

A complete characterization of testable hypotheses

Martin Larsson*

Aaditya Ramdas[†]

Johannes Ruf[‡]

March 3, 2026

Abstract

We revisit a fundamental question in hypothesis testing: given two sets of probability measures $\mathcal{P}$ and $\mathcal{Q}$, when does a nontrivial (i.e. strictly unbiased) test for $\mathcal{P}$ against $\mathcal{Q}$ exist? Le Cam showed that, when $\mathcal{P}$ and $\mathcal{Q}$ have a common dominating measure, a test that has power exceeding its level by more than ε exists if and only if the convex hulls of $\mathcal{P}$ and $\mathcal{Q}$ are separated in total-variation distance by more than ε . The requirement of a dominating measure is frequently violated in nonparametric statistics. In a passing remark, Le Cam described an approach to address more general scenarios, but he stopped short of stating a formal theorem. This work completes Le Cam’s program, by presenting a matching necessary and sufficient condition for testability: for the aforementioned theorem to hold without assumptions, one must take the closures of the convex hulls of $\mathcal{P}$ and $\mathcal{Q}$ in the space of bounded finitely additive measures. We provide simple elucidating examples, and elaborate on various subtle measure theoretic and topological points regarding compactness and achievability.

1 Introduction

Let $\mathcal{M}_1$ denote the set of all (countably additive) probability measures on a given a measurable space $(\Omega, \mathcal{F})$. The null and alternative hypotheses are $\mathcal{P}, \mathcal{Q} \subset \mathcal{M}_1$. A *test* ϕ is a $[0, 1]$ -valued measurable function on this space, and its worst case level (or type-I error) and power are given by $\sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi]$ and $\inf_{\nu \in \mathcal{Q}} \mathbb{E}_\nu[\phi]$. Let $\Phi = \Phi(\mathcal{P}, \mathcal{Q})$ denote the set of all tests. One calls a test ϕ *strictly unbiased* if

$$\sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi] < \inf_{\nu \in \mathcal{Q}} \mathbb{E}_\nu[\phi],$$

meaning that its worst case power is larger than its worst case level. Since this is a rather reasonable minimal requirement for a test, we use the shorter term “nontrivial” to mean strictly unbiased. This paper answers a simple question: when does a nontrivial test exist?

*Department of Mathematical Sciences, Carnegie Mellon University, larsson@cmu.edu

[†]Departments of Statistics and Machine Learning, Carnegie Mellon University aramdas@cmu.edu

[‡]Department of Mathematics, London School of Economics, j.ruf@lse.ac.uk

The risk of a test ϕ is defined as the sum of its worst case type-I and type-II errors,

$$R(\phi) = R(\phi, \mathcal{P}, \mathcal{Q}) = \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi] + \sup_{\nu \in \mathcal{Q}} \mathbb{E}_\nu[1 - \phi].$$

The risk is unaffected if $\mathcal{P}$ and $\mathcal{Q}$ are replaced by their convex hulls $\text{co}(\mathcal{P})$ and $\text{co}(\mathcal{Q})$, which consist of all finite convex combinations of elements of $\mathcal{P}$ and $\mathcal{Q}$. That is,

$$R(\phi, \mathcal{P}, \mathcal{Q}) = R(\phi, \text{co}(\mathcal{P}), \text{co}(\mathcal{Q})).$$

The minimax risk is defined by

$$R(\mathcal{P}, \mathcal{Q}) = \inf_{\phi \in \Phi} R(\phi, \mathcal{P}, \mathcal{Q}).$$

Clearly a test ϕ is nontrivial if and only if $R(\phi) < 1$, and such a test exists if and only if $R(\mathcal{P}, \mathcal{Q}) < 1$. If $R(\mathcal{P}, \mathcal{Q}) = R(\phi^*)$ for some $\phi^* \in \Phi$, we call ϕ^* a minimax optimal test. Because the risk $R(\phi)$ is a supremum of linear functions, it is convex. Thus solving for the minimax risk is a convex minimization problem. However, Φ is not in general compact in any topology for which the risk is continuous (or lower semicontinuous), and so a minimax optimal test may not always exist. It is well-known that a minimax optimal test *does* exist under certain conditions; this is discussed further below.

The minimax risk is closely related to the total-variation (TV) distance between $\text{co}(\mathcal{P})$ and $\text{co}(\mathcal{Q})$. Indeed, for any test ϕ and arbitrary probability measures μ, ν , it straightforwardly holds that

$$\mathbb{E}_\mu[\phi] + \mathbb{E}_\nu[1 - \phi] \geq 1 - d_{\text{TV}}(\mu, \nu), \quad (1.1)$$

where d_{TV} denotes the TV distance; see (2.1) for details. As a result we have

$$R(\phi) \geq 1 - d_{\text{TV}}(\mu, \nu)$$

for any $\phi \in \Phi$, $\mu \in \text{co}(\mathcal{P})$, and $\nu \in \text{co}(\mathcal{Q})$. Minimizing over ϕ and maximizing over μ and ν then yields the “weak duality” inequality

$$R(\mathcal{P}, \mathcal{Q}) \geq 1 - d_{\text{TV}}(\text{co}(\mathcal{P}), \text{co}(\mathcal{Q})), \quad (1.2)$$

where the TV distance between any sets S, T of (finitely or countably additive) probability measures is given by $d_{\text{TV}}(S, T) = \inf_{\mu \in S, \nu \in T} d_{\text{TV}}(\mu, \nu)$. To foreshadow later developments, we note that the definition of TV distance carries over without change to (sets of) *finitely additive* probability measures.

Because the distance between two sets is always the same as the distance between their closures, we could replace $\text{co}(\mathcal{P})$ and $\text{co}(\mathcal{Q})$ by their TV closures $\overline{\text{co}}(\mathcal{P})$ and $\overline{\text{co}}(\mathcal{Q})$ in (1.2). Crediting Lucien Le Cam, Kraft (1955) presented the following fundamental result, which is a sufficient condition for equality to hold in (1.2).

Theorem 1.1. *If $\mathcal{P}$ and $\mathcal{Q}$ have a common dominating measure γ (meaning that $\mu \ll \gamma$ and $\nu \ll \gamma$ for all $\mu \in \mathcal{P}, \nu \in \mathcal{Q}$), then for any $\varepsilon \geq 0$,*

$$\exists \text{ test } \phi \in \Phi: \inf_{\nu \in \mathcal{Q}} \mathbb{E}_\nu[\phi] > \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi] + \varepsilon \iff d_{\text{TV}}(\overline{\text{co}}(\mathcal{P}), \overline{\text{co}}(\mathcal{Q})) > \varepsilon. \quad (1.3)$$

In fact, we have

$$R(\mathcal{P}, \mathcal{Q}) = 1 - d_{\text{TV}}(\overline{\text{co}}(\mathcal{P}), \overline{\text{co}}(\mathcal{Q})) \quad (1.4)$$

where the minimax risk is achieved by some $\phi^ \in \Phi$, but the infimum in d_{TV} is in general not achieved by any $\mu \in \overline{\text{co}}(\mathcal{P}), \nu \in \overline{\text{co}}(\mathcal{Q})$.*

In words, if $\mathcal{P}$ and $\mathcal{Q}$ have a common dominating measure, then a nontrivial test exists if and only if the convex hulls of $\mathcal{P}$ and $\mathcal{Q}$ are separated in TV distance. Actually, [Kraft \(1955\)](#) presented only (1.3) and proved it using the Hahn–Banach separating hyperplane theorem. A statement like (1.4) appears on page 476 of [Le Cam \(2012\)](#) and was proved succinctly using a minimax theorem, but the left hand side there is subtly (yet critically!) different; we return to Le Cam’s claim in Theorem 1.11 and later again in Appendix B. Also using a Hahn–Banach argument, Proposition 1.6 shows that Theorem 1.1 is a corollary of our own more general Theorem 1.5, itself proved using a minimax theorem.

Example 1.2. *The infimum TV distance is not achieved in general in Theorem 1.1. To see this, let $\mathcal{Q} = \{\delta_0\}$ and $\mathcal{P} = \{(\frac{1}{2} - \frac{1}{n})\delta_0 + (\frac{1}{2} + \frac{1}{n})\delta_n: n \in \mathbb{N}\}$. Then $d_{\text{TV}}(\text{co}(\mathcal{P}), \delta_0) = 1/2$, but this is not achieved by any $\mu \in \overline{\text{co}}(\mathcal{P})$ because there is no such μ with mass exactly 1/2 at 0. Said differently, the TV projection of a measure onto $\overline{\text{co}}(\mathcal{P})$ may not exist. (The TV projection will always exist onto $\overline{\text{co}}^*(\mathcal{P})$, an object we introduce later.)*

The assumption of a common dominating measure made in the above theorem appears innocuous but is violated in many standard nonparametric testing problems. For example, here are commonly encountered classes which do not have a dominating measure:

- (i) distributions on a bounded support (like $[0, 1]$) with a specified mean (like 0.5);
- (ii) all symmetric distributions around the origin;
- (iii) all (unbounded) distributions with a specified mean and bound on the variance;
- (iv) all σ -sub-Gaussian distributions for a given $\sigma > 0$;
- (v) the TV, Wasserstein, or Kolmogorov ball around a given distribution.

In all these cases, Theorem 1.1 is silent. The importance of studying such settings has not escaped statisticians. For example, [Huber and Strassen \(1973\)](#) remark at the end of their Section 6 that “...the more interesting sets $\mathcal{P}, \mathcal{Q}$ are not dominated...”. Here is a very simple example from [van der Vaart \(2002\)](#) where the assumption of Theorem 1.1 fails and its conclusion also fails.

Example 1.3. Let $\mathcal{P} = \{\delta_x : x \in [0, 1]\}$, where δ_x is the Dirac delta mass at x , and let $\mathcal{Q} = \{\text{Uni}[0, 1]\}$, where $\text{Uni}[0, 1]$ denotes the uniform distribution on $[0, 1]$. The TV distance between $\text{Uni}[0, 1]$ and $\overline{\text{co}}(\mathcal{P})$ equals 1, the largest possible value. However, as should be intuitively clear, and also follows as a consequence of our main theorem below, no nontrivial test exists for this problem, and the minimax risk is $R(\mathcal{P}, \mathcal{Q}) = 1$.¹ (See also Proposition C.1.)

After seeing the above example, it may be tempting to conjecture that when there is no reference measure, the closure in TV may need to be replaced by the *standard weak closure* (i.e. closure in the standard topology of weak convergence of probability measures). Indeed in the above example, the weak closure of $\mathcal{P}$ contains $\text{Uni}[0, 1]$. However, the following simple counterexample shows that using the weak closure of the convex hull is also in general the wrong answer.

Example 1.4. Take $\Omega = [0, 1]$, $\mathcal{P} = \{\delta_x : x \in [0, 0.5]\}$, and $\mathcal{Q} = \{\delta_y : y \in (0.5, 1]\}$. Then a perfect test exists: $\mathbf{1}_{\{X > 0.5\}}$ has level 0 and power 1. However, the standard weak closures of both $\mathcal{P}$ and $\mathcal{Q}$ include $\delta_{0.5}$. In short, perfect tests exist even though the weak closures of the null and alternatives intersect, so (1.3) could not have used the weak closures of $\text{co}(\mathcal{P})$ and $\text{co}(\mathcal{Q})$, respectively. Once more, our theorem will lead to the right conclusion.

In the first example above, the TV closure was not large enough. In the second example, the standard weak closure was too large. Our main result employs a type of closure that lies in between (in strength) the standard weak and strong closures: we take the closure in the *weak-* topology* on $\mathbf{ba}$, the space of bounded finitely additive measures. The formal definitions are reviewed in Section 2, but let us give some intuition here. Suppose a test ϕ has risk $R(\phi) = r$, so that $\mathbb{E}_\mu[\phi] + \mathbb{E}_\nu[1 - \phi] \leq r$ for all $\mu \in \mathcal{P}$ and $\nu \in \mathcal{Q}$, and r is the smallest number with this property. This inequality remains true for all μ, ν in the convex hulls $\text{co}(\mathcal{P}), \text{co}(\mathcal{Q})$, and then for all μ, ν in the *closures* of the convex hulls—assuming the closures are taken in a topology for which the expectation $\mathbb{E}_\mu[\phi]$ is continuous in μ . The significance of the weak-* topology on $\mathbf{ba}$ is that it is the *weakest* topology with this property for any test ϕ . Thus the weak-* convex closures $\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})$ are the *largest* extensions of $\mathcal{P}, \mathcal{Q}$ that preserve the risk of any test ϕ . It is the distance between these closures that matters for the minimal achievable risk. While this motivates the use of the weak-* topology, it does not explain why the space $\mathbf{ba}$ is needed. This is indeed a more subtle matter, which we address further in Section 2.

We now state our main result giving a quantitative characterization of testable hypotheses in the spirit of Le Cam and Kraft. (We recall later in (2.1) how the TV distance is defined in $\mathbf{ba}$.)

¹Note that, as stated, in this example any test ϕ maps a *single observation* to $[0, 1]$. Letting $\mathcal{P} = \{\delta_x \times \delta_x : x \in [0, 1]\}$ and $\mathcal{Q} = \{\text{Uni}[0, 1] \times \text{Uni}[0, 1]\}$, one can design a perfect test that has zero type-I error and power one (the test rejects if the two observations are different). While our setup allows the underlying Ω and distributions to be arbitrary, and covers the case of $\mathcal{P}$ and $\mathcal{Q}$ consisting of product distributions as a special case, the purpose of the example is to show that Theorem 1.1 can fail in nondominated settings.

Theorem 1.5. *Let $\mathcal{P}$ and $\mathcal{Q}$ be arbitrary nonempty subsets of $\mathcal{M}_1$, and let $\varepsilon \geq 0$. Then we have the following equivalence:*

$$\exists \text{ test } \phi : \inf_{\nu \in \mathcal{Q}} \mathbb{E}_\nu[\phi] > \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi] + \varepsilon \iff d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})) > \varepsilon.$$

In fact, we have

$$R(\mathcal{P}, \mathcal{Q}) = 1 - d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})), \quad (1.5)$$

where the infimum in the d_{TV} is achieved by some $\mu \in \overline{\text{co}}^(\mathcal{P}), \nu \in \overline{\text{co}}^*(\mathcal{Q})$.*

We suspect that outside of special cases, such as the dominated setting of Theorem 1.1, the minimax risk may not in general be achieved by any test ϕ .

The proof of the above result is based on a minimax theorem, as detailed in the following section. We note that the choice of topology is a critical and subtle matter in this regard, because it affects both the compactness of sets and continuity of functions, required for invoking an appropriate minimax theorem. In Appendix A we expand briefly on this, noting that the established literature contains instances where minimax theorems are applied incorrectly.

While Theorem 1.5 holds under weaker assumptions, it is not immediately apparent that it implies Theorem 1.1. The following result clarifies this.

Proposition 1.6. *If $\mathcal{P}$ and $\mathcal{Q}$ admit a common reference measure $\gamma \in \mathcal{M}_1$, then*

$$d_{\text{TV}}(\overline{\text{co}}(\mathcal{P}), \overline{\text{co}}(\mathcal{Q})) = d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})).$$

Moreover, the minimax risk is always attained by a minimax optimal test, so Theorem 1.5 implies Theorem 1.1.

The above result holds despite the fact that $\overline{\text{co}}(\mathcal{P}) \neq \overline{\text{co}}^*(\mathcal{P})$ in general. To see that the two closures can be different, we an example where $\mathcal{P} = \{\delta_n : n \in \mathbb{N}\}$ and note that $\overline{\text{co}}^*(\mathcal{P})$ includes a purely finitely additive measure that is obtained as a limit point of the sequence of point masses δ_n . We denote this limit point by δ_∞ because it is “supported at infinity”, in that it assigns zero mass to any finite subset of $\mathbb{N}$. (We may also note this from Example 1.2, where $\overline{\text{co}}^*(\mathcal{P})$ contains an equal mixture of δ_0 and δ_∞ .)

Thus, Theorem 1.5 appears to be fundamental since it holds without any (statistically relevant) restrictions, and is the main contribution of this paper. We also present another result that is of independent interest, because it involves the usual weak closure and does not involve the weak-* closure in ba .

Let C_b denote the set of continuous bounded functions on Ω , which in the following result is a metric space.

Theorem 1.7. *Assume Ω is a metric space and $\mathcal{F}$ the Borel σ -algebra. If $\mathcal{P}$ and $\mathcal{Q}$ are convex and weakly compact (in the usual sense), then*

$$R(\mathcal{P}, \mathcal{Q}) = \inf_{\phi \in \Phi \cap C_b} R(\phi) = 1 - d_{\text{TV}}(\mathcal{P}, \mathcal{Q}) = 1 - d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})).$$

In particular, (1.3) holds and the test ϕ can be taken to be continuous therein, and the infimum in the total-variation distance d_{TV} is achieved by some $\mu \in \mathcal{P}, \nu \in \mathcal{Q}$.

Note that above, $\mathcal{P}$ and $\mathcal{Q}$ need not have a reference measure, and so Theorem 1.7 itself is incomparable to Theorem 1.1.

Remark 1.8. In the context of Theorem 1.7, let us first note that in general, $\mathcal{P} \subsetneq \overline{\text{co}}^*(\mathcal{P})$. Consider $\Omega = [-1, 1]$ and $\mathcal{P}$ to be the set of all symmetric laws around 0, which is convex and weakly compact. Now, consider the sequence $\mu_n = \text{Uni}[-1/n, 1/n]$, which converges weakly to δ_0 , but does not converge to δ_0 in the weak-* topology. (Indeed, the test function $\phi(x) = 0$ for $x \neq 0$ and $\phi(x) = 1$ for $x = 0$ satisfies $\mathbb{E}_{\mu_n}[\phi] = 0$, but $\mathbb{E}_{\delta_0}[\phi] = 1$, and such discontinuous test functions are permissible in the weak-* topology.) Instead, the sequence (μ_n) has a subnet that converges to an element in ba , showing that $\mathcal{P} \subsetneq \overline{\text{co}}^*(\mathcal{P})$ as claimed. Despite this, the TV distance in the above theorem is unaffected by taking the closure in ba because in a metric space with the Borel sigma algebra, $d_{\text{TV}}(\mu, \nu) = \sup_{\phi \in \Phi \cap C_b} (\mathbb{E}_{\mu}[\phi] - \mathbb{E}_{\nu}[\phi])$; see Lemma 2.3. That is, we only need to consider *continuous* $[0, 1]$ -bounded functions to compute the TV distance. In fact, Lemma 2.3 is the only reason we need a metric, rather than merely a topology, in Theorem 1.7.

The weak-* closure is always at least as large as the TV closure, so the right-hand side of (1.5) is always at least as large as the right-hand side of (1.4). Moreover, there exist problems (like Example 1.3) where

$$R(\mathcal{P}, \mathcal{Q}) > 1 - d_{\text{TV}}(\overline{\text{co}}(\mathcal{P}), \overline{\text{co}}(\mathcal{Q})).$$

Hence, we see that (1.4) is overly optimistic about the minimax risk, in the sense that problems without a reference measure can be harder than what (1.4) may suggest. Even though we have not shown that a minimax optimal test always exists, the following result is still practically useful to verify that a candidate test is, in fact, minimax optimal.

Corollary 1.9. *For any $\mathcal{P}, \mathcal{Q} \subset \mathcal{M}_1$, and any $\phi \in \Phi, \mu \in \overline{\text{co}}^*(\mathcal{P}), \nu \in \overline{\text{co}}^*(\mathcal{Q})$, one has the weak duality inequality*

$$R(\phi) \geq 1 - d_{\text{TV}}(\mu, \nu).$$

Thus, if one can exhibit $\phi^ \in \Phi, \mu^* \in \overline{\text{co}}^*(\mathcal{P}), \nu^* \in \overline{\text{co}}^*(\mathcal{Q})$ such that*

$$R(\phi^*) = 1 - d_{\text{TV}}(\mu^*, \nu^*),$$

then ϕ^ is minimax optimal, and μ^*, ν^* are the TV-closest pair between $\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})$.*

Proof. The first display is an immediate consequence of (1.5). The second display implies

$$R(\mathcal{P}, \mathcal{Q}) \leq R(\phi^*) = 1 - d_{\text{TV}}(\mu^*, \nu^*) \leq 1 - d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})),$$

but since the first and last quantities are equal by (1.5), equality must hold throughout. $\square$

An interesting nonparametric example where the above corollary proves particularly useful is the following.

Example 1.10. *Let $\mathcal{P}$ be all probability measures on $[0, 1]$ whose mean is at most m_1 , and let $\mathcal{Q}$ be those with mean at least $m_2 > m_1$. These sets do not have a dominating reference measure, but are weakly compact and convex, and so Theorem 1.7 applies, but we can say more. Consider the test $\phi^*(x) = x$ and let μ^*, ν^* be Bernoullis with means m_1, m_2 . Then $R(\phi^*)$ is easily seen to equal $m_1 + (1 - m_2)$, which equals $1 - d_{TV}(\mu^*, \nu^*)$. Thus, Corollary 1.9 implies that ϕ^* is minimax optimal and the two Bernoullis are the TV-closest pair. Thus, here, the minimax optimal test exists, despite Theorem 1.7 not guaranteeing its existence.*

The last result that we wish to highlight is that of Huber and Strassen (1973), who prove that if the Choquet capacities induced by $\mathcal{P}, \mathcal{Q}$ are *two-alternating*, then a *least favorable distribution pair* exists, such that a likelihood ratio test between them is minimax optimal. This result applies to both parametric settings which may not be compact (e.g., testing unit variance Gaussians with nonpositive mean against those with mean larger than one), as well as nonparametric settings without a reference measure (e.g., when the null and alternative are each described by a total-variation ball around a simple hypothesis). Nevertheless, it does not yield a necessary and sufficient condition for testability as Theorem 1.5 does, just a (very useful) sufficient one.

As a final note, we highlight that Le Cam appeared to be fully aware of the deficiencies of Theorem 1.1. However, his preferred solution was highly nonstandard. Notice that Theorem 1.5 alters the right hand side of (1.3) but leaves the left hand side unchanged. Le Cam's book (2012, pg. 476) presents a statement that changes the left hand side of (1.3), but leaves the right hand side unchanged. However, while mathematically correct, its statistical implications are limited, as also pointed out by van der Vaart (2002). We nevertheless provide it below for completeness.

Write $\mathcal{M}$ for the set of finite signed countably additive measures on Ω , which is a Banach space when equipped with the TV-norm, and let $\mathcal{M}'$ denote its dual. A *generalized test* is an element $\phi \in \mathcal{M}'$ such that $\phi(\mu) \in [0, 1]$ for all $\mu \in \mathcal{M}_1$. (The reader is warned that Le Cam just called these *tests*, which may lead to confusion unless one very carefully follows his definitions.)

Theorem 1.11 (Le Cam (2012)). *Let $\mathcal{P}$ and $\mathcal{Q}$ be arbitrary nonempty subsets of $\mathcal{M}_1$, and let $\varepsilon \geq 0$. Then we have the following equivalence:*

$$\exists \text{ generalized test } \phi : \inf_{\nu \in \mathcal{Q}} \phi(\nu) > \sup_{\mu \in \mathcal{P}} \phi(\mu) + \varepsilon \iff d_{TV}(\text{co}(\mathcal{P}), \text{co}(\mathcal{Q})) > \varepsilon.$$

As Le Cam had also noted, these generalized tests may not correspond to bounded measurable functions, and thus need not be possible to evaluate using observed samples.

The tests ϕ that we consider also live in $\mathcal{M}'$, but are in fact given by bounded measurable functions (by definition). Thus the above result, while correct, has limited practical applicability, unlike our main Theorem 1.5. We discuss this point in detail in Appendix B.

Once more, Le Cam was aware of this drawback of his mathematically general (but statistically uninterpretable) formulation in Theorem 1.11. In a remark following Lemma 1 on page 476 of Le Cam (2012), he even suggested a way of only working with tests that are bounded measurable functions. Following his sentence “*One case which occurs often in the study of robust procedures is as follows...*”, he goes on to describe a proof strategy for a result similar to Theorem 1.5. However, unfortunately, Le Cam only treats this in a passing remark, that is easily missed (or unappreciated) due to its abstractness. A clearly stated theorem, which actually holds with no unnecessary restrictions on the tests or hypotheses (like Theorem 1.5), was not provided. It remains unclear to us if Le Cam did not realize that what he was proposing in that remark were the seeds of a complete solution, or whether he realized it but simply chose not to emphasize it (despite its apparent importance) for reasons of personal preference. Such ponderings aside, we believe it to be the case that most statisticians, even after parsing his remark, would not be able to clearly answer the question posed at the start of the abstract, motivating us to write this short paper to settle the issue.

We end the introduction by noting that there is much precedent for the use of finitely additive measures in statistics, dating back nearly a century. An early proponent was de Finetti (circa 1930), whose push for finite additivity in subjective Bayesian probability continues to be the topic of much discussion today; see Bingham (2010); Regazzini (2013); Seidenfeld et al. (2025) and references therein. Another prominent book to employ finite additivity is that of Dubins and Savage (2014), who claim that many of their theorems are more general and easier to prove when working with finitely additive measures.

We find it critical to point out that this current paper is very different from the above in its use of finitely additive measures. Our paper does not seek to promote or defend the use of finite additivity as an axiom or mathematical tool of convenience. We do not take a position on this issue. Our mathematical formalism and setup, from questions asked to assumptions made, are about countably additive probability measures, as the currently most prevalent model in probability and statistics. We did make a significant effort stay in the realm of countable additivity, but found that there is no escape: our matching necessary and sufficient condition for the existence of tests arbitrarily close to minimax optimality necessarily (i.e. in general, outside of special cases) must involve finitely additive measures. Examples exist to back up this rather strong claim; see e.g. Example 3.5. Thus finite additivity appears as a fundamental and unavoidable mathematical consequence of providing a complete answer to a question about countable probability, very different from the aforementioned works of de Finetti, Dubins–Savage, and others.

The rest of this paper is organized as follows. Section 2 presents the mathematical background to understand the *weak-* closure in ba*, and the proofs of all the results stated above. Section 3 discusses the interesting case of a singleton $\mathcal{Q} = \{\nu\}$ and its relation to

recent work. Appendix A elaborates on some subtle issues related to the choice of topology. Appendix B discusses Le Cam’s generalized tests in more detail. Appendix C contains some auxiliary technical results. Finally, Appendix D proves Theorem 3.4, which is a result relating the closed convex hull in $\mathbf{ba}$ to the *bipolar* of $\mathcal{P}$, another fundamental object from recent work.

2 Finitely additive measures and proofs of the main results

Our goal in this section is prove the results presented in Section 1 and, in the process of doing so, review the required mathematical machinery and elucidate the role of finitely additive measures. General references on this topic include Yosida and Hewitt (1952); Bhaskara Rao and Bhaskara Rao (1983). Recall that we have fixed a measurable space $(\Omega, \mathcal{F})$, and we let $\mathcal{L}$ denote the set of bounded measurable functions on this space.

Definition 2.1. A *bounded finitely additive measure* is a real-valued set function μ on $\mathcal{F}$ that is (i) *additive*, $\mu(A \cup B) = \mu(A) + \mu(B)$ for all disjoint sets $A, B \in \mathcal{F}$, and (ii) *bounded* in total-variation norm, $\|\mu\|_{\text{TV}} = \sup_{A \in \mathcal{F}} |\mu(A)| < \infty$. The space of all such μ is denoted by $\mathbf{ba}$, the subset taking nonnegative values by $\mathbf{ba}_+$, and the further subset with unit mass by $\mathbf{ba}_1$. Thus $\mathbf{ba}_1$ is the space of finitely additive probability measures, also known as *probability charges*.

Since countable additivity implies finite additivity, one has $\mathcal{M}_1 \subset \mathbf{ba}_1$, and the inclusion is strict as soon as $\mathcal{F}$ is infinite. Every $\mu \in \mathbf{ba}_+$ admits a unique decomposition $\mu = \mu_c + \mu_p$ into a countably additive part $\mu_c \in \mathcal{M}_+$ and a *purely finitely additive* part $\mu_p \in \mathbf{ba}_+$. That μ_p is purely finitely additive means that any countably additive ν such that $0 \leq \nu \leq \mu_p$ in the setwise order must be zero.

The integral $\int_{\Omega} f d\mu$, or $\mathbb{E}_{\mu}[f]$ in probabilistic notation, is defined for any $f \in \mathcal{L}$ and any $\mu \in \mathbf{ba}$, and is linear in f and μ . Limit theorems such as the dominated and monotone convergence theorems do not apply, but one does have the bound $|\mathbb{E}_{\mu}[f]| \leq \|f\|_{\infty} \|\mu\|_{\text{TV}}$, where $\|f\|_{\infty} = \sup_{\omega \in \Omega} |f(\omega)|$ is the supremum norm. This leads to the identification of $\mathbf{ba}$ as the dual of the Banach space $(\mathcal{L}, \|\cdot\|_{\infty})$. Like the dual of any Banach space, $\mathbf{ba}$ can be equipped with the *weak-* topology* $\sigma(\mathbf{ba}, \mathcal{L})$, defined as the weakest topology such that the linear functionals $\mu \mapsto \mathbb{E}_{\mu}[f]$ are continuous for all $f \in \mathcal{L}$. The Banach–Alaoglu theorem implies that the set $\mathbf{ba}_1$ is weak-* compact, and then so is any weak-* closed subset of $\mathbf{ba}_1$. In particular, this is the case for $\overline{\text{co}}^*(\mathcal{P})$ and $\overline{\text{co}}^*(\mathcal{Q})$. It is this compactness property under the weak-* topology that explains the appearance of the space $\mathbf{ba}$ in Theorem 1.5.

The compactness is crucial because the proof of Theorem 1.5 relies on the following special case of a minimax theorem due to Fan (1953, Theorem 2).

Theorem 2.2. *Let X be a compact convex subset of a Hausdorff topological vector space, and Y a convex subset of a vector space. Let F be a real-valued function on $X \times Y$ such*

that $F(x, y)$ is lower semicontinuous and convex in x for each fixed $y \in Y$, and concave in y for each fixed $x \in X$. Then

$$\inf_{x \in X} \sup_{y \in Y} F(x, y) = \sup_{y \in Y} \inf_{x \in X} F(x, y),$$

where the infima are attained thanks to compactness of X and lower semicontinuity in x of $F(x, y)$ and $\sup_{y \in Y} F(x, y)$.

The total-variation norm on $\mathbf{ba}$ has the representation $\|\mu\|_{\text{TV}} = \sup_{f \in B_1} \mathbb{E}_\mu[f]$, where B_1 is the unit ball in $\mathcal{L}$. It is customary to define the total-variation distance $d_{\text{TV}}(\mu, \nu)$ between two probability charges μ, ν as *half* the total-variation norm of $\mu - \nu$. Because μ, ν have unit mass, and because $B_1 = 2\Phi - 1$, meaning that $f \in B_1$ if and only if $f = 2\phi - 1$ for some test ϕ , one obtains the representation

$$d_{\text{TV}}(\mu, \nu) = \frac{1}{2} \|\mu - \nu\|_{\text{TV}} = \sup_{\phi \in \Phi} (\mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi]), \quad \mu, \nu \in \mathbf{ba}_1. \quad (2.1)$$

With these preliminaries in hand, the proof of Theorem 1.5 is straightforward.

Proof of Theorem 1.5. We will apply Fan's minimax theorem with $X = \overline{\text{co}}^*(\mathcal{P}) \times \overline{\text{co}}^*(\mathcal{Q})$, $Y = \Phi$, and $F: ((\mu, \nu), \phi) \mapsto \mathbb{E}_\nu[\phi] - \mathbb{E}_\mu[\phi]$. Indeed, X is a compact convex subset of the product space $\mathbf{ba} \times \mathbf{ba}$, where each factor has the weak-* topology, Y is a convex subset of $\mathcal{L}$, and F has the required concavity and convexity properties and is continuous in (μ, ν) by definition of the weak-* topology. Using first Fan's minimax theorem, then linearity and weak-* continuity of $\mu \mapsto \mathbb{E}_\mu[\phi]$, and finally a simple algebraic manipulation, we have

$$\begin{aligned} \inf_{\substack{\mu \in \overline{\text{co}}^*(\mathcal{P}) \\ \nu \in \overline{\text{co}}^*(\mathcal{Q})}} \sup_{\phi \in \Phi} (\mathbb{E}_\nu[\phi] - \mathbb{E}_\mu[\phi]) &= \sup_{\phi \in \Phi} \inf_{\substack{\mu \in \overline{\text{co}}^*(\mathcal{P}) \\ \nu \in \overline{\text{co}}^*(\mathcal{Q})}} (\mathbb{E}_\nu[\phi] - \mathbb{E}_\mu[\phi]) \\ &= \sup_{\phi \in \Phi} \inf_{\substack{\mu \in \mathcal{P} \\ \nu \in \mathcal{Q}}} (\mathbb{E}_\nu[\phi] - \mathbb{E}_\mu[\phi]) \\ &= 1 - \inf_{\phi \in \Phi} \sup_{\substack{\mu \in \mathcal{P} \\ \nu \in \mathcal{Q}}} (\mathbb{E}_\mu[\phi] + \mathbb{E}_\nu[1 - \phi]), \end{aligned}$$

where the infimum on the left-hand side is attained by some $\mu^* \in \overline{\text{co}}^*(\mathcal{P})$ and $\nu^* \in \overline{\text{co}}^*(\mathcal{Q})$. This is the desired statement, because the left-hand side is equal to $d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q}))$ in view of (2.1), while the right-hand side is equal to $1 - R(\mathcal{P}, \mathcal{Q})$. $\square$

A similar, but different, calculation can be used to obtain Theorem 1.7. This calculation will make use of the assumption that $\mathcal{P}$ and $\mathcal{Q}$ are convex and weakly compact in the usual sense, i.e., in the topology $\sigma(\mathcal{M}_1, C_b)$. By considering *continuous* tests $\phi \in \Phi_c = \Phi \cap C_b$, the map $\mu \mapsto \mathbb{E}_\mu[\phi]$ remains continuous, and the minimax theorem still applies. Restricting to continuous tests raises the question of whether the representation in (2.1) for the total-variation distance remains valid with Φ_c in place of Φ . The following lemma confirms that in the metric space setting of Theorem 1.7 this is indeed so.

Lemma 2.3. *Assume Ω is a metric space with its Borel σ -algebra. Then for any $\mu, \nu \in \mathcal{M}_1$ we have $d_{\text{TV}}(\mu, \nu) = \sup_{\phi \in \Phi_c} (\mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi])$.*

Proof. The Hahn decomposition of the signed measure $\eta = \mu - \nu$ yields a measurable set A such that the positive part η^+ is concentrated on A and the negative part η^- is concentrated on A^c . The total-variation distance is then $d_{\text{TV}}(\mu, \nu) = \mu(A) - \nu(A) = \eta^+(A)$, the total mass of η^+ . Fix any closed sets $F_1 \subset A$ and $F_0 \subset A^c$. Urysohn's lemma, which applies in any metric space, yields a continuous function ϕ with $\phi = 1$ on F_1 , $\phi = 0$ on F_0 , and $0 \leq \phi \leq 1$. Thus

$$\mathbb{E}_\eta[\phi] \geq \eta^+(F_1) - \eta^-(A^c \setminus F_0) = \eta^+(F_1) + \eta^-(F_0) - \eta^-(A^c).$$

We conclude that

$$\sup_{\phi \in \Phi_c} (\mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi]) \geq \eta^+(F_1) + \eta^-(F_0) - \eta^-(A^c)$$

for all closed sets $F_1 \subset A$ and $F_0 \subset A^c$. Taking the supremum of the right-hand side over all such F_1 and F_0 yields $\eta^+(A)$, since positive finite measures on a metric space with the Borel σ -algebra are regular; see Billingsley (1999, Theorem 1.1). We have shown that $\sup_{\phi \in \Phi_c} (\mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi]) \geq d_{\text{TV}}(\mu, \nu)$, and the reverse inequality follows from the fact that $\mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi] \leq d_{\text{TV}}(\mu, \nu)$ for all $\phi \in \Phi$. $\square$

We now proceed with the proof of Theorem 1.7. Although the main calculation is very similar to the one in the proof of Theorem 1.5, we emphasize that it is not subsumed by it. Furthermore, Theorem 1.5 is used in the proof, so the two results are not independent. Recall that $\Phi_c = \Phi \cap C_b$ denotes the set of continuous tests.

Proof of Theorem 1.7. Using first Fan's minimax theorem, then an algebraic manipulation and the definition of $R(\phi, \mathcal{P}, \mathcal{Q})$, and finally relaxing the infimum to a larger set, we obtain

$$\begin{aligned} \inf_{\substack{\mu \in \mathcal{P} \\ \nu \in \mathcal{Q}}} \sup_{\phi \in \Phi_c} (\mathbb{E}_\nu[\phi] - \mathbb{E}_\mu[\phi]) &= \sup_{\phi \in \Phi_c} \inf_{\substack{\mu \in \mathcal{P} \\ \nu \in \mathcal{Q}}} (\mathbb{E}_\nu[\phi] - \mathbb{E}_\mu[\phi]) \\ &= 1 - \inf_{\phi \in \Phi_c} R(\phi, \mathcal{P}, \mathcal{Q}) \\ &\leq 1 - R(\mathcal{P}, \mathcal{Q}). \end{aligned}$$

The infimum on the left-hand side is attained and, in view of Lemma 2.3, is equal to $d_{\text{TV}}(\mathcal{P}, \mathcal{Q})$. By Theorem 1.5, the right-hand side equals $d_{\text{TV}}(\overline{\mathcal{P}}^*, \overline{\mathcal{Q}}^*)$, which is bounded above by $d_{\text{TV}}(\mathcal{P}, \mathcal{Q})$, the distance between two smaller sets. We have thus shown that

$$d_{\text{TV}}(\mathcal{P}, \mathcal{Q}) = 1 - \inf_{\phi \in \Phi_c} R(\phi, \mathcal{P}, \mathcal{Q}) \leq 1 - R(\mathcal{P}, \mathcal{Q}) \leq d_{\text{TV}}(\overline{\mathcal{P}}^*, \overline{\mathcal{Q}}^*) \leq d_{\text{TV}}(\mathcal{P}, \mathcal{Q}).$$

Thus there is equality throughout, and this completes the proof. $\square$

We end with the proof of Proposition 1.6. We could have used the minimax theorem here too, but this does not seem to simplify matters compared to the proof below, which relies on the Hahn–Banach theorem instead.

Proof of Proposition 1.6. We regard $\mathcal{P}$ and $\mathcal{Q}$ as subsets of $L^1 = L^1(\gamma)$ by identifying measures μ with their γ -densities $f_\mu = d\mu/d\gamma$. The total-variation norm is then the L^1 norm of the density, $\|\mu\|_{\text{TV}} = \|f_\mu\|_{L^1}$. Set $r = d_{\text{TV}}(\text{co}(\mathcal{P}), \text{co}(\mathcal{Q})) = \inf \frac{1}{2} \|f_\mu - f_\nu\|_{L^1}$, where the infimum extends over all $\mu \in \mathcal{P}$ and $\nu \in \mathcal{Q}$. Thus the convex set $\text{co}(\mathcal{P}) - \text{co}(\mathcal{Q})$ is disjoint from B_{2r} , the open ball of radius $2r$ in L^1 . The Hahn–Banach theorem (see, e.g., Schaefer and Wolff (1999, Theorem II.9.1)) then yields a unit-norm element $\psi \in L^\infty(\gamma)$ such that $\mathbb{E}_\gamma[\psi(f_\mu - f_\nu)] \geq 2r$ for all $\mu \in \text{co}(\mathcal{P})$, $\nu \in \text{co}(\mathcal{Q})$. We fix a version of ψ such that $|\psi| \leq 1$, and define the test $\phi = (1 + \psi)/2 \in \Phi$. It follows that $\mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi] \geq r$ for all μ, ν in $\text{co}(\mathcal{P})$ and $\text{co}(\mathcal{Q})$, and then in $\overline{\text{co}}^*(\mathcal{P})$ and $\overline{\text{co}}^*(\mathcal{Q})$ by weak-* continuity. We conclude that $d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})) \geq r$. On the other hand, the reverse inequality holds because r is the distance between the two smaller sets $\text{co}(\mathcal{P})$ and $\text{co}(\mathcal{Q})$.

It remains to confirm that the minimax risk is always attained in the dominated setting. To do so, observe that for any $\mu \ll \gamma$, the map $\phi \mapsto \mathbb{E}_\mu[\phi] = \mathbb{E}_\gamma[f_\mu \phi]$ is continuous in the weak-* topology of $L^\infty(\gamma)$. Thus the risk $R(\phi, \mathcal{P}, \mathcal{Q})$ is lower semicontinuous in this topology. Since the set of all (γ -equivalence classes of) tests is weak-* compact in $L^\infty(\gamma)$ thanks to Alaoglu’s theorem, there exists one that minimizes the risk. $\square$

3 The case of a singleton $\mathcal{Q}$

The case of a singleton $\mathcal{Q} = \{\nu\}$ is of particular interest in relation to the recent past work of Larsson et al. (2025). As we formalize below, a corollary of that work (not explicitly stated there, but stated and proved below) is that a nontrivial test exists if and only if ν does not lie in the effective null hypothesis $\mathcal{P}_{\text{eff}}$, defined below. A trivial corollary of our main theorem in this paper is that a nontrivial test exists if and only if ν does not lie in $\overline{\text{co}}^*(\mathcal{P})$. What exactly is the relationship between these two geometric objects appearing in the above statements? This section provides the answer.

We begin with some definitions. Given a set of distributions $\mathcal{P}$, an *e-variable* is a $[0, \infty]$ -valued random variable whose expectation is at most one under every $\mu \in \mathcal{P}$. The set of all e-variables for $\mathcal{P}$ is given by the *polar* set of $\mathcal{P}$,

$$\mathcal{E} = \{Z \geq 0 : \mathbb{E}_\mu[Z] \leq 1 \text{ for all } \mu \in \mathcal{P}\}.$$

Larsson et al. (2025) then define the polar of $\mathcal{E}$ (the *bipolar* of $\mathcal{P}$) as

$$\mathcal{P}_{\text{eff}} = \{\mu \geq 0 : \mathbb{E}_\mu[Z] \leq 1 \text{ for all } Z \in \mathcal{E}\},$$

and call it the *effective null hypothesis*. No test can have nontrivial power against any distribution in $\mathcal{P}_{\text{eff}}$. Indeed, as a corollary of their main theorem, we have the following.

Corollary 3.1. $\nu \in \mathcal{P}_{\text{eff}}$ if and only if for any test ϕ , $\mathbb{E}_\nu[\phi] \leq \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi]$.

Proof. For the forward direction, assume $\nu \in \mathcal{P}_{\text{eff}}$, and consider any test ϕ . If the constant $c = \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi]$ is strictly positive, then ϕ/c is an e-variable, and $\nu \in \mathcal{P}_{\text{eff}}$ implies that $\mathbb{E}_\nu[\phi/c] \leq 1$. If instead $c = 0$, then $\phi + 1$ is an e-variable and $\nu \in \mathcal{P}_{\text{eff}}$ implies $\mathbb{E}_\nu[\phi] = 0$.

For the reverse direction, assume $\nu \notin \mathcal{P}_{\text{eff}}$. Then there exists an e-variable Z such that $\mathbb{E}_\nu[Z] > 1$. By monotone convergence we may assume that $Z \leq n$ for some large $n \in \mathbb{N}$. Then $\phi = Z/n$ is a test, and we have $\mathbb{E}_\nu[\phi] > 1/n \geq \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi]$, a contradiction. $\square$

On the other hand, we also have the following immediate corollary of Theorem 1.5.

Corollary 3.2. $\nu \in \overline{\text{co}}^*(\mathcal{P})$ if and only if for any test ϕ , $\mathbb{E}_\nu[\phi] \leq \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu[\phi]$.

It is immediately clear that $\overline{\text{co}}^*(\mathcal{P}) \neq \mathcal{P}_{\text{eff}}$. For example, the zero measure is always in $\mathcal{P}_{\text{eff}}$ but never in $\overline{\text{co}}^*(\mathcal{P})$. However, it is also not the case that $\mathcal{P}_{\text{eff}}$ contains $\overline{\text{co}}^*(\mathcal{P})$: the latter may contain finitely additive elements that are absent from the former. The following theorem captures the relationship between $\overline{\text{co}}^*(\mathcal{P})$ and $\mathcal{P}_{\text{eff}}$ that makes both the above corollaries true.

Theorem 3.3. $\overline{\text{co}}^*(\mathcal{P}) \cap \mathcal{M}_1 = \mathcal{P}_{\text{eff}} \cap \mathcal{M}_1$.

The above theorem explains why, with their focus on a singleton $\mathcal{Q} = \{\nu\}$, Larsson et al. (2025) managed to avoid **ba**: the qualitative behavior of testability (i.e. the existence of a nontrivial test) does not require **ba**. Further, the quantitative statements of Larsson et al. (2025) in terms of infimum KL divergence from ν to $\mathcal{P}$ also avoided **ba**. However, even for a singleton $\mathcal{Q} = \{\nu\}$, the TV distances in Theorem 1.5 could not have used $\mathcal{P}_{\text{eff}}$ instead of $\overline{\text{co}}^*(\mathcal{P})$; we provide an example soon.

In order to prove the above theorem, we must characterize the relationship between $\mathcal{P}_{\text{eff}}$ and $\overline{\text{co}}^*(\mathcal{P})$. We do this in the following theorem, whose (rather lengthy) proof is presented in Appendix D. It immediately implies Theorem 3.3.

Theorem 3.4. Every $\mu \in \mathcal{P}_{\text{eff}}$ is setwise dominated by some maximal element $\mu' \in \mathcal{P}_{\text{eff}}$. The maximal element μ' is the countably additive part of some $\mu'' \in \overline{\text{co}}^*(\mathcal{P})$. As a consequence, $\mathcal{P}_{\text{eff}}$ is the solid hull of the countably additive parts of $\overline{\text{co}}^*(\mathcal{P})$.

We continue with a simple yet important example, which demonstrates that even in the case of a simple alternative $\mathcal{Q} = \{\nu\}$, in Theorem 1.5 it does not suffice to only consider $\overline{\text{co}}^*(\mathcal{P}) \cap \mathcal{M}_1$ (which restricts attention to countably additive probability measures, such as ν), but we indeed have to consider finitely additive measures in $\overline{\text{co}}^*(\mathcal{P})$.

This example has the subtle feature that it *assumes* that

$$\text{there is no diffuse probability measure on } \text{Pow}(\mathbb{R}), \quad (3.1)$$

where “Pow” denotes the power set (the set of all subsets). This assumption is discussed further in Remark 3.6 below; in particular, it is implied by the continuum hypothesis. The

upshot of that discussion is that the standard setting of probability theory allows for (3.1). Without further undesirable assumptions, situations like the one in the example below can occur and must be covered by the general theory. It is an interesting open question whether similar examples can be constructed without any assumptions such as (3.1).

Example 3.5. Let $\Omega = [-1, 1]$ and define $\mathcal{F} = \{A \cup B : A \in \text{Pow}([-1, 0]), B \in \text{Bor}([0, 1])\}$. Here “Pow” denotes the power set and “Bor” the Borel σ -algebra. Then $\mathcal{F}$ is the σ -algebra generated by the Borel subsets of $[0, 1]$ and all subsets of $[-1, 0]$. We assume that (3.1) holds. Consider

$$\mathcal{P} = \left\{ \frac{1}{2}\delta_x + \frac{1}{2}\delta_{-x} : x \in (0, 1) \right\}, \quad \mathcal{Q} = \{\nu\}, \quad \nu = \text{Uni}[0, 1].$$

We will show that

$$d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}) \cap \mathcal{M}_1, \mathcal{Q}) = 1 \quad \text{but} \quad d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \mathcal{Q}) = \frac{1}{2}. \quad (3.2)$$

This demonstrates that, in general, one cannot restrict attention to the countably additive elements of $\overline{\text{co}}^*(\mathcal{P})$ and $\overline{\text{co}}^*(\mathcal{Q})$ in Theorem 1.5. Put differently, in order for Theorem 1.5 to hold, the closure of $\text{co}(\mathcal{P})$ and $\text{co}(\mathcal{Q})$ cannot be taken in the space $\mathcal{M}_1$ with the topology $\sigma(\mathcal{M}_1, \mathcal{L})$. Finitely additive measures are really needed.

To argue (3.2), we first observe that every $\mu \in \overline{\text{co}}^*(\mathcal{P})$ is symmetric around zero in the sense that $\mu(A \cup (-A)) = 2\mu(A) = 2\mu(-A)$ for every $A \in \text{Bor}([0, 1])$, where $-A = \{-x : x \in A\}$. Indeed, this holds for all $\mu \in \mathcal{P}$, thus for all $\mu \in \text{co}(\mathcal{P})$, and finally for all $\mu \in \overline{\text{co}}^*(\mathcal{P})$ by weak-* continuity.

Consider now any $\mu \in \overline{\text{co}}^*(\mathcal{P}) \cap \mathcal{M}_1$. Its restriction to $[-1, 0]$ is a positive countably additive measure on $\text{Pow}([-1, 0])$, and is therefore concentrated on a countable set $S_0 \subset [-1, 0]$ thanks to Proposition C.3. The symmetry property applied with $A = -S_0$ then shows that μ is concentrated on $S = S_0 \cup (-S_0)$. But since any countable set is a ν -nullset, we have $d_{\text{TV}}(\mu, \nu) \geq \mu(S) - \nu(S) = 1$, and this yields the first part of (3.2).

We now focus on the second part of (3.2). The test $\phi = \mathbf{1}_{[0, 1]}$ certifies that $d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \mathcal{Q})$ is at least $\frac{1}{2}$, and the point is to show that this is achieved for some $\mu \in \overline{\text{co}}^*(\mathcal{P})$. To do this we start with a net (η_α) of purely atomic measures on $(0, 1)$ that converges to ν ; this is possible thanks to Proposition C.1. The symmetrizations $\mu_\alpha(A \cup B) = \frac{1}{2}(\eta_\alpha(-A) + \eta_\alpha(B))$ for $A \in \text{Pow}([-1, 0])$ and $B \in \text{Bor}([0, 1])$ then belong to $\text{co}(\mathcal{P})$. On passing to a convergent subnet, we may assume that (μ_α) converges to some limit $\mu \in \overline{\text{co}}^*(\mathcal{P})$. The restriction of μ to $[0, 1]$ is equal to $\nu/2$ by construction, and it follows that $d_{\text{TV}}(\mu, \nu) = \frac{1}{2}$. This completes the proof of the second part of (3.2).

It is worth noting that the restriction of μ to $[-1, 0]$ must be purely finitely additive. Indeed, its countably additive part is atomless (due to the symmetry property and the fact that on $[0, 1]$, μ is proportional to the uniform distribution) and therefore zero due to Proposition C.3.

Remark 3.6. The assumption (3.1) is known to be consistent with the Zermelo–Fraenkel set theory with the axiom of choice (ZFC). One may thus add (3.1) as an axiom without rendering any theorems proved in ZFC false. Indeed, there are important models of ZFC, such as Gödel’s constructible universe L , where (3.1) can be proved as a theorem. Specifically, Gödel (1940) showed that the continuum hypothesis holds in L , and by a theorem of Banach and Kuratowski (1929, Théorème I), this implies (3.1). See also Theorem 13.20 and Theorem 10.1 in Jech (2003) for a textbook treatment. For these reasons, it seems sensible that our general theory should not depend on any axioms or hypotheses that disallow (3.1).

We end by discussing an important implication of our theorem for the existence of uniformly powered bounded e-variables for composite $\mathcal{Q}$. An e-variable Z for $\mathcal{P}$ is said to be uniformly powered against $\mathcal{Q}$ if $\inf_{\nu \in \mathcal{Q}} \mathbb{E}_{\nu}[Z] > 1$.

Corollary 3.7. *Let $\mathcal{P}$ and $\mathcal{Q}$ be nonempty subsets of $\mathcal{M}_1$. Then there exists a bounded e-variable for $\mathcal{P}$ that is uniformly powered against $\mathcal{Q}$ if and only if $d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})) > 0$.*

Proof. In the forward direction, let Z be the uniformly powered e-variable that lies in $[0, B]$ for some finite B . Then clearly, Z/B is a nontrivial test with type I error at most $1/B$ and worst-case power strictly greater than $1/B$, which implies that $d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})) > 0$ by Theorem 1.5, completing the first direction.

Assume now that $d_{\text{TV}}(\overline{\text{co}}^*(\mathcal{P}), \overline{\text{co}}^*(\mathcal{Q})) > 0$. By Theorem 1.5, there exists a test ϕ such that $\inf_{\nu \in \mathcal{Q}} \mathbb{E}_{\nu}[\phi] > \sup_{\mu \in \mathcal{P}} \mathbb{E}_{\mu}[\phi]$. If $\sup_{\mu \in \mathcal{P}} \mathbb{E}_{\mu}[\phi] = 0$, set $Z = 1 + \phi$. Otherwise, set $Z = \phi / \sup_{\mu \in \mathcal{P}} \mathbb{E}_{\mu}[\phi]$. Then Z is a bounded e-variable that is uniformly powered against $\mathcal{Q}$, completing the proof. $\square$

Following Ramdas and Wang (2025, Chapter 3), we say that an e-variable Z is uniformly *e-powered* against $\mathcal{Q}$ if $\inf_{\nu \in \mathcal{Q}} \mathbb{E}_{\nu}[\log Z] > 0$. Theorem 3.14 in the above book then proves that a uniformly powered bounded e-variable exists if and only if a uniformly e-powered bounded e-variable exists, and so the above theorem could have also been stated with “powered” being replaced by “e-powered”. This extends some results in the above book, and in Zhang et al. (2024).

Acknowledgments

We thank Sivaraman Balakrishnan for several useful discussions regarding Le Cam’s work. ML acknowledges support from the National Science Foundation under grant NSF DMS-2510965. AR acknowledges support from the Sloan Fellowship and the National Science Foundation under grant NSF DMS-2310718.

References

Stefan Banach and Kazimierz Kuratowski. Sur une généralisation du problème de la mesure. *Fundamenta Mathematicae*, 14(1):127–131, 1929. 15

- K. P. S. Bhaskara Rao and M. Bhaskara Rao. *Theory of Charges*, volume 109 of *Pure and Applied Mathematics*. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York, 1983. ISBN 0-12-095780-9. [9](#)
- Patrick Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York, second edition, 1999. ISBN 0-471-19745-9. doi: 10.1002/9780470316962. URL <https://doi.org/10.1002/9780470316962>. A Wiley-Interscience Publication. [11](#)
- NH Bingham. Finite additivity versus countable additivity: De Finetti and Savage. *Electronic Journal History of Probability and Statistics*, 6(1):1–38, 2010. [8](#)
- Lester E Dubins and Leonard J Savage. *How to gamble if you must: Inequalities for stochastic processes*. Courier Corporation, 2014. [8](#)
- Ky Fan. Minimax theorems. *Proc. Nat. Acad. Sci. U.S.A.*, 39:42–47, 1953. ISSN 0027-8424. doi: 10.1073/pnas.39.1.42. URL <https://doi.org/10.1073/pnas.39.1.42>. [9](#)
- D. H. Fremlin. *Measure Theory. Vol. 2*. Torres Fremlin, Colchester, 2003. ISBN 0-9538129-2-8. Broad foundations, Corrected second printing of the 2001 original. [20](#)
- Kurt Gödel. *The consistency of the continuum hypothesis*, volume No. 3 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1940. [15](#)
- Gökhan Gül and Abdelhak M Zoubir. Minimax robust hypothesis testing. *IEEE Transactions on Information Theory*, 63(9):5572–5587, 2017. [18](#)
- Peter J. Huber and Volker Strassen. Minimax tests and the Neyman-Pearson lemma for capacities. *Ann. Statist.*, 1:251–263, 1973. ISSN 0090-5364,2168-8966. URL [http://links.jstor.org/sici?sici=0090-5364\(197303\)1:2<251:MTATNL>2.0.CO;2-R&origin=MSN](http://links.jstor.org/sici?sici=0090-5364(197303)1:2<251:MTATNL>2.0.CO;2-R&origin=MSN). [3](#), [7](#)
- Thomas Jech. *Set Theory*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, millennium edition, 2003. ISBN 3-540-44085-2. [15](#)
- Charles Kraft. Some conditions for consistency and uniform consistency of statistical procedures. *University of California Publication in Statistics*, 2:125–141, 1955. [2](#), [3](#)
- Martin Larsson, Aaditya Ramdas, and Johannes Ruf. The numeraire e-variable and reverse information projection. *Annals of Statistics*, 2025. [12](#), [13](#)
- Lucien Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer Science & Business Media, 2012. [3](#), [7](#), [8](#), [19](#)
- Bernard C Levy. *Principles of Signal Detection and Parameter Estimation*. Springer Science & Business Media, 2008. [18](#)

- Aaditya Ramdas and Ruodu Wang. Hypothesis testing with e-values. *Foundations and Trends® in Statistics*, 1(1-2):1–390, 2025. 15
- Eugenio Regazzini. The origins of de Finetti’s critique of countable additivity. In *Advances in modern statistical theory and applications: a Festschrift in honor of Morris L. Eaton*, volume 10, pages 63–83. Institute of Mathematical Statistics, 2013. 8
- H. H. Schaefer and M. P. Wolff. *Topological Vector Spaces*, volume 3 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1999. ISBN 0-387-98726-6. doi: 10.1007/978-1-4612-1468-7. URL <https://doi.org/10.1007/978-1-4612-1468-7>. 12
- Teddy Seidenfeld, Joseph B Kadane, Mark J Schervish, and Rafael B Stern. Finite additivity, complete additivity, and the comparative principle. *Erkenntnis*, 90(5):1945–1968, 2025. 8
- Aad van der Vaart. The statistical work of Lucien Le Cam. *The Annals of Statistics*, 30(3):631–682, 2002. 3, 7
- Kôsaku Yosida and Edwin Hewitt. Finitely additive measures. *Trans. Amer. Math. Soc.*, 72:46–66, 1952. ISSN 0002-9947,1088-6850. doi: 10.2307/1990654. URL <https://doi.org/10.2307/1990654>. 9, 25
- Zhenyuan Zhang, Aaditya Ramdas, and Ruodu Wang. On the existence of powerful p-values and e-values for composite hypotheses. *The Annals of Statistics*, 52(5):2241–2267, 2024. 15

A Further comments on the choice of topology

In the introduction we showed why the TV topology and the weak convergence topology cannot, in general, be used to obtain characterizations of testability; see Examples 1.3–1.4 and the adjacent discussion. An alternative approach would be to look for topologies on the set Φ of all tests, making this set compact. This could then be leveraged to apply the minimax theorem without compactness of the closed convex hulls of $\mathcal{P}$ and $\mathcal{Q}$. This idea works in the dominated setting of Theorem 1.1, and Le Cam’s approach to embed Φ into the larger, compact, space of generalized tests can be viewed as an extension of this idea to the general case, with important drawbacks as discussed in Appendix B.

There are, however, topologies that make Φ itself compact in the general non-dominated setting. One prominent example is the *topology of pointwise convergence*. This is the weakest topology that makes the evaluation maps $\phi \mapsto \phi(\omega)$ continuous for every $\omega \in \Omega$. Equivalently, this is the product topology on $[0, 1]^\Omega$, which is compact by Tychonoff’s theorem. The trouble with this topology is that the linear functions $\phi \mapsto \mathbb{E}_\mu[\phi]$ fail to be continuous (or semicontinuous) in general.

For example, take $\Omega = [0, 1]$ and μ Lebesgue measure. For every finite subset $A = \{\omega_1, \dots, \omega_n\} \subset [0, 1]$, let $\phi_A = \mathbf{1}_A$ be the indicator function of A , which is a test. These tests form a net that converges to the constant test $\phi = 1$ in the topology of pointwise convergence. However, $\mathbb{E}_\mu[\phi_A] = 0$ for all A , while $\mathbb{E}_\mu[\phi] = 1$, showing that the corresponding expectations do not converge to the expectation of the limiting test. Thus, even though Φ is compact, the semicontinuity hypothesis of the minimax theorem is not satisfied, and the theorem cannot be used.

To restore (semi-)continuity, one has to work with a stronger topology. A very strong topology is the one induced by the supremum norm $\|\phi\|_\infty = \sup_{\omega \in \Omega} |\phi(\omega)|$. The above linear maps $\phi \mapsto \mathbb{E}_\mu[\phi]$ are now continuous thanks to the bounded convergence theorem. However, it is now compactness of Φ that fails, again rendering the minimax theorem inapplicable. Indeed, given any sequence ω_n of mutually distinct points, the tests $\phi_n = \mathbf{1}_{\{\omega_n\}}$ are at a distance $\|\phi_n - \phi_m\|_\infty = 1$ apart, and thus form a sequence with no convergent subsequence.

At this point it may be tempting to consider the topology $\sigma(\Phi, \mathcal{M})$, which is the *weakest* topology on Φ that makes the linear maps $\phi \mapsto \mathbb{E}_\mu[\phi]$ continuous for all $\mu \in \mathcal{M}$. Unfortunately, compactness of Φ still fails, and the net $\phi_A = \mathbf{1}_A$ indexed by finite subsets $A \subset \Omega = [0, 1]$ again furnishes an example. Indeed, suppose for contradiction that compactness holds. Then there is a convergent subnet ϕ_B with some limit ϕ . This limit is the constant function $\phi = 1$. To see this, note that for each $\omega \in [0, 1]$, $\phi_B(\omega) = 1$ eventually, so $\phi(\omega) = \mathbb{E}_{\delta_\omega}[\phi] = \lim_B \mathbb{E}_{\delta_\omega}[\phi_B] = 1$, where we used that $\lim_B \mathbb{E}_\mu[\phi_B] = \mathbb{E}_\mu[\phi]$ for every $\mu \in \mathcal{M}$ by definition of the topology. However, if we now take μ to be Lebesgue measure, we obtain the contradiction $0 = \lim_B \mathbb{E}_\mu[\phi_B] = \mathbb{E}_\mu[\phi] = 1$. This shows that ϕ_A has no convergent subnet, and hence that Φ is not compact.

One thus arrives at the conclusion that either Φ has to be “compactified” (for example by enlarging it to generalized tests as proposed by Le Cam, with its drawbacks discussed in Appendix B), or one has to find a way to compactify the space of measures rather than tests. The latter is the route taken in the present paper (by enlarging the space of measures), which leads us naturally — even necessarily — to working with finitely additive measures.

We hope that this discussion gives a further perspective on the delicate interplay between the choice of ambient space, its topology, and how these must be tailored to fit the confines of the minimax theorem. Unfortunately, we are aware of several published works which incorrectly use minimax theorems to achieve important conclusions. For example [Gül and Zoubir \(2017\)](#) mistakenly invoke Sion’s minimax theorem after using the product topology on $[0, 1]^\Omega$, which as we discussed above destroys continuity. As another example [Levy \(2008, Chapter 6\)](#) invokes von Neumann’s minimax theorem after endowing Φ with $\|\cdot\|_\infty$, which as we discussed above renders Φ non-compact. While we know of several more examples, it is not our goal to point out all such errors in other works, but merely to point to the subtleties involved with the careful choices necessary for everything to work out. It would also be of interest to see whether the corresponding conclusions in these works are incorrect, or if their proofs can be fixed using techniques developed in this paper.

B Generalized tests

In this section we elaborate on the distinction between Theorem 1.5 and Le Cam’s result on generalized tests, Theorem 1.11. After all, the two results lead to different conclusions regarding which hypotheses are testable. While this is not a contradiction mathematically (the two results make use of *tests* and *generalized tests*, respectively), it seems worthwhile to look closer at how they differ in order to understand the implications for applications.

It is useful to return to the basic example in the introduction, where $\Omega = [0, 1]$, $\mathcal{F}$ is the Borel σ -algebra, $\mathcal{P} = \{\delta_x: x \in [0, 1]\}$, and $\mathcal{Q} = \{\nu\}$ with $\nu = \text{Uni}[0, 1]$, the uniform distribution on $[0, 1]$. The weak-* convex closure $\overline{\text{co}}^*(\mathcal{P})$ contains *all* probability measures on $[0, 1]$, in particular ν , so Theorem 1.5 implies that $R(\mathcal{P}, \mathcal{Q}) = 1$. Thus there is no nontrivial test for this pair of hypotheses. On the other hand, we have $d_{\text{TV}}(\text{co}(\mathcal{P}), \nu) = 1$, so Le Cam’s Theorem 1.11 implies that there is a generalized test ϕ^* that perfectly separates $\mathcal{P}$ and $\mathcal{Q}$.

In fact, it is not hard to construct such a generalized test. For any $\mu \in \mathcal{M}$, let μ^a be its absolutely continuous part with respect to ν . The map $\mu \mapsto \mu^a$ is linear and continuous (even a contraction) in the total-variation norm. Now define $\phi^*(\mu) = \mu^a(\Omega)$, the total mass of μ^a . Then ϕ^* is a continuous linear functional on $\mathcal{M}$, and $\phi^*(\mu) \in [0, 1]$ for all $\mu \in \mathcal{M}_1$. That is, ϕ^* is a generalized test. Furthermore, $\phi^*(\nu) = 1$ and $\phi^*(\mu) = 0$ for all $\mu \in \mathcal{P}$.

The key issue here is that $\phi^*(\mu)$ may depend on the entire *distribution* μ , not just on the *sample* $\omega \in [0, 1]$ that is produced by μ . In this example, to compute the value of the generalized test it is not enough to know the realized sample ω —in fact, its value is irrelevant—but rather whether it came from an absolutely continuous distribution or not. This can be thought of as *additional information* that the experimenter would need to acquire in addition to the sample itself.

More generally, Theorem 1.5 shows that whenever no nontrivial test exists, but a nontrivial generalized test does exist, some form of additional information is *always* necessary in order to evaluate the generalized test.

Interestingly, as Le Cam points out in the remark on page 476 of Le Cam (2012), although generalized tests need not be representable by a measurable function ϕ on the original sample space Ω (i.e., by a test), it is always possible to extend $(\Omega, \mathcal{F})$ to a larger sample space $(\tilde{\Omega}, \tilde{\mathcal{F}})$ in such a way that every generalized test *can* be represented by a measurable function on $\tilde{\Omega}$. As such, this extension furnishes the additional data needed to evaluate generalized tests. However, since in applications the original space $(\Omega, \mathcal{F})$ typically models the information available to the experimenter, passing to the extension $(\tilde{\Omega}, \tilde{\mathcal{F}})$ may alter the information structure in an unintended way. The example above illustrates this quite clearly: here the sample ω needs to be accompanied with a “tag” indicating whether or not it was drawn from an absolutely continuous distribution.

Lastly, it may be of interest to know something about the structure of the extended sample space $(\tilde{\Omega}, \tilde{\mathcal{F}})$. The construction outlined in Chapter 1.6 of Le Cam (2012) is very general but also quite abstract and difficult to interpret. In the setting of the present paper,

there is a significantly less abstract construction which we now describe.

We start with the original measurable space $(\Omega, \mathcal{F})$ and its space $\mathcal{M}$ of finite signed measures. Zorn's lemma implies that there exists a set $\{\mu_i : i \in I\}$ of pairwise singular probability measures that is maximal in the collection of all such sets, ordered by inclusion. Every $\mu \in \mathcal{M}$ is then absolutely continuous with respect to some countable combination of the μ_i . Note that the index set I need not be countable. Indeed, I can only be countable if Ω is a countable set, in which case there is a single dominating measure γ and $\mathcal{M}$ can be identified with $L^1(\gamma)$. No extension of Ω is needed in this case, so the following construction is only relevant if Ω is uncountable.

For each $i \in I$ we let $(\Omega_i, \mathcal{F}_i)$ be a copy of $(\Omega, \mathcal{F})$ and consider the disjoint union

$$\tilde{\Omega} = \bigsqcup_{i \in I} \Omega_i,$$

equipped with the σ -algebra $\tilde{\mathcal{F}}$ consisting of all sets $A \subset \tilde{\Omega}$ such that $A \cap \Omega_i \in \mathcal{F}_i$ for all $i \in I$. We further let $\tilde{\mu}$ be the measure on $\tilde{\Omega}$ that restricts to μ_i on each Ω_i ,

$$\tilde{\mu}(A) = \sum_{i \in I} \mu_i(A \cap \Omega_i), \quad A \in \tilde{\mathcal{F}}.$$

If I is uncountable, which is the relevant case, the measure $\tilde{\mu}$ is not σ -finite. Nonetheless, $\tilde{\mathcal{F}}$ -measurable functions f can be integrated provided they are $\tilde{\mu}$ -integrable, meaning that the restriction $f_i = f|_{\Omega_i}$ belongs to $L^1(\mu_i)$ for each $i \in I$ and the norm $\|f\|_{L^1(\tilde{\mu})} = \sum_{i \in I} \|f_i\|_{L^1(\mu_i)}$ is finite. Note that this requires $f_i = 0$ for all but countably many $i \in I$. The space $L^1(\tilde{\mu})$ of (equivalence classes of) $\tilde{\mu}$ -integrable functions with the norm $\|\cdot\|_{L^1(\tilde{\mu})}$ is a Banach space, known as the ℓ^1 -direct sum of $L^1(\mu_i)$, $i \in I$.

The key point is that $L^1(\tilde{\mu})$ is isometrically isomorphic to the Banach space $(\mathcal{M}, \|\cdot\|_{\text{TV}})$. Indeed, any $\mu \in \mathcal{M}$ can be mapped to the unique $f \in L^1(\tilde{\mu})$ whose restriction f_i to Ω_i is the density $d\mu^a/d\mu_i$, where μ^a is the absolutely continuous part of μ with respect to μ_i . Note that μ^a depends on i , although this is not indicated in the notation, and all but countably many f_i are zero. The mapping $\mu \mapsto f$ is furthermore linear and surjective. Thanks to the maximality property of the family $\{\mu_i : i \in I\}$, it preserves the norm $\|\mu\|_{\text{TV}} = \|f\|_{L^1(\tilde{\mu})}$, and is thus an isometry.

In this way the elements of $\mathcal{M}$ can be regarded as $\tilde{\mu}$ -absolutely continuous measures on the (potentially much larger) sample space $\tilde{\Omega}$. Furthermore, the dual space $\mathcal{M}'$ can be identified with the dual of $L^1(\tilde{\mu})$, which is $L^\infty(\tilde{\mu})$. Thus any generalized test can be represented as a measurable function on $\tilde{\Omega}$, which was the goal of the construction. (The L^1 - L^∞ duality holds in the present non- σ -finite case because $\tilde{\mu}$ is *(strictly) localizable*; see [Fremlin \(2003, Chapters 21 and 24\)](#), specifically Definition 211E, Theorem 211L(d), and Theorem 243G.)

Let us clarify how this construction encodes the “additional information” mentioned previously. The observed sample $\tilde{\omega} \in \tilde{\Omega}$ belongs to Ω_i for exactly one i , so it must have

been generated by a distribution μ with nonzero density relative to μ_i . Thus i , which can be inferred from $\tilde{\omega}$, is the “tag” that conveys information about what might be called the “absolute continuity type” of the generating distribution μ . Since the spaces Ω_i are tagged, but otherwise identical, copies of Ω , the sample $\tilde{\omega}$ can indeed be regarded as a point of Ω , supplemented with the tag i . The generalized test, regarded as a bounded measurable function on $\tilde{\Omega}$, can now be evaluated as long this supplemented sample is observed. Furthermore, ordinary (not generalized) tests correspond to functions on $\tilde{\Omega}$ that are the same on every Ω_i , that is, do not depend on the tag i .

Lastly, we note that the considerable complexity of the space $\tilde{\Omega}$ can sometimes be reduced. This happens when the family $\{\mu_i: i \in I\}$ can be replaced by a smaller family that still dominates $\mathcal{P}$ and $\mathcal{Q}$ but perhaps not *every* measure in $\mathcal{M}$. In particular, in the dominated case of Theorem 1.1, the reference measure itself can be used as the single member of the family, and one can take $\tilde{\Omega} = \Omega$ as noted previously.

Remark B.1. Le Cam has shown that the minimax risk can always be attained by a generalized test. The extended space $\tilde{\Omega}$ can be used to construct a net of (non-generalized) tests that converge in the weak-* topology of $L^\infty(\tilde{\mu})$ to a given generalized test. Specifically, let $\phi^* \in L^\infty(\tilde{\mu})$ be a generalized test, and let ϕ_i^* be its restriction to Ω_i for each $i \in I$. For any countable subset of indices $J \subset I$ one can find mutually disjoint sets $A_i \in \mathcal{F}$, $i \in J$, such that $\mu_i(A_i) = 1$ for all $i \in J$, where we regard A_i as a subset of Ω_i . Now define $\phi_J = \sum_{i \in J} \phi_i^* \mathbf{1}_{A_i}$, regarded as a function on Ω . Then ϕ_J is a (non-generalized) test and satisfies $\phi_J = \phi_i^*$, μ_i -a.s., for all $i \in J$. Furthermore, for every $\mu \in \mathcal{M}$ we have $\mathbb{E}_\mu[\phi_J] = \phi^*(\mu)$ whenever J contains all $i \in I$ such that $d\mu/d\mu_i$ is not zero. From this it follows that the ϕ_J form a net that converges to ϕ^* in the weak-* topology of $L^\infty(\tilde{\mu})$.

C Some measure-theoretic results

Proposition C.1. *Let $(\Omega, \mathcal{F})$ be a measurable space and let $\mathcal{P} = \{\delta_\omega: \omega \in \Omega\}$ be the family of point masses. Then $\overline{\text{co}}^*(\mathcal{P}) = \text{ba}_1$.*

Remark C.2. Let us point out that if $\Omega = [0, 1]$ and $\mathcal{P}$ consists of all μ with a Lebesgue density, then $\overline{\text{co}}^*(\mathcal{P})$ is *not* equal to all of ba_1 . This is because every $\mu \in \overline{\text{co}}^*(\mathcal{P})$ satisfies $\mu(A) = 0$ for all Lebesgue nullsets A . In particular, point masses δ_ω are not weak-* limit points of $\text{co}(\mathcal{P})$. For example, the sequence $\mu_n = \text{Uni}[0, 1/n]$ does *not* converge to δ_0 . In fact, this sequence does not converge at all. This is demonstrated, for instance, by the function $f(\omega) = \sin(\log(1/\omega))$ for $\omega \neq 0$ and $f(0) = 0$, which satisfies $\mathbb{E}_{\mu_n}[f] = \sin(\pi/4 + \log(n))/\sqrt{2}$. This does not converge as $n \rightarrow \infty$, and then neither can μ_n . However, weak-* compactness implies that μ_n does have a convergent *subnet*. The limit is a purely finitely additive probability measure μ that assigns unit mass to any interval $[0, \epsilon]$ with $\epsilon > 0$.

Proof of Proposition C.1. Fix any $\mu \in \mathbf{ba}_1$. We construct a net in $\text{co}(\mathcal{P})$ as follows. Let Π be the set of all partitions $\pi = \{A_1, \dots, A_n\}$ of Ω into finitely many nonempty measurable sets, partially ordered by refinement. For each such π , choose points $\omega_i \in A_i$, $i = 1, \dots, n$, and define

$$\mu_\pi = \sum_{i=1}^n \mu(A_i) \delta_{\omega_i} \in \text{co}(\mathcal{P}).$$

Consider now any $B \in \mathcal{F}$. For any $\pi \in \Pi$ and any $\pi' \in \Pi$ that is a common refinement of π and the partition $\{B, B^c\}$, we have $\mu_{\pi'}(B) = \mu(B)$. It follows that the net $(\mu_\pi(B))_{\pi \in \Pi}$ in $\mathbb{R}$ converges to $\mu(B)$. Since this holds for every $B \in \mathcal{F}$, we conclude that the net $(\mu_\pi)_{\pi \in \Pi}$ in $\text{co}(\mathcal{P})$ converges to μ . $\square$

Recall that $\text{Pow}(\mathbb{R})$ refers to the power set of $\mathbb{R}$.

Proposition C.3. *Assume (3.1) and let μ be a positive finite measure on $(\mathbb{R}, \text{Pow}(\mathbb{R}))$. Then μ is purely atomic, i.e., the set $S = \{x \in \mathbb{R} : \mu(\{x\}) > 0\}$ is at most countable and for every $A \subset X$ we have*

$$\mu(A) = \sum_{x \in A \cap S} \mu(\{x\}).$$

Proof. For each $k \in \mathbb{N}$, the set $S_k = \{x \in \mathbb{R} : \mu(\{x\}) > k^{-1}\}$ must be finite. Thus the set $S = \bigcup_{k \in \mathbb{N}} S_k$ is at most countable. Now set $\mu_d(A) = \mu(A) - \mu(A \cap S)$ for all $A \subset \mathbb{R}$. This defines a positive, finite, and diffuse measure on $\text{Pow}(\mathbb{R})$. By (3.1), $\mu_d = 0$, and this completes the proof. $\square$

D Proof of Theorem 3.4

The first part of Theorem 3.4 states that every element of $\mathcal{P}_{\text{eff}}$ is setwise dominated by some maximal element of $\mathcal{P}_{\text{eff}}$. This is shown in Lemma D.4 below, building on several auxiliary results which we now develop.

Lemma D.1. *Let μ_α be a bounded increasing net in $\mathbf{ba}_+$ and define the nonnegative real-valued set function $\mu(A) = \sup_\alpha \mu_\alpha(A)$, $A \in \mathcal{F}$. Then μ_α converges to μ in the weak-* topology. If each μ_α is countably additive, then so is μ .*

Proof. The boundedness assumption means that $\sup_\alpha \mu_\alpha(\Omega) < \infty$, so μ is well-defined. Moreover, a bounded increasing net in $\mathbb{R}$ converges to its supremum, so we have $\mu(A) = \lim_\alpha \mu_\alpha(A)$ for all $A \in \mathcal{F}$. For any disjoint $A, B \in \mathcal{F}$ we pass to the limit in the identity $\mu_\alpha(A \cup B) = \mu_\alpha(A) + \mu_\alpha(B)$ to see that $\mu \in \mathbf{ba}_+$. For bounded nets in $\mathbf{ba}$, setwise convergence is equivalent to weak-* convergence, so $\mu_\alpha \rightarrow \mu$ in the weak-* topology. Lastly, suppose each μ_α is countably additive, and consider an arbitrary countable family of mutually disjoint

sets $A_i \in \mathcal{F}$. Monotonicity and finite additivity of μ yields

$$\mu \left(\bigcup_{i=1}^{\infty} A_i \right) \geq \mu \left(\bigcup_{i=1}^n A_i \right) = \sum_{i=1}^n \mu(A_i)$$

for all n , while countable additivity of μ_α and the fact that μ dominates μ_α yields

$$\mu_\alpha \left(\bigcup_{i=1}^{\infty} A_i \right) = \sum_{i=1}^{\infty} \mu_\alpha(A_i) \leq \sum_{i=1}^{\infty} \mu(A_i)$$

for all α . Passing to the limit in n and α , respectively, shows that μ is countable additive. $\square$

Lemma D.2. *Let μ_α be an increasing net in $\mathcal{P}_{\text{eff}}$. Then μ_α converges in the weak-* topology to an element μ of $\mathcal{P}_{\text{eff}}$.*

Proof. By Lemma D.1, μ_α converges in the weak-* topology to some nonnegative countable additive measure μ . This measure belongs to $\mathcal{P}_{\text{eff}}$ because the property $\int f d\mu_\alpha \leq 1$ for all $f \in \mathcal{E}_b$ is preserved under weak-* convergence. $\square$

Given an element $\mu \in \mathcal{P}_{\text{eff}}$ and a set $A \in \mathcal{F}$, we say that μ is *A-maximal* if any $\mu' \in \mathcal{P}_{\text{eff}}$ with $\mu' \geq \mu$ satisfies $\mu'(A) = \mu(A)$.

Lemma D.3. *Fix $\mu_0 \in \mathcal{P}_{\text{eff}}$ and $A \in \mathcal{F}$. There exists some $\mu' \in \mathcal{P}_{\text{eff}}$ which is A-maximal and satisfies $\mu' \geq \mu_0$.*

Proof. For any $\mu \in \mathcal{P}_{\text{eff}}$ we let $D(\mu) = \{\nu \in \mathcal{P}_{\text{eff}} : \nu \geq \mu\}$ be the set of elements of $\mathcal{P}_{\text{eff}}$ that dominate μ . Starting with μ_0 , we now construct a sequence of measures $\mu_n \in \mathcal{P}_{\text{eff}}$ and numbers a_n as follows. If μ_n has been defined, we let

$$a_n = \sup_{\nu \in D(\mu_n)} \nu(A)$$

be the largest mass assigned to A among all elements of $\mathcal{P}_{\text{eff}}$ that dominate μ_n . We then pick $\mu_{n+1} \in \mathcal{P}_{\text{eff}}$ such that

$$\mu_{n+1} \geq \mu_n \quad \text{and} \quad \mu_{n+1}(A) \geq \frac{1}{2}(a_n + \mu_n(A)),$$

which is possible due to the choice of a_n . We may then define a_{n+1} , μ_{n+2} , etc. Because $\mu_{n+1} \geq \mu_n$, we have $D(\mu_{n+1}) \subset D(\mu_n)$ and hence $a_{n+1} \leq a_n$ for all n . As a result,

$$a_{n+1} - \mu_{n+1}(A) \leq a_n - \mu_{n+1}(A) \leq a_n - \frac{1}{2}(a_n + \mu_n(A)) = \frac{1}{2}(a_n - \mu_n(A)).$$

Iterating this yields

$$a_n - \mu_n(A) \leq 2^{-n}(a_0 - \mu_0(A))$$

for all n . We now apply Lemma D.2 to get that μ_n converges in the weak-* topology to some $\mu' \in \mathcal{P}_{\text{eff}}$ with $\mu' \geq \mu_n$ for all n . Set

$$a' = \sup_{\nu \in D(\mu')} \nu(A).$$

Because $\mu' \geq \mu_n$, we have as above that $a' \leq a_n$, and then

$$a' - \mu'(A) \leq a_n - \mu_n(A) \leq 2^{-n}(a_0 - \mu_0(A))$$

for all n . It follows that $a' = \mu'(A)$, that is, μ' is A -maximal. $\square$

Lemma D.4. *Every $\mu_0 \in \mathcal{P}_{\text{eff}}$ is dominated by some maximal element of $\mathcal{P}_{\text{eff}}$.*

Proof. Let $\{A_\alpha : \alpha < \gamma\}$ be an enumeration of all the sets in $\mathcal{F}$ by ordinals α less than some ordinal γ . The existence of such an enumeration follows from the well-ordering theorem, itself a consequence of the axiom of choice. Note that we do not assume or require that γ be a countable ordinal.

Starting from μ_0 , we construct elements $\mu_\alpha \in \mathcal{P}_{\text{eff}}$ for $\alpha \leq \gamma$ by transfinite recursion as follows. If $\alpha = \beta + 1$ is a successor ordinal with $\alpha < \gamma$, we let μ_α be an A_α -maximal element of $\mathcal{P}_{\text{eff}}$ such that $\mu_\alpha \geq \mu_\beta$. This is possible thanks to Lemma D.3. If $\alpha = \beta + 1 = \gamma$, we simply set $\mu_\alpha = \mu_\beta$. If α is a limit ordinal, we first define ν_α as the weak-* limit of the bounded increasing net $(\mu_\beta)_{\beta < \alpha} \subset \mathcal{P}_{\text{eff}}$, which exists and belongs to $\mathcal{P}_{\text{eff}}$ thanks to Lemma D.2. If $\alpha < \gamma$ we then get an A_α -maximal $\mu_\alpha \in \mathcal{P}_{\text{eff}}$ with $\mu_\alpha \geq \nu_\alpha$ from Lemma D.3. If $\alpha = \gamma$, we simply set $\mu_\alpha = \nu_\alpha$. Now, for each $\alpha < \gamma$ we have $\mu_\alpha \leq \mu_\gamma$, so that μ_γ is also A_α -maximal. Because $\{A_\alpha : \alpha < \gamma\} = \mathcal{F}$, μ_γ is in fact A -maximal for every $A \in \mathcal{F}$. In other words, μ_γ is a maximal element of $\mathcal{P}_{\text{eff}}$ that dominates μ_0 . $\square$

The remaining parts of Theorem 3.4 require the introduction of a third object: one must view the null hypothesis $\mathcal{P}$ as a subset of $\mathbf{ba}$ and consider an appropriate extension of the effective null, namely,

$$\mathcal{P}_{\text{eff}}^{\mathbf{ba}} = \{\mu \in \mathbf{ba}_+ : \int f d\mu \leq 1 \text{ for all } f \in \mathcal{E}_b\}. \quad (\text{D.1})$$

Lemma D.5. $\mathcal{P}_{\text{eff}}^{\mathbf{ba}} = \mathbf{ba}_+ \cap \overline{\text{co}}^*(\mathcal{P} - \mathbf{ba}_+)$.

Proof. The polars and bipolars below are understood with respect to the dual pair $(\mathcal{L}, \mathbf{ba})$. We first show that

$$\mathcal{E}_b = (\mathcal{P} - \mathbf{ba}_+)^\circ. \quad (\text{D.2})$$

The right-hand side equals $\{f \in \mathcal{L} : \int f d(\mu - \nu) \leq 1 \text{ for all } \mu \in \mathcal{P}, \nu \in \mathbf{ba}_+\}$, which certainly contains $\mathcal{E}_b$. Conversely, suppose f belongs to the right-hand side. Taking $\nu = 0$ we see that $\int f d\mu \leq 1$ for all $\mu \in \mathcal{P}$. Moreover, for any $\omega \in \Omega$ we may take an arbitrary $\mu \in \mathcal{P}$ and set $\nu = \mu + t\delta_\omega$ for $t > 0$ to get $f(\omega) \geq -t^{-1}$. Sending t to infinity shows that

$f \in \mathcal{L}_+$, and hence that $f \in \mathcal{E}_b$. This establishes (D.2). Next, the definition (D.1) states that $\mathcal{P}_{\text{eff}}^{\text{ba}} = \text{ba}_+ \cap \mathcal{E}_b^\circ$. Together with (D.2), this yields

$$\mathcal{P}_{\text{eff}}^{\text{ba}} = \text{ba}_+ \cap (\mathcal{P} - \text{ba}_+)^{\circ\circ}.$$

Since $\mathcal{P} - \text{ba}_+$ contains zero, an application of the bipolar theorem completes the proof. $\square$

Lemma D.6. $\mathcal{P}_{\text{eff}}^{\text{ba}}$ is the solid hull of $\overline{\text{co}}^*(\mathcal{P})$, i.e.,

$$\mathcal{P}_{\text{eff}}^{\text{ba}} = \text{solid } \overline{\text{co}}^*(\mathcal{P}). \quad (\text{D.3})$$

As a consequence, the set of maximal elements of $\mathcal{P}_{\text{eff}}^{\text{ba}}$ precisely equals $\overline{\text{co}}^*(\mathcal{P})$, and every $\mu \in \mathcal{P}_{\text{eff}}^{\text{ba}}$ is setwise dominated by some $\mu' \in \overline{\text{co}}^*(\mathcal{P})$.

Proof. We prove (D.3); the remaining statement is then immediate. In view of Lemma D.5, it suffices to show that

$$\text{ba}_+ \cap \overline{\text{co}}^*(\mathcal{P} - \text{ba}_+) \subset \text{solid } \overline{\text{co}}^*(\mathcal{P}) \subset \mathcal{P}_{\text{eff}}^{\text{ba}}. \quad (\text{D.4})$$

We start with the first inclusion, and let η be an element of the left-hand side. Then $\eta \in \text{ba}_+$ and there is a net $\eta_\alpha = \mu_\alpha - \nu_\alpha$ with $\mu_\alpha \in \text{co}(\mathcal{P})$ and $\nu_\alpha \in \text{ba}_+$ such that $\eta_\alpha \rightarrow \eta$. Because each μ_α belongs to ba_1 , which is weak-* compact by the Banach–Alaoglu theorem, on passing to a subnet we have $\mu_\alpha \rightarrow \mu$ for some $\mu \in \text{ba}_1$. In fact, $\mu \in \overline{\text{co}}^*(\mathcal{P})$. We also get $\nu_\alpha = \mu_\alpha - \eta_\alpha \rightarrow \nu = \mu - \eta$ which, like each ν_α , belongs to the weak-* closed set ba_+ . In summary, we have shown that $0 \leq \eta \leq \mu$ with $\mu \in \overline{\text{co}}^*(\mathcal{P})$. That is, $\eta \in \text{solid } \overline{\text{co}}^*(\mathcal{P})$.

We now prove the second inclusion in (D.4). It is clear from the definition (D.1) that $\mathcal{P}_{\text{eff}}^{\text{ba}}$ is convex, weak-* closed, and contains $\mathcal{P}$. Therefore it also contains $\overline{\text{co}}^*(\mathcal{P})$. It is also clear that $\mathcal{P}_{\text{eff}}^{\text{ba}}$ is solid, so it must contain $\text{solid } \overline{\text{co}}^*(\mathcal{P})$ as well. This completes the proof of (D.4), and hence of (D.3). $\square$

We now turn to the second part of Theorem 3.4; this is content of the following lemma.

Lemma D.7. Every maximal element of $\mathcal{P}_{\text{eff}}$ is the countably additive part of some element of $\overline{\text{co}}^*(\mathcal{P})$.

Proof. Let $\mu' \in \mathcal{P}_{\text{eff}}$ be a maximal element. Since $\mathcal{P}_{\text{eff}} \subset \mathcal{P}_{\text{eff}}^{\text{ba}}$, the representation (D.3) of Lemma D.6, which we just established, implies that $\mu' \leq \mu''$ for some $\mu'' \in \overline{\text{co}}^*(\mathcal{P})$. Let $\mu'' = \mu''_c + \mu''_p$ be the Yosida–Hewitt decomposition of μ'' into its countably additive and purely finitely additive parts; see Yosida and Hewitt (1952, Theorem 1.24). The countably additive part μ''_c is the largest countably additive measure that is dominated by μ'' . Indeed, if $\nu \leq \mu''$ is countably additive, then so is $(\nu - \mu''_c)^+ \leq \mu''_p$, implying that $(\nu - \mu''_c)^+ = 0$ since μ''_p is purely finitely additive, and hence $\nu \leq \mu''_c$. We apply this with $\nu = \mu'$ to get $\mu' \leq \mu''_c$. Moreover, because μ''_c is countably additive and dominated by μ'' , it belongs to $\mathcal{P}_{\text{eff}}$. But μ' is maximal in $\mathcal{P}_{\text{eff}}$, so we must actually have $\mu' = \mu''_c$. $\square$