



On admissibility in post-hoc hypothesis testing

Ben Chugg  ^{a,*}, Tyron Lardy ^b, Aaditya Ramdas  ^a, Peter Grünwald  ^{b,c}

^a Carnegie Mellon University, USA

^b University of Leiden, the Netherlands

^c Centrum Wiskunde & Informatica, the Netherlands



ARTICLE INFO

Keywords:

E-values
Hypothesis testing
P-hacking
Significance levels

ABSTRACT

The validity of classical hypothesis testing requires the significance level α be fixed before any statistical analysis takes place. This is a stringent requirement. For instance, it prohibits updating α during (or after) an experiment due to changing concern about the cost of false positives, or to reflect unexpectedly strong evidence against the null. Perhaps most disturbingly, witnessing a p-value $p \ll \alpha$ vs $p = \alpha - \epsilon$ for tiny $\epsilon > 0$ has no (statistical) relevance for any downstream decision-making. Following recent work of Grünwald [1], we develop a theory of *post-hoc* hypothesis testing, enabling α to be chosen after seeing and analyzing the data. To study “good” post-hoc tests we introduce Γ -admissibility, where Γ is a set of adversaries which map the data to a significance level. We classify the set of Γ -admissible rules for various sets Γ , showing they must be based on e-values, and recover the Neyman-Pearson lemma when Γ is the constant map.

1. Introduction

An epidemiologist runs a clinical trial to test the efficacy of a new drug. She does not choose a significance level α beforehand, but waits to see the results. If the p-value is $p = .01$ she rejects at level $\alpha = 0.01$ and if $p = .05$ she rejects at $\alpha = 0.05$. Is this valid statistical practice? In the standard theory of hypothesis testing—dominant since the work of Neyman, Pearson, and Fisher in the 1920s and 30s—it is not. Indeed, if the epidemiologist is prepared to reject at level p for all $p \leq 0.10$ (say), then the true type-I error is 0.10, not p .

In other words, the current paradigm of hypothesis testing requires that the significance level α be chosen independently of (thus without looking at) the data. This is well-known to statisticians, but continues to be a source of frustration and confusion for practitioners, and the epidemiologist’s mistake above is unfortunately all too common [2–4]. This fact is limiting: once the data has been gathered, if the selected value of α gives a vacuous or uninteresting result, nothing more can be done. The same data should not be used in any further analysis of the same question (or, technically, even a separate question motivated by the first¹).

The difficulties do not end there. Requiring that α be chosen independently of the data introduces a fundamental tension between evidence and decision-making. Namely, upon observing a p-value $p \leq \alpha$, it is irrelevant whether p is roughly the same as α or significantly smaller. For example, if $\alpha = 0.05$, then observing $p = .04$ and $p = 10^{-6}$ have the same decision-theoretic consequences. As pointed out recently by [1], this is a frustrating fact: the p-value is ostensibly a measure of evidence against the null hypothesis, yet any p-value below α has no formal relevance for any downstream decision task.

* Corresponding author.

E-mail address: benchugg@cmu.edu (B. Chugg).

¹ This—admittedly philosophically thorny and somewhat controversial—reliance of p-values on counterfactuals is one of their many restrictive features [5,6].

Example 1 (Investment). A funder is interested in investing in a company producing a new drug that looks promising based on an early clinical trial. The funder wants to buy-in at an amount inversely proportional to the risk that the drug is ineffective. The original trial was run at $\alpha = 0.05$ and produced a p-value of less than 0.001. While both the company and the investor would like to use this information, the company can only give guarantees about the false positive rate at the 5% level.

This limitation serves as an uncomfortable reminder that current statistical practice consists of an odd combination of both Ronald Fisher's perspective and that of Jerzy Neyman and Egon Pearson. Whereas Fisher proposed the p-value as measure of evidence [7,8], Neyman and Pearson were focused on decision making [9]. The fact that we both report the p-value *and* make a decision at a prespecified significance level is an awkward methodological hybrid that none of these pioneers promoted.

This hybrid approach reflects (or perhaps introduces) confusion over the proper interpretation of p-values, significance levels, and hypothesis testing in general. And such confusion has substantial practical consequences. “The problem of roving alphas” as Goodman [10] puts it, i.e., choosing α as a function of the data, is well-documented in business and management science [11,12], medicine [13], and psychology [14]. What are statisticians to do in the face of this mistake?

One response to the problem of roving alphas is to continue to educate practitioners on the proper use of statistical tools. This, however, doesn't solve the fundamental tension inherent in the standard theory of hypothesis testing. The epidemiologist's mistake is understandable, after all—there are many situations where it is natural to want to use the p-value to say more than is mathematically warranted. Consider the following example where an analyst might be tempted to incorrectly use p-values in a downstream decision-making task.

Faced with such examples, in this work we consider an alternative response to the problem of roving alphas. Instead of pointing out flaws in the epidemiologist's approach, we focus on building new statistical tools which function both as measures of evidence *and* allow for mathematically sound downstream decision-making. In particular, we develop the theory of *post-hoc hypothesis testing* [1], which allows for data-dependent significance levels. The catch is that the guarantees that an analyst can give on their procedure changes under this theory, moving from error probabilities to expected losses. If α is fixed beforehand, however, then the theory of post-hoc testing recovers the standard theory of hypothesis testing. Thus, in that sense, post-hoc hypothesis testing is a strict generalization of the classical framework. Our focus in this paper is to define and classify “admissible test families”, which are the analogue of uniformly most powerful tests in the post-hoc setting.

One can situate our work as belonging to a recent push to resolve several issues with traditional statistical tools. The inability to handle post-hoc significance levels is just one problem with much of modern statistics—others include the inability to handle optional stopping and optional continuation. The burgeoning area of sequential, anytime-valid inference (SAVI) [15] is focused on such issues. A fundamental tool in this line of work is the *e-value* [16,17], an object which will also play a crucial role in this paper. Overall, we view our work here as an extension of SAVI: broadly, as an attempt to refine existing statistical technology to make life easier for practitioners and to mitigate the risk of statistical malpractice.

1.1. Contributions

Following Grünwald et al. [16], in order to formulate a theory of post-hoc hypothesis testing we begin by following Wald [18] and recast the usual notions of type-I and type-II error probabilities in terms of expectations over data-dependent type-I and type-II loss functions. We require that a test has bounded type-I loss in expectation, where the type-I loss may be chosen adversarially. Such a test is called *type-I risk safe*, a notion which generalizes the property of a test having a bounded type-I error probability.

At this point we depart from the setting studied by Grünwald [1] and introduce a different notion of admissibility. We summarize our contributions below. While the general setting and several of the results are presented for composite hypotheses, we are mainly focused on testing point nulls and alternatives.

Throughout the paper it will often be useful to remark on the differences or similarities between our post-hoc setting and the traditional, non-post-hoc setting. Henceforth, we will often refer to the latter as the “traditional setting” or sometimes as the “Neyman-Pearson paradigm.”

1.1.0.1. Definition of Γ -admissibility. [Definition 2.4](#) introduces the notion of Γ -admissibility, where Γ is a set of functions which map the data to a type-II loss function. This generalizes the definition of a uniformly most powerful test to the post-hoc setting ([Observation 2.1](#) makes this precise). Informally, a test is Γ -admissible if there is no test that has smaller expected type-II loss under each distribution in the alternative and across all mappings in Γ . Our definition of Γ -admissibility departs from the notion of admissibility studied by Grünwald [1], which requires that a test be dominated with probability 1 under each distribution in the alternative to be inadmissible. We require that a test be dominated only in expectation, making fewer tests admissible.

1.1.0.2. Extension of the likelihood ratio test to the post-hoc setting. We give an extension of the likelihood ratio test—shown to be uniformly most powerful in the traditional paradigm [19]—to a post-hoc setting and prove that the extension is Γ -admissible for any Γ . This is in [Section 3](#). This shows that our definition of admissibility is not too strict; the set of Γ -admissible tests is non-empty. This investigation also sheds light on an interesting phenomenon in the post-hoc setting: Even for continuous distributions, if randomization is allowed then admissible tests will always be randomized. More specifically, there will be losses on which the

test is randomized. Since randomized tests may not be kosher for all applications, this motivates explicitly defining and studying the admissibility of *binary* post-hoc tests which reject or accept with probability one.

1.1.0.3. General properties of admissible tests. Moving beyond the likelihood ratio test, we prove several general properties for all Γ -admissible tests for any Γ . We study both randomized and binary tests. Of particular note, we associate with each test an e-variable and show how any admissible test can be written as a function of the associated e-variable. Thus, the e-variable directly determines the behavior of the test. In the opposite direction, given an e-variable, we can associate with it a test obtained by thresholding. See [Lemmas 4.4 and 4.6](#) and [Eqs. \(14\) and \(17\)](#) for the details of these relationships. A second result worth mentioning is a Rao-Blackwell type-argument [\[20,21\]](#), whereby we demonstrate how to (weakly) improve any given test by conditioning on a sufficient statistic. This can be found in [Section 5.2](#).

1.1.0.4. U- and C-admissibility. Finally, we study two specific families of maps Γ and give necessary and sufficient conditions on tests (and e-variables) to be Γ -admissible. In particular, we study both $\Gamma = U$ and $\Gamma = C$ where U is the set of all mappings from the data to the losses, and C is the set of constant mappings. For a point null and alternative, we give a full characterization of all U -admissible tests in terms of e-variables for both randomized and binary tests—see [Theorems 5.1 and 5.2](#). For C , we give several necessary conditions for a randomized test to be admissible, and provide a complete characterization in the binary case—see [Theorem 5.3](#). This final result recovers the Neyman-Pearson lemma [\[9\]](#) when instantiated with a single loss function. The curious reader can jump ahead to [Table 2](#) for a diagrammatic overview of our admissibility results.

To summarize, our main contributions are as follows:

1. We give a definition of admissibility in the post-hoc setting— Γ -admissibility, see [Definition 2.4](#)—which generalizes the notion of admissibility in the traditional Neyman-Pearson setting;
2. We give a Γ -admissible generalization of the likelihood ratio test to the post-hoc setting ([Section 3](#)) for all Γ ;
3. For both randomized and binary tests, we classify the set of U -admissible tests ([Theorems 5.1 and 5.2](#)). For binary tests, we classify the set of C -admissible tests ([Theorem 5.3](#)).

1.2. Related work

The post-hoc framework studied here is not the only proposal for how to resolve the issues with traditional hypothesis testing. Bayesian decision theory [\[22,23\]](#) has a devoted following, where one replaces p-values and significance levels with posterior probabilities and Bayes factors, and makes decisions by minimizing posterior loss. A great benefit of the Bayesian paradigm is that it handles data-dependent loss functions for free, which we show formally in [Appendix C](#) where we explore the connection between our setting and the Bayesian one. The hefty price to pay is that one cannot get prior-free Type-I error probability guarantees.

Another avenue of related work is the subfield of imprecise probability. Here one tries to avoid making claims with false precision by demanding that decisions be defensible over a set of models and/or priors (as done recently in, e.g., Chau et al. [\[24\]](#) for two-sample testing). E-values may already be considered as a tool of imprecise probability, as they are closed under convex combinations of the null (sometimes called “credal sets”). Moreover, the e-posterior [\[25\]](#) was recently connected to credal sets [\[26, Proposition 10\]](#). While we do not pursue further connections here, we believe this to be an interesting direction.

Besides imprecise probability and Bayesian methods, other approaches to navigate around the drawbacks of traditional hypothesis testing involve changing what researchers report, e.g., switching p-values with confidence intervals [\[27\]](#) or likelihood ratios [\[28\]](#). Post-hoc hypothesis testing has more in common with some of these than with others (it plays nicely with the law of likelihood for instance, since the likelihood ratio is an e-value), but importantly it is the only approach to seek a strict generalization of the traditional Neyman-Pearson paradigm. Given that most hypothesis testing in practice still relies on this paradigm, we believe it worthwhile to study how to expand the scope of its tools instead of casting them aside.

Our work can be viewed as part of a more general focus on e-variables/e-values, which are simple yet surprisingly fruitful mathematical objects that have proven useful across a wide variety of statistical problems. While some authors had previously used them implicitly, it was not until recently (roughly 2020) that excitement grew concerning their applications and they were given a name. Since then, there has been growing interest in both developing the underlying theory of e-variables and deploying them in various problems. We refer to Grünwald et al. [\[16\]](#), Ramdas and Wang [\[17\]](#), and Vovk and Wang [\[29\]](#) for introductions to e-values and their use in hypothesis testing.

That e-values can play nicely with data-dependent parameters, thus overcoming some of the drawbacks of p-values, has been noted previously by several authors. The first work in this vein came from the multiple testing literature and worked with e-values only implicitly [\[30\]](#). As far as we are aware, the first result to explicitly link e-values and data-dependent thresholds is [Lemma 4.1](#) of Wang and Ramdas [\[31\]](#) in the context of false discovery rate control. Then Grünwald [\[1\]](#) put these connection on firm decision-theoretic ground and launched the original investigation into post-hoc hypothesis testing. In contradistinction to this paper, Grünwald considers multiple actions, whereas we consider only accept/reject decisions. Working under a different notion of admissibility, he gives a partial classification of all admissible tests in terms of e-variables. We further discuss his results and his notion of admissibility in [Appendix B](#).

As part of this work, Grünwald [\[1\]](#) notes that e-value based confidence intervals and confidence distributions (using the so-called e-posterior [\[25\]](#)) are safe under data-dependent selections of α . Since then, a number of authors have noted the promise of e-values in enabling testing and estimating when significance levels are data-dependent [\[32–36\]](#). In particular, Koning [\[33\]](#) also studies post-hoc

hypothesis testing and defines post-hoc p-values. He further relates post-hoc p-values to e-values and shows that any post-hoc test defines an e-value. An overview of post-hoc testing and decision-making with e-values was recently given in Ramdas and Wang [17, Chapter 4].

Besides Grünwald [1], several other authors have studied questions concerning the admissibility of e-variables. These include Wang [37], Vovk and Wang [38], and Ramdas et al. [39]. However, their notions of admissibility either differ from ours or the settings are incomparable. Wang [37] and Vovk and Wang [38] are concerned with designing e-values that are as large as possible under the alternative, and Ramdas et al. [39] are concerned with admissible anytime-valid inference.

1.3. Outline

[Section 2](#) introduces the general framework of post-hoc hypothesis testing, recasting the usual notions of type-I and type-II error probabilities as type-I and type-II risk. Here we formalize a post-hoc hypothesis test as a “test family,” introducing both randomized and binary test families ([Definitions 2.1](#) and [2.2](#)). [Section 2.1](#) introduces Γ -admissibility and discusses its relationship with traditional hypothesis testing. [Section 2.2](#) discusses various assumption we make in the remainder of the paper and gives relevant background on e-values. With this machinery in hand, [Section 3](#) generalizes the likelihood ratio test to a post-hoc setting. [Section 4](#) then gives general properties of Γ -admissible tests for any Γ . [Section 4.1](#) studies tests which allow randomization and [Section 4.2](#) studies tests that are restricted to be binary. [Sections 5.1](#) and [5.3](#) study U-admissibility and C-admissibility respectively. [Section 6](#) concludes. All longer proofs are delegated to [Appendix A](#).

2. Post-hoc hypothesis testing

We begin with a standard hypothesis testing setup. Let \mathcal{P} and \mathcal{Q} be two disjoint families of probability measures on some measurable space (Ω, \mathcal{F}) . We observe data encoded by a random variable $X : \Omega \rightarrow \mathcal{X}$ taking values in some space \mathcal{X} whose behavior is governed by some unknown distribution $R \in \mathcal{P} \cup \mathcal{Q}$. We are testing \mathcal{P} against \mathcal{Q} , meaning that we are tasked with deciding between the null hypothesis $H_0 : R \in \mathcal{P}$ and the alternative $H_1 : R \in \mathcal{Q}$. We either *reject* or *sustain* the null hypothesis H_0 . If the former, we are implicitly adopting the alternative.

We will introduce the framework of post-hoc hypothesis testing in terms of possibly composite hypotheses \mathcal{P} and \mathcal{Q} , but thereafter switch to considering only point hypothesis $\mathcal{P} = \{P\}$ and $\mathcal{Q} = \{Q\}$. While some of our results extend to the composite case, not all do, and considering point hypotheses keeps the presentation cleaner.

In order to formalize the notion of post-hoc significance levels, we follow Wald [18] and Grünwald [1] and frame hypothesis testing in terms of expected losses instead of error probabilities. For an index set $\mathcal{B} \subset \mathfrak{N}_{\geq 0}$ which we call a set of *scenarios*, let $\{L_b\}_{b \in \mathcal{B}}$ be a set of *loss functions*, where $L_b : \{0, 1\} \times \{0, 1\} \rightarrow \mathfrak{N}_{>0}$. The first argument denotes whether we are considering the loss under the null or alternative: $L_b(0, \cdot)$ is loss under the former and $L_b(1, \cdot)$ that under the latter. The second argument denotes whether we sustain (0) or reject (1) the null. We will assume that $L_b(0, 0) = L_b(1, 1) = 0$ for each $b \in \mathcal{B}$. In other words, making the correct decision (sustaining when $R \in \mathcal{P}$ and rejecting when $R \in \mathcal{Q}$) incurs zero loss.

We choose the letters \mathcal{B} and b to accord with [Example 1](#). For instance, we might consider $L_b(0, 1)$ to be the buy in of the funding agency—that is, the amount they stand to lose if the drug turns out to be ineffective. We do not assume that \mathcal{B} is finite (or even countable). For instance, we might have $L_b(0, 1) = b$ for all $b \geq 1$. (Indeed, this is arguably the most natural choice; see the discussion at the end of this section). Finally, we assume without loss of generality that the losses $L_b(0, 1)$ are increasing in b (if not, we may simply relabel accordingly). Note this does not imply any ordering of the losses under the alternative.

Next we define the analogue of a hypothesis test in this post-hoc setting. Our decision to sustain or reject the null takes the form of a test family, which sees both the data and the loss, and then makes a (possibly randomized) decision. We emphasize that while we will allow the loss to be data-dependent, the test takes this loss as an argument. If it were otherwise, a test would have to be needlessly conservative. This is also the case in the traditional setting: the test knows the parameter α .

In what follows, let $\text{Ber}(p)$ denote a Bernoulli distribution with parameter $p \in [0, 1]$.

Definition 2.1 (Test family). *A test family δ is a map $\delta : \mathcal{X} \times \mathcal{B} \rightarrow [0, 1]$. We call δ a family because it represents a hypothesis test for each $b \in \mathcal{B}$. Given $X \in \mathcal{X}$ and $b \in \mathcal{B}$, the decision to reject is drawn as $D \sim \text{Ber}(\delta(X, b))$ where $D = 1$ means reject and $D = 0$ means sustain.*

We also define *binary* test families, which do not allow randomization.

Definition 2.2 (Binary test family). *A binary test family φ is a map $\varphi : \mathcal{X} \times \mathcal{B} \rightarrow \{0, 1\}$. Given $X \in \mathcal{X}$ and $b \in \mathcal{B}$, we reject if $\varphi(X, b) = 1$ and sustain otherwise.*

When studying binary test families, we will assume that \mathcal{P} and \mathcal{Q} consist of continuous distributions. Even in the traditional setting, we know that powerful hypothesis tests must generally be randomized for discrete distributions. Further, in the case of discrete distributions and binary tests, there is no neat classification of UMP tests as far as we are aware. Thus, studying admissibility in this case becomes quite challenging.

We will typically denote binary test families with the Greek letters φ or ψ —see [Table 1](#). Unless binary test families are mentioned explicitly, it should be assumed in the sequel that we are discussing test families that could be either randomized or binary. We will denote such families with letters δ and ϕ . We may sometimes refer to test families that allow randomization ([Definition 2.1](#)) as *randomized* test families to emphasize the distinction from binary test families.

For notational simplicity, for a test family δ and $a \in \{0, 1\}$ define

$$L_b(a, \delta(X, b)) \equiv \mathbb{E}_{D \sim \text{Ber}(\delta(X, b))} L_b(a, D), \quad (1)$$

which is the expected loss (type-I or type-II) when playing $\delta(X, b)$ on loss L_b . Note that we can write $L_b(0, \delta(X, b)) = L_b(0, 1)\delta(X, b)$ and $L_b(1, \delta(X, b)) = L_b(1, 0)(1 - \delta(X, b))$, since we are assuming that $L_b(0, 0) = L_b(1, 1) = 0$. These two identities will be used extensively. Next we define the risk of a test family.

Definition 2.3 (Type-I Risk). *For a test family δ , set*

$$\mathcal{R}_P(\delta) \equiv \sup_{P \in \mathcal{P}} \mathbb{E}_{X \sim P} \sup_{b \in \mathcal{B}} L_b(0, \delta(X, b)). \quad (2)$$

We call $\mathcal{R}_P(\delta)$ the type-I risk of δ . If $\mathcal{R}_P(\delta) \leq 1$, we say that δ is type-I risk safe.

Definition 2.3 implicitly assumes that $\sup_b L_b(0, \delta(X, b))$ is measurable. In this work we will make certain modest assumptions that ensure this is the case (see the discussion in [Section B.1](#)).

Risk is the equivalent of type-I error probability in the standard Neyman-Pearson paradigm. Indeed, consider the case of a single loss $\mathcal{B} = \{b_0\}$. Then $\mathcal{R}_P(\delta) = \sup_{P \in \mathcal{P}} \mathbb{E}_P L_{b_0}(0, \delta(X, b_0)) = L_{b_0}(0, 1) \sup_{P \in \mathcal{P}} \mathbb{E}_P[\delta(X, b_0)]$. A risk bounded by 1 is thus identical to a type-I error probability bounded by $\alpha > 0$ if we take $L_{b_0}(0, 1) = 1/\alpha$. The upper bound of 1 on the risk is arbitrary. It can be replaced by any positive constant and our results will scale accordingly. See for instance Ramdas and Wang [[17](#), Section 4] which replaces 1 with a constant $\beta > 0$.

While risk recovers type-I error probability when there is a single loss, risk is of course a different metric than an error probability in general. A framework built on replacing error probabilities with risk may therefore make some readers uncomfortable. We provide a longer defense of risk in [Appendix D](#), but let us make two points here. First, risk is actually the norm in decision theory. Second, type-I risk seems to us the most straightforward generalization of type-I error control.²

In the definition of type-I risk, the supremum over losses is inside the expectation. This allows the loss to be data-dependent, which can be seen explicitly by writing type-I risk in terms of the supremum over mappings from the data to losses. In particular, under reasonable measure-theoretic assumptions on the data and the losses, we have the equality:

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{X \sim P} \sup_{b \in \mathcal{B}} L_b(0, \delta(X, b)) = \sup_{P \in \mathcal{P}} \sup_{B: \mathcal{X} \rightarrow \mathcal{B}} \mathbb{E}_{X \sim P} L_{B(X)}(0, \delta(X, B(X))), \quad (3)$$

Each function $B: \mathcal{X} \rightarrow \mathcal{B}$ in (3) can be thought of as an adversary who chooses a loss based on the data. Thus, if a decision δ is type-I risk safe, it is *post-hoc* type-I risk safe in the sense that it guarantees a type-I loss of at most 1 on any data-dependent loss. While intuitively clear, making this fully mathematically precise requires some non-trivial work. The formal statement and its proof can be found in [Appendix B.1](#).

We define the following fundamental object:

$$E_\delta(X) \equiv \sup_{b \in \mathcal{B}} L_b(0, \delta(X, b)). \quad (4)$$

Using the letter “E” is no coincidence: If δ is type-I risk safe, then E_δ is an e-value (see [Section 2.2](#)). There is also a natural way to define test families from e-values which we will see in [Section 4](#). This creates a bidirectional mapping between e-values and tests, a relationship which will be convenient both for classifying the admissible test families (e.g., [Theorem 5.1](#)) and in proving properties of admissible tests.

2.1. Admissible test families

In the standard theory of hypothesis testing, two classes of tests are of interest: uniformly most powerful (UMP) tests and admissible tests. A level- α test (i.e., having type-I error bounded by α) is UMP if it has greater or equal power than every other level- α test for every $Q \in \mathcal{Q}$. If \mathcal{Q} is composite, then UMP tests only exist in special settings. These include practically relevant cases, however, such as the one-sided t -test and χ^2 -tests, and more generally when $\mathcal{P} \cup \mathcal{Q}$ admit a monotone likelihood ratio. An admissible test, meanwhile, cannot be uniformly improved across all $Q \in \mathcal{Q}$ (meaning that it cannot be strictly improved for some Q and not made worse on any other Q). In this work we generalize admissibility to a post-hoc setting.

Let $\mathbf{U} \equiv \mathbf{U}(\mathcal{X}, \mathcal{B})$ be the set of all measurable functions (“adversaries”) $B: \mathcal{X} \rightarrow \mathcal{B}$. Intuitively, we call a test family δ Γ -inadmissible for some $\Gamma \subset \mathbf{U}$ if there is another test which is type-I risk safe and never has higher, but sometimes lower, expected type-II loss under some adversary $B \in \Gamma$. More formally:

Definition 2.4 (Γ -admissibility). *We say the decision family ϕ is weakly preferable to δ relative to $\Gamma \subset \mathbf{U}$ if it is type-I risk safe and for all $B \in \Gamma$ and all $Q \in \mathcal{Q}$,*

$$\mathbb{E}_{X \sim Q}[L_{B(X)}(1, \phi(X, B(X)))] \leq \mathbb{E}_{X \sim Q}[L_{B(X)}(1, \delta(X, B(X)))] \quad (5)$$

We call ϕ strictly preferable to δ if the inequality is strict for at least one $B \in \Gamma$ and $Q \in \mathcal{Q}$. If there exists some type-I risk safe ϕ that is strictly preferable to δ relatively to Γ we say that δ is Γ -inadmissible. If no decision family ϕ is strictly preferable to δ relative to Γ then we say that δ is Γ -admissible. (Here is an alternative definition: δ is Γ -admissible if for every ϕ that is weakly preferable to δ , (5) actually holds with equality for all $B \in \Gamma$ and $Q \in \mathcal{Q}$.)

² For the interested reader, Koning [[33](#)] explores two other options, so-called “geometric” and “arithmetic” post-hoc validity.

Remark 1. Unless we consider $\Gamma = \mathbf{U}$ there is an asymmetry between the adversaries considered in the definition of type-I risk and the definition of Γ -admissibility. [Lemma B.2](#) shows that a type-I risk safe rule has expected type-I loss bounded by 1 across any mapping $B : \mathcal{X} \rightarrow \mathcal{B}$, while in Γ -admissibility we are deliberately restricting the mappings under consideration, but only for the alternative (i.e. for type II risk). This is by design. In the worst case, the loss may be data-dependent—we want to ensure we remain in type-I risk safe in this case. However, we might also be cautiously optimistic that it won't be data-dependent, and allow ourselves to choose a test that is \mathbf{C} -admissible based on this hope, all the while knowing that if we are wrong then at least we are protected under the null. Therefore, our first goal is to find tests which have bounded type-I risk under even the worst case adversary. Only then do we consider type-II risk. This is analogous to how Neyman and Pearson prioritize type-I error bounded by α and only then look for tests with higher power.

We leave implicit in [Definition 2.4](#) whether we are considering binary or randomized tests; the definition holds for both. We will always make clear in the text which kinds of tests are under consideration.

In what follows we will always assume that any collection Γ of adversaries includes the constant mappings. This will simplify many of the proofs. It is also intuitive: it would be odd if an adversary could play $L_b(1, 0)$ for only a subset of the observations. To formalize this assumption, we will consider only those $\Gamma \supset \mathbf{C}$ where

$$\mathbf{C} \equiv \{B \in \mathbf{U} : B(X) = b^* \text{ for all } X \in \mathcal{X} \text{ and some } b^*\}. \quad (6)$$

In the sequel, we will often shorthand $B(X)$ to B in equations to save space. Recall that for any $b \in \mathcal{B}$, $L_b(1, \delta(X, b)) = L_b(1, 0)(1 - \delta(X, b))$ (we are assuming that $L_b(1, 1) = L_b(0, 0) = 0$), so (5) can be equivalently written as $\mathbb{E}_{X \sim Q} L_B(1, 0)\delta(X, B) \leq \mathbb{E}_{X \sim Q} L_B(1, 0)\phi(X, B)$. In particular, this implies that if $\phi(X, b) \geq \delta(X, b)$ for some $b \in \mathcal{B}$ and all $X \in \mathcal{F}$ for some $\mathcal{F} \subset \mathcal{X}$, then $\mathbb{E}_Q [L_b(1, \phi(X, b))\mathbf{1}\{X \in \mathcal{F}\}] \leq \mathbb{E}_Q [L_b(1, \delta(X, b))\mathbf{1}\{X \in \mathcal{F}\}]$, a fact we will use often.

Despite the differences between being UMP and admissible, it's worth investigating when the two notions coincide. Suppose that there is only one loss function L_{b_0} . If ϕ is UMP at level $L_{b_0}^{-1}(0, 1)$ then it satisfies $\mathbb{E}_Q [\delta(X, b_0)] \leq \mathbb{E}_Q [\phi(X, b_0)]$ for all $Q \in \mathcal{Q}$ and all other tests δ . Thus, in particular, these inequalities are satisfied when δ is admissible (which notion of admissibility does not matter, since there is only one loss), thus implying that the inequalities must be equalities. Hence δ has the same power as ϕ on each Q , and is also UMP. We may therefore conclude:

Observation 2.1. For a single loss L (the traditional Neyman-Pearson setting) with $\alpha = L^{-1}(0, 1)$, if a uniformly most powerful level- α test exists then any admissible test family is such a test.

One could consider other notions of admissibility. Grünwald [1] considers a family ϕ to be preferable to δ if, for any adversary, ϕ rejects whenever δ rejects with probability 1. We compare our notion of admissibility with that of Grünwald [1] further in [Appendix B](#), but suffice it to say here that requiring that ϕ be better than δ with probability 1 is a strong requirement, and does not recover the notion of a uniformly most powerful test.

A second natural idea is to move the supremum over $B \in \Gamma$ inside the expectation and replace (5) with $\mathbb{E}_Q [\sup_{B \in \Gamma} L_B(1, \phi(X, B))] \leq \mathbb{E}_{X \sim Q} [\sup_{B \in \Gamma} L_B(1, \delta(X, B))]$. That is, we require that the worst case adversary for ϕ is no worse than the worst case of δ , in expectation. However, such a definition of admissibility is too strong: the set of admissible families is empty. This is also shown in [Appendix B](#).

One might expect that if δ is Γ_1 -admissible, then it is Γ_2 -admissible for any $\Gamma_1 \subset \Gamma_2$. But this does not necessarily hold. Consider $\Gamma_1 = \{B : B(x) = b_0\}$ for some fixed b_0 . The family $\delta(X, b) = \phi_{b_0}^{\text{NP}}(X)$ if $b = b_0$ and $\delta(x, b) = 0$ otherwise is Γ_1 -admissible, but is not \mathbf{U} -admissible by [Lemma 4.4](#) if \mathcal{B} contains more than b_0 . To put it another way, by introducing more adversaries we open ourselves to the possibility of a test family having strictly better type-II risk than δ on one of those adversaries.

Remark 2. Related to the previous remark, one could also consider different sets of adversaries in the definition of type-I risk safety. In particular, we could consider the notion of (Γ_1, Γ_2) -admissibility, where Γ_1 is the set of mappings considered in type-I risk (i.e., the supremum on the right hand side of (3)), and Γ_2 is the set of adversaries considered in [Definition 2.4](#). This paper would then be concerned with (\mathbf{U}, Γ) -admissibility while standard Neyman-Pearson theory focuses on (\mathbf{C}, \mathbf{C}) -admissibility. While this nicely emphasizes the relationship between our notion of type-I risk safety and that of traditional hypothesis testing, it is burdensome to continuously write (\mathbf{U}, Γ) -admissibility. Therefore, to spare ourselves notational pain, in this paper we shorten (\mathbf{U}, Γ) -admissibility to Γ -admissibility.

Throughout this work, we will be primarily concerned with testing a simple null P against simple alternative Q . [Section 4](#) will give general properties of admissible test families for all Γ , and then [Sections 5.1](#) and [5.3](#) will provide full or partial classifications of Γ -admissible test families for two particular classes Γ :

1. $\Gamma = \mathbf{U}$. The first natural set of adversaries to consider is all mappings $B : \mathcal{X} \rightarrow \mathcal{B}$. Under such a liberal notion of admissibility, a great many test families are admissible. Indeed, under mild assumptions on the losses, every test family δ which satisfies $\mathcal{R}_P(\delta) = 1$ is admissible. In other words, every test which is defined by a sharp e-value is admissible.
2. $\Gamma = \mathbf{C}$ for \mathbf{C} as in (6). We discuss the interpretation of \mathbf{C} further in [Section 5.3](#), and provide a concrete motivation in [Example 2](#), but let us mention here that it may be helpful to interpret \mathbf{C} -admissibility in terms of counterfactuals. Namely, if δ is \mathbf{C} -admissible then for every type-I risk safe ϕ , there exists some loss b such that if the adversary had played L_b , we would not regret playing δ instead of ϕ . We will give several necessary conditions for a randomized test family to be \mathbf{C} -admissible and provide a full characterization of \mathbf{C} -admissible binary test families.

Note that \mathbf{U} and \mathbf{C} are, respectively, the largest and smallest sets of adversaries we can consider under the assumption that $\Gamma \supset \mathbf{C}$: we are studying the two extrema of [Definition 2.4](#). While we do not explicitly study other sets of adversaries here (apart from when

Example 2. (Example 1, Continued) Consider the following idealized instantiation(s) of Example 1. Suppose b represents the cost of investment expressed in some monetary unit; for concreteness say the cost is $\$b \cdot 10^5$. A funder is given the opportunity to invest b in the drug's production.

If the drug turns out to be ineffective yet an investment is made then the investor loses their capital. We formalize this as $L_b(0, 0) = 0$, $L_b(0, 1) = b$. When the drug is effective, we consider two pay-off variations:

1. The pay-off is a fixed constant C^* . This reflects, for example, a situation in which the funder is the sole investor, and the investment is used to build a factory that can produce a fixed quantity of the drug. We view the loss as the regret incurred if the funder does not invest in the successful drug, so $L_b(1, 1) = 0$ and $L_b(1, 0) = C^*$ for all b .
2. The pay-off is a strictly increasing function $f(b)$ of b . This could for example be the case if larger b 's represent a larger factory (more productive capacity) or if there are many funders who each get a share of profits proportional to their investment (in that case f would be linear). In this case, again viewing $L_b(1, \cdot)$ as the regret of not investing in a successful drug, we set $L_b(1, 0) = f(b)$ and $L_b(1, 1) = 0$.

In the first variation, $L_b(1, 0) = C^*$ is independent of b , showing that $\Gamma = \mathbf{C}$ is an important special case to consider. In the second variation, $L_b(1, 0)$ is increasing in b . It then seems reasonable to assume that the investment b is increasing in the amount of evidence for effectiveness—in practice the ‘adversaries’ might be the data-analysts in a company, who suggest to management to make a larger investment the more evidence they see in the data. This reflects how p-hacking typically occurs in practice: an analyst is tempted to lower α (increase $L_b(0, 1)$), which, in the second case also increases $L_b(1, 0)$) after seeing promising results. This suggests an analysis in which Γ is restricted to those $B(X)$ which are (monotonically) increasing in a sufficient statistic of the data (e.g., the likelihood ratio); thus $\mathbf{C} \subsetneq \Gamma \subsetneq \mathbf{U}$. On the other hand, if we allowed $f(b)$ to be all functions of b , not just increasing functions, then $\Gamma = \mathbf{U}$.

Table 1

Some common notation used throughout the paper.

δ, ϕ	Test families, either binary or randomized	Definition 2.1
φ, ψ	Binary test families	Definition 2.2
t_φ	Decision curve of (admissible) binary test φ	Corollary 5.3
E_δ	e-variable associated with δ	Eq. (4)
δ_E, ϕ_E	Randomized test families defined by E	Eq. (14)
φ_E, ψ_E	Binary test families defined by E	Eq. (17)
\mathcal{L}	Set of all type-I losses, $\{L_b(0, 1) : b \in \mathcal{B}\}$	Eq. (12)
\mathbf{U}	All measurable maps from \mathcal{X} to \mathcal{B}	—
\mathbf{C}	All constant maps from \mathcal{X} to \mathcal{B}	Eq. (6)

our results easily generalize to handle all families Γ ; cf. Section 4), we believe that other sets could be of interest, as illustrated by Example 2 below, which also provides further motivation for studying \mathbf{C} -admissibility.

Table 2 gives an overview of our admissibility results.

Finally, it is very natural to take $L_b(0, 1) = b$ for $b \geq 1$. In this case,

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\sup_{b \in \mathcal{B}} L_b(0, \delta(X, b)) \right] = \sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\sup_{\alpha \in (0, 1]} \frac{\delta(X, 1/\alpha)}{\alpha} \right]. \quad (7)$$

That is, letting the losses take on any value in $[1, \infty)$ lets us consider any post-hoc significance level $\alpha \in (0, 1)$. Arguably, the right hand side of (7) is what one would expect of a definition of post-hoc hypothesis testing (see, e.g., Wang and Ramdas [31, Section 6.1], Koning [33] and Gauthier et al. [36]).

2.2. Notation, assumptions, and further background

As discussed previously, our results rely heavily on e-variables/e-values. An e-variable for a set of distributions \mathcal{P} is a nonnegative random variable E with expected value at most 1 under the null: $\sup_{P \in \mathcal{P}} \mathbb{E}_P[E] \leq 1$. The realized value of an e-variable is called an e-value. E-variables are typically functions of the underlying data X and we will often write $E = E(X)$. We call an e-variable E sharp if $\sup_{P \in \mathcal{P}} \mathbb{E}_P[E] = 1$.

We use $\mathbf{1}\{\cdot\}$ as an indicator, i.e., $\mathbf{1}\{F\}$ is 1 if F occurs and 0 otherwise. Throughout, we synonymize “increasing” with “non-decreasing,” and not with “strictly increasing.” Given a function $h : \mathfrak{R} \rightarrow \mathfrak{R}$, let $h^-(y) = \inf\{z : h(z) \geq y\}$ denote its (lower) generalized inverse. If h is strictly increasing and h^{-1} is well-defined, then $h^- = h^{-1}$. If h is increasing, then h^- is left-continuous and $\mathbf{1}\{h(x) \geq y\} = \mathbf{1}\{x \geq h^-(y)\}$. The generalized inverse will appear mainly in Section 5.3.1.

For a point null P and point alternative Q with $Q \ll P$ we use Λ to refer to the likelihood ratio $\Lambda(X) = \frac{Q}{P}(X)$ where Q/P is the Radon-Nikodym derivative between Q and P . We let $\mathcal{L} = \{L_b(0, 1) : b \in \mathcal{B}\}$ be the “span” of the type-I losses, to appear in Section 4.

Let us now discuss some assumptions. Suppose we observe X which has no support under any $P \in \mathcal{P}$. Any reasonable test family must reject the null in this case. If not, then a test family which does so will have lower type-II risk on all losses. Similarly, a test family which observes X with no support under any $Q \in \mathcal{Q}$ must sustain. This leads to the following observation.

Observation 2.2. *If δ is a Γ -admissible test family for any Γ , then for any $A \subset \mathcal{X}$ with $\sup_{P \in \mathcal{P}} P(A) = 0$, $\delta(X, b) = 1$ for all $X \in A$ Q -almost surely and all $b \in \mathcal{B}$. Likewise, for any $C \subset \mathcal{X}$ with $\sup_{Q \in \mathcal{Q}} Q(C) = 0$, we have $\delta(X, b) = 0$ for all $X \in C$ P -almost surely and $b \in \mathcal{B}$.*

Thus, going forward, we may safely assume that $\bigcup_{Q \in \mathcal{Q}} \text{supp}(Q) = \bigcup_{P \in \mathcal{P}} \text{supp}(P)$ (where $\text{supp}(P)$ is the support of P). In the case of testing a simple null P against a simple alternative Q , this implies that both Q and P are absolutely continuous with respect to one another ($Q \ll P$ and $P \ll Q$). Because of this assumption, we will usually say P -almost surely, instead of P - and Q -almost surely.

We will also assume that there exists some b such that $L_b(0, 1) > 1$. Otherwise a test family may always reject with probability 1 and remain type-I risk safe, making the problem uninteresting. This assumption is akin to not studying $\alpha = 1$ in the traditional Neyman-Pearson paradigm, which is uninteresting for the same reason. We also make the necessary (quite modest) measure-theoretic assumptions so that (3) holds. See [Appendix B.1](#) for precise details. Finally, going forward, we assume that we are testing a point null $\mathcal{P} = \{P\}$ against a point alternative $\mathcal{Q} = \{Q\}$.

We end this section by noting how one might modify the notation we use throughout the paper.

Remark 3 (*Koning representation* of post-hoc hypothesis tests). For the choice $L_b(0, 1) = b$, Koning [\[40\]](#) recently suggested merging the loss into the decision space of the test. Instead of a test mapping to $[0, 1]$, he rescales the range to $[0, L_b(0, 1)] \supseteq [0, 1]$. One could do the same thing throughout this paper. Such a transformation amounts to rescaling our usual test family δ and defining a new test family f given by $f(X, b) = L_b(0, 1)\delta(X, b)$. In this case, type-I risk safety is the condition $\mathbb{E}_P[\sup_b f(X, b)] \leq 1$. That is, the test itself becomes an e-value, as opposed to the quantity $\sup_b L_b(0, \delta(X, b))$. While Koning's representation definitely has a strong appeal, in this paper we opt to work with the standard range of $[0, 1]$ in the hope that it makes the framework more familiar to those new to post-hoc testing and e-values.

3. Warmup: likelihood ratio tests

In the case of a single loss function, the Neyman-Pearson lemma [\[19\]](#) implies that the likelihood ratio test is most powerful. What can we say about likelihood ratio-style tests in a setting with multiple losses?

For each fixed $b \in \mathcal{B}$, let $\phi^{\text{NP}}(X, b)$ denote the likelihood ratio test on loss $L_b(0, 1)$. Let Λ be the likelihood ratio between Q and P . The likelihood ratio test has the form

$$\phi^{\text{NP}}(X, b) = \mathbf{1}\{\Lambda(X) > \kappa(b)\} + \gamma \mathbf{1}\{\Lambda(X) = \kappa(b)\}, \quad (8)$$

where $\kappa(b)$ and γ are chosen such that

$$\mathbb{E}_P[\phi^{\text{NP}}(X, b)] = L_b^{-1}(0, 1).$$

We use the superscript "NP" to refer to Neyman-Pearson. Note that the second term $\gamma \mathbf{1}\{\Lambda(X) = \kappa(b)\}$ is only relevant for discrete distributions. We refer to Shao [\[41\]](#) for a modern treatment of the Neyman-Pearson lemma and the optimality of the likelihood ratio test.

Let δ be an arbitrary Γ -admissible test family and suppose that it acts as $\phi^{\text{NP}}(\cdot, b^*)$ on some b^* (that is, $\delta(X, b^*) = \phi^{\text{NP}}(X, b^*)$). How does it behave for $b \neq b^*$? Naively, one would like to set $\delta(\cdot, b) = \phi^{\text{NP}}(\cdot, b)$ on each b . Such a test family would not be type-I risk safe, however. This is easy to see, even in the case of two losses. Suppose that $\delta(X, b_1) = \phi^{\text{NP}}(X, b_1)$ and $\delta(X, b_2) = \phi^{\text{NP}}(X, b_2)$ with $b_1 < b_2$. Since $L_{b_2}(0, 1) > L_{b_1}(0, 1)$ we have $\kappa(b_2) > \kappa(b_1)$. If an adversary B plays loss b_2 for X with $\Lambda(X) \geq \kappa(b_2)$ and b_1 for X with $\kappa(b_1) \leq \Lambda(X) < \kappa(b_2)$, then $\mathbb{E}_P[L_B(0, \delta(X, B))] = L_{b_2}(0, 1)P(\Lambda(X) \geq \kappa(b_2)) + L_{b_1}(0, 1)Pr(\kappa(b_1) \leq \Lambda(X) < \kappa(b_2)) > 1$, so $\mathcal{R}_P(\delta) > 1$.

However, for $b \leq b^*$, $\delta(\cdot, b)$ can act as $\phi^{\text{NP}}(\cdot, b^*)$ and remain type-I risk safe. (In fact, in the discrete case, for those X such that $\Lambda(X) = \kappa(b)$ it can even raise the probability of rejection due to the extra freedom afforded by the fact that $L_b(0, 1) \leq L_{b^*}(0, 1)$.) For $b > b^*$ meanwhile, $\delta(\cdot, b)$ must lower the probability of rejection, but may still reject. That is, even in the continuous case, in contrast to the NP setting with a single loss function, if we allow randomized test families then admissible test families will always be randomized. Formally, [Proposition 3.1](#) demonstrates that, for any $\Gamma \supset \mathcal{C}$, the only admissible test family which plays as $\phi^{\text{NP}}(\cdot, b^*)$ on b^* has the following form:

$$\delta(X, b) = \min \left\{ 1, \frac{L_{b^*}(0, 1)}{L_b(0, 1)} \phi^{\text{NP}}(X, b^*) \right\}. \quad (9)$$

The proof is in [Appendix A.2](#). One may prove the result directly, but instead we leverage some of the properties of general admissible rules proven in [Section 4](#). Even so, the proof is more involved than one might initially expect.

Proposition 3.1. *The test family δ defined by (9) is Γ -admissible for any $\Gamma \supset \mathcal{C}$. Conversely, if ϕ is Γ -admissible for any $\Gamma \supset \mathcal{C}$ and there exists some b^* such that $\phi(X, b^*) = \phi^{\text{NP}}(X, b^*)$ P -almost surely then ϕ acts as (9) P -almost surely.*

The representation of δ in (9) foreshadows a more general pattern that we explore further in [Section 4](#). In particular, after noting that $E_\delta(X) = L_{b^*}(0, 1)\phi^{\text{NP}}(X, b^*)$, we see that we can write $\delta(X, b) = \min\{1, E_\delta(X)/L_b(0, 1)\}$ which is a representation that holds for all admissible tests (for any Γ) and illustrates a fundamental connection between e-values and test families.

It is worth noting that when P and Q are continuous, (9) can be rewritten as:

$$\delta(X, b) = \begin{cases} \frac{L_{b^*}(0,1)}{L_b(0,1)} \phi^{\text{NP}}(X, b^*), & b > b^*, \\ \phi^{\text{NP}}(X, b^*), & b \leq b^*. \end{cases} \quad (10)$$

This representation is perhaps more intuitive to digest. If an adversary B plays $B(X) = b^*$ for all X , the expected type-I risk $\mathbb{E}_P[L_B(1, \delta(X, B))]$ is precisely 1 by definition of the likelihood ratio test. For $b \leq b^*$ we have $L_b(0,1) \leq L_{b^*}(0,1)$ so an adversary playing such a b can only lower the type-I risk and δ can play as $\phi^{\text{NP}}(X, b^*)$. For $b > b^*$, δ may still reject, but it must lower its rejection probability to make up for the increase in type-I risk. We set δ such that $L_b(0,1)\delta(X, b) = L_{b^*}(0,1)\phi^{\text{NP}}(X, b^*)$ so that the type-I risk stays the same. We note that Koning [33] also studied (10) (see his example 9). He cites it as an example of a post-hoc test, though he does not prove that it is admissible.

The test family in (10) illustrates a general phenomenon we discussed previously, namely that admissible rules will always be randomized (or, more precisely, there will exist losses on which the test will be randomized), even for continuous distributions. Because randomization may not be desirable, it is reasonable to search for binary test families. In this case, we can show that

$$\varphi(X, b) = \begin{cases} 0, & b > b^*, \\ \phi^{\text{NP}}(X, b^*), & b \leq b^*, \end{cases} \quad (11)$$

is admissible when P and Q are assumed to be continuous. When P and Q are discrete, then $\phi^{\text{NP}}(\cdot, b^*)$ is randomized—see (8)—so this case is not of interest. The proof of the following proposition is in [Appendix A.3](#).

Proposition 3.2. *Fix any $\Gamma \supset \mathbf{C}$ and suppose P and Q are continuous. If we restrict our attention to binary test families, then (11) is Γ -admissible. Moreover, if ψ is any binary Γ -admissible test family and there exists some b^* such that $\psi(X, b^*) = \phi^{\text{NP}}(X, b^*)$, then ψ acts as (11) P -almost surely.*

Just as (9) was illustrative of a more general representation of test families when randomization is allowed, (11) is illustrative of a general representation in the binary case. Indeed, [Lemma 4.6](#) in [Section 4](#) will show that, for any $\Gamma \supset \mathbf{C}$, a Γ -admissible binary test φ can be written as $\varphi(X, b) = \mathbf{1}\{E_\delta(X) \geq L_b(0,1)\}$.

4. Properties of admissible rules

Here we prove properties of Γ -admissible families that hold for general $\Gamma \supset \mathbf{C}$, the most important results being [lemmas 4.4](#) and [4.6](#), which provide explicit representations of any admissible test δ in terms of its associated e-value E_δ (defined in (4)), generalizing [Eqs. \(9\)](#) and [\(10\)](#). All omitted proofs can be found in [Appendix A](#).

We begin by proving an intuitive property of any test family: If it can be modified to reject with strictly higher probability on some loss without sacrificing type-I risk safety, then it cannot be admissible. This fact is immediate in the traditional Neyman-Pearson framework for any uniformly most powerful test (indeed, it is the definition of being uniformly most powerful). But it requires proof in the post-hoc setting.

Lemma 4.1. *Let ϕ, δ be two type-I risk safe test families such that $\phi(X, b) \geq \delta(X, b)$ P -almost surely for all $b \in \mathcal{B}$. If ϕ is anywhere strictly greater than δ , i.e., there exists some $b^* \in \mathcal{B}$ such that $P(\phi(X, b^*) > \delta(X, b^*)) > 0$, then δ is Γ -inadmissible for any $\Gamma \supset \mathbf{C}$. The same statement holds if ϕ and δ are binary test families.*

One property which follows from [Lemma 4.1](#) is that any admissible test must be decreasing in the loss. In other words, as the loss increases, a test should only become more conservative. This fact was also noted by Koning [33, Remark 4], though not related to admissibility.

Lemma 4.2. *Fix any $\Gamma \supset \mathbf{C}$ and let δ be Γ -admissible. Then $\delta(X, b_1) \geq \delta(X, b_2)$ P -almost surely if and only if $b_1 \leq b_2$. The same statement holds for binary test families.*

A second consequence of [Lemma 4.1](#) is that the e-value of an admissible test cannot be strictly dominated by the e-value of another admissible test. Consequently, admissible tests define sharp e-values (recall that $E_\delta(X) = \sup_b L_b(0, \delta(X, b))$). The proof is postponed to [Appendix A.6](#). Note that the properties of Γ -admissible tests given by [Lemmas 4.2](#) and [4.3](#) are both illustrated by the Neyman-Pearson style tests of [Section 3](#).

Lemma 4.3. *Let ϕ, δ be two type-I risk safe test families. If $E_\phi(X) \geq E_\delta(X)$ P -almost surely and $P(E_\phi(X) > E_\delta(X)) > 0$ then δ is Γ -inadmissible for any $\Gamma \supset \mathbf{C}$. Consequently, if E_δ is not sharp then it is Γ -inadmissible for any $\Gamma \supset \mathbf{C}$. The same statement holds if ϕ and δ are binary test families and P and Q are continuous.*

Given a test family δ we can study the properties of E_δ , sometimes giving necessary and sufficient conditions on E_δ for δ to be admissible. But we will also be interested in the other direction. That is, given an e-variable E for P , the next sections will define an associated test family δ_E —[Section 4.1](#) in the randomized case and [Section 4.2](#) in the binary case. We may then ask about the properties of E that lead to admissible tests. A helpful concept when reasoning about the behavior of e-variables is that of *compatibility*. Recall that

$$\mathcal{L} \equiv \{L_b(0,1) : b \in \mathcal{B}\}. \quad (12)$$

Definition 4.1. We say an e-variable E for P is compatible with \mathcal{L} if $E(X) \in \mathcal{L} \cup \{0\}$ P -almost surely. If \mathcal{L} is understood from context, we will often simply say that E is compatible.

[Lemma 4.7](#) will show that if E defines a Γ -admissible binary test family for any Γ then E is compatible. Such a general result does not hold in the randomized case, which can be seen by considering the e-values of Neyman-Pearson tests when the distributions are discrete. However, [Lemma 5.2](#) proves that when P and Q are continuous and the randomized test δ_E is C-admissible, then E is compatible. Interestingly, the issue of compatibility also arose for Grünwald [1], even when studying a different notion of admissibility. He called the condition “richness.”

Beyond this point, the properties of randomized test families and binary test families begin to diverge so we give a separate treatment of each.

4.1. Randomized test families

Throughout this subsection we assume that all test families allow for randomization. The distributions P and Q can be continuous or discrete, unless stated otherwise. As usual, we assume they have the same support. Our first result provides an explicit form for any test family in terms of its associated e-value.

Lemma 4.4. Fix any $\Gamma \supset \mathbf{C}$. If δ is Γ -admissible then for all $b \in \mathcal{B}$,

$$\delta(X, b) = \min \left\{ 1, \frac{E_\delta(X)}{L_b(0, 1)} \right\}, \quad P\text{-almost surely.} \quad (13)$$

We will sometimes call (13) the *canonical* representation of δ , or simply say that δ is canonical if written as such. [Lemma 4.4](#) will be particularly useful in the upcoming theorems on admissibility, since it lets us assume that all test families we deal with are canonical. We note that Koning [33, Corollary 2] shows that the test in (13) is type-I risk safe.

Just as test families define e-variables, it will be convenient to enable e-variables to define test families. In view of [Lemma 4.4](#), given an e-variable E for P , we define the corresponding test family as

$$\delta_E(X, b) \equiv \min \left\{ 1, \frac{E(X)}{L_b(0, 1)} \right\}. \quad (14)$$

It is easy to verify that this definition is consistent in the sense that $\delta_{E_\delta}(X, b) = \delta(X, b)$ for a Γ -admissible test δ (this follows immediately from [Lemma 4.4](#)). Further, if we begin with an e-variable E , define δ_E and then compute E_{δ_E} , we recover E under weak assumptions. To wit,

$$E_{\delta_E}(X) = \sup_{b \in \mathcal{B}} L_b(0, \delta_E(X, b)) = \sup_{b \in \mathcal{B}} \min \{L_b(0, 1), E(X)\} = E(X) \wedge \sup_b L_b(0, 1). \quad (15)$$

If $\sup_b L_b(0, 1) = \infty$ then the final quantity is $E(X)$. In fact, we show next that if δ_E is admissible, then $E(X) \wedge \sup_b L_b(0, 1) = E(X)$ P -almost surely, which suffices for us to conclude that the final expression in (15) equals E .

Lemma 4.5. Let E be an e-variable for P . If $P(E(X) > \sup_b L_b(0, 1)) > 0$ then δ_E is Γ -inadmissible for any $\Gamma \supset \mathbf{C}$. Otherwise, $E_{\delta_E} = E$ P -almost surely, and in particular this holds if δ_E is Γ -admissible.

The proof can be found in [Appendix A.8](#). It is helpful to view [Lemma 4.5](#) as a weaker form of compatibility. It doesn't guarantee that E takes values in $\mathcal{L} \cup \{0\}$, but it does guarantee that it won't take values strictly larger than those in \mathcal{L} . As we discussed earlier, [Lemma 4.5](#) can be strengthened to bona fide compatibility when we consider admissibility with respect to constant adversaries. We now move onto studying binary test families.

4.2. Binary test families

Throughout this section we assume that P and Q are continuous. We begin with the analogue of [Lemma 4.4](#) but for binary test families. The same representation was given recently in Ramdas and Wang [17, Chapter 4], though stated as a definition, not as the unique form of any admissible binary test family. Let us note that [Koning [33] [Theorem 5.2](#) and Corollary 2] proves a very similar result. In particular, after translating his result to our setting, he shows that given an e-variable E , we can obtain a type-I risk safe test via $(X, b) \mapsto \mathbf{1}\{L_b(0, 1) \leq E(X)\}$. Conversely, given a type-I risk safe test, then $E_\phi = \sup\{L_b(0, 1) \leq E(X)\}$ is an e-value. We show that, starting with ϕ , one can write $\phi(X, b) = \mathbf{1}\{L_b(0, 1) \leq E_\phi(X)\}$ if ϕ is admissible. This also follows from Koning's result after only slightly more work.

Lemma 4.6. Fix any $\Gamma \supset \mathbf{C}$. If ϕ is a Γ -admissible binary test family then, for all $b \in \mathcal{B}$,

$$\phi(X, b) = \mathbf{1}\{L_b(0, 1) \leq E_\phi(X)\} \quad P\text{-almost surely.} \quad (16)$$

Similarly to the randomized case, let us say that a test family with the representation in (16) is *canonical*, or has canonical representation. And again as we did in the randomized case, let us associate test families to e-variables. Given an e-variable E for P , we define its associated binary test family as

$$\phi_E(X, b) \equiv \mathbf{1}\{L_b(0, 1) \leq E(X)\}. \quad (17)$$

Table 2

An overview of our admissibility results and how they relate to the traditional Neyman-Pearson paradigm ($|B| = 1$). E always refers to an e-variable above. Note the one-way implication (highlighted in red) for C-admissibility for randomized test families, which differs from the biconditionals (if and only if) in the other cases. Recall that E is sharp if $\mathbb{E}_P[E] = 1$ and E is compatible with \mathcal{L} if $E \in \mathcal{L} \cup \{0\}$ P -almost surely (Definition 4.1).

Necessary and sufficient conditions for admissibility		
Γ	$ B > 1$	$ B = 1$
Randomized test families (Continuous or discrete distributions)	U	If (C1) holds: δ is U-admissible \Leftrightarrow $\delta = \delta_E$ for E sharp and $E \leq \sup_b L_b(0, 1)$
		δ is C-admissible \Rightarrow $\delta = \delta_E$ for E sharp, increasing in Λ , and compatible with \mathcal{L}
	C	If (C2) holds: φ is U-admissible \Leftrightarrow $\varphi = \varphi_E$ for E sharp and compatible with \mathcal{L}
		φ is C-admissible \Leftrightarrow $\varphi = \varphi_E$ for E sharp, increasing in Λ , and compatible with \mathcal{L}
Binary test families (Continuous distributions)	U	
	C	

As usual, we use φ_E or ψ_E to denote the binary test family associated with E , as opposed to the randomized test defined by (14). We will always explicitly mention whether we are considering randomized or binary tests.

Both (16) and (17) are satisfying in that they are precisely what one would obtain from naively applying the results in the previous section and rounding all non-integer values taken by a test family to zero. Similarly to the randomized case, it is easy to check that $\varphi_{E_\varphi}(X, b) = \varphi(X, b)$ P -almost surely. Moreover, given an e-variable E , we have

$$E_{\varphi_E}(X) = \sup_b L_b(0, 1) \mathbf{1}\{E(X) \geq L_b(0, 1)\}, \quad (18)$$

which is equal to $E(X)$ if $E(X) = L_b(0, 1)$ for some b , or $E(X) = 0$. In other words, if $E(X)$ is compatible with \mathcal{L} then $E_{\varphi_E} = E$ for any e-variable E .

If E is not compatible with \mathcal{L} , then the test family δ_E is “wasting” some of its risk budget for no reduction in type-II loss. Indeed, suppose that for some X $L_{b_1}(0, 1) < E(X) < L_{b_2}(0, 1)$ where there is no other loss between $L_{b_1}(0, 1)$ and $L_{b_2}(0, 1)$. Then $E(X)$ defines the same test as taking $\tilde{E}(X) = L_{b_1}(0, 1)$. Moreover, $\tilde{E}(X)$ can take advantage of the reduction in type-I risk in order to take larger values elsewhere and improve its type-II risk. The following result formalizes this argument. Its proof is provided in Appendix A.10.

Lemma 4.7. *Let E be an e-variable for P and fix any $\Gamma \supset \mathbf{C}$. If φ_E is Γ -admissible then E is compatible with \mathcal{L} . Consequently, if φ_E is Γ -admissible then $E = E_{\varphi_E}$.*

5. Classifying the admissible rules

The following two sections study U-admissibility and C-admissibility. Table 2 provides a summary of the results.

5.1. U-admissibility

The first natural set of adversaries Γ to consider in the definition of Γ -admissibility is all mappings from the data to the losses; i.e., $\Gamma = \mathbf{U}$. This section classifies the set of U-admissible test families and binary test families for a point null P versus a point alternative Q . For randomized test families we will let P and Q to be either discrete or continuous. As usual, for binary test families we assume that P and Q are continuous.

When discussing randomized test families we will impose the following conditions on the losses, which restricts the behavior the ratio of type-II loss to type-I loss in the limit. In particular, we will assume that

$$\liminf_{b \rightarrow \sup(\mathcal{B})} \frac{L_b(1, 0)}{L_b(0, 1)} = 0. \quad (\text{C1})$$

Condition (C1) implies that an adversary can make type-II loss arbitrarily unimportant compared to type-I loss. In other words, as the cost of false positives increases ($L_b(0, 1)$ grows) we care more about false positives than false negatives.

Intuitively, given two test families ϕ and δ , if condition (C1) holds then an adversary B can make ϕ arbitrarily small on regions where $\phi > \delta$, thus ensuring a lower type-II risk for δ . Without such a condition, the set of admissible rules appear to depend

in unsatisfactory ways on the precise null and alternative being tested. Condition (C1) can be replaced by the assumption that $\liminf_{b \rightarrow \inf(\mathcal{B})} L_b(1, 0)/L_b(0, 1) = 0$, but this requires that $L_b(1, 0) \approx 0$ for small b , which is unsatisfying for practical applications. If the type-II loss is zero, then one should always accept the null.

With that, we state the main result for U-admissibility of randomized test families. Recall that δ is canonical if it is written as in (13).

Theorem 5.1 (U-admissibility, randomized tests). *Let δ be a randomized test family and suppose condition (C1) holds. Then δ is U-admissible if and only if E_δ is sharp and δ is canonical. Furthermore, given an e-variable E for P , δ_E is U-admissible if and only if E is sharp and $E \leq \sup_b L_b(0, 1)$ P -almost surely.*

The proof is in [Appendix A.11](#). The idea for the first part of the theorem is straightforward. Given two type-I risk safe test families δ and ϕ , we split \mathcal{X} into regions depending on whether E_δ is greater than E_ϕ or vice versa. When $E_\delta(X) > E_\phi(X)$, we can pick the adversary $B(X)$ such that $\phi(X, B) < \delta(X, B)$, ensuring that δ has smaller type-II risk in such regions. When $E_\phi(X) > E_\delta(X)$, we choose $B(X) \rightarrow \sup(\mathcal{B})$ which drives the contribution of $\phi(X, B)$ on such regions to 0.

Remark 4. The invocation of Condition (C1) in [Theorem 5.1](#) prohibits us from translating the results into the standard Neyman-Pearson setting. That is, [Observation 2.1](#) does not apply. Indeed, for a singleton $\mathcal{B} = \{b\}$, (C1) implies that $L_b(1, 0) = 0$, so type-II risk loses its meaning and does not translate into the power of a test in the traditional paradigm.

An immediate consequence of [Theorem 5.1](#) is that any U-admissible test family can be written as the test family corresponding to some e-variable, highlighting their central role in post-hoc testing:

Corollary 5.1 (U-admissibility, randomized tests). *Suppose condition (C1) holds. Then a randomized test family δ is U-admissible if and only if $\delta = \delta_E$ for some sharp e-variable E such that $E \leq \sup_b L_b(0, 1)$ P -almost surely.*

Intriguingly, [Theorem 5.1](#) and [Corollary 5.1](#) recover the result of Grünwald [1] up to some assumptions, even though he studied a weaker form of admissibility. In particular, [Theorem 5.1](#) of Grünwald [1] essentially shows that any G-admissible rule is defined by a sharp e-variable. Roughly speaking, he considers a test family to be G-admissible if it is not dominated with probability 1 by another test. (We use G-admissibility to refer to Grünwald.)

Upon reflection, the convergence of our two results is unsurprising. Given two sharp but distinct e-variables, each must be greater than the other in some region of \mathcal{X} . Neither will dominate the other with probability 1, hence neither will the corresponding tests. Thus both are G-admissible. In our setting meanwhile, Condition (C1) ensures that we can choose the adversary to take advantage of the difference between the two e-values.

Let us now provide the analogue of [Theorem 5.1](#) for binary test families. In this case we replace condition (C1) with an assumption on the density of the losses in the reals. In particular:

$$\mathcal{L} \text{ is contained in and is dense in } \mathfrak{R}_{\geq M} \text{ for some } M > 0. \quad (\text{C2})$$

Theorem 5.2 (U-admissibility, binary tests). *Let P and Q be continuous and suppose that (C2) holds. A binary test family φ is U-admissible if and only if E_φ is sharp and φ is canonical. Moreover, given an e-variable E for P , φ_E is U-admissible if and only if E is sharp and compatible with \mathcal{L} .*

The proof is in [Appendix A.12](#). It follows a similar logic to the proof of [Theorem 5.1](#). Like Condition (C1), the density assumption in [Theorem 5.2](#) prohibits us from applying the result to the standard Neyman-Pearson paradigm. That is, the conclusion of [Remark 4](#) applies in this case as well.

To end this section, we state the analogue of [Corollary 5.1](#) for binary test families.

Corollary 5.2 (U-admissibility, binary tests). *Let P and Q be continuous and suppose that \mathcal{L} is contained in and is dense in $\mathfrak{R}_{\geq 1}$. A binary test family φ is U-admissible if and only if $\varphi = \varphi_E$ for some sharp e-variable E for P which is compatible with \mathcal{L} .*

5.2. Improving tests via Rao-Blackwellization

What kinds of e-variables are available to us when testing point nulls and alternatives? For the Neyman-Pearson style tests in [Section 3](#), the corresponding e-variable is easily seen to be $E^{\text{NP}}(X) = L_{b^*}(0, 1)\mathbf{1}\{\Lambda(X) \geq \kappa(b^*)\}$ for some fixed b^* . Another e-variable is of course the likelihood ratio itself which, in the binary case, defines the satisfyingly intuitive test family $\varphi_\Lambda(X, b) = \mathbf{1}\{\Lambda(X) \geq L_b(0, 1)\}$.

These examples might make us wonder whether all “good” e-variables for post-hoc hypothesis testing for point hypotheses are functions of the likelihood ratio. In general, of course, there are e-variables which are not functions of Λ . Indeed, $E = R/P$ for any distribution R is a (sharp) e-variable. However, unless $R = Q$, such an e-variable is not log-optimal [42]. We might therefore expect that it can be improved as a post-hoc test.

Here we show that Rao-Blackwellization [20,21] can weakly improve the type-II risk of any test family based on an e-variable, with respect to a particular set of adversaries. Recall that the Rao-Blackwell theorem says that the expected loss of an estimator can be weakly improved (with respect to a convex loss) by conditioning on a sufficient statistic. See Wasserman [43, Theorem 9.42] for a modern statement of the result.

Let E be an arbitrary e-variable for P and let T be any sufficient statistic for testing P vs Q . More formally, we assume that testing P vs Q can be written as $H_0 : \theta = \theta_0$ and $H_1 : \theta = \theta_1$ and that T is sufficient for θ . Define

$$S_T(X) = \mathbb{E}_P[E(X)|\sigma(T(X))], \quad (19)$$

where $\sigma(T(X))$ is the σ -field defined by $T(X)$. Note that S is an e-variable by the law of total expectation: $\mathbb{E}_P[S] = \mathbb{E}_P[\mathbb{E}_P[E|\sigma(T)] = \mathbb{E}_P[E] \leq 1$. Further note that S_T is *observable*, meaning that it does not depend on the unknown parameter. This is by definition of sufficiency.

The following result shows that S_T defines a test family that is weakly-preferable to that defined by E . However, we must restrict ourselves to adversaries that are $\sigma(T)$ -measurable.

Proposition 5.1. *Let T be a sufficient statistic for P vs Q . Given any e-variable E for P , the test family δ_{S_T} for S_T in (19) is weakly-preferable to δ_E with respect to any $\Gamma \subset \Gamma_{\sigma(T)}$ where $\Gamma_{\sigma(T)}$ is the set of all adversaries which are $\sigma(T)$ -measurable. The same statement holds for the binary test families φ_{S_T} and φ_E .*

Proof. Let

$$f(E, B) = \min \left\{ L_B(1, 0), \frac{L_B(1, 0)}{L_B(0, 1)} E \right\}.$$

Note that f is concave in E . Therefore, by Jensen's inequality, if B is $\sigma(T)$ -measurable, we have

$$\mathbb{E}_P[f(E, B)|\sigma(T)] \leq f(\mathbb{E}_P[E|\sigma(T)], B) = f(S_T, B).$$

Since Λ is *minimally* sufficient, it follows that Λ is $\sigma(T)$ measurable (apply the Fisher-Neyman characterization to write Λ as a function of T). Therefore,

$$\begin{aligned} \mathbb{E}_Q[f(S_T, B)] &= \mathbb{E}_P[\Lambda f(S_T, B)] \geq \mathbb{E}_P[\Lambda \mathbb{E}_P[f(E, B)|\sigma(T)]] \\ &= \mathbb{E}_P[\mathbb{E}_P[\Lambda f(E, B)|\sigma(T)]] = \mathbb{E}_P[\Lambda f(E, B)] = \mathbb{E}_Q[f(E, B)]. \end{aligned}$$

We have thus shown that

$$\mathbb{E}_Q[L_B(1, 0)\delta_{S_T}(X, B)] = \mathbb{E}_Q[f(S_T, B)] \geq \mathbb{E}_Q[f(E, B)] = \mathbb{E}_Q[L_B(1, 0)\delta_E(X, B)],$$

thus showing that δ_{S_T} is weakly-preferable to δ_E with respect to all adversaries B that are $\sigma(T)$ -measurable, as desired. When considering binary test families, we instead take $f(E, B) = L_B(1, 0)\mathbf{1}\{E \geq L_B(0, 1)\}$, which is also concave in E . The rest of the argument remains unchanged. \square

For **C**-admissibility, Lemma 5.1 will show that, not only is δ_E a function of the likelihood ratio, it must be an *increasing* function of the likelihood ratio.

5.3. C-admissibility

The results of Section 5.1 were, in some sense, disappointingly liberal. As Remark 4 points out, Condition (C1) ensures that Theorems 5.1 and 5.2 cannot be mapped back to the standard hypothesis testing setting. However, even if this were not the case, the result itself is unsatisfying: There are many sharp e-variables that do not correspond to the likelihood ratio test. Ideally, we would like a result which recovers the Neyman-Pearson lemma when instantiated with a single loss function. This motivates considering a different set Γ of adversaries. In this section we study **C**-admissibility, where we recall that $\mathbf{C} = \{B : B(X) = b \text{ for some } b \text{ and all } X\}$.

It is worth revisiting Remark 1 at this point. We stress that considering a restricted family Γ of adversaries in the definition of admissibility does not change our notion of type-I risk safety, which continues to provide a guarantee with respect to all adversaries. In other words, the set of type-I risk safe test families stays the same, but the subset which are **C**-admissible is smaller (or so it will turn out) than the set which are **U**-admissible.

For test families which are permitted to randomize, we give two necessary conditions to be **C**-admissible. The first is that the test families must be increasing functions of the likelihood ratio. The second is that, when P and Q are continuous, the e-variable associated to the test must be compatible with \mathcal{L} . For binary test families we give a complete characterization of all **C**-admissible rules in terms of e-variables. As in Section 5.1 we consider a point null P and a point alternative Q .

Studying **C**-admissibility has an interesting consequence: The type-II losses do not matter. Indeed, instantiating Definition 2.4 when $\Gamma = \mathbf{C}$, we see that δ is **C**-inadmissible if, there exists a ϕ such that for all $b \in \mathcal{B}$, $\mathbb{E}_Q[L_b(1, \delta(X, b))] \leq \mathbb{E}_Q[L_b(1, \phi(X, b))]$, i.e., $\mathbb{E}_Q[\delta(X, b)] \leq \mathbb{E}_Q[\phi(X, b)]$ for all $b \in \mathcal{B}$, with a strict inequality for some b . That is, the precise values of $L_b(1, 0)$, assuming they are not zero, do not affect the admissibility of a test.

In other words, subject to type-I risk safety, we are looking for tests which maximize their power on each fixed b . This leads us to the counterfactual interpretation of **C**-admissibility of δ : For every other test ϕ , there exists a $b \in \mathcal{B}$ such that had we set our significance level to be $L_b^{-1}(0, 1)$ at the beginning of the experiment, we would not regret using δ instead of ϕ (in terms of power, i.e., $\mathbb{E}_Q[\delta(X, b)]$).

We begin by proving that any randomized **C**-admissible test must be increasing in the likelihood ratio. The result is, of course, satisfyingly intuitive—a larger likelihood ratio indicates more evidence against the null. A test that rejects the null at $\Lambda(X)$ but fails to reject when $\Lambda(Y) > \Lambda(X)$ is doing so despite increased odds in favor of the alternative. That such a result does not hold for richer classes of adversaries Γ is yet another reason to consider **C**. Indeed, tests which don't satisfy this property are flaunting the law of likelihood [28].

Lemma 5.1. *Let δ be a randomized test family. If a randomized test family δ is **C**-admissible then E_δ is an increasing function of the likelihood ratio P -almost surely. Consequently, for each fixed $b \in \mathcal{B}$, $\delta(X, b)$ is an increasing function of the likelihood ratio P -almost surely.*

The proof may be found in [Appendix A.13](#). This result mirrors a result by Koning [40] (see [Theorem 5.1](#)) with a similar flavor: he shows that e-variables which maximize a fixed expected concave utility under a (simple) alternative must also be increasing in the likelihood ratio. Connecting his result with [Lemma 5.1](#) may be an interesting avenue for future research.

Our second result shows that when P and Q are continuous, [Lemma 4.7](#) holds for randomized tests as well as binary ones. The proof may be found in [Appendix A.14](#).

Lemma 5.2. *Let P and Q be continuous. If a randomized test family δ is C-admissible then E_δ is compatible with \mathcal{L} .*

It is worth noting that the properties given in [Lemmas 5.1](#) and [5.2](#) are both on display for the likelihood ratio test. For continuous P and Q , its associated e-value is $E^{\text{NP}}(X) = L_{b^*}(0, 1)$ if $\Lambda(X) \geq \kappa(b^*)$ and 0 otherwise (see [Section 3](#)), which is clearly compatible with \mathcal{L} . However, in the discrete case, the e-value obeys $E^{\text{NP}}(X) = L_{b^*}(0, 1)(1 + \gamma)\mathbf{1}\{\Lambda(X) \geq \kappa(b^*)\}$, which may not belong to \mathcal{L} . Hence, the assumption of continuity in [Lemma 5.2](#) cannot be removed.

We leave it as an open question whether the conditions specified by [Lemmas 5.1](#) and [5.2](#) (in addition, of course, to the canonical representation in [Lemma 4.4](#)) are also sufficient to specify C-admissible tests. We expect the answer is no. We do, however, conjecture that in addition to our extensions of the likelihood ratio test, *mixtures* of likelihood ratio tests are C-admissible, subject to compatibility. More formally, if ϕ_k^{NP} is the post-hoc likelihood ratio test on loss b_k (i.e., the test defined by [\(9\)](#)) then we suspect that the test defined by the (sharp) e-variable

$$E_{\text{mix}} = \frac{1}{K} \sum_{k=1}^K E_{\phi_k}, \quad (20)$$

is C-admissible, as long as $\frac{1}{K} \sum_{k \leq j} L_{b_k}(0, 1) \in \mathcal{B}$ for each $1 \leq j \leq K$. Note that ϕ_k^{NP} is C-admissible by [Proposition 3.1](#).

5.3.1. Binary test families

We now consider C-admissible binary test families. In this case we can provide a complete characterization of admissibility. As usual for binary tests, throughout this section we assume that P and Q are continuous. We begin with the analogue of [Lemma 5.1](#) in the binary case. The proof is in [Appendix A.15](#).

Lemma 5.3. *Let P and Q be continuous and let φ be a binary test family. If φ is C-admissible then E_φ is increasing in the likelihood ratio P -almost surely. Consequently, for each fixed $b \in \mathcal{B}$, $\varphi(X, b)$ is an increasing function of the likelihood ratio P -almost surely.*

One might initially suspect that [Lemma 5.3](#) follows from [Lemma 5.1](#). In that proof, however, we tweaked the given randomized test slightly, and there was no guarantee that the resulting test family was binary (in fact it almost certainly would not be). While the spirit of the proof of [Lemma 5.3](#) is the same, the mechanics are quite different and use the fact that we begin with a binary test.

Recall from [Lemma 4.6](#) that we can represent any C-admissible binary test as $\varphi(X, b) = \mathbf{1}\{E_\varphi(X) \geq L_b(0, 1)\}$. At a fixed b , this allows us to compare two binary rules by comparing their e-values. [Lemma 5.3](#) implies that there exists some threshold $t(b)$ such that $\varphi(X, b) = \mathbf{1}\{\Lambda(X) \geq t(b)\}$ and, moreover, t is increasing in b by [Lemma 4.2](#). This allows us to compare two tests at the same X by comparing these thresholds. Moving forward, our analysis will make heavy use of these thresholds. Let us state this formally. The proof is in [Appendix A.16](#).

Corollary 5.3. *Let P and Q be continuous and let φ be a C-admissible binary test family. Then there exists an increasing function, henceforth called a decision curve, $t : \mathcal{B} \rightarrow \mathfrak{R}_{\geq 0} \cup \{\infty\}$ such that*

$$\varphi(X, b) = \mathbf{1}\{\Lambda(X) \geq t(b)\}. \quad (21)$$

We write the decision curve associated with φ as t_φ .

Remark 5. There can be multiple decision curves for a given test, but they can differ only outside the range of Λ under P and Q . That is, if t and ℓ are two decision curves for φ , then $P(t(b) \leq \Lambda(X) < \ell(b)) = 0$ for all b , otherwise φ is not a well-defined function. What happens outside the support of P and Q does not affect admissibility, so going forward we will refer to the unique decision curve t_φ of a C-admissible test φ .

As an example, if φ is the binary NP-type decision rule defined by [\(11\)](#) in [Section 3](#), then $t_\varphi(b) = \infty$ for $b > b^*$ and $t_\varphi(b) = \kappa(b^*)$ for $b \leq b^*$.

A test can be shown to be better than a second if the decision curve of the first is never above that of the second and is sometimes strictly lower. More formally, we say t_ϕ *Q*-strictly dominates t_φ if $t_\phi(b) \leq t_\varphi(b)$ for all $b \in \mathcal{B}$ and there exists some b^* such that $t_\phi(b^*) < t_\varphi(b^*)$ and $Q(t_\phi(b^*) \leq \Lambda(X) < t_\varphi(b^*)) > 0$.

It turns out that we can characterize the C-admissible by appealing to decision curves. [Lemma A.4](#) in the appendix does precisely this. But such a classification is not practically useful. To move beyond decision curves and state C-admissibility in terms of e-variables, we require the following technical definition.

Definition 5.1. *We say that a binary test family φ has a minimal decision curve if for any $b^* \in \mathcal{B}$, whenever $\sup_{b < b^*} L_b(0, 1) = L_{b^*}(0, 1)$, we either have $\sup_{b < b^*} t_\varphi(b) = t_\varphi(b^*)$ or $P(\sup_{b < b^*} t_\varphi(b) \leq \Lambda(X) < t_\varphi(b^*)) = 0$.*

[Definition 5.1](#) is vacuous whenever \mathcal{L} is discrete, meaning for every $L_b(0, 1) \in \mathcal{L}$ there exists some $\epsilon > 0$ such that $(L_b(0, 1) - \epsilon, L_b(0, 1) + \epsilon)$ does not contain any other element of \mathcal{L} . It is relevant only for \mathcal{L} that is dense (or has a subset which is dense). When P

puts mass everywhere, the definition is equivalent to requiring that the decision curve be left continuous whenever the type-I losses are left continuous.

Lemma A.5 shows that a C-admissible test must have a minimal decision curve. The intuition is as follows. Suppose $\sup_{b < b^*} L_b(0, 1) = L_{b^*}(0, 1)$ but $\sup_{b < b^*} t_\varphi(b) < t_\varphi(b^*)$. Then, for X such that $\sup_{b < b^*} t_\varphi(b) \leq \Lambda(X) \leq t_\varphi(b^*)$, $\sup_b L_b(0, 1) \mathbf{1}\{\Lambda(X) \geq t_\varphi(b)\} = L_{b^*}(0, 1)$. However, $\varphi(X, b^*) = 0$ for such X . That is, φ is suffering a type-I loss of $L_{b^*}(0, 1)$ on such X but not receiving the benefits of rejecting on loss b^* . In this case, one can define a second test family which rejects on b^* and suffers the same type-I loss, thus making φ C-inadmissible.

We can now state the main result of this section. The proof is in [Appendix A.17](#).

Theorem 5.3. *Let P and Q be continuous. A canonical binary test family φ is C-admissible if and only if it has a minimal decision curve and E_φ is sharp and an increasing function of the likelihood ratio. Moreover, given an e-variable E for P , φ_E is C-admissible if and only if E is sharp, compatible, and increasing in the likelihood ratio.*

As was done in [Section 5.1](#), it is worth stating the following immediate corollary of [Theorem 5.3](#) which highlights the direct correspondence between admissible tests and e-variables.

Corollary 5.4. *Let P and Q be continuous. A binary test family φ is C-admissible if and only if $\varphi = \varphi_E$ for some sharp e-variable E for P which is compatible with \mathcal{L} and increasing in the likelihood ratio.*

[Theorem 5.3](#) makes no assumptions on the losses, hence can recover the Neyman-Pearson lemma for continuous distributions. Indeed, as we discussed previously, for a single loss $L(0, 1) = 1/\alpha$, the e-variable associated to the likelihood ratio $\varphi^{\text{NP}}(X) = \mathbf{1}\{\Lambda(X) \geq \kappa\}$ is $E^{\text{NP}}(X) = \alpha^{-1} \mathbf{1}\{\Lambda(X) \geq \kappa\}$ and is sharp, increasing in Λ , and compatible by construction. We can also see that φ^{NP} is the unique test which satisfies [Corollary 5.4](#). Consider any other C-admissible rule, which can be written as $\varphi_E(X) = \mathbf{1}\{E(X) \geq 1/\alpha\}$ for some E . Since E is compatible with $\mathcal{L} = \{1/\alpha\}$ and $\mathbb{E}_P[E] = 1$, we have $P(E(X) = 1/\alpha) = \alpha$. Additionally, $E = h(\Lambda)$ for some increasing h by [Lemma 5.3](#), so $\varphi_E(X) = \mathbf{1}\{\Lambda(X) \geq h^{-1}(1/\alpha)\}$. Therefore $\alpha = P(E(X) \geq 1/\alpha) = P(\Lambda(X) \geq h^{-1}(1/\alpha))$, meaning that $h^{-1}(1/\alpha) = \kappa$ and $\varphi_E = \varphi^{\text{NP}}$.

Now consider $L_b(0, 1) = b$ for all $b > 0$. In this case the test

$$\varphi_\Lambda(X, b) = \mathbf{1}\{\Lambda(X) \geq L_b(0, 1)\}, \quad (22)$$

is C-admissible. Indeed, Λ is sharp, compatible with $\mathfrak{R}_{\geq 0}$ and obviously increasing in itself. This test is satisfying in that it accords well with calls to report the likelihood ratio instead of the p-value [44–47]. This nicely bridges the Fisherian and Neyman-Pearson perspectives on hypothesis testing: If a study summarizes its findings via $\Lambda(X)$ as a measure of evidence, downstream analysts can reject at whatever level they like, maintaining type-I risk safety.

6. Summary

We studied post-hoc hypothesis testing, a framework which allows the significance level to be chosen as a function of the data. The notions of type-I and type-II error are replaced by type-I and type-II risk, which take the form of expectations over data-dependent type-I and type-II loss functions. In order to generalize the notion of a uniformly most powerful test to this setting, we introduced the notion of Γ -admissibility, where Γ is a set of functions mapping the data to a type-II loss function.

Different families Γ result in different classes of admissible tests. For point hypotheses, when Γ is all mappings from the data to losses we show that all admissible tests correspond to sharp e-variables. When Γ is the set of constant mappings and the tests are binary, we give a classification of all admissible tests which recovers the Neyman-Pearson lemma [19].

This work suggests several open questions. For one, in the case of randomized tests, the set of necessary and sufficient conditions for C-admissibility do not match. Can we close the gap and give a classification of all admissible tests in this case? As discussed in [Section 5.3](#), we expect that mixtures of likelihood ratio tests are admissible but we have been unable to prove it. Second, while we have focused mainly on point hypotheses, hypothesis testing in practice typically involves composite hypotheses. While some of our results extend straightforwardly to composite settings,³ can we obtain composite versions of our main classification theorems? Third, we have adopted the typical hypothesis testing setup of two actions: either sustaining or rejecting the null. Grünwald [1], however, studies multiple actions. How must our results be modified in this case?

Overall, post-hoc hypothesis testing is in its nascent stage. We hope our work provides a foundation that will inspire further research. The connections we have established between admissible tests and e-variables connect post-hoc hypothesis testing to exciting developments in modern statistics, and our generalization of classical results suggest that post-hoc testing may serve as a principled bridge between traditional hypothesis testing and more adaptive approaches to statistical inference. While our work here has been primarily theoretical, we also hope that post-hoc hypothesis testing may offer solutions to some of the problems plaguing various empirical fields as they grapple with issues of statistical validity and reproducibility.

³ In particular, [Lemmas 4.1–4.3](#) can be extended to composite settings if handled correctly. For instance, a composite version of [Lemma 4.1](#) would posit that $\varphi(X, b) \geq \delta(X, b)$ P -almost surely for all $P \in \mathcal{P}$ and that there exists some $Q \in \mathcal{Q}$ such that $Q(\varphi(X, b) > \delta(X, b)) > 0$. The representations in [Lemmas 4.4](#) and [4.6](#) likewise hold in composite settings.

CRediT authorship contribution statement

Ben Chugg: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization; **Tyron Lardy:** Writing – review & editing, Methodology, Investigation, Formal analysis, Conceptualization; **Aaditya Ramdas:** Writing – review & editing; **Peter Grünwald:** Writing – review & editing, Formal analysis, Conceptualization.

Data availability

No data was used for the research described in the article.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ben Chugg reports financial support was provided by National Science Foundation. Peter Grunwald reports financial support was provided by European Research Council. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Auxiliary results and omitted proofs

A.1. Technical Lemmas

Lemma A.1. Suppose there exist sets A_1, A_2 such that $\sup_{X_1 \in A_1} \Lambda(X_1) < \inf_{X_2 \in A_2} \Lambda(X_2)$. Then both of the following hold: (a) if $P(A_1) \leq P(A_2)$ then $Q(A_1) < Q(A_2)$, and (b) $Q(A_1)/Q(A_2) < P(A_1)/P(A_2)$.

Proof. Define

$$\bar{c} = \sup_{X \in A_1} \frac{Q}{P}(X), \text{ and } \underline{c} = \inf_{X \in A_2} \frac{Q}{P}(X).$$

Observe that

$$Q(A_1) = \int_{A_1} Q(x) \frac{P}{P}(x) = \int_{A_1} \Lambda(x) \frac{P}{P}(x) \leq \bar{c} \int_{A_1} \frac{P}{P}(x) = \bar{c} P(A_1).$$

Similarly,

$$Q(A_2) = \int_{A_2} Q(x) \frac{P}{P}(x) \geq \underline{c} \int_{A_2} \frac{P}{P}(x) = \underline{c} P(A_2).$$

To prove (b), use the above two displays and note that $\bar{c} < \underline{c}$. Then, by assumption, we have

$$\frac{Q(A_1)}{Q(A_2)} \leq \frac{\bar{c} P(A_1)}{\underline{c} P(A_2)} \leq \frac{\bar{c} P(A_1)}{\underline{c} P(A_2)} < \frac{P(A_1)}{P(A_2)}.$$

As for (a), suppose that $P(A_1) \leq P(A_2)$. Then,

$$Q(A_1) \leq \bar{c}, P(A_1) \leq \bar{c} \Pr(A_2) \leq \underline{c} \Pr(A_2) \leq Q(A_2),$$

as desired. \square

Lemma A.2. Let δ be Γ -admissible for any $\Gamma \supset \mathbf{C}$. Then $\delta(\cdot, b)$ is increasing in the likelihood ratio for each $b \in \mathcal{B}$ if and only if E_δ is increasing in the likelihood ratio. The same statement holds if $\delta = \varphi$ is a binary test family.

Proof. Suppose $\Lambda(X_1) \leq \Lambda(X_2)$. If $E_\delta(X_1) \leq E_\delta(X_2)$ then $\min\{1, E_\delta(X_1)/L_b(0, 1)\} \leq \min\{1, E_\delta(X_2)/L_b(0, 1)\}$ for each b , so applying Lemma 4.4 demonstrates that $\delta(X_1, b) \leq \delta(X_2, b)$. Conversely, suppose that $E_\delta(X_1) > E_\delta(X_2)$. Consider any b such that $L_b(0, 1) \geq E_\delta(X_1)$ (this must exist, otherwise $E_\delta(X_1) > \sup_b L_b(0, 1)$, which is impossible by definition). For such b , $\min\{L_b(0, 1), E_\delta(X_1)\} = E_\delta(X_1) > E_\delta(X_2) = \min\{L_b(0, 1), E_\delta(X_2)\}$, implying that $\phi(X_1, b) > \phi(X_2, b)$. This proves the result in the randomized case. The argument in the binary case is proved similarly. \square

Lemma A.3. Let P, Q be continuous and let δ be Γ_{const} -admissible. Let $Y \subset \mathcal{X}$ satisfy $Q(Y) > 0$ and $E_\delta(X) > 0$ for all $X \in Y$. Then there exists an adversary $B^* : \mathcal{X} \rightarrow \mathcal{B}$ such that $E_\delta(X) \geq L_{B^*(X)}(0, 1)$ for all $X \in Y$ Q -almost surely.

Proof. Suppose not. Then

$$\sup_{X \in Y} E_\delta(X) < \inf_{b \in \mathcal{B}} L_b(0, 1).$$

Since $E_\delta > 0$ on Y , we can find some $\gamma > 0$ and $A \subset Y$ with $Q(A) > 0$ such that

$$\gamma \leq \inf_{X \in A} E_\delta(X).$$

Let $A = A_1 \cup A_2$ where $P(A_1) = P(A_2)$ and $\Lambda(X_1) < \Lambda(X_2)$ for all $X_1 \in A_1$ and $X_2 \in A_2$. (This is possible by continuity.) Define a new decision rule $\hat{\delta}$ such that, for all $b \in \mathcal{B}$, $\hat{\delta}(X, b) = \delta(X, b) - \epsilon/L_b(0, 1)$ for all $X \in A_1$ and $\hat{\delta}(X, b) = \delta(X, b) + \epsilon/L_b(0, 1)$ for all $X \in A_2$. For $X \notin A$, $\hat{\delta}$ is the same as δ . Note that by [Lemma 4.4](#), $\delta(X, b) = E_\delta(X)/L_b(0, 1) \in (0, 1)$ for $X \in A$ and $b \in \mathcal{B}$. Therefore,

$$\hat{\delta}(X, b) = \frac{E_\delta(X) - \epsilon}{L_b(0, 1)} \text{ for } X \in A_1 \text{ and } \hat{\delta}(X, b) = \frac{E_\delta(X) + \epsilon}{L_b(0, 1)} \text{ for } X \in A_2.$$

Therefore, we may find ϵ small enough such that $0 < \hat{\delta}(X, b) \leq 1$ for all $X \in A$. Indeed, to ensure that $\hat{\delta}(X, b) > 0$ we need to ensure that $\epsilon < \gamma$ and for $\hat{\delta}(X, b) \leq 1$ we need $\epsilon < \inf_{b \in \mathcal{B}} L_b(0, 1) - \sup_{X \in A} E_\delta(X)$, where the final quantity is greater than zero by assumption. Fix any such $\epsilon > 0$.

To see that $\hat{\delta}$ is type-I risk safe, write

$$\begin{aligned} \int_A E_{\hat{\delta}}(x) \mathbb{P} &= \int_{A_1} \sup_b L_b(0, 1) \left(\delta(x, b) - \frac{\epsilon}{L_b(0, 1)} \right) \mathbb{P} \\ &\quad + \int_{A_1} \sup_b L_b(0, 1) \left(\delta(x, b) + \frac{\epsilon}{L_b(0, 1)} \right) \mathbb{P} \\ &= \int_{A_1} E_\delta(x) \mathbb{P} - \epsilon P(A_1) + \int_{A_2} E_\delta(x) \mathbb{P} + \epsilon P(A_2) \\ &= \int_A \sup_b L_b(0, 1) \delta(x, b) \mathbb{P}, \end{aligned}$$

since $P(A_1) = P(A_2)$ by assumption. From this it follows that $\mathcal{R}_P(\hat{\delta}) \leq 1$ since

$$\mathbb{E}_P[E_{\hat{\delta}}(X)] = \int_A E_{\hat{\delta}}(x) \mathbb{P} + \int_{\mathcal{X} \setminus A} E_\delta(x) \mathbb{P} = \int_{\mathcal{X}} E_\delta(x) \mathbb{P} \leq 1.$$

For type-II risk, fix any $b \in \mathcal{B}$ and notice that

$$\int_A \hat{\delta}(x, b) \mathbb{Q} = \int_{A_1 \cup A_2} \delta(x, b) \mathbb{Q} - \frac{\epsilon}{L_b(0, 1)} Q(A_1) + \frac{\epsilon}{L_b(0, 1)} Q(A_2).$$

[Lemma A.1](#) implies that $Q(A_2) > Q(A_1)$ by our choice of A_1 and A_2 . Therefore $-\epsilon Q(A_1) + \epsilon Q(A_2) > 0$ and we conclude

$$\int_A \hat{\delta}(x, b) \mathbb{Q} > \int_A \delta(x, b) \mathbb{Q},$$

implying that $\mathbb{E}_Q[L_B(1, \delta(X, B))] < \mathbb{E}_Q[L_B(1, \hat{\delta}(X, B))]$. Hence δ is Γ_{all} -inadmissible. \square

Recall that a decision curve t_ϕ Q -strictly dominates t_φ if $t_\phi(b) \leq t_\varphi(b)$ for all $b \in \mathcal{B}$ and there exists some b^* such that $t_\phi(b^*) < t_\varphi(b^*)$ and $Q(t_\phi(b^*)) \leq \Lambda(X) < t_\varphi(b^*)) > 0$.

Lemma A.4. *Let P and Q be continuous and let φ be a type-I risk safe binary test family. Then φ is C-inadmissible if and only if there exists some binary type-I risk safe ψ such that t_ψ Q -strictly dominates t_φ .*

Proof. Let ψ be a binary test family which is type-I risk safe and satisfies $t_\psi(b) \leq t_\varphi(b)$ for all $b \in \mathcal{B}$ and $t_\psi(b^*) < t_\varphi(b^*)$ where $Q(t_\psi(b^*)) < \Lambda(X) < t_\varphi(b^*)) > 0$. Then for all $b \in \mathcal{B}$,

$$\varphi(X, b) = \mathbf{1}\{\Lambda(X) \geq t_\varphi(b)\} \leq \mathbf{1}\{\Lambda(X) \geq t_\psi(b)\} = \psi(X, b),$$

and the inequality is strict for b^* on a set of positive measure under Q . Apply [Lemma 4.1](#) to see that φ is C-inadmissible. Conversely, suppose that φ is C-inadmissible, so there exists some ψ satisfying $\mathbb{E}_Q[\psi(X, b)] \geq \mathbb{E}_Q[\varphi(X, b)]$ for all b , with a strict inequality for at least one b^* . Using the representation (21) of both test families, this implies that

$$Q(\Lambda(X) \geq t_\varphi(b)) \leq Q(\Lambda(X) \geq t_\psi(b)),$$

for all b , so $t_\varphi(b) \geq t_\psi(b)$. For b^* , $0 < \mathbb{E}_Q[\mathbf{1}\{\Lambda(X) \geq t_\psi(b^*)\} - \mathbf{1}\{\Lambda(X) \geq t_\varphi(b^*)\}] = Q(t_\psi(b^*) \leq \Lambda(X) < t_\varphi(b^*))$, completing the proof. \square

Lemma A.5. *Let P and Q be continuous and let φ be a type-I risk safe binary test family. If φ is C-admissible then it has a minimal decision curve.*

Proof. We prove the contrapositive. Suppose that φ does not have a minimal decision curve, so there exists some b^* such that $\sup_{b < b^*} t_\delta(b) < t_\delta(b^*)$ and $P(\sup_{b < b^*} t_\delta(b) \leq \Lambda(X) < t_\delta(b^*)) > 0$. Define a new function $\ell : \mathcal{B} \rightarrow \mathfrak{R}_{\geq 0} \cup \{\infty\}$ as

$$\ell : b \mapsto \begin{cases} t_\delta(b), & b \neq b^*, \\ \sup_{b < b^*} t_\delta(b), & b = b^*. \end{cases}$$

Observe that ℓ Q -strictly dominates t_δ . We claim that it also defines a type-I risk safe test family, which will show that δ is Γ_{const} -inadmissible by [Lemma A.4](#). Note that since t_δ is increasing in b by [Corollary 5.3](#), so too is ℓ .

Partition \mathcal{X} into $C_1 = \{X : \Lambda(X) < \ell(b^*)\}$, $C_2 = G(b^*) = \{X : \ell(b^*) \leq \Lambda(X) < t_\delta(b^*)\}$, $C_3 = \{X : \Lambda(X) \geq t_\delta(b^*)\}$. On C_1 and C_3 , we have $\sup_b L_b(0, 1) \mathbf{1}\{\Lambda(X) \geq t_\delta(b)\} = \sup_b L_b(0, 1) \mathbf{1}\{\Lambda(X) \geq \ell(b)\}$ since, in both cases, the supremum does not involve b^* . Now, for $X \in C_2$, $\mathbf{1}\{\Lambda(X) \geq t_\delta(b)\} = 1$ if and only if $b < b^*$, so

$$\sup_b L_b(0, \varphi(X, b)) = \sup_b L_b(0, 1) \mathbf{1}\{\Lambda(X) \geq t_\delta(b)\} = \sup_{b < b^*} L_b(0, 1) = L_{b^*}(0, 1).$$

Moreover, again for $X \in C_2$, $\mathbf{1}\{\Lambda(X) \geq \ell(b)\} = 1$ if and only if $b \leq b^*$, giving that $\sup_b L_b(0, 1)\mathbf{1}\{\Lambda(X) \geq \ell(b)\} = \sup_{b \leq b^*} L_b(0, 1) = L_{b^*}(0, 1)$. Therefore,

$$\int_{C_2} \sup_b L_b(0, 1)\mathbf{1}\{\Lambda(x) \geq \ell(b)\} \mathbb{P} = \int_{C_2} \sup_b L_b(0, \varphi(x, b)) \mathbb{P},$$

which, combined with equality on C_1 and C_3 , implies that the risk of the test family defined by ℓ is equal to $\mathcal{R}(\varphi)$, completing the argument. \square

A.2. Proof of Proposition 3.1

First let us show that δ is type-I risk safe. Since $\delta(X, b) = 0$ on $\{\Lambda(X) < \kappa(b^*)\}$,

$$\begin{aligned} \mathcal{R}_P(\delta) &= \int_{\{\Lambda(x) \geq \kappa(b^*)\}} \sup_b L_b(0, \delta(x, b)) \mathbb{P} + \int_{\{\Lambda(x) < \kappa(b^*)\}} \sup_b L_b(0, \delta(x, b)) \mathbb{P} \\ &= \int_{\{\Lambda(x) > \kappa(b^*)\}} \sup_b L_b(0, 1) \frac{L_{b^*}(0, 1)}{L_b(0, 1)} \mathbb{P} + \int_{\{\Lambda(x) = \kappa(b^*)\}} \sup_b L_b(0, 1) \frac{L_{b^*}(0, 1)}{L_b(0, 1)} \gamma \mathbb{P} \\ &= \int_{\{\Lambda(x) > \kappa(b^*)\}} L_{b^*}(0, 1) \mathbb{P} + \int_{\{\Lambda(x) = \kappa(b^*)\}} \gamma L_{b^*}(0, 1) \mathbb{P} \\ &= L_{b^*}(0, 1) \mathbb{E}_P[\phi^{NP}(X, b^*)] = 1. \end{aligned}$$

Now we show it is admissible. Suppose there is a type-I risk safe test family ϕ (distinct from ϕ^{NP}) which is weakly preferable to δ (with respect to Γ). That is, $\mathbb{E}_Q[L_B(1, \delta(X, b))] \leq \mathbb{E}_Q[L_B(1, \phi(X, b))]$ for all $B \in \Gamma$. If we consider the constant map $B(X) = b^*$, which is in Γ by assumption, we obtain $\mathbb{E}_Q[\phi^{NP}(X, b^*)] \leq \mathbb{E}_Q[\phi(X, b^*)]$. We claim this implies that $\phi(\cdot, b^*) = \phi^{NP}(\cdot, b^*)$ P -almost-surely. Indeed, this follows from the uniqueness in the Neyman-Pearson lemma, but let us prove it directly since it may not be a priori obvious that the same logic transfers to the post-hoc setting. Consider the integral

$$I = \int_{\mathcal{X}} (\underbrace{\phi^{NP}(x, b^*) - \phi(x, b^*)}_{:= t_1}) (\underbrace{Q - \kappa(b^*) \mathbb{P}}_{:= t_2}).$$

We claim that $I = 0$. We begin by showing that the integrand is nonnegative. This is clear if $\phi^{NP}(x, b^*) = \phi(x, b^*)$. If $\phi^{NP}(x, b^*) = 1$ then $Q(x)/P(x) > \kappa(b^*)$ so both t_1 and t_2 are nonnegative. If $\phi^{NP}(x, b^*) = 0$ then $Q(x)/P(x) < \kappa(b^*)$ so both terms are nonpositive. If $0 < \phi^{NP}(x, b^*) < 1$ then $Q(x) = \kappa P(x)$ so t_1 and t_2 multiply to zero. This shows that $I \geq 0$. Hence,

$$\begin{aligned} \int_{\mathcal{X}} (\phi^{NP}(x, b^*) - \phi(x, b^*)) Q &\geq \kappa(b^*) \int_{\mathcal{X}} (\phi^{NP}(x, b^*) - \phi(x, b^*)) \mathbb{P} \\ &= \kappa(b^*) (L_{b^*}^{-1}(0, 1) - \mathbb{E}_P[\phi(X, b^*)]). \end{aligned}$$

Now, we must have $\mathbb{E}_P[\phi(X, b^*)] \leq L_{b^*}(0, 1)$, otherwise $\mathcal{R}_P(\phi) \geq \mathbb{E}_P[L_{b^*}(0, 1)\phi(X, b^*)] > 1$, contradicting that ϕ is type-I risk safe. The last term in the above display is thus nonnegative, and we have that

$$\mathbb{E}_Q[\phi^{NP}(X, b^*)] - \mathbb{E}_Q[\phi(X, b^*)] = \int_{\mathcal{X}} (\phi^{NP}(x, b^*) - \phi(x, b^*)) Q \geq 0.$$

Since we know from above that $\mathbb{E}_Q[\phi^{NP}(X, b^*)] \leq \mathbb{E}_Q[\phi(X, b^*)]$, we have shown that $\mathbb{E}_Q[\phi^{NP}(X, b^*)] = \mathbb{E}_Q[\phi(X, b^*)]$. Finally, we can rewrite I as $I = -\kappa(b^*)(L_{b^*}(0, 1) - \mathbb{E}_P[\phi(X, b^*)])$ which we now see must equal 0, since $I \geq 0$ but we argued above that $\mathbb{E}_P[\phi(X, b^*)] \leq L_{b^*}(0, 1)$. Since $I = 0$ with a positive integrand, the integrand must be 0 almost surely. This implies that $\phi^{NP}(X, b^*) = \phi(X, b^*)$ except on the set $\{x : Q(x) = \kappa(b^*)P(x)\}$. If this set has positive measure under Q (as it does in the discrete case) and $\phi^{NP}(X, b^*) \neq \phi(X, b^*)$, then one has higher power than the other, implying that $\mathbb{E}_Q[\phi^{NP}(X, b^*)] \neq \mathbb{E}_Q[\phi(X, b^*)]$, a contradiction. We conclude that $\phi(X, b^*) = \phi^{NP}(X, b^*)$ P -almost surely.

Returning to the main proof, write

$$E_{\phi}(X) = \sup_b L_b(0, \phi(X, b)) \geq L_{b^*}(0, \phi(X, b^*)) = L_{b^*}(0, \phi^{NP}(X, b^*)).$$

Hence $\mathbb{E}_P[E_{\phi}(X)] \geq \mathbb{E}_P[L_{b^*}(0, \phi^{NP}(X, b^*))] = 1$ and we may conclude that

$$E_{\phi}(X) = L_{b^*}(0, \phi^{NP}(X, b^*)),$$

since ϕ is assumed to be type-I risk safe. By Lemma 4.4 we can write

$$\phi(X, b) = \min \left\{ 1, \frac{E_{\phi}(X)}{L_b(0, 1)} \right\} = \min \left\{ 1, \frac{L_{b^*}(0, 1)}{L_b(0, 1)} \phi^{NP}(X, b^*) \right\}, \quad (\text{A.1})$$

as desired. Now we prove the second half of the proposition. Suppose δ is Γ -admissible and satisfies $\delta(X, b^*) = \phi^{NP}(X, b^*)$ for some b^* . We are tasked with showing that $\delta = \hat{\phi}$ P -almost surely, where $\hat{\phi}$ is defined as the right hand side of (A.1). First we show that δ cannot

be greater than $\hat{\phi}$. Suppose that $\delta(X, b_0) > \hat{\phi}(X, b_0)$ for all $X \in A \subset \mathcal{X}$ and some $b_0 \in \mathcal{B}$ where $P(A) > 0$. We will show that δ is not type-I risk safe. Observe that since $\hat{\phi}(X, b_0) < 1$ for $X \in A$ (since $1 \geq \delta(X, b_0) > \hat{\phi}(X, b_0)$), we have $\hat{\phi}(X, b_0) = L_{b^*}(0, 1)\phi^{\text{NP}}(X, b^*)/L_{b_0}(0, 1)$. Consider an adversary B defined as $B(X) = b^*$ for $X \in \mathcal{X} \setminus A$ and $B(X) = b_0$ for $X \in A$. Then

$$\begin{aligned} \mathcal{R}_P(\delta) &\geq \int_A L_{B(x)}(0, \delta(x, B(x)))P + \int_{\mathcal{X} \setminus A} L_{B(x)}(0, \delta(x, B(x)))P \\ &> \int_A L_{b_0}(0, 1)\hat{\phi}(X, b_0)P + \int_{\mathcal{X} \setminus A} L_{b^*}(0, 1)\delta(x, b^*)P \\ &= \int_A L_{b^*}(0, 1)\phi^{\text{NP}}(X, b^*)P + \int_{\mathcal{X} \setminus A} L_{b^*}(0, 1)\phi^{\text{NP}}(X, b^*)P \\ &= L_{b^*}(0, 1)\mathbb{E}_P[\phi^{\text{NP}}(X, b^*)] = 1, \end{aligned}$$

so δ is not type-I risk safe. We conclude that we must have $\delta(X, b) \leq \hat{\phi}(X, b)$ P -almost surely for all $b \in \mathcal{B}$. Now suppose that for some b_0 , $\delta(X, b_0) < \hat{\phi}(X, b_0)$ for all $X \in A$ where again $P(A) > 0$. Consider defining a new test family $\hat{\delta}$ such that $\hat{\delta} = \delta$, except on A and b_0 in which case we set $\hat{\delta}(X, b_0) = \hat{\phi}(X, b_0)$. Since we've already shown that $\delta \leq \hat{\phi}$ P -almost surely, it follows that $\mathcal{R}_P(\hat{\delta}) \leq \mathcal{R}_P(\delta) \leq 1$, so $\hat{\delta}$ is type-I risk safe. Moreover, it has greater type-II risk than δ . To see this, consider any adversary B and let $A_0 \subset A$ be the (possibly empty) set of X such that $B(X) = b_0$. Then

$$\begin{aligned} \mathbb{E}_Q[L_B(1, 0)\hat{\delta}(X, B)] &= \int_{A_0} L_{b_0}(1, 0)\hat{\delta}(x, b_0)Q + \int_{\mathcal{X} \setminus A_0} L_{B(x)}(1, 0)\delta(x, B(x))Q \\ &\geq \int_{A_0} L_{b_0}(1, 0)\hat{\delta}(x, b_0)Q + \int_{\mathcal{X} \setminus A_0} L_{B(x)}(1, 0)\delta(x, B(x))Q \\ &= \mathbb{E}_Q[L_B(1, 0)\delta(X, B)], \end{aligned}$$

hence $\mathbb{E}_Q[L_B(1, \hat{\delta}(X, B))] \leq \mathbb{E}_Q[L_B(1, \delta(X, B))]$. If A_0 has positive measure under Q then the inequality above becomes a strictly inequality, showing that $\hat{\delta}$ sometimes has strictly lower type-II risk than δ (such an adversary exists in Γ by assumption) making it strictly preferable to δ with respect to Γ , completing the proof.

A.3. Proof of Proposition 3.2

The proof has similar mechanics to that of [Proposition 3.2](#). First note that φ is indeed type-I risk safe. By definition, $\varphi(X, b) = 0$ whenever $\Lambda(X) < \kappa(b^*)$ and $\varphi(X, b)$ iff $b \leq b^*$ otherwise. Therefore,

$$\begin{aligned} \mathcal{R}_P(\delta) &= \int_{\Lambda \geq \kappa(b^*)} \sup_b L_b(0, \varphi(X, b))P + \int_{\Lambda(X) < \kappa(b^*)} \sup_b L_b(0, 0)P \\ &= \int_{\Lambda \geq \kappa(b^*)} L_{b^*}(0, 1)P = L_{b^*}(0, 1)P(\Lambda(X) \geq \kappa(b^*)) = 1, \end{aligned}$$

by definition of $\kappa(b^*)$. Now suppose that ψ is any other type-I risk safe binary test family that is weakly preferable to φ with respect to any Γ . For the adversary $B(X) = b^*$ we have $\mathbb{E}_Q[\psi(X, b^*)] \geq \mathbb{E}_Q[\varphi(X, b^*)] = \mathbb{E}_Q[\phi^{\text{NP}}(X, b^*)]$. As in [Appendix A.2](#), this implies that $\psi(X, b^*) = \phi^{\text{NP}}(X, b^*)$ Q -almost surely, which in turn implies that $E_\psi(X) = \sup_b L_b(0, \psi(X, b)) = L_{b^*}(0, 1)\phi^{\text{NP}}(X, b^*)$ P -almost surely. Therefore, by the canonical representation given in [Lemma 4.6](#), we have

$$\psi(X, b) = \mathbf{1}\{E_\psi(X) \geq L_b(0, 1)\} = \mathbf{1}\{L_{b^*}(0, 1)\phi^{\text{NP}}(X, b^*) \geq L_b(0, 1)\}, \quad (\text{A.2})$$

P -almost surely, which is precisely $\varphi(X, b)$. Next, suppose that φ is Γ -admissible and there exists some b^* such that $\varphi(X, b^*) = \phi^{\text{NP}}(X, b^*)$. Then we cannot have $\varphi(X, b) > 0$ for any $b > b^*$ on any set A of positive measure under P , otherwise an adversary can ensure the type-I risk is higher than 1 by playing b on A and b^* everywhere else. (We omit the precise calculation because it is similar to the proof in [Appendix A.2](#)). Similarly, we cannot have $\delta(X, b) < \phi^{\text{NP}}(X, b^*)$ for $b \leq b^*$ and $\Lambda(X) \geq \kappa(b^*)$ otherwise φ is Γ -inadmissible by [Lemma 4.1](#). This completes the proof.

A.4. Proof of Lemma 4.1

By assumption we may assume that anything with positive probability under P has positive probability under Q and vice versa. For any $b \in \mathcal{B}$, $L_b(1, \phi(X, b)) = L_b(1, 0)(1 - \phi(X, b)) \leq L_b(1, 0)(1 - \delta(X, b)) = L_b(1, \delta(X, b))$. Therefore, for any $B \in \Gamma$,

$$\mathbb{E}_Q[L_B(1, \phi(X, B))] \leq \mathbb{E}_Q[L_B(1, \delta(X, B))].$$

Suppose $A \subset \mathcal{X}$ is such that $\phi(X, b^*) > \delta(X, b^*)$ for all $X \in A$ and $Q(A) > 0$. Then ϕ has strictly lower type-II risk on the map $B(X) = b^*$ since

$$\mathbb{E}_Q[L_{b^*}(1, \phi(X, b^*))\mathbf{1}\{X \in A\}] < \mathbb{E}_Q[L_{b^*}(1, \delta(X, b^*))\mathbf{1}\{X \in A\}],$$

making ϕ strictly preferable to δ with respect to any Γ ($B \in \Gamma$ by assumption).

A.5. Proof of Lemma 4.2

Let δ be any test family, binary or otherwise, such that $\delta(X, b_1) < \delta(X, b_2)$ for some $b_1 < b_2$ and all X belonging to some set A of positive measure on Q . Define a new test family $\hat{\delta}$ such that $\hat{\delta}(X, b_1) = \delta(X, b_2)$ for all $X \in A$. Otherwise $\hat{\delta}$ is the same as δ . Since $L_{b_1}(0, 1)\hat{\delta}(X, b_1) = L_{b_1}(0, 1)\delta(X, b_2) \leq L_{b_2}(0, 1)\delta(X, b_2)$, it follows that $\sup_b L_b(0, 1)\hat{\delta}(X, b) = \sup_b L_b(0, 1)\delta(X, b)$. This implies that $\hat{\delta}$ is type-I risk safe and we conclude that δ is Γ -inadmissible by Lemma 4.1.

A.6. Proof of Lemma 4.3

If $E_\delta(X) < E_\phi(X)$ then there exists some b such that $\delta(X, b) < \phi(X, b)$. If not, then $L_{b_0}(0, 1)\phi(X, b_0) \leq L_{b_0}(0, 1)\delta(X, b_0) \leq \sup_b L_b(0, 1)\delta(X, b) = E_\delta(X)$ for all b_0 , so taking the supremum over b_0 gives $E_\phi(X) \leq E_\delta(X)$, a contradiction. Now apply Lemma 4.1 to see that δ is Γ -inadmissible for any Γ . As for the second part of the result, suppose that $\mathbb{E}_P[E_\delta] < 1$. We will define a new test family $\hat{\delta}$ such that $E_{\hat{\delta}} > E_\delta$ in some region A while remaining type-I risk safe. First note that there exists some region $A \subset \mathcal{X}$ with $P(A) > 0$ and some b_0 such that $\delta(X, b_0)$ for all $X \in A$. If not, then $\delta(X, b) = 1$ for all $b \in \mathcal{B}$ P -almost surely. Since δ is type-I risk safe, we have $\mathbb{E}_P[\sup_b L_b(0, \delta(X, b))] = \mathbb{E}_P[\sup_b L_b(0, 1)] \leq 1$, hence $L_b(0, 1) \leq 1$ for all b , contradicting our assumption on the losses. For $X \in A$, set

$$\hat{\delta}(X, b_0) = \delta(X, b_0) + \frac{\epsilon}{L_{b_0}(0, 1)},$$

for some $\epsilon > 0$ to be determined. For $X \notin A$ or any $b \neq b_0$, let $\hat{\delta}(X, b) = \delta(X, b)$. In order to ensure that $\hat{\delta}(X, b_0) \leq 1$, we require that

$$\frac{\epsilon}{L_{b_0}(0, 1)} \leq 1 - \delta(X, b_0), \quad (\text{A.3})$$

for all $X \in A$. We can choose $\epsilon > 0$ small enough such that (A.3) is satisfied since $\delta(X, b_0) < 1$. If $\hat{\delta}$ is type-I risk safe then it is strongly preferable to δ for any Γ by Lemma 4.1, thus rendering δ Γ -inadmissible. To see that ϵ can be chosen to make $\hat{\delta}$ type-I risk safe, first write

$$\begin{aligned} \int_A \sup_b L_b(0, \hat{\delta}(x, b))P &\leq \int_A \sup_b L_b(0, 1) \left(\delta(x, b) + \frac{\epsilon}{L_b(0, 1)} \right)P \\ &= \int_A \sup_b L_b(0, \delta(x, b))P + \epsilon P(A). \end{aligned}$$

Therefore,

$$\begin{aligned} \mathcal{R}_P(\hat{\delta}) &= \int_A \sup_b L_b(0, \hat{\delta}(x, b))P + \int_{\mathcal{X} \setminus A} \sup_b L_b(0, \delta(x, b))P \\ &\leq \int_{\mathcal{X}} \sup_b L_b(0, \delta(x, b))P + \epsilon P(A) = \mathcal{R}_P(\delta) + \epsilon P(A). \end{aligned}$$

Since $\mathcal{R}_P(\delta) = \mathbb{E}_P[E_\delta] < 1$ by assumption, choosing any

$$0 < \epsilon \leq \frac{1 - \mathbb{E}_P[E_\delta]}{P(A)}, \quad (\text{A.4})$$

ensures that $\hat{\delta}$ remains type-I risk safe. Thus, choosing any ϵ that satisfies both (A.4) and (A.3) completes the proof when the test families are allowed to be randomized. Now let us consider the case of binary test families. Let φ be such that $\mathbb{E}_P[E_\varphi] < 1$. As before, this implies the existence of some $A \subset \mathcal{X}$ and b_0 such that $\varphi(X, b_0) = 0$. Otherwise φ rejects with probability 1 everywhere and on all losses, implying that $L_b(0, 1) \leq 1$ for all b . We define $\hat{\varphi}$ which acts as δ everywhere except that $\hat{\varphi}(X, b_0) = 1 > 0 = \varphi(X, b_0)$ for all X in some $A_0 \subset A$ with $Q(A_0) > 0$. In particular, we choose A_0 to satisfy

$$P(A_0) = \frac{1 - \mathbb{E}_P[E_\delta]}{L_{b_0}(0, 1)}, \quad (\text{A.5})$$

which is possible by continuity of P . Next, notice that by Lemma 4.2, $\delta(X, b) = 0$ for all $b \geq b_0$ and $X \in A_0$, otherwise δ is already inadmissible. Therefore, $\sup_b L_b(0, \hat{\varphi}(X, b)) = L_{b_0}(0, 1)$ in A_0 . This gives

$$\begin{aligned} \mathcal{R}_P(\hat{\delta}) &= \int_{\mathcal{X} \setminus A_0} \sup_b L_b(0, \varphi(x, b))P + \int_{A_0} \sup_b L_b(0, \hat{\varphi}(x, b))P \\ &= \int_{\mathcal{X} \setminus A_0} \sup_b L_b(0, \varphi(x, b))P + L_{b_0}(0, 1)P(A_0) \\ &\leq \mathcal{R}_P(\delta) + 1 - \mathbb{E}_P[E_\delta] = 1, \end{aligned}$$

which proves that $\hat{\varphi}$ is type-I risk safe. To complete the proof, apply Lemma 4.1 to see that $\hat{\delta}$ is strictly preferable to δ with respect to any Γ .

A.7. Proof of Lemma 4.4

Suppose there exists some A satisfying $P(A), Q(A) > 0$ such that for some $b \in \mathcal{B}$ and all $X \in A$, $\delta(X, b) > \min\{1, E_\delta(X)/L_b(0, 1)\}$. That is,

$$L_b(0, 1)\delta(X, b) > \min\{L_b(0, 1), E_\delta(X)\}, \text{ for all } X \in A. \quad (\text{A.6})$$

Observe that we must have $\min\{L_b(0, 1), E_\delta(X)\} = E_\delta(X)$. Otherwise (A.6) implies that $L_b(0, 1)\delta(X, b) > L_b(0, 1)$, which is impossible since $\delta(X, b) \leq 1$ by definition. Therefore the minimum is achieved by $E_\delta(X)$ and we have $L_b(0, 1)\delta(X, b) > E_\delta(X)$, contradicting the definition of E_δ . Thus no such b can exist, and we conclude that the left hand side of (13) is at most the right hand side.

Next, suppose there exists some b with $\delta(X, b) < \min\{1, E_\delta(x)/L_b(0, 1)\}$ for all $X \in A$, where again $P(A), Q(A) > 0$. Consider a new decision family $\hat{\delta}$ defined as $\hat{\delta}(X, b) = \min\{1, E_\delta(X)/L_b(0, 1)\}$ for all $X \in A$ and which acts as δ everywhere else. Note that for $X \in A$, $L_b(0, 1)\hat{\delta}(X, b) = \min\{L_b(0, 1), E_\delta(X)\} \leq E_\delta(X)$, hence $E_{\hat{\delta}}(X) = \sup_b L_b(0, 1)\hat{\delta}(X, b) = E_\delta(X)$ (using that $\hat{\delta}(\cdot, c) = \delta(\cdot, c)$ for all $c \neq b$). Therefore,

$$\mathcal{R}_P(\hat{\delta}) = \int_{\mathcal{X} \setminus A} E_\delta(x)\mathbb{P} + \int_A E_{\hat{\delta}}(x)\mathbb{P} = \mathcal{R}_P(\delta) \leq 1, \quad (\text{A.7})$$

so $\hat{\delta}$ remains type-I risk safe. Moreover, it strictly improves type-II risk by Lemma 4.1 ($\hat{\delta}(\cdot, b) > \delta(\cdot, b)$ on a set of positive measure under Q), implying that δ is Γ -inadmissible.

A.8. Proof of Lemma 4.5

Suppose there exists a set $A \subset \mathcal{X}$ such that $E(X) > \sup_b L_b(0, 1)$ for all $X \in A$ and $P(A) > 0$. We will define a new e-variable F for P such that $F(X) \leq \sup_b L_b(0, 1)$ for all $X \in A$ and δ_F is strictly preferable to δ with respect to any Γ . First we observe that there exists some $Y \subset \mathcal{X}$ such that $E(X) < \sup_b L_b(0, 1)$ for all $X \in Y$ and $P(Y) > 0$. If not, then $E(X) \geq \sup_b L_b(0, 1)$ P -almost surely, implying by definition (see (14)) that δ_E rejects with probability 1 P -almost surely, contradicting our assumption that $L_b(0, 1) > 1$ for some b . For some fixed $\epsilon > 0$ to be specified later, define F as follows:

$$F(X) = \begin{cases} \sup_b L_b(0, 1), & \text{if } X \in A, \\ E(X) + \epsilon, & \text{if } X \in Y, \\ E(X), & \text{otherwise.} \end{cases} \quad (\text{A.8})$$

First notice that on Y there exists some b such that $\delta_F(X, b) = \min\{1, E(X)/L_b(0, 1)\} > E(X)/L_b(0, 1) = \delta_E(X)$. Moreover, on A , $\delta_F(X) = 1 = \delta_E(X, b)$ for all b since $E(X) > F(X) = \sup_b L_b(0, 1)$. Finally, everywhere else we have $\delta_E(X, b) = \delta_F(X, b)$ for all b since $F(X) = E(X)$. Therefore, it follows from Lemma 4.1 that δ_E is Γ -inadmissible if δ_F is type-I risk safe. Compute

$$\begin{aligned} \mathcal{R}_P(\delta_F) &= \int_A F\mathbb{P} + \int_Y F\mathbb{P} + \int_{\mathcal{X} \setminus (A \cup Y)} E\mathbb{P} \\ &= \int_A (\sup_b L_b(0, 1) - E + E)\mathbb{P} + \int_E F\mathbb{P} + \epsilon P(Y) + \int_{\mathcal{X} \setminus (A \cup Y)} E\mathbb{P} \\ &= \int_{\mathcal{X}} E\mathbb{P} + \int_A (\sup_b L_b(0, 1) - E)\mathbb{P} + \epsilon P(Y). \end{aligned}$$

Taking

$$\epsilon = \frac{1}{P(Y)} \int_A (\sup_b L_b(0, 1) - E)\mathbb{P}, \quad (\text{A.9})$$

thus suffices to ensure that $\mathcal{R}_P(\delta_F) = \mathcal{R}_P(\delta_E) \leq 1$. We have thus proven that if $P(E(X) > \sup_b L_b(0, 1)) > 0$ then δ is Γ -inadmissible. Then, if $P(E(X) > \sup_b L_b(0, 1)) = 0$, $E_{\delta_E} = E$ P -almost surely by (15).

A.9. Proof of Lemma 4.6

The proof is similar to that of Lemma 4.4. First note that we can't have $\varphi(X, b) = 1$ if $L_b(0, 1) > E_\varphi(X)$ by definition of E_φ . Therefore, $\varphi(X, b) \leq 1 \{L_b(0, 1) \leq E_\varphi(X)\}$. Conversely, suppose that, for some b^* , $\varphi(X, b^*) = 0$ when $L_{b^*}(0, 1) < E_\varphi(X)$ for all $X \in A$ where $P(A), Q(A) > 0$. Define a new rule $\hat{\varphi}$ such that $\hat{\varphi}(X, b^*) = 1$ for $X \in A$ and otherwise $\hat{\varphi}$ is the same as φ . Since $E_\varphi = E_{\hat{\varphi}}$, $\hat{\varphi}$ is type-I risk safe. Now apply Lemma 4.1 with $\phi = \hat{\varphi}$ to see that φ is Γ -inadmissible.

A.10. Proof of Lemma 4.7

We prove the contrapositive. Suppose that E is not compatible with \mathcal{L}_0 . Then there exists some set $Y \subset \mathcal{X}$ with $P(Y) > 0$ such that, for all $X \in Y$, either

1. $L_{b_1}(0, 1) < E(X) < L_{b_2}(0, 1)$ for some $b_1 < b_2$ such that there is no $b \in \mathcal{B}$ with $b_1 < b < b_2$, or
2. $E(X) > \sup_b L_b(0, 1)$. Of course, in this case $\sup_b L_b(0, 1) < \infty$.

We consider these two cases separately. First suppose that (1) holds. For some *disjoint* subsets $Y_0, Y_1 \subset Y$ to be determined, we define a new random variable as follows:

$$F(X) = \begin{cases} L_{b_1}(0, 1), & \text{if } X \in Y_0, \\ L_{b_2}(0, 1), & \text{if } X \in Y_1, \\ E(X), & \text{otherwise.} \end{cases} \quad (\text{A.10})$$

Notice that for all $X \in Y_0$ and all $b \in \mathcal{B}$, $\varphi_F(X, b) = \varphi_E(X, b)$; the fact that $E(X) > L_{b_1}(0, 1)$ makes no difference since it is strictly smaller than the next loss, $L_{b_2}(0, 1)$ and thus cannot influence the value of φ_E . Meanwhile, $\varphi_F(X, b_2) = 1 > 0 = \varphi_E(X, b_2)$ for $X \in Y_1$, and $\varphi_F(X, b) = \varphi_E(X, b)$ for all other b . For $X \notin Y_0 \cup Y_1$, clearly $\varphi_E(X, b) = \varphi_F(X, b)$. Therefore, as long as Y_1 has positive measure under P and φ_E is type-I risk safe, then φ_E is Γ -inadmissible by [Lemma 4.1](#). Let us compute

$$\begin{aligned} \mathcal{R}_P(\delta_F) &= \int_{Y_0} F \mathbb{P} + \int_{Y_1} F \mathbb{P} + \int_{\mathcal{X} \setminus (Y_0 \cup Y_1)} F \mathbb{P} \\ &= \int_{Y_0} (F - E + E) \mathbb{P} + \int_{Y_1} (F - E + E) \mathbb{P} + \int_{\mathcal{X} \setminus (Y_0 \cup Y_1)} E \mathbb{P} \\ &= \int_{\mathcal{X}} E \mathbb{P} + \int_{Y_0} (L_{b_1}(0, 1) - E) \mathbb{P} + \int_{Y_1} (L_{b_2}(0, 1) - E) \mathbb{P} \\ &\leq 1 + \int_{Y_0} (L_{b_1}(0, 1) - E) \mathbb{P} + (L_{b_2}(0, 1) - L_{b_1}(0, 1)) P(Y_1). \end{aligned}$$

Thus, for a fixed Y_0 with positive measure, taking Y_1 small enough such that

$$0 < P(Y_1) = (L_{b_2}(0, 1) - L_{b_1}(0, 1))^{-1} \int_{Y_0} (E - L_{b_1}(0, 1)) \mathbb{P}, \quad (\text{A.11})$$

ensures that $\mathcal{R}_P(\delta_F) \leq 1$. Such choices of Y_0 and Y_1 are possible by continuity of P . This completes the proof of case (1).

Let us now consider case (2). Here the proof is similar to that of [Lemma 4.5](#), but it must be handled somewhat differently to account for binary tests. Since $E(X) > \sup_b L_b(0, 1)$ on Y , there exists some U such that $E(X) < \sup_b L_b(0, 1)$ for all $X \in U$. The reasoning is the same as in previous proofs.

For some $U_0 \subset U$, define

$$F(X) = \begin{cases} \sup_b L_b(0, 1), & \text{if } X \in Y \cup U_0, \\ E(X), & \text{otherwise.} \end{cases} \quad (\text{A.12})$$

Following similar reasoning as above, $\varphi_F(X, b) \geq \varphi_E(X, b)$ with strict inequality for some b when $x \in U_0$. We thus need only show that φ_F is type-I risk safe to complete the proof. Let $s = \sup_b L_b(0, 1)$ and write

$$\begin{aligned} \mathcal{R}_P(\varphi_F) &= \int_Y (s - E + E) \mathbb{P} + \int_{U_0} (s - E + E) \mathbb{P} + \int_{\mathcal{X} \setminus (Y \cup U_0)} E \mathbb{P} \\ &\leq \int_{\mathcal{X}} E \mathbb{P} + \int_Y (s - E) \mathbb{P} + \int_{U_0} s \mathbb{P} \\ &\leq 1 + \int_Y (s - E) \mathbb{P} + s P(U_0). \end{aligned}$$

Since $s < \infty$ and P is continuous, we can choose U_0 such that $s < P(U_0) = \int_Y (E - s) \mathbb{P}$, which gives $\mathcal{R}_P(\varphi_F)$, completing the argument.

A.11. Proof of [Theorem 5.1](#)

[Lemma 4.3](#) implies that if E_δ is not sharp then δ is Γ_{all} -inadmissible. Meanwhile, [Lemma 4.4](#) implies that if δ is not canonical then it is not Γ_{all} -admissible. This proves one direction. To prove the other direction, suppose that δ is canonical and E_δ is sharp. Let ϕ be any other type-I risk safe test family. We may assume that ϕ is canonical and E_ϕ sharp, otherwise we may improve it and consider the resulting test family.

Define the following subsets of \mathcal{X} . Let

1. $R_1 \subset \mathcal{X}$ be the region where $E_\phi(X) = E_\delta(X) = 0$,
2. $R_2 \subset \mathcal{X}$ be the region where $0 = E_\delta(X) < E_\phi(X)$,
3. $R_3 \subset \mathcal{X}$ be the region where $0 < E_\delta(X) \leq E_\phi(X)$,
4. and $R_4 \subset \mathcal{X}$ be the region where $E_\delta(X) > E_\phi(X)$.

Observe that R_4 cannot have Q -measure zero. Indeed, if so, then $E_\delta(X) \leq E_\phi(X)$ Q -almost surely. Since both E_δ and E_ϕ have the same expected value and both are nonnegative, this implies that $E_\phi(X) = E_\delta(X)$ Q -almost surely, implying by [Lemma 4.4](#) that $\delta = \phi$ Q -almost surely.

It suffices to find a map B such that $\mathbb{E}_Q[L_B(1, \phi(X, B))] > \mathbb{E}_Q[L_B(1, \delta(X, B))]$, thus demonstrating that ϕ cannot be strongly Γ_{all} -preferable to δ . To this end, write

$$\mathbb{E}_Q[L_B(1, \phi(X, B)) - L_B(1, \delta(X, B))] = \sum_{1 \leq k \leq 4} \Delta_k, \quad (\text{A.13})$$

where

$$\begin{aligned} \Delta_k &\equiv \int_{R_k} L_B(1, \phi(x, B)) - L_B(1, \delta(x, B)) \mathbb{Q} \\ &= \int_{R_k} L_B(1, 0)(\delta(x, B) - \phi(x, B)) \mathbb{Q}. \end{aligned}$$

Define the adversary B as follows. First consider R_4 which, as we argued above, has positive measure under Q . For $X \in R_4$, we set $B(X)$ to be any b such that $L_b(0, 1) > E_\phi(X)$. Such a b must exist otherwise, $E_\delta(X) > E_\phi(X) \geq \sup_b L_b(0, 1)$, contradicting the definition of $E_\delta(X)$. Therefore, on R_4 , ϕ satisfies $\phi(X, B(X)) = E_\phi(X)/L_{B(X)}(0, 1) < 1$, by Lemma 4.4, so

$$\delta(X, B(X)) = \min \left\{ 1, \frac{E_\delta(X)}{L_{B(X)}(0, 1)} \right\} > \frac{E_\phi(X)}{L_{B(X)}(0, 1)} = \phi(X, B(X)), \quad Q\text{-almost surely.}$$

Therefore,

$$\Delta_4 = \int_{R_4} L_B(1, 0) \left(\frac{E_\delta(X)}{L_B(0, 1)} - \frac{E_\phi(X)}{L_B(0, 1)} \right) \mathbb{Q} > 0,$$

Now let us choose how B acts on the rest of \mathcal{X} . For $X \in R_1$, the choice is arbitrary, since $\phi(X, b) = \delta(X, b) = 0$ so $\Delta_1 = 0$. For R_2 , write

$$\Delta_2 = - \int_{R_2} L_B(1, 0) \min \left\{ 1, \frac{E_\phi(X)}{L_B(0, 1)} \right\} \mathbb{Q}.$$

For all $X \in R_2$, take $B(X) = b$ where b is large enough such that

$$\frac{L_b(1, 0)}{L_b(0, 1)} \leq \frac{\Delta_4}{4 \int_{R_2} E_\phi(X) \mathbb{Q}} \quad \text{and} \quad L_b(0, 1) \geq \sup_{X \in R_2} E_\phi(X). \quad (\text{A.14})$$

Such a b must exist by condition (C1) and by definition of E_ϕ . In this case,

$$\Delta_2 \geq - \int_{R_2} \frac{L_b(1, 0)}{L_b(0, 1)} E_\phi(X) \mathbb{Q} \geq - \frac{\Delta_4}{4}.$$

For R_3 we make a similar argument. Write

$$\Delta_3 = - \int_{R_3} L_B(1, 0)(\phi(x, B) - \delta(x, B)) \mathbb{Q} \geq - \int_{R_3} L_B(1, 0) \min \left\{ 1, \frac{E_\phi(X)}{L_B(0, 1)} \right\} \mathbb{Q}.$$

Take $B(X) = b$ where b defined as in (A.14) but with R_3 in place of R_2 . Then $\Delta_3 \geq -\Delta_4/4$. Therefore, overall we have $\sum_{k \leq 4} \Delta_k \geq -\Delta_4/4 - \Delta_4/4 + \Delta_4 > 0$ since $\Delta_4 > 0$. This completes the first part of the argument.

For the second part, let E be an e-variable for P . If $E \leq \sup_b L_b(0, 1)$ P -almost surely, then $E_{\delta_E} = E$ P -almost surely by (15). Thus if E is sharp then E_{δ_E} is sharp, and we may apply the first part of the theorem to see that E_{δ_E} is Γ_{all} -admissible. Conversely, if δ_E is Γ_{all} -admissible then E_{δ_E} is sharp, again by applying the first part of the theorem. This in turn implies that E is sharp, since $E_{\delta_E}(X) \leq E(X)$ by (15). Next we claim that $E = E_{\delta_E}$, which will complete the proof since $E_{\delta_E}(X) \leq \sup_b L_b(0, 1)$ for all X by construction. Let $W = \{X : E(X) \neq E_{\delta_E}(X)\}$ be the subset on which E and E_{δ_E} disagree. By (15), we know they disagree if and only if $E(X) > \sup_b L_b(0, 1)$. Therefore, we can write

$$W = \{X : E(X) > \sup_b L_b(0, 1)\}.$$

Since E is sharp,

$$\begin{aligned} 1 &= \mathbb{E}_P[E] = \mathbb{E}_P[E \mathbf{1}\{X \notin W\}] + \mathbb{E}_P[E \mathbf{1}\{X \in W\}] \\ &= \mathbb{E}_P[E_{\delta_E} \mathbf{1}\{X \notin W\}] + \mathbb{E}_P[E \mathbf{1}\{X \in W\}] \\ &> \mathbb{E}_P[E_{\delta_E} \mathbf{1}\{X \notin W\}] + \sup_b L_b(0, 1)P(W). \end{aligned}$$

Therefore,

$$\begin{aligned} 1 &= \mathbb{E}_P[E_{\delta_E}] = \mathbb{E}_P[E \mathbf{1}\{X \notin W\}] + \mathbb{E}_P[E_{\delta_E} \mathbf{1}\{X \in W\}] \\ &< 1 - \sup_b L_b(0, 1)P(W) \mathbb{E}_P[E_{\delta_E} \mathbf{1}\{X \in W\}] \\ &\leq 1 - \sup_b L_b(0, 1)P(W) \sup_b L_b(0, 1)P(W) = 1, \end{aligned}$$

a contradiction. Hence we must have $P(W) = 0$, meaning $E = E_{\delta_E}$ P -almost surely. This completes the proof.

A.12. Proof of Theorem 5.2

The proof is similar to that of [Theorem 5.1](#), so we provide only a sketch. [Lemma 4.3](#) implies that E_φ must be sharp; [Lemma 4.6](#) implies that φ must be canonical. For the converse, we partition the sample space the same way as in the proof of [Theorem 5.1](#). On R_4 we have $E_\delta(X) > E_\phi(X)$. Because the losses are dense in $\mathfrak{R} \cap [M, \infty)$, we can find some b such that $E_\phi(X) < L_b(0, 1) \leq E_\delta(X)$. (Note that $E_\delta(X) \geq \inf b L_b(0, 1) \geq M$ for binary δ .) We take $B(X)$ to be this value of b . Then $\Delta_4 > 0$. For R_1 , the choice of B is again arbitrary. For R_2 , choose $B(X)$ such that $L_{B(X)}(0, 1) > E_\phi(X)$ which is possible $\sup_b L_b(0, 1) = \infty$ (E_ϕ cannot be infinite on a set of positive measure under P otherwise ϕ is not type-I risk safe). Then $\Delta_2 = 0$. A similar argument applies to R_3 . This shows that $\mathbb{E}_Q[L_B(1, \phi(X, B))] - \mathbb{E}_Q[L_B(1, \delta(X, B))] = \sum_i \Delta_i > 0$.

Now consider an e-variable E for P . If E is not sharp then neither is E_{δ_E} since $E_{\delta_E}(X) = \sup_b L_b(0, 1) \mathbf{1}\{E(X) \geq L_b(0, 1)\} \leq E(X)$. Therefore δ_E is Γ_{all} -inadmissible by the first part of the theorem. Further, if E is not compatible then δ_E is Γ_{all} -inadmissible by [Lemma 4.7](#). Now we prove the other direction. If E is compatible, then $E_{\delta_E} = E$ P -almost surely. Therefore if E is sharp so is E_{δ_E} so δ_E is Γ_{all} -admissible, again by applying the first part of the theorem.

A.13. Proof of Lemma 5.1

Suppose not, so there exists A_1, A_2 and b such that

$$\inf_{X_1 \in A_1} E_\delta(X_1) > \sup_{X_2 \in A_2} E_\delta(X_2), \quad \text{and} \quad \sup_{X_1 \in A_1} \Lambda(X_1) < \sup_{X_2 \in A_2} \Lambda(X_2).$$

For convenience, define

$$\gamma_1 = \inf_{X_1 \in A_1} E_\delta(X_1), \quad \gamma_2 = \sup_{X_2 \in A_2} E_\delta(X_2).$$

We will define a new decision rule $\hat{\delta}$ which is strictly preferable to δ with respect to Γ_{const} . Pick any $\Delta \in \mathfrak{R}$ such that $0 < \Delta < \gamma_1 - \gamma_2$, which exists by the density of the reals. Define

$$\mathcal{B}_0 \equiv \{b : L_b(0, 1) \leq \gamma_2\}$$

$$\mathcal{B}_1 \equiv \{b : \gamma_2 < L_b(0, 1) \leq \gamma_2 + \Delta\}$$

$$\mathcal{B}_2 \equiv \mathcal{B} \setminus (\mathcal{B}_0 \cup \mathcal{B}_1).$$

Notice that $\mathcal{B}_0, \mathcal{B}_1$, and \mathcal{B}_2 are all disjoint and form a partition of \mathcal{B} . It may be helpful to draw a plot with $L_b(0, 1)$ on the y-axis and $\Lambda(X)$ on the x-axis. \mathcal{B}_0 are those b at the bottom, \mathcal{B}_2 those at the top. Further notice that \mathcal{B}_2 is non-empty. Otherwise

$$\sup_{b \in \mathcal{B}} L_b(0, 1) \leq \gamma_2 + \Delta < \gamma_1 \leq \sup_{X \in A_1} \sup_{b \in \mathcal{B}} L_b(0, 1) \delta(X, b) \leq \sup_{b \in \mathcal{B}} L_b(0, 1),$$

a contradiction. Finally, note that since $L_b(0, 1)$ is increasing in b , if $b_1 \notin \mathcal{B}_0$ and $b_2 > b_1$, then $b_2 \notin \mathcal{B}_0$. Now, for all $b \in \mathcal{B}_0 \cup \mathcal{B}_1$, we let $\hat{\delta}(\cdot, b) = \delta(\cdot, b)$. For all $b \in \mathcal{B}_2$, we set

$$\hat{\delta}(X_1, b) = \delta(X_1, b) - \frac{\epsilon_1}{L_b(0, 1)}, \quad \text{and} \quad \hat{\delta}(X_2, b) = \delta(X_2, b) + \frac{\epsilon_2}{L_b(0, 1)}, \quad (\text{A.15})$$

for all $X_1 \in A_1$ and $X_2 \in A_2$. Outside of A_1 and A_2 , $\hat{\delta}$ will be defined as δ . Throughout this proof, to save ourselves from constantly writing quantifiers, we will always assume that $X_1 \in A_1$ and $X_2 \in A_2$.

We must first ensure that $\hat{\delta}$ is well-defined, i.e., $\hat{\delta}(X_1, b) \geq 0$ and $\hat{\delta}(X_2, b) \leq 1$ (it is clear that $\hat{\delta}(X_1, b) \leq \delta(X_1, b) \leq 1$ and $\hat{\delta}(A_2, b) \geq \delta(A_2, b) \geq 0$). By [Lemma 4.4](#), we may assume that $\delta(X_1, b) = \min\{1, E_\delta(X_1)/L_b(0, 1)\}$ and $\delta(X_2, b) = \min\{1, E_\delta(X_2)/L_b(0, 1)\}$. Therefore, to ensure that $\hat{\delta}(X_1, b) = \min\{1, E_\delta(X_1)/L_b(0, 1)\} - \epsilon_1/L_b(0, 1) \geq 0$, it suffices that (a) $\epsilon_1 \leq L_b(0, 1)$ and (b) $\epsilon_1 \leq \inf_{X_1 \in A_1} E_\delta(X_1) = \gamma_1$ (consider the two cases in the minimum). Since $L_b(0, 1)$ is increasing in b and $L_b(0, 1) > \gamma_2 + \Delta$ for all $b \in \mathcal{B}_2$, we can ensure that (a) and (b) are both met by taking

$$\epsilon_1 \leq \gamma_2 + \Delta < \gamma_1. \quad (\text{A.16})$$

As for $\hat{\delta}$ on A_2 , notice that $\hat{\delta}(X_2, b) = (E_\delta(X_2) + \epsilon_2)/L_b(0, 1)$ for all $b \in \mathcal{B}_2$, which is at most 1 if $E_\delta(X_2) + \epsilon_2 \leq L_b(0, 1)$. Since $\gamma_2 + \Delta \leq L_b(0, 1)$ for $b \in \mathcal{B}_2$ and $E_\delta(X_2) \leq \gamma_2$, it suffices to take $\epsilon_2 \leq \gamma_2 + \Delta - \gamma_2$, i.e.,

$$0 < \epsilon_2 \leq \Delta. \quad (\text{A.17})$$

We have shown that there exist $\epsilon_1, \epsilon_2 > 0$ such that $\hat{\delta}$ is well-defined. Next we show that they can be chosen to ensure that $\hat{\delta}$ remains type-I risk safe and is Γ_{const} -preferable to δ . We begin with type-I safety. Since $\hat{\delta}$ changes only on A_1 and A_2 , it suffices to show that

$$\int_{A_1 \cup A_2} \sup_b L_b(0, \hat{\delta}(x, b)) \mathbb{P} \leq \int_{A_1 \cup A_2} \sup_b L_b(0, \delta(x, b)) \mathbb{P}. \quad (\text{A.18})$$

For all $b \in \mathcal{B}_0 \cup \mathcal{B}_1$, we have $\delta(X_1, b) = \min\{1, E_\delta(X_1)/L_b(0, 1)\} = 1$ since $L_b(0, 1) \leq \gamma_2 + \Delta < \gamma_1 \leq E_\delta(X_1)$. Moreover, $\hat{\delta}(X_1, b) = \delta(X_1, b)$ for such b . For $b \in \mathcal{B}_2$ meanwhile, $\hat{\delta}(X_1, b) = \min\{1, E_\delta(X_1)/L_b(0, 1)\} - \epsilon_1/L_b(0, 1) \leq (E_\delta(X_1) - \epsilon_1)/L_b(0, 1)$ so

$$\sup_{b \in \mathcal{B}_2} L_b(0, \hat{\delta}(X_1, b)) \leq \sup_{b \in \mathcal{B}_2} L_b(0, 1) \left(\frac{E_\delta(X_1) - \epsilon_1}{L_b(0, 1)} \right) = E_\delta(X_1) - \epsilon_1,$$

and

$$\begin{aligned} \sup_{b \in \mathcal{B}} L_b(0, \hat{\delta}(X_1, b)) &= \max \left\{ \sup_{b \in \mathcal{B}_0 \cup \mathcal{B}_1} L_b(0, \delta(X_1, b)), \sup_{b \in \mathcal{B}_2} L_b(0, \hat{\delta}(X_1, b)) \right\} \\ &\leq \max \left\{ \sup_{b \in \mathcal{B}_0 \cup \mathcal{B}_1} L_b(0, 1), E_\delta(X_1) - \epsilon_1 \right\} \\ &\leq \max \{ \gamma_2 + \Delta, E_\delta(X_1) - \epsilon_1 \}. \end{aligned}$$

We want this final quantity to equal $E_\delta(X_1) - \epsilon_1$, so we choose ϵ_1 such that

$$0 < \epsilon_1 \leq \gamma_1 - (\gamma_2 + \Delta) \leq E_\delta(X_1) - (\gamma_2 + \Delta), \quad (\text{A.19})$$

which is possible by our choice of Δ . This ensures the maximum in the above display is achieved by $E_\delta(X_1) - \epsilon_1$ and we have

$$\sup_b L_b(0, \hat{\delta}(X_1, b)) \leq E_\delta(X_1) - \epsilon_1. \quad (\text{A.20})$$

Now let us consider A_2 . For $b \in \mathcal{B}_0$ we have $\hat{\delta}(X_2, b) = \delta(X_2, b) \leq 1$ and for $b \in \mathcal{B}_1$, $\hat{\delta}(X_2, b) = \delta(X_2, b) = \min\{1, \gamma_2 / L_b(0, 1)\} = E_\delta(X_2) / L_b(0, 1)$. For $b \in \mathcal{B}_2$, $\hat{\delta}(X_2, b) = \delta(X_2, b) = E_\delta(X_2) / L_b(0, 1) + \epsilon_2 / L_b(0, 1)$, so

$$\begin{aligned} \sup_{b \in \mathcal{B}} L_b(0, \hat{\delta}(X_2, b)) &\leq \max \left\{ \sup_{b \in \mathcal{B}_0} L_b(0, 1), \sup_{b \in \mathcal{B}_1} L_b(0, \hat{\delta}(X_2, b)), \sup_{b \in \mathcal{B}_2} L_b(0, \hat{\delta}(X_2, b)) \right\} \\ &= \max \left\{ \sup_{b \in \mathcal{B}_0} L_b(0, 1), \sup_{b \in \mathcal{B}_1} E_\delta(X_2), \sup_{b \in \mathcal{B}_2} E_\delta(X_2) + \epsilon_2 \right\} \\ &\leq \max \{ \gamma_2, \gamma_2, \gamma_2 + \epsilon_2 \} = \gamma_2 + \epsilon_2. \end{aligned}$$

Combining this and (A.20), the left hand side of (A.18) can be upper bounded as

$$\begin{aligned} \int_{A_1 \cup A_2} \sup_{b \in \mathcal{B}} L_b(0, \hat{\delta}(x, b)) \mathbb{P} &\leq \int_{A_1} E_\delta(x) \mathbb{P} - \epsilon_1 P(A_1) + \int_{A_2} E_\delta(x) \mathbb{P} - \epsilon_2 P(A_2) \\ &= \int_{A_1 \cup A_2} \sup_b L_b(0, \delta(x, b)) \mathbb{P} - \epsilon_1 P(A_1) + \epsilon_2 P(A_2). \end{aligned}$$

For (A.18) to hold we thus require that $-\epsilon_1 P(A_1) + \epsilon_2 P(A_2) \leq 0$, i.e.,

$$\frac{\epsilon_2}{\epsilon_1} \leq \frac{P(A_1)}{P(A_2)}. \quad (\text{A.21})$$

Still leaving the precise choice of ϵ_1 and ϵ_2 unspecified for the moment, let us move on to type-II risk. We claim that for all $b \in \mathcal{B}_2$, $\mathbb{E}_Q[\hat{\delta}(X, b)] \geq \mathbb{E}_Q[\delta(X, b)]$. Since type-II risk remains unchanged for $b \in \mathcal{B}_0 \cup \mathcal{B}_1$ (since $\hat{\delta}(\cdot, b) = \delta(\cdot, b)$ for such b), this will show that $\hat{\delta}$ is strictly preferable to δ with respect to Γ_{const} .

As with type-I risk, because $\hat{\delta}$ is the same as δ outside of A , it suffices to focus on A_1 and A_2 and show that

$$\mathbb{E}_Q[\hat{\delta}(X, b) \mathbf{1}\{X \in A\}] \geq \mathbb{E}_Q[\delta(X, b) \mathbf{1}\{X \in A\}], \quad (\text{A.22})$$

for all $b \in \mathcal{B}_2$. For these b , write

$$\int_{A_1} \hat{\delta}(x, b) \mathbb{Q} = \int_{A_1} \left(\delta(x, b) - \frac{\epsilon_1}{L_b(0, 1)} \right) \mathbb{Q} = \int_{A_1} \delta(x, b) \mathbb{Q} - \frac{\epsilon_1}{L_b(0, 1)} Q(A_1).$$

Similarly,

$$\int_{A_2} \hat{\delta}(x, b) \mathbb{Q} = \int_{A_2} \delta(x, b) \mathbb{Q} + \frac{\epsilon_2}{L_b(0, 1)} Q(A_2).$$

Combining these two displays, we see that (A.22) holds iff

$$\frac{\epsilon_2}{\epsilon_1} > \frac{Q(A_1)}{Q(A_2)}. \quad (\text{A.23})$$

We choose $\epsilon_1 = wP(A_1)$ and $\epsilon_2 = wP(A_2)$, where $w > 0$ is small enough to ensure that conditions (A.16), (A.17), and (A.19) are met (note that none of these conditions are in conflict with one another; all are upper bounds on ϵ_1 and ϵ_2). Then,

$$\frac{\epsilon_2}{\epsilon_1} = \frac{P(A_1)}{P(A_2)} > \frac{Q(A_1)}{Q(A_2)},$$

where the final inequality follows from Lemma A.1. Inequalities (A.21) and (A.23) are thus met, demonstrating that $\hat{\delta}$ is strongly preferable to δ with respect to Γ_{const} . This contradicts that δ is Γ_{const} -admissible, so we conclude that $E_\delta(X)$ is indeed increasing in the likelihood ratio P -almost surely. The second part of the lemma follows by applying Lemma A.2.

A.14. Proof of Lemma 5.2

Suppose not. By Lemma A.3 we have that $E_\delta(X) \geq \inf_{b \in \mathcal{B}} L_b(0, 1)$ Q -almost surely. Moreover, we cannot have $E_\delta(X) > \sup_{b \in \mathcal{B}} L_b(0, 1)$ by definition. Therefore, there exists some A with $Q(A) > 0$ and some partition of $\mathcal{B} = \mathcal{B}_0 \cup \mathcal{B}_1$ such that

$$\sup_{b \in \mathcal{B}_0} L_b(0, 1) < E_\delta(X) < \inf_{b \in \mathcal{B}_1} L_b(0, 1).$$

For simplicity we will assume that there exist b_1, b_2 such that $b_1 = \sup_{b \in \mathcal{B}_0} L_b(0, 1)$ and $b_2 = \inf_{b \in \mathcal{B}_1} L_b(0, 1)$, but the proof can be amended if these supremum and infimum are not in \mathcal{B} . We can thus write $\mathcal{B}_1 = \{b \geq b_2\}$ since the losses are assumed to be increasing in b .

Let $A = A_1 \cup A_2$ where $P(A_1) = P(A_2)$ and $\Lambda(X_1) < \Lambda(X_2)$ for all $X_1 \in A_1$ and $X_2 \in A_2$ (this is possible by continuity). Define a new decision rule $\hat{\delta}$ such that, for all $b \in \mathcal{B}_0$, $\hat{\delta}(X, b) = \delta(X, b) - \epsilon/L_b(0, 1)$ for all $X \in A_1$ and $\hat{\delta}(X, b) = \delta(X, b) + \epsilon/L_b(0, 1)$ for all $X \in A_2$. Otherwise $\hat{\delta}$ is the same as δ . We have $0 < \delta(X, b) < 1$ for all $X \in A$ and $b \in \mathcal{B}_1$ (the lower bound of 0 is implied by Lemma 4.4). Therefore, we may find ϵ small enough such that $0 < \hat{\delta}(X, b) \leq 1$ for all $X \in A$. In particular, let us ensure that ϵ is small enough such that

$$L_{b_1}(0, 1) + \epsilon < E_\delta(X). \quad (\text{A.24})$$

To see that $\hat{\delta}$ is type-I risk safe, observe that for $X \in A_1$,

$$\begin{aligned} E_{\hat{\delta}}(X) &= \sup_b L_b(0, \hat{\delta}(X, b)) \\ &= \max \left\{ \sup_{b < b_2} L_b(0, \delta(X, b)), \sup_{b \geq b_2} L_b(0, \hat{\delta}(X, b)) \right\} \\ &= \max \left\{ L_{b_1}(0, 1), \left(\sup_{b \geq b_2} L_b(0, 1) \delta(X, b) \right) - \epsilon \right\} \\ &= \max \left\{ L_{b_1}(0, 1), E_\delta(X) - \epsilon \right\} \\ &= E_\delta(X) - \epsilon, \end{aligned}$$

where the final equality follows by (A.24). Similarly, for $X \in A_2$,

$$\begin{aligned} E_{\hat{\delta}}(X) &= \max \left\{ L_{b_1}(0, 1), \left(\sup_{b \geq b_2} L_b(0, 1) \delta(X, b) \right) + \epsilon \right\} \\ &= \max \left\{ L_{b_1}(0, 1), E_\delta(X) + \epsilon \right\} \\ &= E_\delta(X) + \epsilon. \end{aligned}$$

Therefore,

$$\begin{aligned} \int_A E_{\hat{\delta}}(x) \mathbb{P} &= \int_{A_1} (E_\delta(x) - \epsilon) \mathbb{P} + \int_{A_2} (E_\delta(x) + \epsilon) \mathbb{P} \\ &= \int_A E_\delta(x) - \epsilon P(A_1) + \epsilon P(A_2) = \int_A E_\delta(x) \mathbb{P}, \end{aligned}$$

since $P(A_1) = P(A_2)$ by assumption. From this it follows that $\mathcal{R}_P(\hat{\delta}) \leq 1$ since

$$\mathbb{E}_P[E_{\hat{\delta}}(X)] = \int_{A'} E_{\hat{\delta}}(x) \mathbb{P} + \int_A E_\delta(x) \mathbb{P} = \int_{A'} E_\delta(x) \mathbb{P} \leq 1.$$

So much for type-I risk. For type-II risk, the argument is nearly identical to that in Lemma A.3, so we omit it here.

A.15. Proof of Lemma 5.3

Throughout the proof we drop the superscript “bin” on δ . By Lemma A.2 $\delta(X, b)$ is increasing in Λ iff E_δ is increasing in the Λ , so we focus on the test family itself.

Suppose that for some $b^* \in \mathcal{B}$, $\delta(\cdot, b^*)$ is not an increasing function of Λ . Then we can find two sets, $A_1, A_2 \subset \mathcal{X}$ with positive measure under both P and Q such that, for all $X_1 \in A_1$ and $X_2 \in A_2$,

$$\delta(X_1, b^*) = 1, \quad \delta(X_2, b^*) = 0, \quad \text{and} \quad \sup_{X \in A_1} \Lambda(X_1) < \inf_{X_2 \in A_2} \Lambda(X_2) \quad (\text{A.25})$$

Define a set $G = [t_-, t_+] \subset \mathcal{B} \cup \{\inf(\mathcal{B})\} \cup \{\sup(\mathcal{B})\}$ as follows.⁴ Suppose we can find some $A'_1 \subset A_1$ and $t > b^*$ such that $\delta(X_1, b) = 0$ for all $b > t$ and all $X_1 \in A'_1$. Then take $t^+ = t$ and redefine $A_1 \leftarrow A'_1$. If no such A'_1 and t exist, then $\delta(X_1, b) = 1$ for all $X_1 \in A_1$ all

⁴ Here $\sup(\mathcal{B}) = \sup_{b \in \mathcal{B}} b$. Similarly for $\inf(\mathcal{B})$.

$b \geq b^*$, since δ is nonincreasing in b by [Lemma 4.2](#) (if $\delta(X, b) = 0$ for some $b \geq b^*$ then $\delta(X, b') = 0$ for all $b' \geq b$). In this case we set $t^+ = \sup(\mathcal{B})$.

As for t_- , suppose we can find some $A'_2 \subset A_2$ and some $t \leq b^*$ such that $\delta(X, b) = 1$ for all $x \in A'_2$ and $b < t$. Then we take $t = t_-$ and redefine $A_2 \leftarrow A'_2$. Otherwise, $\delta(X, b) = 0$ for all $b \leq b^*$ for all $X \in A_2$. In this case take $t_- = \inf(\mathcal{B})$.

Finally, if $P(A_2) > P(A_1)$, then take a subset of $A'_2 \subset A_2$ such that $P(A'_2) = P(A_1)$ (which is possible by continuity). We may therefore assume that $P(A_1) = P(A_2)$.

The upshot of this construction is contained in the following three facts:

- F1. $P(A_1) = P(A_2)$
- F2. If $b > t_+$ then $\delta(X_1, b) = 0$ for all $X_1 \in A_1$.
- F3. If $b < t_-$ then $\delta(X_2, b) = 1$ for all $X_2 \in A_2$.

As we did in the proof of [Lemma 5.1](#), we will always assume that $X_1 \in A_1$ and $X_2 \in A_2$ to save ourselves from constantly writing quantifiers.

Now, let $G^+ := \{b \in \mathcal{B} : b > t_+\}$ and $G^- := \{b \in \mathcal{B} : b < t_-\}$. Either (or both) of these sets may be empty. Note that $G \cup G^+ \cup G^- = \mathcal{B}$. If G^+ and G^- are non-empty, facts F1 and F2 above can be rephrased as

$$\sup_{X_1 \in A_1} \sup_{b \in G^+} \delta(X_1, b) = 0, \quad \inf_{X_2 \in A_2} \inf_{b \in G^-} \delta(X_2, b) = 1. \quad (\text{A.26})$$

We define a new decision family, $\hat{\delta}$, where

$$\hat{\delta}(X, b) = \begin{cases} \delta(X, b) & X \notin A_1 \cup A_2 \text{ or } b \notin G, \\ 0, & X \in A_1 \text{ and } b \in G, \\ 1, & X \in A_2 \text{ and } b \in G. \end{cases} \quad (\text{A.27})$$

We claim that $\hat{\delta}$ remains type-I risk safe and has (sometimes strictly) lower type-II risk than δ , thus demonstrating that δ is $\Gamma_{\text{const-inadmissible}}$. To see type-I risk safety, write

$$\mathcal{R}_P(\hat{\delta}) = \int_{\mathcal{X} \setminus A_1 \cup A_2} \sup_b L_b(0, \delta(x, b)) \mathbb{P} + \int_{A_1 \cup A_2} \sup_b L_b(0, \hat{\delta}(x, b)) \mathbb{P}, \quad (\text{A.28})$$

where we've used that on $\mathcal{X} \setminus A_1 \cup A_2$, $\hat{\delta}(\cdot, b) = \delta(\cdot, b)$ for all b . Meanwhile for all $X_1 \in A_1$, $L_b(0, 1) \hat{\delta}(X_1, b) = 0$ for all $b \in G \cup G^+$ by F2 and definition of $\hat{\delta}$ (since $\hat{\delta}(X_1, b) \equiv 0$ by construction for $b \in G$ and $\hat{\delta}(X_1, b) = \delta(X_1, b) \equiv 0$ for $b \in G^+$ by F2. If $G^+ = \emptyset$ this holds trivially). Hence, for any $X_1 \in A_1$,

$$\sup_{b \in \mathcal{B}} L_b(0, 1) \hat{\delta}(x_1, b) = \sup_{b \in G^-} L_b(0, 1) \hat{\delta}(x_1, b) = \sup_{b \in G^-} L_b(0, 1) \delta(x_1, b) = \sup_{b \in G^-} L_b(0, 1),$$

where we interpret the supremum as 0 if $G^- = \emptyset$. Therefore,

$$\int_{A_1} \sup_b L_b(0, 1) \hat{\delta}(x, b) \mathbb{P} \leq \int_{A_1} \sup_{b \in G^-} L_b(0, 1) \mathbb{P} = \int_{A_2} \sup_{b \in G^-} L_b(0, 1) \mathbb{P},$$

since $P(A_1) = P(A_2)$. Moreover,

$$\int_{A_2} \sup_{b \in G^-} L_b(0, 1) \mathbb{P} = \int_{A_2} \sup_{b \in G^-} L_b(0, 1) \delta(x, b) \mathbb{P} \leq \int_{A_2} \sup_{b \in \mathcal{B}} L_b(0, 1) \delta(x, b) \mathbb{P}.$$

To summarize, the previous two displays imply that:

$$\int_{A_1} \sup_{b \in \mathcal{B}} L_b(0, 1) \hat{\delta}(x, b) \mathbb{P} \leq \int_{A_2} \sup_{b \in \mathcal{B}} L_b(0, 1) \delta(x, b) \mathbb{P}. \quad (\text{A.29})$$

Next, note that for all $X_2 \in A_2$, $\hat{\delta}(X_2, b) = 0$ for any $b \in G^+$, $\hat{\delta}(X_2, b) = 1$ for $b \in G$ by definition, and $\hat{\delta}(X_2, b) = \delta(X_2, b) = 1$ for $b \in G^-$ by F3. Therefore, since $L_b(0, 1)$ is increasing in b ,

$$\sup_{b \in \mathcal{B}} L_b(0, 1) \hat{\delta}(X_2, b) \leq \sup_{b \in G \cup G^-} L_b(0, 1) = \sup_{b \in G} L_b(0, 1),$$

and so, again using that $P(A_1) = P(A_2)$,

$$\begin{aligned} \int_{A_2} \sup_{b \in \mathcal{B}} L_b(0, 1) \hat{\delta}(x, b) \mathbb{P} &\leq \int_{A_2} \sup_{b \in G} L_b(0, 1) \mathbb{P} \\ &= \int_{A_1} \sup_{b \in G} L_b(0, 1) \mathbb{P} \leq \int_{A_1} \sup_{b \in \mathcal{B}} L_b(0, 1) \delta(x, b) \mathbb{P}. \end{aligned} \quad (\text{A.30})$$

Here the final inequality uses that

$$\sup_{b \in G} L_b(0, 1) = \sup_{b \in G} L_b(0, 1) \delta(X_1, b) \leq \sup_{b \in \mathcal{B}} L_b(0, 1) \delta(X_1, b),$$

for all $X_1 \in A_1$. Combining (A.29) and (A.30) we have

$$\int_{A_1 \cup A_2} \sup_{b \in B} L_b(0, 1) \hat{\delta}(x, b) \mathbb{P} \leq \int_{A_1 \cup A_2} \sup_{b \in B} L_b(0, 1) \delta(x, b) \mathbb{P}. \quad (\text{A.31})$$

Hence, by (A.28), $\mathcal{R}_P(\hat{\delta}) \leq \int_{\mathcal{X}} \sup_{b \in B} L_b(0, 1) \delta_b(X) \mathbb{P} = \mathcal{R}_P(\delta) \leq 1$.

Next we show that $\hat{\delta}$ has strictly lower type-II risk than δ for all $b \in G$. Since the type-II risk of $\hat{\delta}$ is the same as that of δ for all $b \notin G$ (by construction), this will show that δ is inadmissible. Since $P(A_1) = P(A_2)$ by construction, we have $Q(A_1) < Q(A_2)$ by Lemma A.1. Therefore, for any $b \in G$, since $\hat{\delta}(X_2, b) = 1$ for all $X_2 \in A_2$ and $\hat{\delta}(X_1, b) = 0$ for all $X_1 \in A_1$, we have

$$\begin{aligned} \mathbb{E}_Q[\hat{\delta}_b(X)] &= \int_{\mathcal{X} \setminus A_1 \cup A_2} \hat{\delta}(x, b) Q + \int_{A_1 \cup A_2} \hat{\delta}(x, b) Q(A_2) \\ &= \int_{\mathcal{X} \setminus A_1 \cup A_2} \delta(x, b) dQ + Q(A_2) \\ &> \int_{\mathcal{X} \setminus A_1 \cup A_2} \delta(x, b) Q + Q(A_1) = \mathbb{E}_Q[\delta(X, b)], \end{aligned}$$

which proves that δ is Γ_{const} -inadmissible, completing the argument.

A.16. Proof of Corollary 5.3

Lemma 5.3 implies that a C-admissible binary test φ is increasing in Λ , hence so too is E_φ by Lemma A.2 in the appendix. Therefore, there exists an increasing function h such that $E_\varphi(X) = h(\Lambda(X))$ and we may write $\varphi(X, b) = \mathbf{1}\{E_\varphi(X) \geq L_b(0, 1)\} = \mathbf{1}\{h(\Lambda(X)) \geq L_b(0, 1)\} = \mathbf{1}\{\Lambda(X) \geq h^-(L_b(0, 1))\}$, where h^- is the generalized inverse of h and the final equality holds since h is increasing. The function $t(b) = h^-(L_b(0, 1))$ defines the decision curve.

A.17. Proof of Theorem 5.3

If φ is Γ_{const} -admissible then E_φ is sharp by Lemma 4.3, increasing in Λ by Lemma 5.3, and has a minimal decision curve by Lemma A.5. This proves the forward direction. Conversely, suppose that E_φ is sharp and increasing in Λ and that φ has a minimal decision curve but that φ is Γ_{const} -inadmissible. Then, by Lemma A.4 there exists some ψ such that $t_\varphi(b) \leq t_\delta(b)$ for all $b \in B$ and $t_\varphi(b^*) < t_\delta(b^*)$ for some b^* , with $Q(t_\varphi(b^*) \leq \Lambda(X) < t_\delta(b^*)) > 0$. Let $I(b^*) = \{X : t_\varphi(b^*) \leq \Lambda(X) < t_\delta(b^*)\}$. We claim that

$$\int_{I(b^*)} \sup_b L_b(0, \varphi(x, b)) \mathbb{P} < \int_{I(b^*)} \sup_b L_b(0, \psi(x, b)) \mathbb{P}. \quad (\text{A.32})$$

For $X \in I(b^*)$, $\varphi(X, b^*) = \mathbf{1}\{\Lambda(X) \geq t_\varphi(b^*)\} = 1$, so $\sup_b L_b(0, \varphi(X, b)) \geq L_{b^*}(0, 1)$ and the right hand side of (A.32) is at least $L_{b^*}(0, 1)P(I(b^*))$. We claim that the integrand on the left hand side is less than $L_{b^*}(0, 1)$. We consider two cases.

1. Suppose that we can find some b_0 such that $t_\varphi(b^*) < t_\delta(b_0) < t_\delta(b^*)$ and $P(t_\varphi(b^*) \leq \Lambda(X) < t_\delta(b_0)) > 0$. Let $I(b^*) = I_1 \cup I_2$ where $I_1 = \{X : t_\varphi(b^*) \leq \Lambda(X) < t_\delta(b_0)\}$ and $I_2 = \{X : t_\delta(b_0) \leq \Lambda(X) < t_\delta(b^*)\}$. Since t_δ is increasing, for $X \in I_1$, we have $\mathbf{1}\{\Lambda(X) \geq t_\delta(b)\} = 0$ for all $b > b_0$ so

$$\sup_b L_b(0, \varphi(X, b)) = \sup_b L_b(0, 1) \mathbf{1}\{\Lambda(X) \geq t_\delta(b)\} \leq L_{b_0}(0, 1) < L_{b^*}(0, 1).$$

Similarly, for $X \in I_2$, $\sup_b L_b(0, \varphi(X, b)) \leq L_{b^*}(0, 1)$ and it follows that

$$\int_{I(b^*)} \sup_b L_b(0, \varphi(x, b)) \mathbb{P} = L_{b_0}(0, 1)P(I_1) + L_{b^*}(0, 1)P(I_2) < L_{b^*}(0, 1)P(I(b^*)),$$

using since $b_0 < b^*$ (again since t_δ is increasing). Thus, in this case, (A.32) holds.

2. Next suppose that no such b_0 exists. Therefore $\sup_{b < b^*} t_\delta(b) \leq t_\varphi(b^*) < t_\delta(b^*)$. Since $Q(\sup_{b < b^*} t_\delta(b) \leq \Lambda(X) < t_\delta(b^*)) \geq Q(t_\varphi(b^*) \leq \Lambda(X) < t_\delta(b^*)) > 0$ by assumption of φ , it follows that $\sup_{b < b^*} L_b(0, 1) < L_{b^*}(0, 1)$ since φ has a minimal decision curve. Therefore, since $\varphi(X, b) = \mathbf{1}\{\Lambda(X) \geq t_\delta(b)\}$ for all $b > b^*$ for $X \in I(b^*)$, we have

$$\begin{aligned} \int_{I(b^*)} \sup_b L_b(0, \varphi(x, b)) \mathbb{P} &= \int_{I(b^*)} \sup_{b < b^*} L_b(0, 1) \mathbf{1}\{\Lambda(x) \geq t_\delta(b)\} \mathbb{P} \\ &\leq \int_{I(b^*)} \sup_{b < b^*} L_b(0, 1) \mathbb{P} < L_{b^*}(0, 1)P(I(b^*)), \end{aligned}$$

and the final quantity is a lower bound on the right hand side of (A.32).

We have thus proved (A.32). Note all for all $X \in \mathcal{X}$, $\varphi(X, b) \geq \delta(X, b)$ since $t_\varphi(b) \leq t_\delta(b)$. Therefore,

$$\begin{aligned} \mathcal{R}_P(\varphi) &= \int_{\mathcal{X} \setminus I(b^*)} \sup_b L_b(0, \varphi(x, b)) \mathbb{P} + \int_{I(b^*)} \sup_b L_b(0, \varphi(x, b)) \mathbb{P} \\ &< \int_{\mathcal{X} \setminus I(b^*)} \sup_b L_b(0, \psi(x, b)) \mathbb{P} + \int_{I(b^*)} \sup_b L_b(0, \psi(x, b)) \mathbb{P} \end{aligned}$$

$$= \mathcal{R}_P(\psi) \leq 1,$$

contradicting that E_φ is sharp. This proves the first part of theorem.

For the second half of the theorem, let E be an e-variable for P . Suppose φ_E is Γ_{const} -admissible. By Lemma 4.7, $E = E_{\varphi_E}$. Then, by the first part of the theorem, E is sharp, compatible, and increasing in Λ . Conversely, suppose E is sharp, compatible, and increasing in Λ . By compatibility, $E = E_{\varphi_E}$ so it remains only to show that φ_E has a minimal decision-curve, in which case Γ_{const} -admissibility will again follow from applying the first part of the theorem. Let h be an increasing function such that $E(X) + h(\Lambda(X))$ and let $h^-(y) = \inf\{x \in \mathfrak{R} : h(x) \geq y\}$ by its generalized inverse. Since h is increasing, h^- is left continuous. Further,

$$\varphi_E(X, b) = \mathbf{1}\{E(X) \geq L_b(0, 1)\} = \mathbf{1}\{\Lambda(X) \geq h^-(L_b(0, 1))\},$$

so $h^-(L_b(0, 1))$ is equal to $t_\delta(b)$ P -almost surely. Therefore, if $\sup_{b < b^*} L_b(0, 1) = L_{b^*}(0, 1)$ then $h^-(\sup_{b < b^*} L_b(0, 1)) = h^-(L_{b^*}(0, 1))$, hence $t_\delta(b)$ is a minimal decision curve. This completes the proof.

Appendix B. On the definition of admissibility

First let us discuss the notion of admissibility used by Grünwald [1], which we henceforth refer to as G-admissibility. We say a test family δ is G-inadmissible if there exists a type-I risk safe test family ϕ such that, for all $b \in \mathcal{B}$ and all $Q \in \mathcal{Q}$,

$$Q(L_b(1, \phi(X, b)) > L_b(1, \delta(X, b))) = 0, \quad (\text{B.1})$$

and there exists some $b \in \mathcal{B}$ and some $Q \in \mathcal{Q}$ such that

$$Q(L_b(1, \phi(X, b)) < L_b(1, \delta(X, b))) > 0. \quad (\text{B.2})$$

In other words, ϕ has loss at most that of δ with Q -probability 1 on all losses, and has lower loss than δ with positive probability. If no such ϕ exists then we say that δ is G-admissible. We note that Grünwald [1] defines G-admissibility in terms of type-I losses instead of type-II, after preprocessing the losses to remove any such that $L_b(1, a) > L_b(1, a)$ and $L_b(0, a) > L_b(0, a)$ for some a , which can be done without loss of generality.

Now, if we translate Grünwald's results to our setting—in particular, we consider binary actions only—then his main result can be formulated as follows. Consider testing \mathcal{P} against \mathcal{Q} . Then:

1. If δ is G-admissible, then there exists some e-variable E for \mathcal{P} such that $\delta(X, b) = \mathbf{1}\{E(X) \geq L_b(0, 1)\}$.
2. If $\delta = \mathbf{1}\{E(X) \geq L_b(0, 1)\}$ for some some e-variable E for \mathcal{P} that is sharp, compatible with \mathcal{L}_0 , and all $P \in \mathcal{P}$ are absolutely continuous with one another, then δ is G-admissible.

Let us now discuss some relationships between Γ -admissibility and G-admissibility. First, C-admissibility implies G-admissibility. This is intuitive: If a test family δ cannot be beat in expectation for fixed b , then it cannot be beat with probability 1 either.

Lemma B.1. *C-admissibility implies G-admissibility.*

Proof. We prove the contrapositive. Suppose that δ is G-inadmissible. So there exists some ϕ such that, for all $b \in \mathcal{B}$,

$$\phi(X, b) \geq \delta(X, b) \quad Q\text{-almost everywhere for all } Q \in \mathcal{Q},$$

and $\phi(X, b_0) > \delta(X, b_0)$ for some b_0 and all $X \in \mathcal{A}$ with $Q_0(A) > 0$ for some Q_0 . Therefore, $\mathbb{E}_{Q_0}[\delta(X, b_0)\mathbf{1}\{A\}] < \mathbb{E}_{Q_0}[\phi(X, b_0)\mathbf{1}\{A\}]$ and $\mathbb{E}_{Q_0}[\delta(X, b_0)\mathbf{1}\{A^c\}] \leq \mathbb{E}_{Q_0}[\phi(X, b_0)\mathbf{1}\{A^c\}]$. Thus $\mathbb{E}_{Q_0}[\delta(X, b_0)] < \mathbb{E}_{Q_0}[\delta(X, b_0)]$. Moreover, since $\delta(X, b) \geq \phi(X, b)$ Q -almost surely for all b , we have $\mathbb{E}_Q[\delta(X, b)] \leq \mathbb{E}_Q[\phi(X, b)]$ for all b , hence ϕ is strictly preferable to δ with respect to U . Thus G-inadmissibility implies C-inadmissibility. \square

Note, however, that this proof does not work if we replace C-admissibility by U-admissibility. Consider the following example: Set $\mathcal{B} = [1, \infty)$ and let X take values on $[1, \infty)$. Suppose that $\delta(X, b) > \phi(X, b)$ for all $X \neq b$, and for $X = b$ we have $\delta(X, b) < \phi(X, b)$. Since $X = b$ has measure zero, δ is strictly preferable to ϕ with respect to G-admissibility. But for $B(X) = x$ we have $\mathbb{E}_Q[L_B(1, \delta(X, B))] = \mathbb{E}_Q[L_X(1, 0)(1 - \delta(X, X))] > \mathbb{E}_Q[L_X(1, 0)(1 - \phi(X, X))] = \mathbb{E}_Q[L_B(1, \phi(X, B))]$, so δ is not preferable to ϕ with respect to U-admissibility. This shows that we cannot conclude U-inadmissibility from G-inadmissibility.

Of course, this does not provide a counterexample to the claim that U-admissibility implies G-admissibility per se. But it sheds light on the fact that one would need a proof totally unlike the one above. This makes us doubt that the claim is true.

Next let us return to a question raised in Section 2: Why not define admissibility with a supremum over losses inside the expectation? That is, we could say that δ is admissible if it is type-I risk safe and for all other type-I risk safe tests ϕ ,

$$\mathbb{E} \left[\sup_{b \in \mathcal{B}} (L_b(1, \delta_b(X)) - L_b(1, \phi_b(X))) \right] \leq 0. \quad (\text{B.3})$$

Such a definition, however, leads to an empty set of admissible test families. Intuitively, no test can be as good as any other on every data-dependent loss. To do so, it must be the most powerful test on every loss. That is, for a point null versus a point alternative, it must be the Neyman-Pearson test. But it is invalid to play the Neyman-Pearson on each loss, since such a test family is not type-I risk safe. This is formalized in the following argument.

Suppose we are testing a simple null vs simple alternative, $\mathcal{P} = \{P\}$ vs $\mathcal{Q} = \{Q\}$. Consider two loss functions, $\mathcal{B} = \{0, 1\}$. Suppose δ is admissible according to (B.3) and consider the Neyman-Pearson test on each loss. That is, for $i \in \{0, 1\}$, define

$$\phi^{(i)}(X, b) = \begin{cases} \mathbf{1}\{\Lambda(X) \geq \kappa_i\}, & \text{if } b = i \\ 0, & \text{otherwise,} \end{cases}$$

where, as usual, κ_i is chosen such that $\mathbb{E}_P[\phi^{(i)}(X, b)] = 1/L_i(0, 1)$. Note that $\phi^{(i)}$ is type-I risk safe. Also,

$$\begin{aligned} & \mathbb{E}_Q \left[\sup_b (L_b(1, \delta(X, b)) - L_b(1, \phi^{(i)}(X, b))) \right] \\ & \geq \mathbb{E}_Q [L_i(1, \delta(X, i)) - L_i(1, \phi^{(i)}(X, i))] \\ & = L_i(1, 0) \mathbb{E}_Q [\phi^{(i)}(X, i) - \delta(X, i)] \geq 0, \end{aligned}$$

where the last inequality follows since $\phi^{(i)}$ is the uniformly most powerful test by the Neyman-Pearson lemma. Therefore, by admissibility we conclude that $\mathbb{E}_Q[\phi^{(i)}(X, i) - \delta(X, i)] = 0$. That is, δ has the same power as $\phi^{(i)}$ under loss b_i and thus, by the uniqueness of the Neyman-Pearson lemma, we have $\delta = \phi^{(i)}$ almost surely. But i was arbitrary, so $\phi^{(1)} = \delta = \phi^{(2)}$, a contradiction.

B.1. Type-I risk as the supremum over adversaries

Let us now discuss (3), i.e., the fact that we can write type-I risk as the supremum over maps from the data to losses.

Recall that (\mathcal{A}, Σ) is a standard Borel space if there exists a complete, separable metric ρ on X such that Σ becomes the Borel σ -field under ρ . A probability space (Ω, \mathcal{F}, P) is Borel if (Ω, \mathcal{F}) is Borel. If we call a function $Y : (\mathcal{A}, \Sigma_{\mathcal{A}}) \rightarrow (\mathfrak{R}, \Sigma_{\mathfrak{R}})$ $\Sigma_{\mathcal{A}}$ -measurable, we mean that it is $\Sigma_{\mathcal{A}}/\Sigma_{\mathfrak{R}}$ -measurable where $\Sigma_{\mathfrak{R}}$ is understood to be the Borel σ -field; likewise if \mathfrak{R} is replaced with any subset of \mathfrak{R} .

Lemma B.2. *Let (Ω, \mathcal{F}, P) be a Borel probability space. Equip \mathcal{X} and \mathcal{B} with σ -fields $\Sigma_{\mathcal{X}}$ and $\Sigma_{\mathcal{B}}$ such that $(\mathcal{X}, \Sigma_{\mathcal{X}})$ and $(\mathcal{B}, \Sigma_{\mathcal{B}})$ are standard Borel spaces. Consider any set of type-I loss functions $\{L_b(0, 1) : b \in \mathcal{B}\}$ such that $b \mapsto L_b(0, 1)$ is $\Sigma_{\mathcal{B}}$ -measurable. Let δ be a test family that is $(\Sigma_{\mathcal{X}} \otimes \Sigma_{\mathcal{B}})$ -measurable. If δ has finite type-I risk, then:*

$$\mathbb{E}_{X \sim P} \sup_{b \in \mathcal{B}} L_b(0, \delta(X, b)) = \sup_{B: \mathcal{X} \rightarrow \mathcal{B}} \mathbb{E}_{X \sim P} L_{B(X)}(0, \delta(X, B(X))), \quad (\text{B.4})$$

where the supremum on the right is over all maps B that are $\Sigma_{\mathcal{X}}/\Sigma_{\mathcal{B}}$ -measurable.

Proof. First we claim the left hand side is well-defined. Define the function $g : \Omega \rightarrow \mathfrak{R}$ by $g(\omega) = \sup_b L_b(0, \delta(X(\omega), b))$. We claim that g is \mathcal{F} -measurable. Since $L_b = L_b(0, 1)$ is $\Sigma_{\mathcal{B}}$ -measurable and δ is $(\Sigma_{\mathcal{X}} \otimes \Sigma_{\mathcal{B}})$ -measurable, $L_b(0, \delta(x, b)) = L_b(0, 1)\delta(x, b)$ is $(\Sigma_{\mathcal{X}} \otimes \Sigma_{\mathcal{B}})$ -measurable. Consequently, the map $(\omega, b) \mapsto L_b(0, \delta(X(\omega), b))$ is $(\mathcal{F} \otimes \Sigma_{\mathcal{B}})$ -measurable. Consider the super-level sets

$$\{\omega : g(\omega) > c\} = \Gamma_{\Omega}(\underbrace{\{\omega, b : L_b(0, 1)\delta(X(\omega), b) > c\}}_{=: A_c})$$

where Γ_{Ω} represents the projection of a set in $\Omega \times \mathcal{B}$ onto Ω . The set A_c is Borel (i.e., lies in $\Sigma_{\mathcal{X}} \otimes \mathcal{B}$) by the measurability of $L_b(0, 1)\delta(X(\omega), b)$. The projection of a Borel set on a product space is analytic by Suslin's theorem, hence $\{\omega : g(\omega) > c\}$ is analytic [48, Exercise 14.3]. Moreover, every analytic set is universally measurable [48, Theorem 21.10], meaning it is measurable for any finite Borel measure on Ω , and P is finite and Borel by assumption. Finally, if the super-level sets of a real-valued function are measurable, then the function itself measurable. We conclude that g is \mathcal{F} -measurable, so the left hand side of (3) exists. Now, for any X and B , $L_{B(X)}(0, \delta(X, B(X))) \leq \sup_b L_b(0, \delta(X, b))$. Taking expectations and then taking the supremum over all maps B , we see that the right hand side of (3) is at most the left hand side.

To prove the converse fix $\epsilon > 0$ and let ω be given. Define the multifunction $\Xi_{\epsilon} : \Omega \rightrightarrows \mathcal{B}$

$$\Xi_{\epsilon}(\omega) = \left\{ b^* \in \mathcal{B} : L_{b^*}(0, \delta(X(\omega), b^*)) > \sup_b L_b(0, \delta(X(\omega), b)) - \epsilon \right\},$$

which is \mathcal{F} -measurable since $(\omega, b) \mapsto L_b(0, \delta(X(\omega), b))$ is measurable. The Kuratowski-Ryll-Nardzewski measurable selection theorem [49] (see also Kechris 48, Theorem 12.13) guarantees the existence of a measurable selector $B_{\epsilon}^*(X(\omega))$ such that $B_{\epsilon}^*(X(\omega)) \in \Xi_{\epsilon}(\omega)$, hence $L_{B_{\epsilon}^*(X)}(0, \delta(X(\omega), B_{\epsilon}^*(X))) + \epsilon > \sup_b L_b(0, \delta(X, b))$. Taking expectations and letting $\epsilon \downarrow 0$ shows that the left hand side of (3) is at most the right, which proves the result. \square

The following corollary extends the result to composite null hypotheses. The proof follows by taking the supremum over $P \in \mathcal{P}$ on both sides of (B.4).

Corollary B.1. *Let (Ω, \mathcal{F}) be a measurable space and let \mathcal{P} be a collection of probability measures such that (Ω, \mathcal{F}, P) is a Borel probability space for each $P \in \mathcal{P}$. Let $(\mathcal{X}, \Sigma_{\mathcal{X}})$, $(\mathcal{B}, \Sigma_{\mathcal{B}})$, and $\{L_b(0, 1) : b \in \mathcal{B}\}$ be as in Lemma B.2. Then*

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{X \sim P} \sup_{b \in \mathcal{B}} L_b(0, \delta(X, b)) = \sup_{P \in \mathcal{P}} \sup_{B: \mathcal{X} \rightarrow \mathcal{B}} \mathbb{E}_{X \sim P} L_{B(X)}(0, \delta(X, B(X))), \quad (\text{B.5})$$

where the rightmost supremum is over all maps B that are $\Sigma_{\mathcal{X}}/\Sigma_{\mathcal{B}}$ -measurable.

Remark 6. There exist pathological choices of $\{L_b\}_{b \in \mathcal{B}}$ which ensure that $\sup_b L_b(0, \delta(X, b))$ is not measurable, in which case Definition 2.3 breaks down and Lemma B.2 doesn't hold. Instead of making various measurability assumptions as we do in this work, one could circumvent this problem by following Grünwald [1] and calling δ type-I risk safe if there exists some measurable function $\xi : \mathcal{X} \rightarrow \mathfrak{R}_{\geq 0}$ such that

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{X \sim P} \xi(X) \leq 1 \text{ and } \sup_{b \in \mathcal{B}} L_b(0, \delta(X, b)) \leq \xi(X) \text{ } \mathcal{P}\text{-almost surely.} \quad (\text{B.6})$$

Appendix C. On the relationship to Bayesian decision theory

A common response to the drawbacks of classical Neyman-Pearson theory that we've highlighted is to pivot to Bayesian decision theory [22]. While the Bayesian often/typically omits the notion of type-I error control and focuses solely on loss minimization, we show below that he gains the ability to handle post-hoc loss functions for free. We also discuss what a Bayesian notion of type-I risk control might look like, and how to achieve it.

Moving away from the setting considered in the main paper for a moment, let Θ be a parameter space, \mathcal{A} a set of actions, and $L : \Theta \times \mathcal{A} \rightarrow \mathfrak{R}_{\geq 0}$ a loss function. Associated to each $\theta \in \Theta$ is a distribution P_θ . Let π be a prior on Θ . Given observations X , Bayesian decision theory tells us to minimize the expected posterior loss. That is, we take action

$$\delta_\pi(X) \equiv \arg \min_{a \in \mathcal{A}} \mathbb{E}_{\theta \sim \pi(\cdot|X)} L(\theta, a). \quad (\text{C.1})$$

A fundamental result is that the *Bayes optimal decision rule* [50, Theorem 4.1.1], henceforth 'Bayes decision' for brevity, is always given by δ_π , meaning that it satisfies $B_\pi(\delta_\pi) = \inf_\delta B_\pi(\delta)$ where

$$B_\pi(\delta) = \mathbb{E}_{\theta \sim \pi} \mathbb{E}_{X \sim P_\theta} [L(\theta, \delta(X))]. \quad (\text{C.2})$$

(Note that θ is drawn from the prior π in (C.2), not the posterior $\pi(\cdot|X)$ as in (C.1).) Bayes estimators are, as the name suggests, the Bayesian equivalent to frequentist minimax estimators, so a unified solution for the Bayes decision regardless of the loss function is a big benefit of Bayesian decision theory.

To apply this to our setting, we take $\mathcal{A} = \{0, 1\}$ and we assume that Θ is naturally partitioned into Θ_0 and Θ_1 , where Θ_0 corresponds to the null and Θ_1 to the alternative. This induces the conditional prior distributions $\pi_i(\theta) = \pi_i(\theta | \theta \in \Theta_i)$ and conditional posterior distributions $\pi_i(\theta | X) = \pi_i(\theta | X, \theta \in \Theta_i)$.

Assuming as usual that $L(\theta, i) = 0$ for $\theta \in \Theta_i$ and $i \in \{0, 1\}$, the Bayes decision becomes

$$\begin{aligned} \delta_\pi(X) &= \arg \min \{ \mathbb{E}_{\theta \sim \pi(\cdot|X)} [L(\theta, 0)], \mathbb{E}_{\theta \sim \pi(\cdot|X)} [L(\theta, 1)] \} \\ &= \arg \min \{ \mathbb{E}_{\theta \sim \pi_1(\cdot|X)} [L(\theta, 0)] \pi(\Theta_1 | X), \mathbb{E}_{\theta \sim \pi_0(\cdot|X)} [L(\theta, 1)] \pi(\Theta_0 | X) \} \\ &= 1 \left\{ \mathbb{E}_{\theta \sim \pi_1(\cdot|X)} [L(\theta, 0)] \pi(\Theta_1 | X) \geq \mathbb{E}_{\theta \sim \pi_0(\cdot|X)} [L(\theta, 1)] \pi(\Theta_0 | X) \right\}. \end{aligned} \quad (\text{C.3})$$

In words, we play whichever action has the lowest cost if wrong (i.e., results in a false positive or negative), weighted by the probability under the posterior.

To move to a post-hoc setting, let's consider a set of losses $\{L_b\}_{b \in \mathcal{B}}$. As before, we allow an adversary to choose b after seeing the data and we encode the adversary as a measurable map $B : \mathcal{X} \rightarrow \mathcal{B}$. As usual, we let the decision rule see $B(X)$ before making a decision.

Can Bayesian decision theory accommodate a post-hoc selection of the loss? Yes, in the following sense: If we minimize the expected posterior loss as before but on loss $L_{B(X)}$ (which we observe), then the resulting decision rule minimizes Bayes risk, where the Bayes risk is now defined in terms of the loss selected by the adversary. More formally:

Proposition C.1. The decision rule

$$\phi_\pi(X) = \arg \min_{a \in \mathcal{A}} \mathbb{E}_{\theta \sim \pi(\cdot|X)} L_{B(X)}(\theta, a),$$

satisfies

$$\mathbb{E}_{\theta \sim \pi} \mathbb{E}_{X \sim P_\theta} L_{B(X)}(\theta, \phi_\pi(X)) = \inf_{\phi} \mathbb{E}_{\theta \sim \pi} \mathbb{E}_{X \sim P_\theta} L_{B(X)}(\theta, \phi(X)),$$

for all $B : \mathcal{X} \rightarrow \mathcal{B}$.

We emphasize that ϕ_π is well-defined since, as mentioned earlier, ϕ is allowed to see $B(X)$. We also note that this result is not new. For example, Grünwald [25] notes that that Bayesian decision can accommodate post-hoc loss functions, but we hope the formal treatment above clarifies in what sense this is true, and allows for easier comparison to our frequentist results.

Proof. For any decision rule ϕ , let $\ell(\phi(x)|x) = \mathbb{E}_{\theta \sim \pi(\cdot|X)} L_{B(X)}(\theta, \phi(x))$ be its expected posterior loss at a given x on loss $L_{B(X)}$. For simplicity, assume that P_θ has density $p(\cdot|\theta)$. Let $p(x) = \int p(x|\theta) \pi(\theta) d\theta$ be the marginal and note that $p(x|\theta) \pi(\theta) = \pi(\theta|x) p(x)$. Then, by Fubini's theorem,

$$\begin{aligned} \mathbb{E}_{\theta \sim \pi} \mathbb{E}_{X \sim P_\theta} L_{B(X)}(\theta, \phi(X)) &= \int_{\Theta} \int_{\mathcal{X}} L_{B(X)}(\theta, \phi(x)) p(x|\theta) \pi(\theta) dx d\theta \\ &= \int_{\mathcal{X}} \int_{\Theta} L_{B(X)}(\theta, \phi(x)) \pi(\theta|x) p(x) d\theta dx \end{aligned}$$

$$= \int_X \ell(\phi(x)|x)p(x)d\theta.$$

By minimizing the expected posterior $\ell(\phi(x)|x)$ we're minimizing the above integrand at every x , thus minimizing the value of the integral. \square

The Bayesian paradigm is thus post-hoc valid in the sense of [Proposition C.1](#). But it does not take into account the notion of type-I error—it is solely concerned with minimizing the loss. In other words, Bayesian theory treats type-I and type-II loss as symmetric, whereas Neyman-Pearson theory, and our generalized version thereof, views type-I risk as taking precedence. A natural question is whether one can incorporate a notion of type-I risk safety into Bayesian decision theory; we provide one answer below.

We might consider the following definition of type-I risk in the Bayesian context:

$$S_\pi(\delta) = \sup_{B: \mathcal{X} \rightarrow \mathcal{B}} \mathbb{E}_{\theta \sim \pi_0} \mathbb{E}_{X \sim P_\theta} L_{B(X)}(\theta, \delta(X)), \quad (\text{C.4})$$

where we emphasize that θ is being drawn from the induced prior π_0 over Θ_0 , not from π itself. It's helpful to compare [\(C.4\)](#) to the frequentist counterpart in the main paper, [Definition 2.3](#). First, following the Bayesian philosophy, we've replaced the supremum over Θ_0 with an expectation (the supremum over Θ_0 appears as a supremum over \mathcal{P} in [Definition 2.3](#)). But we have also moved the supremum over B to the outside; we leave moving the supremum inside $\mathbb{E}_{\theta \sim \pi_0}$ as a possible future direction.

It turns out that we can solve a constrained optimization problem in order provide an alternative estimator which satisfies $S_\pi(\delta) \leq 1$. This alternative estimator is the Bayes decision for a new problem with a different loss function. Only the general form of this estimator can be given, however. It involves a parameter λ which, to calculate explicitly, would require knowledge of the particular distributions involved (and thus can be calculated for any given problem, just not in a closed form manner that applies to all problems).

Proposition C.2. *There exists some $\lambda \geq 0$ such that the Bayes decision δ_π with loss*

$$\hat{L}(\theta, a) = \begin{cases} L(\theta, a), & \theta \in \Theta_1, \\ \frac{\pi(\Theta_0) + \lambda}{\pi(\Theta_0)} L(\theta, a), & \theta \in \Theta_0, \end{cases} \quad (\text{C.5})$$

satisfies $S_\pi(\delta_\pi) \leq 1$ on the original problem.

Proof. In order to enforce the constraint $S_\pi(\delta_\pi) \leq 1$ we introduce the Lagrange multiplier λ and solve the optimization problem

$$\min_{\delta} \{ B_\pi(\delta) + \lambda(S_\pi(\delta) - 1) \}. \quad (\text{C.6})$$

Expanding, we can collect terms and obtain

$$\begin{aligned} B_\pi(\delta) + \lambda(S_\pi(\delta) - 1) &= (\pi(\Theta_0) + \lambda) \mathbb{E}_{\theta \sim \pi_0} \mathbb{E}_{X \sim P_\theta} L(\theta, \delta(X)) \\ &\quad + \mathbb{E}_{\theta \sim \pi_1} \mathbb{E}_{X \sim P_\theta} L(\theta, \delta(X)) \pi(\Theta_1) - \lambda \\ &= \mathbb{E}_{\theta \sim \pi} \mathbb{E}_{X \sim P_\theta} \hat{L}(\theta, \delta(X)) - \lambda. \end{aligned}$$

This is minimized by the Bayes estimator in [\(C.1\)](#) with L replaced by \hat{L} . The value of λ can then be adjusted to ensure that $S_\pi(\delta) \leq 1$. \square

Appendix D. On risk control vs error probabilities

The formulation of post-hoc hypothesis testing we've advanced here changes the kinds of guarantees that an analyst can give concerning the performance of their tests. The currency of our framework is risk control ([Definition 2.3](#)) instead of type-I error control—one obtains a bound on the expected loss instead of the error probability. Given that modern hypothesis testing uses error probabilities, it is reasonable to be skeptical of changing the metric. Here we consider four brief arguments in favor of risk control. To provide a full-throated defense would require its own paper, but it's worth at least sketching these considerations here.

The first argument is historical. As we've discussed, Wald [18] noted that standard Neyman-Pearson hypothesis testing is a part of classical statistical decision theory once it is reformulated in terms of loss functions. And decision theory has always been concerned with expected loss as the error metric, due to compelling arguments such as representation theorems [51], coherence and Dutch books arguments [52], complete class theorems [53,54], amenability to central limit theorems and laws of large numbers, and so on. Thus, our emphasis on risk control can be seen as a return to the roots of hypothesis testing as decision theory, where expected risk is the norm.

The second is a pragmatic argument. As we mentioned in the introduction, there are alternatives to Neyman-Pearson hypothesis testing, the most successful of which is the Bayesian agenda. But if one is fond of the tools of the Neyman-Pearson paradigm—such as significance levels, power, and accept/reject procedures—and seeks to generalize those tools to handle post-hoc analysis, then the formulation of post-hoc hypothesis testing studied here (and in 1) appears to be the only option currently on offer. (See [Appendix C](#) for a longer discussion on Bayesian decision theory and how it relates to our framework.)

The third argument is pedagogical. It is well known that p-values and error probabilities are misunderstood by practitioners, necessitating a near-constant stream of mildly condescending educational articles on their correct use and interpretation. Changing the metric to risk control (and reporting e-values instead of p-values) may be prophylactic, allowing us to leave some of this confusion behind.

Last is a normative argument. Traditional hypothesis testing, with its focus on controlling error probabilities, can mislead practitioners into treating the results as definitive solely based on whether $p \leq \alpha$. This obscures the fundamental uncertainty inherent in any single study. Risk control, by contrast, emphasizes expected loss across many applications of a testing procedure. The guarantee is only that the average loss will be controlled in the long run. It emphasizes that any individual study may be wrong, and one should remain humble when asserting conclusions.

References

- [1] P. Grünwald, Beyond Neyman–Pearson: E-values enable hypothesis testing with a data-driven alpha, *Proc. Natl. Acad. Sci.* 121 (39) (2024) e2302098121.
- [2] D. Bakan, The test of significance in psychological research, *Psychol. Bull.* 66 (6) (1966) 423.
- [3] R. Dar, R.C. Serlin, H. Omer, Misuse of statistical tests in three decades of psychotherapy research, *J. Consult. Clin. Psychol.* 62 (1) (1994) 75.
- [4] G. Gigerenzer, Mindless statistics, *J. Socio Econ.* 33 (5) (2004) 587–606.
- [5] J.W. Pratt, The foundations of statistical inference, *Q. Appl. Math.* (1964) 170–172.
- [6] E.-J. Wagenmakers, A practical solution to the pervasive problems of p values, *Psychon. Bull. Rev.* 14 (5) (2007) 779–804.
- [7] R. Fisher, Theory of statistical estimation, in: *Mathematical Proceedings of the Cambridge Philosophical Society*, 22, Cambridge University Press, 1925, pp. 700–725.
- [8] R. Fisher, The logic of inductive inference, *J. R. Stat. Soc.* 98 (1) (1935) 39–82.
- [9] J. Neyman, E. Pearson, On the use and interpretation of certain test criteria for purposes of statistical inference part i, *Biometrika* 20 (1–2) (1928) 175–240.
- [10] S.N. Goodman, P values, hypothesis tests, and likelihood: implications for epidemiology of a neglected historical debate, *Am. J. Epidemiol.* 137 (5) (1993) 485–496.
- [11] R. Hubbard, M.J. Bayarri, P Values are not Error Probabilities, *Working Paper 03–26*, Institute of Statistics and Decision Sciences, Duke University, Durham, NC, 2003.
- [12] R. Hubbard, The widespread misinterpretation of p-values as error probabilities, *J. Appl. Stat.* 38 (11) (2011) 2617–2626.
- [13] S.N. Goodman, Toward evidence-based medical statistics. 1: the p value fallacy, *Ann. Intern. Med.* 130 (12) (1999) 995–1004.
- [14] G. Gigerenzer, The superego, the ego, and the id in statistical reasoning, in: *A Handbook for Data Analysis in the Behavioral Sciences*, Psychology Press, 2014, pp. 311–340.
- [15] A. Ramdas, P. Grünwald, V. Vovk, G. Shafer, Game-theoretic statistics and safe anytime-valid inference, *Stat. Sci.* 38 (4) (2023) 576–601.
- [16] P. Grünwald, R. de Heide, W. Koolen, Safe testing, *J. R. Stat. Soc. Series B Stat. Methodol.* 86 (5) (2024) 1091–1128.
- [17] A. Ramdas, R. Wang, Hypothesis testing with e-values, *Foundations and Trends in Statistics* (2025).
- [18] A. Wald, Contributions to the theory of statistical estimation and testing hypotheses, *Ann. Math. Stat.* 10 (4) (1939) 299–326.
- [19] J. Neyman, E. Pearson, On the problem of the most efficient tests of statistical hypotheses, *Philos. Trans. R. Soc. London. Series A Contain. Pap. Math. Phys. Char.* 231 (694–706) (1933) 289–337.
- [20] C.R. Rao, Information and accuracy attainable in the estimation of statistical parameters, *Bull. Calcutta Math. Soc.* 37 (3) (1945) 81–91.
- [21] D. Blackwell, Conditional expectation and unbiased sequential estimation, *Ann. Math. Stat.* (1947) 105–110.
- [22] J.M. Bernardo, A.F.M. Smith, M. Berliner, *Bayesian Theory*, 586, Wiley Online Library, 1994.
- [23] A. Gelman, J.B. Carlin, H.S. Stern, D.B. Rubin, *Bayesian Data Analysis*, Chapman and Hall/CRC, 1995.
- [24] S.L. Chau, A. Schrab, A. Gretton, D. Sejdinovic, K. Muandet, Credal two-sample tests of epistemic uncertainty, in: *International Conference on Artificial Intelligence and Statistics*, PMLR, 2025, pp. 127–135.
- [25] P. Grünwald, The e-posterior, *Philos. Trans. R. Soc. A* 381 (2247) (2023) 20220146.
- [26] M. Caprio, The joys of categorical conformal prediction, (2025) *arXiv preprint arXiv:2507.04441* .
- [27] M.J. Gardner, D.G. Altman, Confidence intervals rather than p values: estimation rather than hypothesis testing, *Br. Med. J. (Clin. Res. Ed)* 292 (6522) (1986) 746–750.
- [28] R. Royall, *Statistical Evidence: A Likelihood Paradigm*, Routledge, 2017.
- [29] V. Vovk, R. Wang, E-values: calibration, combination and applications, *Ann. Stat.* 49 (3) (2021) 1736–1754.
- [30] E. Katsevich, A. Ramdas, Simultaneous high-probability bounds on the false discovery proportion in structured, regression and online settings, *Ann. Stat.* 48 (6) (2020) 3465–3487.
- [31] R. Wang, A. Ramdas, False discovery rate control with e-values, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 84 (3) (2022) 822–852.
- [32] Z. Xu, R. Wang, A. Ramdas, Post-selection inference for e-value based confidence intervals, *Electron. J. Stat.* 18 (1) (2024) 2292–2338.
- [33] N.W. Koning, Post-hoc α hypothesis testing and the Post-hoc p-value, (2023) *arXiv preprint arXiv:2312.08040* .
- [34] J. Hemerik, N.W. Koning, Choosing alpha post hoc: the danger of multiple standard significance thresholds, (2024) *arXiv preprint arXiv:2410.02306* .
- [35] Z. Xu, A. Solari, L. Fischer, R. de Heide, A. Ramdas, J. Goeman, Bringing closure to false discovery rate control: a general principle for multiple testing, (2025) *arXiv preprint arXiv:2509.02517* .
- [36] E. Gauthier, F. Bach, M.I. Jordan, E-values expand the scope of conformal prediction, (2025) *arXiv preprint arXiv:2503.13050* .
- [37] R. Wang, The only admissible way of merging arbitrary e-values, *Biometrika* 112 (2) (2025) 0–20.
- [38] V. Vovk, R. Wang, Combining p-values via averaging, *Biometrika* 107 (4) (2020) 791–808.
- [39] A. Ramdas, J. Ruf, M. Larsson, W. Koolen, Admissible anytime-valid sequential inference must rely on nonnegative martingales, (2020) *arXiv preprint arXiv:2009.03167* .
- [40] N.W. Koning, Continuous testing: unifying tests and e-values, (2024) *arXiv preprint arXiv:2409.05654* .
- [41] J. Shao, *Mathematical Statistics*, Springer Science & Business Media, 2008.
- [42] G. Shafer, Testing by betting: a strategy for statistical and scientific communication, *J. R. Stat. Soc. Ser. A Stat. Soc.* 184 (2) (2021) 407–431.
- [43] L. Wasserman, *All of Statistics: A Concise Course in Statistical Inference*, Springer Science & Business Media, 2004.
- [44] R. Royall, The effect of sample size on the meaning of significance tests, *Am. Stat.* 40 (4) (1986) 313–315.
- [45] R. Royall, On the probability of observing misleading statistical evidence, *J. Am. Stat. Assoc.* 95 (451) (2000) 760–768.
- [46] J.O. Berger, T. Sellke, Testing a point null hypothesis: the irreconcilability of p values and evidence, *J. Am. Stat. Assoc.* 82 (397) (1987) 112–122.
- [47] T.V. Perneger, Sifting the evidence: likelihood ratios are alternatives to p values, *BMJ Br. Med. J.* 322 (7295) (2001) 1184.
- [48] A. Kechris, *Classical Descriptive Set Theory*, 156, Springer Science & Business Media, 2012.
- [49] K. Kuratowski, C. Ryll-Nardzewski, A general theorem on selectors, *Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys.* 13 (6) (1965) 397–403.
- [50] E.L. Lehmann, G. Casella, *Theory of Point Estimation*, Springer, 1998.
- [51] L.J. Savage, *The Foundations of Statistics*, John Wiley & Sons, 1954.
- [52] B. De Finetti, La prévision: ses lois logiques, ses sources subjectives, in: *Annales de l'institut Henri Poincaré*, 7, 1937, pp. 1–68.
- [53] A. Wald, *Statistical decision functions*, in: *Breakthroughs in Statistics: Foundations and Basic Theory*, Springer, 1950, pp. 342–357.
- [54] D. Blackwell, M.A. Girshick, *Theory of Games and Statistical Decisions*, Wiley, 1954.