\{P_1 + P_2 \leq \epsilon\}}] \\ &\leq \mathbb{E}\left[\frac{1}{\epsilon}(2\epsilon - P_1 - P_2)_+\right] \\ &\leq \mathbb{E}\left[\frac{1}{\epsilon}(\epsilon - P_1)_+\right] + \mathbb{E}\left[\frac{1}{\epsilon}(\epsilon - P_2)_+\right] \\ &\leq 2\mathbb{E}\left[\frac{1}{\epsilon}(\epsilon - U)_+\right] = \epsilon, \end{aligned}$$

where the last inequality is because $U \leq_2 P_1, P_2$ and $u \mapsto (\epsilon - u)_+$ is convex and decreasing. Therefore, $P_1 + P_2$ is a p -variable. \square

The following lemma is needed in the proof of Theorem 5.2.

Lemma A.1. *Let M be an increasing Borel function on $[0, \infty)^K$. Then $M(P_1, \dots, P_K)$ is a p -variable for all p^* -variables P_1, \dots, P_K if and only if*

$$\inf\{q_1(M(\alpha P_1, \dots, \alpha P_K)) : P_1, \dots, P_K \geq_2 U\} \geq \alpha \text{ for all } \alpha \in (0, 1), \quad (20)$$

where U is uniform on $[0, 1]$ and $q_1(X)$ is the essential supremum of a random variable X .

Proof. Let $g(\alpha)$ be the critical value for testing with $M(P_1, \dots, P_K)$, that is, the largest value such that

$$\mathbb{P}(M(P_1, \dots, P_K) < g(\alpha)) \leq \alpha \text{ for all } p^*\text{-variables } P_1, \dots, P_K.$$

Converting between the distribution function and the quantile function, this means

$$g(\alpha) = \inf\{q_\alpha(M(P_1, \dots, P_K)) : P_1, \dots, P_K \geq_2 U\},$$

where $q_\alpha(X)$ be the left α -quantile of a random variable X . For an increasing function M , its infimum α -quantile can be converted to the essential supremum of random variables with conditional distributions on their lower α -tail; see the proof of [15, Theorem 3] which deals with the case of additive functions (see also the proof of [4, Proposition 1] where this technique is used). Note that for $P \geq 2U$, its lower α -tail conditional distribution dominates αU . This argument leads to

$$g(\alpha) = \inf\{q_1(M(P_1, \dots, P_K)) : P_1, \dots, P_K \geq_2 \alpha U\}.$$

Noting that $P \geq_2 \alpha U$ is equivalent to $P/\alpha \geq_2 U$, we obtain (20). \square

Proof of Theorem 5.2. First, suppose that w_1, \dots, w_K are non-negative constants adding up to 1. Let $g(\alpha)$ be the critical value for testing with \tilde{P} , that is, the smallest value such that

$$\mathbb{P}\left(\prod_{k=1}^K P_k^{w_k} \leq g(\alpha)\right) \leq \alpha \text{ for all } p^*\text{-variables } P_1, \dots, P_K.$$

We will show that $g(\alpha) \geq \alpha/\epsilon$. Using Lemma A.1, we get

$$g(\alpha) = \alpha \inf \left\{ q_1 \left(\prod_{k=1}^K P_k^{w_k} \right) : P_1, \dots, P_K \geq_2 U \right\}.$$

For any p^* -variables P_1, \dots, P_K , since $\log(\cdot)$ is an increasing concave function, we have

$$q_1 \left(\sum_{k=1}^K w_k \log(P_k) \right) \geq \mathbb{E} \left[\sum_{k=1}^K w_k \log(P_k) \right] \geq \sum_{k=1}^K w_k \mathbb{E}[\log(U)] = -1.$$

Therefore, $\log(g(\alpha)/\alpha) \geq -1$, leading to the desired bound $g(\alpha) \geq \alpha/e$, and thus

$$\mathbb{P} \left(\prod_{k=1}^K P_k^{w_k} \leq \alpha/e \right) \leq \mathbb{P} \left(\prod_{k=1}^K P_k^{w_k} \leq g(\alpha) \right) \leq \alpha.$$

For random w_1, \dots, w_k , taking an expectation leads to (7). \square

Proof of Proposition 5.3. The validity of M_K as a p^* -merging function is implied by Theorem 4.1. To show its admissibility, suppose that there exists a p^* -merging function M that strictly dominates M_K . Let P_1, \dots, P_K be iid uniform random variables on $[0, 1]$. The strict domination implies $M \leq M_K$ and $\mathbb{P}(M(P_1, \dots, P_K) < M_K(P_1, \dots, P_K)) > 0$. We have

$$\mathbb{E}[M(P_1, \dots, P_K)] < \mathbb{E}[M_K(P_1, \dots, P_K)] = \frac{1}{2}.$$

This means that $M(P_1, \dots, P_K)$ is not a p^* -variable, a contradiction. \square

Proof of Proposition 5.4. Let P_1, \dots, P_K be p^* -variables, and P be a random variable such that the distribution of P is the equally weighted mixture of those of P_1, \dots, P_K . Note that P a p^* -variable by Proposition 3.3. Let $P_{(1)} = \bigwedge_{k=1}^K P_k$. Using the Bonferroni inequality, we have, for any $\epsilon \in (0, 1)$,

$$\mathbb{P}(P_{(1)} \leq \epsilon) \leq \sum_{k=1}^K \mathbb{P}(P_k \leq \epsilon) = K\mathbb{P}(P \leq \epsilon). \quad (21)$$

Let G_1 be the left-quantile function of $P_{(1)}$ and G_2 be that of P . By (21), we have $G_1(Kt) \geq G_2(t)$ for all $t \in (0, 1/K)$. Hence, for each $y \in (0, 1/K)$, using the equivalent condition (1), we have

$$\int_0^y K G_1(t) dt \geq \int_0^y K G_2(t/K) dt = K^2 \int_0^{y/K} G_2(t) dt \geq K^2 \frac{y^2}{2K^2} = \frac{y^2}{2}.$$

This implies, via the equivalent condition (1) again, that $KP_{(1)}$ is a p^* -variable.

Next we show the admissibility of M_B for $K \geq 2$, since the case $K = 1$ is trivial. Suppose that there is a p^* -merging function M which strictly dominates M_B . There exists $(p_1, \dots, p_K) \in [0, 1]^K$ such that $M(p_1, \dots, p_K) < M_B(p_1, \dots, p_K) = KP_{(1)}$, where $p_{(1)} = \bigwedge_{k=1}^K p_k$. Suppose, without loss of generality, that $p_1 = p_{(1)}$. We can assume $p_1 \leq 1/K$, because otherwise

$$M(1/K, p_2, \dots, p_K) \leq M(p_1, p_2, \dots, p_K) < M_B(p_1, p_2, \dots, p_K) \leq 1 = M_B(1/K, p_2, \dots, p_K);$$

that is, we can replace p_1 by $1/K$. Since M_B is continuous, there exists $\epsilon > 0$ such that $M < M_B$ on the set

$$S := [p_1 - \epsilon, p_1] \times \cdots \times [p_K - \epsilon, p_K] \subset [0, 1]^K.$$

Let $p = p_1 - \epsilon/2$. Define identically distributed random variables P_1, \dots, P_K by

$$P_k = pV_k \mathbb{1}_{A_k} + (p + (1-p)V_k) \mathbb{1}_{A_k^c}, \quad k = 1, \dots, K,$$

where A_1, \dots, A_K are disjoint events with $\mathbb{P}(A_k) = p$ for each k , and V_1, \dots, V_K are iid standard uniform random variables, independent of (A_1, \dots, A_K) . It is easy to check that each of P_1, \dots, P_K has a standard uniform distribution. Moreover, using independence of A_1 and V_1, \dots, V_K , we have

$$\begin{aligned} \mathbb{P}((P_1, \dots, P_K) \in S \mid A_1) &\geq \mathbb{P}\left(\{pV_1 \in [p_1 - \epsilon, p_1]\} \cap \bigcap_{k=2}^K \{(p + (1-p)V_k) \in [p_k - \epsilon, p_k]\} \mid A_1\right) \\ &= \mathbb{P}(pV_1 \in [p_1 - \epsilon, p_1]) \prod_{k=2}^K \mathbb{P}((p + (1-p)V_k) \in [p_k - \epsilon, p_k]) > 0, \end{aligned}$$

where the last inequality uses the fact that $p_1 - \epsilon < p < p_1 \leq p_k$ for each k . Therefore, we obtain $\mathbb{P}((P_1, \dots, P_K) \in S \mid A) > 0$.

For simplicity write $X = M(P_1, \dots, P_K)$ and $X_B = M_B(P_1, \dots, P_K)$. Let G be the left-quantile function of X and G_B be the left-quantile function of X_B . On each set A_k , $X_B = KpV_k$ which is uniformly distributed on $[0, Kp]$. Therefore, on $A := \bigcup_{k=1}^K A_k$, X_B is uniformly distributed on $[0, Kp]$. Moreover, by construction, X_B takes its smallest values on A , i.e., $X(\omega) \leq X(\omega')$ for almost every $\omega \in A$ and $\omega' \in A^c$. Hence, $G_B(t) = t$ for every $t \in (0, 2Kp)$. Since $M \leq M_B$, $M < M_B$ on S , and $\mathbb{P}((P_1, \dots, P_K) \in S \mid A) > 0$, we obtain $X \leq X_B$ on A and $\mathbb{P}(X < X_B \mid A) > 0$. Since X_B takes its smallest values on A , this implies $G(t) \leq G_B(t)$ for every $t \in (0, 2Kp)$ and $G(t) < G_B(t)$ for some $t \in (0, 2Kp)$. With left continuity of G_B and monotonicity of G , this implies $G(t) < G_B(t)$ on a subset of $(0, 2Kp)$ with positive Lebesgue measure. Therefore,

$$\int_0^{2Kp} G(t) dt < \int_0^{2Kp} G_B(t) dt = \frac{(2Kp)^2}{2},$$

showing that $M(P_1, \dots, P_K)$ is not a p^* -variable by (1), a contradiction. \square

A.4. Proofs of results in Section 6

Proof of Theorem 6.1. Let U be a uniform random variable on $[0, 1]$.

- (i) The validity of the calibrator $u \mapsto (2u) \wedge 1$ is implied by (i), and below we show that it dominates all others. For any function g on $[0, \infty)$, suppose that $g(u) < 2u$ for some $u \in (0, 1/2]$. Consider the random variable V defined by $V = U1_{\{U > 2u\}} + u1_{\{U \leq 2u\}}$. Clearly, V is a p^* -variable. Note that

$$\mathbb{P}(g(V) \leq g(u)) \geq \mathbb{P}(U \leq 2u) = 2u > g(u),$$

implying that g is not a p^* -to- p calibrator. Hence, any p^* -to- p calibrator g satisfies $g(u) \geq 2u$ for all $u \in (0, 1/2]$, thus showing that $u \mapsto (2u) \wedge 1$ dominates all p^* -to- p calibrators.

- (ii) By (1), we know that $f(U) \geq_2 U$, and thus f is a valid p-to-p* calibrator. To show its admissibility, it suffices to notice that f is left-continuous (lower semi-continuous) on $[0, 1]$, and if $g \leq f$ and $g \neq f$, then $\int_0^1 g(u) du < 1/2$, implying that $g(U)$ cannot be a p*-variable. \square

Proof of Theorem 6.2. (i) Let f be a convex p-to-e calibrator. Note that $-f$ is increasing and concave. For any $[0, 1]$ -valued p*-variable P , by definition, we have $\mathbb{E}[-f(P)] \geq \mathbb{E}[-f(U)]$. Hence,

$$\mathbb{E}[f(P)] = -\mathbb{E}[-f(P)] \leq -\mathbb{E}[-f(U)] = \mathbb{E}[f(U)] \leq 1.$$

Since a $[0, \infty)$ -valued p*-variable is first-order stochastically larger than some $[0, 1]$ -valued p*-variable (e.g., [29, Theorem 4.A.6]), we know that $\mathbb{E}[f(P)] \leq 1$ for all p*-variables P . Thus, f is a p*-to-e calibrator.

Next, we show the statement on admissibility. An admissible convex p-to-e calibrator f is a p*-to-e calibrator as shown above. Since the class of p-to-e calibrators is larger than the class of p*-to-e calibrators, f is not strictly dominated by any p*-to-e calibrator.

- (ii) We only need to show the “only if” direction, since the “if” direction is implied by (i). Suppose that a non-convex function f is an admissible p-to-e calibrator. Since f is not convex, there exist two points $t, s \in [0, 1]$ such that

$$f(t) + f(s) < 2f\left(\frac{t+s}{2}\right).$$

Left-continuity of f implies that there exists $\epsilon \in (0, |t-s|)$ such that

$$f(t-\epsilon) + f(s-\epsilon) < 2f\left(\frac{t+s}{2}\right).$$

Note that

$$\int_{t-\epsilon}^t \left(f\left(\frac{t+s}{2}\right) - f(u)\right) du \geq \epsilon \left(f\left(\frac{t+s}{2}\right) - f(t-\epsilon)\right),$$

and the inequality also holds if the positions of s and t are flipped. Hence, by letting $A = [t-\epsilon, t] \cup [s-\epsilon, t]$, we have

$$\begin{aligned} & \int_A \left(f\left(\frac{t+s}{2}\right) - f(u)\right) du \\ & \geq \epsilon \left(2f\left(\frac{t+s}{2}\right) - f(t-\epsilon) - f(s-\epsilon)\right) > 0. \end{aligned} \quad (22)$$

Let U be a uniform random variable on $[0, 1]$ and P be given by

$$P = U \mathbb{1}_{\{U \notin A\}} + \frac{t+s}{2} \mathbb{1}_{\{U \in A\}}.$$

For any increasing concave function g and $x \in [t-\epsilon, t]$ and $y \in [s-\epsilon, s]$, we have

$$2g\left(\frac{t+s}{2}\right) \geq g(t) + g(s) \geq g(x) + g(y).$$

Therefore, $\mathbb{E}[g(P)] \geq \mathbb{E}[g(U)]$, and hence $U \leq_2 P$. Thus, P is a p*-variable. Moreover, using (22), we have

$$\mathbb{E}[f(P)] = \int_0^1 f(u) du + \int_A \left(f\left(\frac{t+s}{2}\right) - f(u)\right) du > \int_0^1 f(u) du = 1.$$

Hence, f is not a p^* -to-e calibrator. Thus, f has to be convex if it is both an admissible p -to-e calibrator and a p^* -to-e calibrator.

- (iii) First, we show that $f : e \mapsto (2e)^{-1} \wedge 1$ is an e-to- p^* calibrator. Clearly, it suffices to show that $1/(2E)$ is a p^* -variable for any e-variable E with mean 1, since any e-variable with mean less than 1 is dominated by an e-variable with mean 1. Let δ_x be the point-mass at x .

Assume that E has a two-point distribution (including the point-mass δ_1 as a special case). With $\mathbb{E}[E] = 1$, the distribution F_E of E can be characterized with two parameters $p \in (0, 1)$ and $a \in (0, 1/p]$ via

$$F_E = p\delta_{1+(1-p)a} + (1-p)\delta_{1-pa}.$$

The distribution F_P of $P := 1/(2E)$ (we allow P to take the value ∞ in case $a = 1/p$) is given by

$$F_P = p\delta_{1/(2+2(1-p)a)} + (1-p)\delta_{1/(2-2pa)}.$$

Let G_P be the left-quantile function of P on $(0, 1)$. We have

$$G_P(t) = \frac{1}{2+2(1-p)a} \mathbb{1}_{\{t \in (0, p]\}} + \frac{1}{2+2pa} \mathbb{1}_{\{t \in (p, 1)\}}.$$

Define two functions g and h on $[0, 1]$ by $g(v) := \int_0^v G_P(u) du$ and $h(v) := v^2/2$. For $v \in (0, p]$, we have, using $a \leq 1/p$,

$$\begin{aligned} g(v) &= \int_0^v G_P(u) du \\ &= \frac{v}{2+2(1-p)a} \geq \frac{v}{2+2(1-p)/p} = \frac{vp}{2} \geq \frac{v^2}{2} = h(v). \end{aligned}$$

Moreover, Jensen's inequality gives

$$g(1) = \int_0^1 G_P(u) du = \mathbb{E}[P] = \mathbb{E}\left[\frac{1}{2E}\right] \geq \frac{1}{2\mathbb{E}[E]} = \frac{1}{2} = h(1).$$

Since g is linear on $[p, 1]$, and h is convex, $g(p) \leq h(p)$ and $g(1) \leq h(1)$ imply $g(v) \leq h(v)$ for all $v \in [p, 1]$. Therefore, we conclude that $g \leq h$ on $[0, 1]$, namely

$$\int_0^v G_P(u) du \leq \frac{v^2}{2} \quad \text{for all } v \in [0, 1].$$

Using (1), we have that P is a p^* -variable.

For a general e-variable E with mean 1, its distribution can be rewritten as a mixture of two-point distributions with mean 1 by Lemma 2.7 of [38]. Since the set of p^* -variables is closed under distribution mixtures (Proposition 3.3), we know that $f(E)$ is a p^* -variable. Hence, f is an e-to- p^* calibrator.

To show that f essentially dominates all other e-to- p^* calibrators, we take any e-to- p^* calibrator f' . Using Theorem 6.1, the function $e \mapsto (2f'(e)) \wedge 1$ is an e-to- p calibrator. Using Proposition 2.2 of [36], any e-to- p calibrator is dominated by $e \mapsto e^{-1} \wedge 1$, and hence

$$(2f'(e)) \wedge 1 \geq e^{-1} \wedge 1 \quad \text{for } e \in [0, \infty),$$

which in term gives $f'(e) \geq (2e)^{-1}$ for $e \geq 1$. Since f' is decreasing, we know that $f'(e) < 1/2$ implies $e > 1$. For any $e \geq 0$ with $f'(e) < 1/2$, we have $f'(e) \geq f(e)$, and thus f essentially dominates f' . \square

Proof of Proposition 6.4. The first statement is implied by Theorem 6.2 (iii). For the second statement, we note that $1/2$ is a p^* -variable. For $\alpha \in (0, 1)$, applying (23) with $P = 1/2$ and $V = \alpha E$, we obtained a p^* -test with a random threshold with mean at most α . Using Proposition B.5, this test has size at most α , that is, $\mathbb{P}(2E \geq 1/\alpha) \leq \alpha$. Hence, P is a p -variable. \square

Appendix B: Randomized p^* -test and applications

In this appendix we discuss several applications of testing with p^* -values and randomization.

B.1. Randomized p^* -test

We first introduce the a generalized version of the randomized p^* -test in Proposition 4.5. The following density condition (DP) for a $(0, 1)$ -valued random variable V will be useful.

(DP) V has a decreasing density function on $(0, 1)$.

The canonical choice of V satisfying (DP) is a uniform random variable on $[0, 2\alpha]$ for $\alpha \in (0, 1/2]$, which we will explain later. For a $(0, 1)$ -valued random variable V with mean α and a p^* -variable P independent of V , we consider the test

$$\text{rejecting the null hypothesis} \iff P \leq V. \quad (23)$$

The following theorem justifies the validity of the test (23) with the necessary and sufficient condition (DP).

Proposition B.1. *Suppose that V is a $(0, 1)$ -valued random variable with mean α .*

- (i) *The test (23) has size at most α for all p^* -variables P independent of V if and only if V satisfies (DP).*
- (ii) *The test (23) has size at most α for all p -variables P independent of V , and the size is precisely α if P is uniformly distributed on $[0, 1]$.*

The proof of Proposition B.1 relies on the following lemma.

Lemma B.2. *For any non-negative random variable V with a decreasing density function on $(0, \infty)$ (with possibly a probability mass at 0) and any p^* -variable P independent of V , we have $\mathbb{P}(P \leq V) \leq \mathbb{E}[V]$.*

Proof. Let F_V be the distribution function of V , which is an increasing concave function on $[0, \infty)$ because of the decreasing density. Since P is a p^* -variable, we have $\mathbb{E}[F_V(P)] \geq \int_0^1 F_V(u) du$. Therefore,

$$\mathbb{P}(P \leq V) = \mathbb{E}[\mathbb{P}(P \leq V|P)] = \mathbb{E}[1 - F_V(P)] \leq \int_0^1 (1 - F_V(u)) du = \mathbb{E}[V].$$

Hence, the statement in the lemma holds. \square

Proof of Proposition B.1. We first note that (ii) is straightforward: for a uniform random variable U on $[0, 1]$ independent of V , then $\mathbb{P}(U \leq V) = \mathbb{E}[V]$. If $P \geq_1 U$, then $\mathbb{P}(P \leq V) \leq \mathbb{P}(U \leq V) \leq \mathbb{E}[V]$.

The “if” statement of point (i) directly follows from Lemma B.2, noting that condition (DP) is stronger than the condition in Lemma B.2. Below, we show the “only if” statement of point (i).

Let F_V be the distribution function of V and U be a uniform random variable on $[0, 1]$. Suppose that F_V is not concave on $(0, 1)$. It follows that there exists $x, y \in (0, 1)$ such that $F_V(x) + F_V(y) > 2F_V((x+y)/2)$. By the right-continuity of F_V , there exists $\epsilon \in (0, |x-y|)$ such that

$$F_V(x) + F_V(y) > 2F_V\left(\frac{x+y+\epsilon}{2}\right). \quad (24)$$

Let $A = [x, x+\epsilon]$ and $B = [y, y+\epsilon]$, which are disjoint intervals. Define a random variable P by

$$P = U\mathbb{1}_{\{U \notin A \cup B\}} + \frac{x+y+\epsilon}{2}\mathbb{1}_{\{U \in A \cup B\}}. \quad (25)$$

We check that P defined by (25) is a p^* -variable. For any concave function g , Jensen’s inequality gives

$$\begin{aligned} \mathbb{E}[g(P)] &= \int_{[0,1] \setminus (A \cup B)} g(u) du + 2\epsilon g\left(\frac{x+y+\epsilon}{2}\right) \\ &\geq \int_{[0,1] \setminus (A \cup B)} g(u) du + \int_{A \cup B} g(u) du = \mathbb{E}[g(U)]. \end{aligned}$$

Hence, P is a p^* -variable. It follows from (24) and (25) that

$$\begin{aligned} \mathbb{E}[F_V(P)] &= \int_{[0,1] \setminus (A \cup B)} F_V(u) du + \int_{A \cup B} F_V\left(\frac{x+y+\epsilon}{2}\right) du \\ &< \int_{[0,1] \setminus (A \cup B)} F_V(u) du + \int_{A \cup B} \frac{F_V(x) + F_V(y)}{2} du \\ &= \int_{[0,1] \setminus (A \cup B)} F_V(u) du + \epsilon(F_V(x) + F_V(y)) \\ &= \int_{[0,1] \setminus (A \cup B)} F_V(u) du + \int_A F_V(x) du + \int_B F_V(y) du \\ &\leq \int_0^1 F_V(u) du = 1 - \mathbb{E}[V] = 1 - \alpha. \end{aligned}$$

Therefore,

$$\mathbb{P}(P \leq V) = 1 - \mathbb{E}[F_V(P)] > 1 - (1 - \alpha) = \alpha.$$

Since this contracts the validity requirement, we know that V has to have a concave distribution function, and hence a decreasing density on $(0, 1)$. \square

Lemma B.2 gives $\mathbb{P}(P \leq V) \leq \mathbb{E}[V]$ for V possibly taking values larger than 1 and possibly having a probability mass at 0. We are not interested designing a random threshold with positive probability to be 0 or larger than 1, but this result will become helpful in Section 7. Since condition (DP) implies $\mathbb{E}[V] \leq 1/2$, we will assume $\alpha \in (0, 1/2]$, which is certainly harmless for practice.

With the help of Proposition B.1, we formally define α -random thresholds and the randomized p^* -test.

Definition B.3. For a significance level $\alpha \in (0, 1/2]$, an α -random threshold V is a $(0, 1)$ -valued random variable independent of the test statistics (a p^* -variable P in this section) with mean α satisfying (DP). For an α -random threshold V and a p^* -variable P , the randomized p^* -test is given by (23), i.e., rejecting the null hypothesis $\iff P \leq V$.

Proposition B.1 implies that the randomized p^* -test always has size at most α , just like the classic p -test (5). Since the randomized p^* -test (23) has size equal to α if P is uniformly distributed on $[0, 1]$, the size α of the randomized p^* -test cannot be improved in general.

As mentioned in Section 4.4, randomization is generally undesirable in testing. Like any other randomized methods, different scientists may arrive at different statistical conclusions by the randomized p^* -test for the same data set generating the p^* -value. Because of assumption (DP), which is necessary for validity by Proposition B.1, we cannot reduce the α -random threshold V to a deterministic α . This undesirable feature is the price one has to pay when a p -variable is weakened to a p^* -variable.

If one needs to test with a deterministic threshold, then $\alpha/2$ needs to be used instead of α . In other words, the test

$$\text{rejecting the null hypothesis} \iff P \leq \alpha/2 \tag{26}$$

has size α for all p^* -variable P . The validity of (26) was noted by Rüschendorf [24], and it is a direct consequence of Proposition 5.1. Unlike the random threshold $U[0, 2\alpha]$ which gives a size precisely α in realistic situations, the deterministic threshold $\alpha/2$ is often overly conservative in practice (see discussions in [19, Section 5]), but it cannot be improved in general when testing with the average of p -variables ([35, Proposition 3]); recall that the average of p -variables is a p^* -variable.

We will see an application of the randomized p^* -test in Section B.2, leading to new tests on the weighted average of p -values, which can be made deterministic if one of the p -values is independent of the others. Moreover, the randomized p^* -test can be used to improve the power of tests with e -values and martingales in Section 7.

As mentioned above, the extra randomness introduced by the random threshold V is often considered undesirable. One may wish to choose V such that the randomness is minimized. The next result shows that $V \sim U[0, 2\alpha]$ is the optimal choice if the randomness is measured by variance or convex order.

Proposition B.4. For any α -random threshold V , we have $\text{var}(V) \geq \alpha^2/3$, and this smallest variance is attained by $V^* \sim U[0, 2\alpha]$. Moreover, $\mathbb{E}[f(V)] \geq \mathbb{E}[f(V^*)]$ for any convex function f (hence $V \leq_2 V^*$ holds).

Proof. We directly show $\mathbb{E}[f(V)] \geq \mathbb{E}[f(V^*)]$ for all convex functions f , which implies the statement on variance as a special case since the mean of V is fixed as α . Note that V has a concave distribution function F_V on $[0, 1]$, and V^* has a linear distribution function F_{V^*} on $[0, 2\alpha]$. Moreover, they have the same mean. Hence, there exists $t \in [0, 2\alpha]$ such that $F_V(x) \geq F_{V^*}(x)$ for $x \leq t$ and $F_V(x) \leq F_{V^*}(x)$ for $x \geq t$. This condition is sufficient for $\mathbb{E}[f(V)] \geq \mathbb{E}[f(V^*)]$ by Theorem 3.A.44 of [29]. \square

Combining Propositions B.1 and B.4, the canonical choice of the threshold in the randomized p^* -test has a uniform distribution on $[0, 2\alpha]$.

We note that it is also possible to use some V with mean less than α and variance less than $\alpha^2/3$. This reflects a tradeoff between power and variance. Such a random threshold does not necessarily have a decreasing density. For instance, the point-mass at $\alpha/2$ is a valid choice; the next proposition gives some other choices.

Proposition B.5. *Let V be an α -random threshold and V' is a random variable satisfying $V' \leq_1 V$. We have $\mathbb{P}(P \leq V') \leq \alpha$ for arbitrary p*-variable P independent of V' .*

Proof. Let F be the distribution function of \bar{P} . For any increasing function f , we have $\mathbb{E}[f(V')] \leq \mathbb{E}[f(V)]$, which follows from $V' \leq_1 V$. Hence, we have

$$\mathbb{P}(P \leq V') = \mathbb{E}[F(V')] \leq \mathbb{E}[F(V)] = \mathbb{P}(P \leq V) \leq \alpha,$$

where the last inequality follows from Proposition B.1. \square

Proposition B.5 can be applied to a special situation where a p-variable P and an independent p*-variable P^* are available for the same null hypothesis. Note that in this case $2\alpha(1 - P) \leq_1 V \sim U[0, 2\alpha]$. Hence, by Proposition B.5, the test

$$\text{rejecting the null hypothesis} \iff \frac{P^*}{2(1 - P)} \leq \alpha. \quad (27)$$

has size at most α . Alternatively, using the fact that $\mathbb{P}(P \leq 2\alpha(1 - P^*)) \leq \alpha$ implied by Proposition B.1 (ii), we can design a test

$$\text{rejecting the null hypothesis} \iff \frac{P}{2(1 - P^*)} \leq \alpha. \quad (28)$$

The tests (27) and (28) both have a deterministic threshold of α . This observation is useful in Section B.2.

B.2. Testing with averages of p-values

In this section we illustrate applications of the randomized p*-test to tests with averages of dependent p-values.

Let P_1, \dots, P_K be K p-variables for a global null hypothesis and they are generally not independent. Vovk and Wang [35] proposed testing using generalized means of the p-values, so that the type-I error is controlled at a level α under arbitrary dependence. We focus on the weighted (arithmetic) average $\bar{P} := \sum_{k=1}^K w_k P_k$ for some weights $w_1, \dots, w_K \geq 0$ with $\sum_{k=1}^K w_k = 1$. In case $w_1 = \dots = w_K = 1/K$, we speak of the arithmetic average.

The method of [35] on arithmetic average is given by

$$\text{rejecting the null hypothesis} \iff \bar{P} \leq \alpha/2. \quad (29)$$

We will call (29) the *arithmetic averaging test*. The extra factor of $1/2$ is needed to compensate for arbitrary dependence among p-values. Since \bar{P} is a p*-variable by Theorem 4.1, the test (29) is a special case of (26). This method is quite conservative, and it often has relatively low power compared to the Bonferroni correction and other similar methods unless p-values are very highly correlated, as illustrated by the numerical experiments in [35].

To enhance the power of the test (29), we apply the randomized p*-test in Section 4.4 to design the *randomized averaging test* by

$$\text{rejecting the null hypothesis} \iff \bar{P} \leq V, \quad (30)$$

where V is an α -random threshold independent of (P_1, \dots, P_K) . Comparing the fixed-threshold test (29) and the randomized averaging test (30) with $V \sim U[0, 2\alpha]$, there is a 3/4 probability that the randomized averaging test has a better power, with the price of randomization.

Next, we consider a special situation, where a p-variable among P_1, \dots, P_K is independent of the others under the null hypothesis. In this case, we can apply (27), and the resulting test is no longer randomized, as it is determined by the observed p-values.

Without loss of generality, assume that P_1 is independent of (P_2, \dots, P_K) . Let $\bar{P}_{(-1)} := \sum_{k=2}^K w_k P_k$ be a weighted average of (P_2, \dots, P_K) . Using $\bar{P}_{(-1)}$ as the p*-variable, the test (27) becomes

$$\text{rejecting the null hypothesis} \iff \bar{P}_{(-1)} + 2\alpha P_1 \leq 2\alpha. \quad (31)$$

Following directly from the validity of (27), for any p-variables P_1, \dots, P_K with P_1 independent of (P_2, \dots, P_K) , the test (31) has size at most α .

Comparing (31) with (29), we rewrite (31) as

$$\text{rejecting the null hypothesis} \iff \bar{P} := \sum_{k=1}^n w'_k P_k \leq \frac{2\alpha}{1+2\alpha},$$

where

$$w'_1 = \frac{2\alpha}{1+2\alpha} \quad \text{and} \quad w'_k = \frac{w_k}{1+2\alpha}, \quad k = 2, \dots, K.$$

Note that \bar{P} is a weighted average of (P_1, \dots, P_K) . Since α is small, the rejection threshold is increased by almost three times, compared to the test (29) applied to $\sum_{k=1}^n w'_k P_k$ using the threshold $\alpha/2$. For this reason, we will call (31) the *enhanced averaging test*.

In particular, if $2\alpha(K-1) = 1$, and $\bar{P}_{(-1)}$ is the arithmetic average of (P_2, \dots, P_K) , then \bar{P} is the arithmetic average of (P_1, \dots, P_K) . For instance, if $K = 51$ and $\alpha = 0.01$, then the rejection condition for test (31) is $\bar{P} \leq 1/51$, and the rejection condition for (29) is $\bar{P} \leq 1/200$.

B.3. Simulation experiments

We compare by simulation the performance of a few tests via merging p-values. For the purpose of illustration, we conduct correlated z-tests for the mean of normal samples with variance 1. More precisely, the null hypothesis H_0 is $N(0, 1)$ and the alternative is $N(\delta, 1)$ for some $\delta > 0$. The p-variables P_1, \dots, P_K are specified as $P_k = 1 - \Phi(X_k)$ from the Neyman-Pearson lemma, where Φ is the standard normal distribution function, and X_1, \dots, X_K are generated from $N(\delta, 1)$ with pair-wise correlation ρ . As illustrated by the numerical studies in [35], the arithmetic average test performs poorly unless p-values are strongly correlated. Therefore, we consider the cases where p-values are highly correlated, e.g., parallel experiments with shared data or scientific objects. We set $\rho = 0.9$ in our simulation studies; this choice is harmless as we are interested in the relative performance of the averaging methods in this section, instead of their performance against other methods (such as method of Simes [30]) that are known work well for lightly correlated or independent p-values.

The significance level α is set to be 0.01. For a comparison, we consider the following tests:

- (a) the arithmetic averaging test (29): reject H_0 if $\bar{P} \leq \alpha/2$;
- (b) the randomized averaging test (30): reject H_0 if $\bar{P} \leq V$ where $V \sim U[0, 2\alpha]$ independent of \bar{P} ;
- (c) the Bonferroni method: reject H_0 if $\min(P_1, \dots, P_K) \leq \alpha/K$;
- (d) the Simes method: reject H_0 if $\min_k (K P_{(k)}/k) \leq \alpha$ where $P_{(k)}$ is the k -th smallest p-value;

- (e) the harmonic averaging test of [35]: reject H_0 if $(\sum_{k=1}^K P_k^{-1})^{-1} \leq \alpha/c_K$ where $c_K > 1$ is a constant in [35, Proposition 6].

The validity (size no larger than α) of the Simes method is guaranteed under some dependence conditions on the p-values; see [26,3]. Moreover, as shown recently by Vovk et al. [34, Theorem 6], the Simes method dominates any symmetric and deterministic p-merging method valid for arbitrary dependence (such as the (a), (c) and (e); the Simes method itself is not valid for arbitrary dependence).

In the second setting, we assume that one of the p-variables (P_1 without loss of generality) is independent of the rest, and the rest p-variables have a pair-wise correlation of $\rho = 0.9$. For this setting, we further include

- (f) the enhanced averaging test (31): reject H_0 if $\bar{P}_{(-1)} + 2\alpha P_1 \leq 2\alpha$.

The power (i.e., the probability of rejection) of each test is computed from the average of 10,000 replications for varying signal strength δ and for $K \in \{20, 100, 500\}$. Results are reported in Figure 5.

In the first setting of correlated p-values, the randomized averaging test (b) improves the performance of (a) uniformly, at the price of randomization. The Bonferroni method (d) and the harmonic averaging test (e) perform poorly and are both penalized significantly as K increases. None of these methods visibly outperforms the Simes method, although in some situations the test (b) performs comparably to the Simes method.

In the second setting where an independent p-value exists, the enhanced arithmetic test (f) performs quite well; it outperforms the Simes method for most parameter values especially for small signal strength δ . This illustrates the significant improvement via incorporating an independent p-value.

We remark that the averaging methods (a), (b) and (f) should not be used in situations in which correlation among p-values is known to be not very strong. This is because the arithmetic mean does not benefit from an increasing number of independent p-values of similar strength, unlike the methods of Bonferroni and Simes.

Acknowledgments

The author thanks an Editor, an Associate Editor, three anonymous referees, Ilmun Kim, Tiantian Mao, Marcel Nutz, Aaditya Ramdas, Vladimir Vovk, Qiuqi Wang, Qinyu Wu and Zhenyuan Zhang for helpful comments and technical discussions on an earlier version of the paper.

Funding

The author acknowledges financial support from and the Natural Sciences and Engineering Research Council of Canada (RGPIN-2018-03823, RGPAS-2018-522590).

References

- [1] Bates, S., Candès, E., Lei, L., Romano, Y. and Sesia, M. (2021). Testing for outliers with conformal p-values. *arXiv*: 2104.08279.
- [2] Benjamini, Y. and Hochberg, Y. (1997). Multiple hypotheses testing with weights. *Scand. J. Stat.* **24**(3) 407–418.
- [3] Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.* **29**(4) 1165–1188.

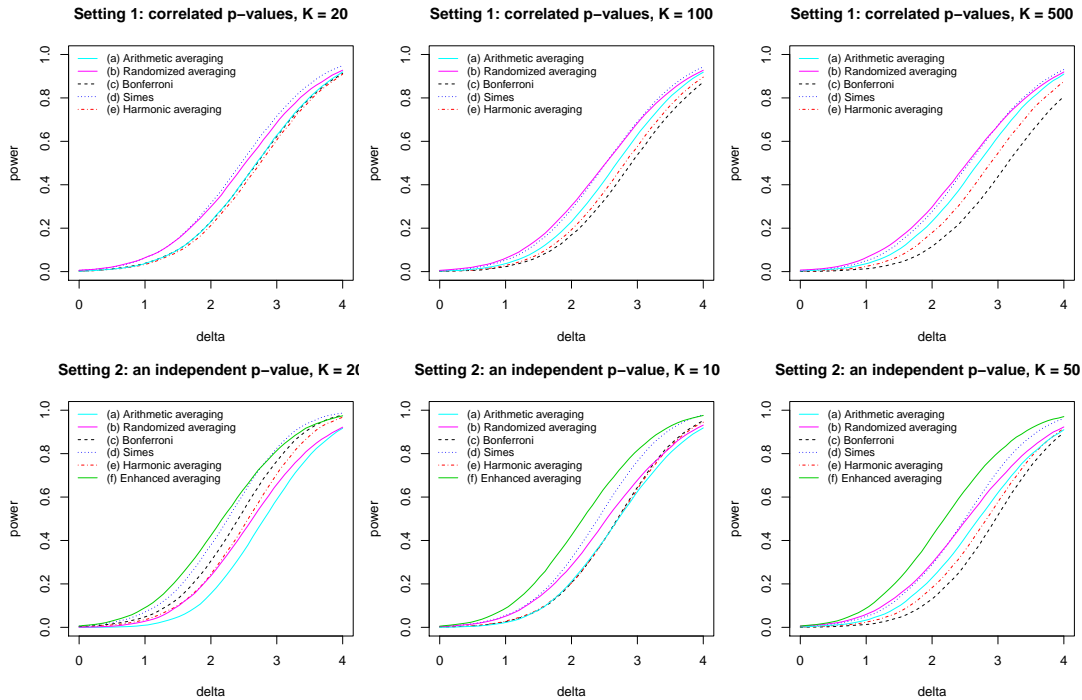


Figure 5: Tests based on combining p-values

- [4] Chen, Y., Liu, P. and Tan, K. S. and Wang, R. (2023). Trade-off between validity and efficiency of merging p-values under arbitrary dependence. *Statist. Sinica*, **33**, 851–872.
- [5] Howard, S. R., Ramdas, A., McAuliffe, J. and Sekhon, J. (2021). Time-uniform, nonparametric, nonasymptotic confidence sequences. *Ann. Stat.* **49**(2) 1055–1080.
- [6] Döhler, S., Durand, G. and Roquain, E. (2018). New FDR bounds for discrete and heterogeneous tests. *Electron. J. Stat.* **12**(1) 1867–1900.
- [7] Duan, B., Ramdas, A., Balakrishnan, S. and Wasserman, L. (2020). Interactive martingale tests for the global null. *Electron. J. Stat.* **14**(2) 4489–4551.
- [8] Efron, B. (2010). *Large-scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press.
- [9] Genovese, R. and Wasserman, L. (2004). A stochastic process approach to false discovery control. *Ann. Stat.* **32** 1035–1061.
- [10] Goeman, J. J. and Solari, A. (2011). Multiple testing for exploratory research. *Stat. Sci.* **26**(4) 584–597.
- [11] Grünwald, P., de Heide, R. and Koolen, W. M. (2020). Safe testing. *arXiv*: 1906.07801v2.
- [12] Habiger, J. D. (2015). Multiple test functions and adjusted p-values for test statistics with discrete distributions. *J. Stat. Plann. Inference* **167** 1–13.
- [13] Huber, M. (2019). Halving the bounds for the Markov, Chebyshev, and Chernoff Inequalities using smoothing. *Am. Math. Mon.* **126**(10) 915–927.
- [14] Lancaster, H. O. (1952). Statistical control of counting experiments. *Biometrika* **39** 419–422.
- [15] Liu, F. and Wang, R. (2021). A theory for measures of tail risk. *Math. Oper. Res.* **46**(3) 1109–1128.
- [16] Liu, Y. and Xie, J. (2020). Cauchy combination test: A powerful test with analytic p-value calculation under arbitrary dependency structures. *J. Am. Stat. Assoc.* **115**(529) 393–402.
- [17] Mao, T., Wang, B. and Wang, R. (2019). Sums of uniform random variables. *J. Appl. Probab.* **56**(3) 918–936.
- [18] Nutz, M., Wang, R., Zhang, Z. (2022). Martingale transports and Monge maps. *arXiv*: 2209.14432.

- [19] Meng, X. L. (1994). Posterior predictive p -values. *Ann. Stat.* **22**(3) 1142–1160.
- [20] Müller, A. and Stoyan, D. (2002). *Comparison Methods for Stochastic Models and Risks*. Wiley, England.
- [21] Ramdas, A. K., Barber, R. F., Wainwright, M. J. and Jordan, M. I. (2019). A unified treatment of multiple testing with prior knowledge using the p -filter. *Ann. Stat.* **47**(5) 2790–2821.
- [22] Ramdas, A., Grünwald, P., Vovk, V. and Shafer, G. (2022). Game-theoretic statistics and safe anytime-valid inference. *arXiv*: 2210.01948.
- [23] Rubin-Delanchy, P., Heard, N. A. and Lawson, D. J. (2019). Meta-analysis of mid- p -values: Some new results based on the convex order. *J. Am. Stat. Assoc.* **114**(527) 1105–1112.
- [24] Rüschendorf, L. (1982). Random variables with maximum sums. *Adv. Appl. Probab.* **14**(3) 623–632.
- [25] Rüschendorf, L. (2013). *Mathematical Risk Analysis. Dependence, Risk Bounds, Optimal Allocations and Portfolios*. Springer, Heidelberg.
- [26] Sarkar, S. K. (1998). Some probability inequalities for ordered MTP2 random variables: a proof of the Simes conjecture. *Ann. Stat.* **26**(2) 494–504.
- [27] Shafer, G. (2021). The language of betting as a strategy for statistical and scientific communication. *J. R. Stat. Soc. A* **184**(2) 407–431.
- [28] Shafer, G., Shen, A., Vereshchagin, N. and Vovk, V. (2011). Test martingales, Bayes factors, and p -values. *Stat. Sci.* **26** 84–101.
- [29] Shaked, M. and Shanthikumar, J. G. (2007). *Stochastic Orders*. Springer Series in Statistics.
- [30] Simes, R. J. (1986). An improved Bonferroni procedure for multiple tests of significance. *Biometrika* **73** 751–754.
- [31] Strassen, V. (1965). The existence of probability measures with given marginals. *Ann. Math. Stat.* **36**(2) 423–439.
- [32] Vovk, V. (2021). Testing randomness online. *Stat. Sci.* **36**(4) 595–611.
- [33] Vovk, V., Gammerman, A. and Shafer, G. (2005). *Algorithmic Learning in a Random World*. Springer, New York.
- [34] Vovk, V., Wang, B. and Wang, R. (2022). Admissible ways of merging p -values under arbitrary dependence. *Ann. Stat.* **50**(1) 351–375.
- [35] Vovk, V. and Wang, R. (2020). Combining p -values via averaging. *Biometrika* **107**(4) 791–808.
- [36] Vovk, V. and Wang, R. (2021). E-values: Calibration, combination, and applications. *Ann. Stat.* **49**(3) 1736–1754.
- [37] Vovk, V. and Wang, R. (2023). Confidence and discoveries with e-values. *Stat. Sci.*, forthcoming.
- [38] Wang, R. (2014). Sum of arbitrarily dependent random variables. *Electron. J. Probab.* **19**(84) 1–18.
- [39] Wang, R. and Ramdas, A. (2022). False discovery rate control with e-values. *J. R. Stat. Soc. B* **84**(3) 822–852.
- [40] Wasserman, L., Ramdas, A. and Balakrishnan, S. (2020). Universal inference. *Proc. Natl. Acad. Sci. U.S.A.* **117**(29) 16880–16890.
- [41] Wilson, D. J. (2019). The harmonic mean p -value for combining dependent tests. *Proc. Natl. Acad. Sci. U.S.A.* **116**(4) 1195–1200.