

FUNDAMENTALS OF P-VALUES: INTRODUCTION

YURI GUREVICH AND VLADIMIR VOVK

ABSTRACT. We explain the concept of p-values presupposing only rudimentary probability theory. We also use the occasion to introduce the notion of p-function, so that p-values are values of a p-function.

The explanation is restricted to the discrete case with no outcomes of zero probability. We are going to address the general case elsewhere.

1. PRELUDE

Q¹: The Cournot Principle that one of you wrote about [3] caught my attention: “It is practically certain that a predicted event of small probability does not happen.” How small should the probability be?

A: Traditionally there have been two camps of naive Cournotians, the lax ones using the 5% threshold and strict ones using the 1% threshold.

Q: Why do you call them naive?

A: A better way is to report an appropriate p-value, without committing yourself to a threshold.

Q: Oh yes, I heard about p-values and even tried to read about them but got confused. Do you have a good reference?

A: Cox and Hinkley define p-values, even though they do not use the term, in §3 of their 1974 book “Theoretical Statistics” [2]. That is the best reference that we have.

Q: How about explaining the concept to me right now? But please take into account that, while I have been exposed to probability theory and logic, I have not studied statistics.

A: The general case is somewhat involved [4], but the discrete case is important and simple enough to explain the concept, especially if one presumes that every outcome has positive probability.

¹Q and A are Quisani (a former student of the first author) and the authors respectively.

2. PROBABILITY TRIALS

Recall that a discrete probability space is given by a nonempty set Ω and a function \mathbf{P} from Ω to the real interval $[0, 1]$ satisfying the following three requirements.

- (1) Ω is countable, i.e. finite or infinite countable.
- (2) Every $\mathbf{P}(x) \geq 0$.
- (3) $\sum_{x \in \Omega} \mathbf{P}(x) = 1$.

To simplify our exposition, we replace (2) with a stronger and relatively innocuous requirement

- (2') Every $\mathbf{P}(x) > 0$.

The function \mathbf{P} naturally extends to subsets E of Ω : $\mathbf{P}(E) = \sum \{P(x) : x \in E\}$.

Think of the probability space (Ω, \mathbf{P}) as a probability trial. Elements of Ω are outcomes, subsets of Ω are events, and \mathbf{P} assigns probabilities to outcomes and events. The probability distribution \mathbf{P} is the *null hypothesis* of the trial. Think of it as an alleged probability distribution. The purpose of the trial is to test \mathbf{P} .

Q: I guess, a predicted event of small probability, if and when it occurs, provides falsifying evidence against the null hypothesis.

A: Yes, but we do not need to define what probabilities count as small. For now, it suffices to say this: The smaller the probability of the event in question, the stronger the impugning power of the event. Suppose for example that $\Omega = \{a, b, c\}$ and that $\mathbf{P}(a) = 990/1000$, $\mathbf{P}(b) = 9/1000$ and $\mathbf{P}(c) = 1/1000$. We can order the nonempty events in the order of increasing impugning power: $\{a, b, c\}$, $\{a, b\}$, $\{a, c\}$, $\{a\}$, $\{b, c\}$, $\{b\}$, $\{c\}$.

Q: But the event $\{a, b, c\}$ has zero impugning power.

A: True, and we don't say that it has any. We just say that its impugning power is less than that of the other 6 events. Also, we ignore the empty event because it does not occur.

3. THE LOGIC ANGLE

Q: For logicians, a discrete probabilistic trial (Ω, \mathbf{P}) is a relatively simple logic structure, an extension of real arithmetic with a separate sort Outcome and a function \mathbf{P} of type Outcome \rightarrow Real, subject to requirements (1)–(3) of §2.

A: In applications, a trial comes with additional information. Here are two examples, in a simplified form, from the paper you cited.

Coin Example. A coin is tossed 42 times. The null hypothesis is that all 2^{42} outcomes are equally probable. The actual outcome has 41 heads and just one tail. Clearly the impugning power of that outcome exceeds that of a random outcome.

Lottery Example. 4,000,000 people bought 10,000,000 tickets. Any of them could be the winner; thus these are 4,000,000 potential outcomes. The null hypothesis is that the probability of a person to win is proportional to the number of his or her tickets. The wife Donna of the lottery organizer John bought 3 tickets and won the lottery. Clearly the impugning power of that outcome exceeds that of a random outcome.

Q: Additional information complicates the matter. Clearly there are countless different kinds of additional information. How do statisticians deal with them?

A: They simplify the situation. They replace any additional structure with a linear preorder of the outcomes and restrict attention to the downward closed events.

Q: Wow, this is quite a simplification.

A: Often the linear preorder is given implicitly, by means of a test statistic or a nested family of events.

Q: Please remind me the necessary definitions.

4. TEST TOOLS

Let (Ω, \mathbf{P}) be a probability trial. A *preorder* of a nonempty set Ω is a binary relation \leq on Ω that is reflexive and transitive so that

- $x \leq x$ for every x , and
- if $x \leq y$ and $y \leq z$ then $x \leq z$.

A preorder \leq of Ω is *linear* if

- $x \leq y$ or $y \leq x$ for all x, y .

Definition 1.

- A *test order* of Ω is a linear preorder on Ω .
- A *test pyramid* over Ω is a family of events linearly ordered by inclusion.
- A *test statistic* on Ω is a function from Ω to the reals.

Coin Example, cont. Let $\text{Heads}(x)$ and $\text{Tails}(x)$ be the numbers of the heads and tails in an outcome x respectively. The function $M(x) = \text{Min}\{\text{Heads}(x), \text{Tails}(x)\}$ is a test statistic. The relation $M(x) \leq M(y)$ is a linear preorder of the outcomes. Events $\{x : M(x) \leq n\}$, where n ranges from 0 to 42 form a test pyramid.

Given a linear preorder \leq on Ω , an event $E \subseteq \Omega$ is *downward closed* if $x \leq y \in E$ implies $x \in E$ for all x, y . The downward closed events are nested: if E_1, E_2 are downward closed then either $E_1 \subseteq E_2$ or $E_2 \subseteq E_1$. Indeed if $E_2 \not\subseteq E_1$ and $x \in E_2 - E_1$ then $E_1 \subseteq \{y : y \leq x\} \subseteq E_2$.

Note that there may be distinct outcomes x, y with $x \leq y$ and $y \leq x$; such outcomes x, y called *quasi-equivalent*. Linear preorders are also known as linear quasi-orders.

Test orders, test pyramids and test statistics are three tools to measure the relative strength of the impugning evidence provided by a given outcome or event. We adopt the following convention.

- For a given test order, smaller outcomes provide stronger impugning evidence.
- For a given test statistic, smaller values provide stronger impugning evidence.

And of course, for a given test pyramid of events, smaller events provide stronger impugning evidence.

In the following two sections we show that the three tools are equivalent in an appropriate sense.

5. TEST ORDERS AND TEST PYRAMIDS

If \leq is a test order, let $[\leq x]$ be the downward closure $\{y : y \leq x\}$ of x .

Definition 2.

- Every test order \leq *induces* a test pyramid, namely the family of events $[\leq x]$.
- Every test pyramid \mathcal{F} *induces* a test order

$$\{(x, y) : (\forall E \in \mathcal{F})[y \in E \implies x \in E]\}$$

In words, the formula $(\forall E \in \mathcal{F})[y \in E \implies x \in E]$ says that every member of \mathcal{F} that contains y contains x as well.

Lemma 3. *If a test order \leq induces a test pyramid \mathcal{F} then \mathcal{F} induces \leq .*

Proof. Let \preceq be the test order induced by \mathcal{F} .

$$\begin{aligned} x \preceq y &\iff (\forall E \in \mathcal{F})[y \in E \implies x \in E] && \text{by the definition of } x \preceq y \\ &\iff (\forall z \in \Omega)[y \in [\leq z] \implies x \in [\leq z]] && \text{by the definition of } \mathcal{F} \\ &\iff x \in [\leq y] \\ &\iff x \leq y \end{aligned} \quad \square$$

Definition 4.

- Two test pyramids are *equivalent* if they induce the same test order.
- The *canonic version* $\hat{\mathcal{F}}$ of a test pyramid \mathcal{F} is the test pyramid induced by the test order induced by \mathcal{F} .
- A test pyramid \mathcal{F} is *canonic* if it is the canonic version of itself.

Q: Show me a non-canonic test pyramid.

A: Let Ω be the set of numbers $\pm\frac{1}{n}$ where $n = 1, 2, \dots$. The precise definition of the distribution \mathbf{P} on Ω is not important provided that every element of Ω has a positive probability. Let \leq be the standard order on the numbers $\pm\frac{1}{n}$. The desired test pyramid \mathcal{F} is the collection of all subsets of Ω downward closed with respect to \leq . Clearly \mathcal{F} induces the test order \leq . The pyramid $\hat{\mathcal{F}}$, induced by \leq , consists of sets $[\leq x]$ where $x \in \Omega$. It is a proper subcollection of \mathcal{F} . The difference $\mathcal{F} - \hat{\mathcal{F}}$ consists of three sets: \emptyset , $[\leq 0]$ and Ω .

Q: You could have require that a test order \leq induces the pyramid of all the events downward closed with respect to \leq . Then your counter-example would not work.

A: That is true. But then a non-canonic pyramid would be obtained from the pyramid of downward-closed events by omitting one or more of the sets \emptyset , $[\leq 0]$, Ω .

Q: What about canonic test orders?

A: In our current setup every test order can be viewed as canonic. We saw already that every test order is induced by a test pyramid. We'll see below that every test order is induced by a test statistic. This is not so in the general case. In fact, it is not so even in the discrete case with outcomes of zero probability.

Lemma 5. *Every pyramid \mathcal{F} is equivalent to a unique canonic test pyramid, namely to $\hat{\mathcal{F}}$.*

Proof. Let \leq be the test order induced by \mathcal{F} . By the definition, $\hat{\mathcal{F}}$ is the collection of events $[\leq x]$. Clearly $\hat{\mathcal{F}}$ induces \leq , and so it is equivalent to \mathcal{F} .

It remains to check that $\hat{\mathcal{F}}$ is the only canonic test pyramid equivalent to \mathcal{F} . Let \mathcal{F}' be a canonic test pyramid that is equivalent to \mathcal{F} and therefore induces \leq . Since \mathcal{F}' is canonic, it is induced by \leq . But then $\mathcal{F}' = \hat{\mathcal{F}}$. \square

Corollary 6.

- (1) *If a test pyramid \mathcal{F} induces a test order \leq then \leq induces the canonic version $\hat{\mathcal{F}}$ of \mathcal{F} .*
- (2) *If a test order \leq induces a test pyramid \mathcal{F} then \mathcal{F} is canonic.*

Proof. (1) By the definition of $\hat{\mathcal{F}}$.

(2) By Lemma 3, \mathcal{F} induces \leq . By (1), $\mathcal{F} = \hat{\mathcal{F}}$. □

6. TEST ORDERS AND TEST STATISTICS

Definition 7.

- Every test order \leq induces a test statistic $f(x) = \mathbf{P}[\leq x]$.
- Every test statistic f induces a test order $\{(x, y) : f(x) \leq f(y)\}$.

For a test statistics f and a real number r , let $[f \leq r]$ be the event $\{x : f(x) \leq r\}$. In the following lemma, we take advantage of having no zero-probability outcomes.

Lemma 8. *For any test order \leq , we have $x \leq y$ if and only if $\mathbf{P}[\leq x] \leq \mathbf{P}[\leq y]$.*

Proof. If $x \leq y$ then $[\leq x] \subseteq [\leq y]$ and therefore $\mathbf{P}[\leq x] \leq \mathbf{P}[\leq y]$. This establishes the only-if implication. We prove the if implication by contrapositive. Suppose that $x > y$. Then x belongs to $[\leq x]$ but not to $[\leq y]$ so that $[\leq x] \not\subseteq [\leq y]$. Since $\mathbf{P}(x) > 0$, we have $\mathbf{P}[\leq x] > \mathbf{P}[\leq y]$. □

Lemma 9. *For any test statistic f , we have:*

- (1) *If f induces a test order \leq then every $[\leq x] = [f \leq f(x)]$.*
- (2) *$f(x) \leq f(y) \iff \mathbf{P}[f \leq f(x)] \leq \mathbf{P}[f \leq f(y)]$.*

Proof.

$$(1) [\leq x] = \{y : y \leq x\} = \{y : f(y) \leq f(x)\} = [f \leq f(x)].$$

(2) Follows from (1) and Lemma 8. □

Lemma 10. *If a test order \leq induces a test statistic f then f induces \leq .*

Proof. Let \preceq be the test order induced by f .

$$\begin{aligned} x \leq y &\iff \mathbf{P}[\leq x] \leq \mathbf{P}[\leq y] && \text{by Lemma 8,} \\ &\iff f(x) \leq f(y) && \text{by the definition of } f, \\ &\iff x \preceq y && \text{by the definition of } \preceq. \end{aligned} \quad \square$$

Definition 11.

- Two test statistics are *equivalent* if they induce the same test order.
- The *canonic version* of a test statistic f is the test statistic $\hat{f}(x) = \mathbf{P}[f \leq f(x)]$.
- A test statistic is *canonic* if it is the canonic version of itself.

Lemma 12. *Every test statistic f is equivalent to a unique canonic test statistics, namely to \hat{f} .*

Proof. First we check that f is equivalent to \hat{f} .

$$\begin{aligned} f(x) \leq f(y) &\iff \mathbf{P}[f \leq f(x)] \leq \mathbf{P}[f \leq f(y)] && \text{by Lemma 9,} \\ &\iff \hat{f}(x) \leq \hat{f}(y) && \text{by the definition of } \hat{f}. \end{aligned}$$

Second we check that \hat{f} is the only canonic test statistic equivalent to f . If f' is a canonic test statistic equivalent to f then

$$\begin{aligned} f'(x) \leq f'(y) &\iff f(x) \leq f(y) && \text{because } f \text{ and } f' \text{ are equivalent,} \\ &\iff \hat{f}(x) \leq \hat{f}(y) && \text{because } f \text{ and } \hat{f} \text{ are equivalent.} \end{aligned}$$

Thus f' , \hat{f} and f are all equivalent and therefore induce the same order \leq .

$$\begin{aligned} f'(x) &= \mathbf{P}[f' \leq f'(x)] && \text{as } f' \text{ is canonic} \\ &= \mathbf{P}[\leq x] && \text{by Lemma 8, part 1} \\ &= \mathbf{P}[\hat{f} \leq \hat{f}(x)] && \text{by Lemma 8, part 1} \\ &= \hat{f}(x) && \text{as } \hat{f} \text{ is canonic.} \quad \square \end{aligned}$$

Corollary 13.

- If a test statistic f induces a test order \leq then \leq induces the canonic version \hat{f} of f .
- The test statistic induced by any test order is canonic.

Proof.

(1) Let f' be the test statistic induced by \leq . We have

$$\begin{aligned} f'(x) &= \mathbf{P}[\leq x] && \text{by the definition of } f', \\ &= \mathbf{P}[f \leq f(x)] && \text{by Lemma 9,} \\ &= \hat{f}(x) && \text{by the definition of } \hat{f}. \end{aligned}$$

(2) Let f be the test statistic induced by a given test order \leq . By Lemma 10, f induces \leq . By (1), $f = \hat{f}$. □

7. EXACT P-VALUES

In this section we define exact p-values. A more general notion of p-value will be given in the next section.

Definition 14. Given a probabilistic trial (Ω, \mathbf{P}) furnished with a test order \leq , the *exact p-value* of an outcome x is the probability $\mathbf{P}[\leq x]$.

Notice the dependence on the test order. Of course, the test order can be given by means a test pyramid or test statistic. Test statistics are especially popular in applications. In this connection, let us give an alternative definition of exact p-values.

Definition 15. Given a probabilistic trial (Ω, \mathbf{P}) furnished with a test statistic f , the *exact p-value* of an outcome x is the number $\hat{f}(x) = \mathbf{P}[f \leq f(x)]$.

Of course the two definitions are consistent.

Lemma 16.

- (1) *If a test statistic f induces a test order \leq then $\mathbf{P}[\leq x] = \mathbf{P}[f \leq f(x)]$.*
- (2) *If a test order \leq induces a test statistic f then $\mathbf{P}[f \leq f(x)] = \mathbf{P}[\leq x]$.*

Proof.

(1) Use Lemma 9, part 1.

(2) By Lemma 10, f induces \leq . Now use (1). □

Q: Let's go back to the two examples of §3, compute the exact p-values of the actual outcomes and compare them.

A: The examples do not have sufficient information. We need to furnish them with reasonable test tools. For the Coin Example, we can use the test statistic $M(x) = \text{Min}\{\text{Heads}(x), \text{Tails}(x)\}$ defined in the continuation of the Coin Example in §4. With respect to that test statistic, the exact p-value p_1 of the actual outcome is as follows.

$$p_1 = \mathbf{P}[f \leq 1] = \frac{2(1 + 42)}{2^{42}} < \frac{2 \times 2^6}{2^{42}} = 2^{-35}$$

Concerning the Lottery Example, one possible approach is given by the following graph G .

- The vertices of G are the 4,000,000 lottery participants plus John, the lottery organizer, whether he bought any lottery tickets or not.
- Two distinct vertices X and Y are connected by an edge in G if and only if at least one of the following conditions holds:
 - X is a parent of Y or the other way round.
 - X, Y are spouses.
 - X, Y are close friends.

Given the graph G , we have a natural test statistic $\delta(X)$ on the lottery participants, namely the length of the minimal chain of edges from John to X . If John bought at least one ticket then the minimal value of δ is 0; otherwise it is one. Donna, John's

wife, is at distance 1 from John, and she bought 4 tickets. We can estimate the exact p-value p_2 of the actual outcome from below.

$$p_2 = \mathbf{P}(\delta \leq 1) \geq \frac{4}{10,000,000} > \frac{2^2}{2^4 \times 2^{10} \times 2^{10}} = 2^{-22}.$$

It follows that

$$p_1 < 2^{-35} < 2^{-22} < p_2.$$

Q: This is interesting. The win of the lottery organizer's wife seemed more striking to me than getting 41 heads in 42 coin tosses. I wonder whether p_2 is small enough to impugn the hypothesis that the lottery was fair.

A: Well, it seems safe to assume that the lottery participants at distance ≤ 1 from John bought less than 10,000 tickets. Under this assumption,

$$p_2 < \frac{10,000}{10,000,000} = \frac{1}{1000} = 0.1\%$$

which is small enough for the lax as well as strict Cournotians.

8. P-FUNCTIONS

In this section ε ranges over the non-negative reals.

Lemma 17. *Let f be a canonic test statistic.*

- (1) $\mathbf{P}[f \leq \varepsilon] = \varepsilon$ if $\varepsilon \in \text{Range}(f)$.
- (2) $\mathbf{P}[f \leq \varepsilon] = \varepsilon$ if $\varepsilon = \sup(S)$ for some $S \subseteq \text{Range}(f)$.
- (3) $\mathbf{P}[f \leq \varepsilon] \leq \varepsilon$ for every ε .

Proof. (1) If $\varepsilon = f(x)$ then

$$\begin{aligned} \mathbf{P}[f \leq \varepsilon] &= \mathbf{P}[f \leq f(x)] \\ &= f(x) && \text{as } f \text{ is canonic} \\ &= \varepsilon \end{aligned}$$

(2) If $\varepsilon = \sup(S)$ for some $S \subseteq \text{Range}(f)$ then there is a sequence $s_1 < s_2 < \dots$ of elements of S converging to ε . Then

$$\mathbf{P}[f \leq \varepsilon] = \mathbf{P}[f \leq s_1] + \mathbf{P}[s_1 < f \leq s_2] + \mathbf{P}[s_2 < f \leq s_3] + \dots$$

so that the probability

$$\mathbf{P}[f \leq s_n] = \mathbf{P}[f \leq s_1] + \mathbf{P}[s_1 < f \leq s_2] + \dots + \mathbf{P}[s_{n-1} < f \leq s_n]$$

converges to $\mathbf{P}[f \leq \varepsilon]$ as $n \rightarrow \infty$. Therefore

$$\begin{aligned} \mathbf{P}[f \leq \varepsilon] &= \lim_{n \rightarrow \infty} \mathbf{P}[f \leq s_n] \\ &= \lim_{n \rightarrow \infty} s_n && \text{by (2)} \\ &= \varepsilon && \text{by the choice of } s_1 < s_2 < \dots \end{aligned}$$

(3) Let $\varepsilon_0 = \sup\{s \in \text{Range}(f) : s \leq \varepsilon\}$. Then $[f \leq \varepsilon] = [f \leq \varepsilon_0]$ and therefore $\mathbf{P}[f \leq \varepsilon] = \mathbf{P}[f \leq \varepsilon_0] = \varepsilon_0 \leq \varepsilon$. \square

Definition 18.

- A *p*-function is a test statistic f such that $\mathbf{P}[f \leq \varepsilon] \leq \varepsilon$ for every ε .
- A *p*-function f is *exact* if $\mathbf{P}[f \leq \varepsilon] = \varepsilon$ for every $\varepsilon \in \text{Range}(f)$; otherwise f is *conservative*.
- Values of a *p*-function are *p-values*.

If f is a *p*-function then cf is a *p*-function for every $c \geq 1$. Indeed

$$\mathbf{P}[cf \leq \varepsilon] = \mathbf{P}[f \leq \varepsilon/c] \leq \varepsilon/c \leq \varepsilon.$$

If $c < 1$ then cf may not be a *p*-function. In particular if $\mathbf{P}[f \leq \varepsilon] = \varepsilon$ for at least one ε then cf is not a *p*-function because, for that ε , we have

$$\mathbf{P}[cf \leq c\varepsilon] = \mathbf{P}[f \leq \varepsilon] = \varepsilon > c\varepsilon.$$

Theorem 19. *For any test statistic f , the following two claims are equivalent:*

- (1) f is *canonic*.
- (2) f is an *exact p*-function.

Proof. The implication (1) \rightarrow (2) is proven in Lemma 17. To prove the implication (2) \rightarrow (1), assume that f is an exact *p*-function, x is an arbitrary outcome and $\varepsilon = f(x)$. We have

$$\begin{aligned} \mathbf{P}[f \leq f(x)] &= \mathbf{P}[f \leq \varepsilon] \\ &= \varepsilon && \text{as } f \text{ is exact} \\ &= f(x). && \square \end{aligned}$$

Corollary 20.

- *Exact p-values are values of a exact p-function.*
- *Every test order induces a unique exact p-function.*
- *Every test statistic f is equivalent to a unique exact p-function.*

9. DATA SNOOPING

Q: What about conservative p-functions? Do you really have any use for them?

A: Yes. A good example is the so-called Bonferroni correction. There is a good explanation in [1].

Q: But what is the idea?

A: Consider a version of our coin example where, for simplicity, you toss a coin only 15 times. Again the null hypothesis is that all outcomes are equally probable, and again we will use the test statistic $M(x) = \text{Min}\{\text{Heads}(x), \text{Tails}(x)\}$. Suppose that you really want to impugn the null hypothesis. So you repeat the trial over and over until you encounter an outcome where M is 1. At this point, you report that you performed a trial and got an outcome with p-value

$$\mathbf{P}[M \leq 1] = \frac{2(1 + 15)}{2^{15}} = \frac{2^5}{2^{15}} = 2^{-10} = \frac{1}{1024}.$$

Q: But this is cheating.

A: Exactly. Yet, such data snooping occurs in science though rarely in such a blatant form; a scientist may genuinely forget about some steps in the preliminary analysis of data. In this connection, you may enjoy a relevant issue [6] of the XKCD webcomic.

Q: I remember hearing about publication bias. What is it?

A: It is a rather insidious kind of data snooping. Imagine that a large number of groups of scientists are testing an important null hypothesis. Eventually some group gets a significant p-value and submits their findings for publication while the other groups don't publish anything on the issue. Some influential statisticians argue that most published research findings are wrong and that publication bias may be the main reason for that [5].

Q: You cannot avoid multiple testing of a null hypothesis. What do you do?

A: Consider a case where the same probabilistic trial is performed n times, using test statistics f_1, \dots, f_n and getting p-values p_1, \dots, p_n . Of course one would love to report the p-value $\min\{p_1, \dots, p_n\}$ but it needs to be adjusted for the multiple testing of the null hypothesis. The Bonferroni correction is one way to adjust that p-value, by multiplying it by n . The p-value $n \cdot \min\{p_1, \dots, p_n\}$ is valid though usually conservative. More exactly the test statistic $f = n \min(f_1, \dots, f_n)$ is a p-function, typically conservative. Indeed,

$$\mathbf{P}[f \leq \varepsilon] \leq \mathbf{P}\left(\cup_{i=1}^n [f_i \leq \varepsilon/n]\right) \leq \sum_{i=1}^n \mathbf{P}[f_i \leq \varepsilon/n] \leq \sum_{i=1}^n \varepsilon/n = \varepsilon.$$

Acknowledgments. We thank Andreas Blass for useful comments on a draft of this paper.

REFERENCES

- [1] Hervé Abdi, “Bonferroni and Šidák corrections for multiple comparisons,” <http://www.utdallas.edu/~herve/Abdi-Bonferroni2007-pretty.pdf>, accessed January 8, 2016. Appeared in N. J. Salkind (ed.), *Encyclopedia of Measurement and Statistics*, Sage Publications, 2007.
- [2] D. R. Cox and D. V. Hinkley, “Theoretical Statistics,” Chapman & Hall, 1974.
- [3] Yuri Gurevich and Grant O. Passmore, “Impugning Randomness, Convincingly,” *Bulletin of EATCS* 104, June 2011; also — with an added Prologue — in *Studia Logica* 82 (2012), 1–31.
- [4] Yuri Gurevich and Vladimir Vovk, “Fundamentals of p-values,” unpublished manuscript.
- [5] John P. A. Ioannidis, “Why Most Published Research Findings Are False”, *PLoS Medicine*, Volume 2, Issue 8, e124, August 2005.
- [6] XKCD, <http://xkcd.com/882/>, accessed on Jan. 8, 2016.