

The numeraire e-variable and reverse information projection

Martin Larsson* Aaditya Ramdas† Johannes Ruf‡

December 30, 2024

Abstract

We consider testing a composite null hypothesis \mathcal{P} against a point alternative \mathbb{Q} using e-variables, which are nonnegative random variables X such that $\mathbb{E}_{\mathbb{P}}[X] \leq 1$ for every $\mathbb{P} \in \mathcal{P}$. This paper establishes a fundamental result: under no conditions whatsoever on \mathcal{P} or \mathbb{Q} , there exists a special e-variable X^* that we call the numeraire, which is strictly positive and satisfies $\mathbb{E}_{\mathbb{Q}}[X/X^*] \leq 1$ for every other e-variable X . In particular, X^* is log-optimal in the sense that $\mathbb{E}_{\mathbb{Q}}[\log(X/X^*)] \leq 0$. Moreover, X^* identifies a particular sub-probability measure \mathbb{P}^* via the density $d\mathbb{P}^*/d\mathbb{Q} = 1/X^*$. As a result, X^* can be seen as a generalized likelihood ratio of \mathbb{Q} against \mathcal{P} . We show that \mathbb{P}^* coincides with the reverse information projection (RIPr) when additional assumptions are made that are required for the latter to exist. Thus \mathbb{P}^* is a natural definition of the RIPr in the absence of any assumptions on \mathcal{P} or \mathbb{Q} . In addition to the abstract theory, we provide several tools for finding the numeraire and RIPr in concrete cases. We discuss several nonparametric examples where we can indeed identify the numeraire and RIPr, despite not having a reference measure. Our results have interpretations outside of testing in that they yield the optimal Kelly bet against \mathcal{P} if we believe reality follows \mathbb{Q} . We end with a more general optimality theory that goes beyond the ubiquitous logarithmic utility. We focus on certain power utilities, leading to reverse Rényi projections in place of the RIPr, which also always exist.

1 Introduction

Suppose we observe data from an unknown distribution on a measurable space (Ω, \mathcal{F}) . This paper is concerned with testing the composite null hypothesis H_0 that the data distribution belongs to a given set \mathcal{P} of distributions on this space, against the point alternative H_1 that the data was generated by a given distribution \mathbb{Q} . Here \mathcal{P} and \mathbb{Q} are completely arbitrary, and our general theory will not require any further assumptions whatsoever.

We approach the testing problem by designing *e-variables*: a $[0, \infty]$ -valued random variable X is called an e-variable for \mathcal{P} if $\mathbb{E}_{\mathbb{P}}[X] \leq 1$ for all $\mathbb{P} \in \mathcal{P}$. Specifically, we will look for ‘optimal’ e-variables, in senses to be formalized later, that tend to be large under \mathbb{Q} .

E-variables measure evidence against the null: the larger its realized value, the stronger the evidence. One can also use e-variables to make hard accept/reject decisions. In particular,

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

†Departments of Statistics and ML, Carnegie Mellon University; aramdas@cmu.edu

‡Department of Mathematics, London School of Economics; School of Data Science, Chinese University Hong Kong Shenzhen j.ruf@lse.ac.uk

Markov’s inequality implies that we can reject H_0 at level α if an e-variable X exceeds $1/\alpha$, since $\mathbb{P}(X \geq 1/\alpha) \leq \alpha$ for each $\mathbb{P} \in \mathcal{P}$. This paper will not be concerned with such hard decisions, but just in the measure of evidence X itself, which requires no pre-specification of any level α .

E-variables have been employed since at least [Robbins \(1970\)](#). More recently, in 2018–2020, a number of papers appeared considering new applications and constructions of e-variables ([Wasserman et al., 2020](#); [Howard et al., 2020, 2021](#); [Shafer, 2021](#); [Vovk and Wang, 2021](#); [Waudby-Smith and Ramdas, 2023](#); [Grünwald et al., 2024](#)), but all using different (or no) terminology. In 2020, these authors agreed to use the term ‘e-variable’ (and ‘e-value’ for its realization). See [Ramdas et al. \(2023\)](#) for a recent survey.

Our main result establishes that *under absolutely no restrictions or conditions on \mathcal{P} or \mathbb{Q}* , there exists a special \mathbb{Q} -almost surely unique and strictly positive e-variable X^* that we call the *numeraire e-variable*, or just the *numeraire*. It is characterized by the property that $\mathbb{E}_{\mathbb{Q}}[X/X^*] \leq 1$ for every other e-variable X . This in turn implies two other interpretable properties $\mathbb{E}_{\mathbb{Q}}[X^*/X] \geq 1$ and $\mathbb{E}_{\mathbb{Q}}[\log(X/X^*)] \leq 0$, the second of which is log-optimality. In fact, X^* is the numeraire if and only if it is log-optimal in this sense. The terminology derives from mathematical finance, where the *numeraire portfolio* is a central object analogous to the numeraire e-variable; see e.g. [Long \(1990\)](#); [Becherer \(2001\)](#); [Platen \(2006\)](#); [Karatzas and Kardaras \(2007, 2021\)](#). In the standard Neyman-Pearson framework with a composite null and alternative, already [Cvitanic and Karatzas \(2001\)](#) observed how techniques from financial mathematics can be transferred to obtain optimal statistical tests.

Although no assumptions on \mathcal{P} or \mathbb{Q} are needed in general, the theory simplifies if the alternative \mathbb{Q} does not assign positive probability to any event that is null under every $\mathbb{P} \in \mathcal{P}$. In this case we say that \mathbb{Q} is *absolutely continuous with respect to \mathcal{P}* , written $\mathbb{Q} \ll \mathcal{P}$. This natural generalization of absolute continuity for pairs of measures goes back to [Halmos and Savage \(1949\)](#) and is satisfied in most statistically relevant settings, while being significantly weaker than requiring a common dominating reference measure. [Huber and Strassen \(1973\)](#) provide early insights on statistical testing in the absence of a reference measure.

The condition $\mathbb{Q} \ll \mathcal{P}$ has several consequences: (a) the numeraire is \mathbb{Q} -almost surely finite; (b) every e-variable is \mathbb{Q} -almost surely finite; and (c) the set of all e-variables for \mathcal{P} is bounded in probability under \mathbb{Q} , meaning that $\lim_{a \rightarrow \infty} \sup_X \mathbb{Q}(X > a) = 0$ where the supremum extends over all e-variables for \mathcal{P} . In fact, we show that the above properties are equivalent in that they imply each other, and they also imply that $\mathbb{Q} \ll \mathcal{P}$.

Next, we associate to the numeraire a special sub-probability measure $\mathbb{P}^* \ll \mathbb{Q}$ defined via its density $d\mathbb{P}^*/d\mathbb{Q} = 1/X^*$. This measure belongs to the *effective null hypothesis* associated with \mathcal{P} , denoted by \mathcal{P}_{eff} . The effective null is the largest family of sub-probability measures that has the same set of e-variables as \mathcal{P} . These are the distributions against which e-variables for \mathcal{P} do not provide any evidence.

We show that \mathbb{P}^* can be characterized in terms of properties normally associated with the *reverse information projection* (RIPr), first used by [Csiszár and Tusnády \(1984\)](#). Earlier work on the RIPr involve assumptions related to absolute continuity and/or finiteness of the minimum relative entropy between \mathbb{Q} and \mathcal{P} , and in such settings the above \mathbb{P}^* coincides with the RIPr. This includes, in particular, work by [Grünwald et al. \(2024\)](#) and [Lardy et al. \(2024\)](#), which in turn build on results in [Li \(1999\)](#). This leads us to argue that \mathbb{P}^* is the natural definition of the RIPr in the absence of any assumptions on \mathcal{P} or \mathbb{Q} .

One of our main contributions is to significantly relax the conditions required for a log-optimality theory of e-variables to be built, and for a RIPr to exist. Furthermore, in addition to the abstract theory, we develop tools to aid in the computation of the numeraire and the RIPr in concrete examples. We establish several sufficient conditions to verify that an e-variable at hand is indeed the numeraire, and we provide several examples of nonparametric nulls \mathcal{P} (such as bounded distributions, symmetric distributions and sub-Gaussian distributions) that illustrate how these sufficient conditions can be used. These examples are out of the reach of previous works due to their lack of a reference measure on which the aforementioned works centrally rely.

An important message of our work is that in concrete examples, the duality between the set of e-variables and the null hypothesis can be exploited in different ways depending on the nature of the null. For parametric nulls, the simplest approach is usually to first find the RIPr (and then the numeraire). For non-parametric nulls, it is often easier to first find the numeraire (and then the RIPr).

We end the paper with a more general optimality theory that goes beyond the ubiquitous logarithmic utility, by studying other concave utilities and associated divergences. Without aiming for full generality, we focus on certain power utilities in place of the logarithm, which leads to *reverse Rényi projections* in place of the RIPr. These developments, again, take place under absolutely no conditions on \mathcal{P} or \mathcal{Q} .

Finally, we note that if $\mathcal{P} = \{\mathbb{P}\}$ is a singleton and $\mathcal{Q} \ll \mathbb{P}$, we obtain as a consequence that the numeraire is the likelihood ratio $d\mathcal{Q}/d\mathbb{P}$, which is well known to be log-optimal (Shafer, 2021). Thus, our work can be viewed as deriving certain generalized likelihood ratios for composite null hypotheses, and extending the corresponding optimality theory.

Paper outline. Section 2 defines the numeraire, establishes existence and uniqueness under no assumptions, characterizes the absolute continuity property $\mathcal{Q} \ll \mathcal{P}$, and establishes various other useful facts such as log-optimality of the numeraire. Section 3 introduces the RIPr and characterizes it via various optimality properties dual to the properties and log-optimality of the numeraire. Section 4 establishes a number of tools for finding the numeraire and the RIPr in practice. These tools are then applied to specific examples in Section 5. Section 6 goes beyond log-optimality by considering certain power utilities and the associated reverse Rényi projections. Section 7 shows how our results imply that the method of universal inference is inadmissible. Section 8 concludes. Appendix A contains a generalized Lebesgue decomposition that we rely on in certain proofs, and Appendix B gives an overview of the relation between absolute continuity properties and finiteness of the numeraire and relative entropies.

Notation. Given the measurable space (Ω, \mathcal{F}) , we let F_+ , M_+ , and M_1 denote the set of all $[0, \infty]$ -valued measurable functions, nonnegative measures, and probability measures on \mathcal{F} , respectively. Given a null hypothesis \mathcal{P} , which is an arbitrary nonempty subset of M_1 , we write \mathcal{E} for the associated set of e-variables, i.e., the family of all $X \in F_+$ that satisfy $\mathbb{E}_{\mathbb{P}}[X] \leq 1$ for all $\mathbb{P} \in \mathcal{P}$. Infinite quantities play a key role in this paper. We generally rely on standard measure theoretic conventions such as $0 \cdot x = 0$ for $x \in [-\infty, \infty]$ and $x \cdot \infty = \infty$ for $x \in (0, \infty]$, as well as $\log(0) = -\infty$ and $\log(\infty) = \infty$. When evaluating ratios of $[0, \infty]$ -valued random variables we often use the following convention, the last part of which is somewhat

unusual but simplifies notation:

$$x/\infty = 0 \text{ and } \infty/x = \infty \text{ for } x \in [0, \infty), \text{ and } \infty/\infty = 1. \quad (1.1)$$

Clarifying explanations are given in potentially ambiguous cases. Finally, we write $\text{conv}(\dots)$ for convex hull, that is, the set of (finite) convex combinations of the indicated objects.

2 The numeraire e-variable

Here is the key definition of this paper. The convention (1.1) is used to evaluate the ratio inside the expectation.

Definition 2.1. A *numeraire e-variable* (or just *numeraire* for short) is a \mathbb{Q} -almost surely strictly positive e-variable X^* such that $\mathbb{E}_{\mathbb{Q}}[X/X^*] \leq 1$ for every e-variable X .

Our choice of terminology is inspired by the *numeraire portfolio* in the mathematical finance literature; see e.g. Long (1990); Becherer (2001); Platen (2006); Karatzas and Kardaras (2007, 2021). In that literature, the defining property of the numeraire portfolio is that any positive wealth process divided by the numeraire becomes a supermartingale. In view of the gambling interpretation of e-variables (Shafer, 2021), this is precisely what is expressed by Definition 2.1. A subtle distinction is that we allow the numeraire to take the value infinity, whereas this is excluded in the finance literature.

Numeraires are unique up to \mathbb{Q} -nullsets. Indeed, if X_1^* and X_2^* are numeraires, then the ratio $Y = X_2^*/X_1^*$ satisfies $1 \leq 1/\mathbb{E}_{\mathbb{Q}}[Y] \leq \mathbb{E}_{\mathbb{Q}}[1/Y] \leq 1$ thanks to the numeraire property of X_1^* , Jensen's inequality, and the numeraire property of X_2^* . Thus Jensen's inequality holds with equality, so Y is \mathbb{Q} -almost surely equal to a constant which must be one. It follows that X_1^* and X_2^* are \mathbb{Q} -almost surely equal. In view of this uniqueness, we often speak of 'the' numeraire.

Example 2.2. In the case of a simple null $\mathcal{P} = \{\mathbb{P}_0\}$ with \mathbb{Q} absolutely continuous with respect to \mathbb{P}_0 , the numeraire is just the likelihood ratio $X^* = d\mathbb{Q}/d\mathbb{P}_0$. Indeed, X^* is \mathbb{Q} -almost surely strictly positive, it is an e-variable, and for any other e-variable X we have $\mathbb{E}_{\mathbb{Q}}[X/X^*] = \mathbb{E}_{\mathbb{P}_0}[X] \leq 1$. In particular settings, such as when a reference measure exists and certain additional assumptions are satisfied, we will show that the numeraire is a likelihood ratio involving the 'reverse information projection' (Grünwald et al., 2024; Lardy et al., 2024). Even without these restrictions, X^* can still be interpreted as a likelihood ratio. One of our goals is to establish this in full generality. This is the subject of Section 3.

The following general property of the numeraire can sometimes be useful.

Lemma 2.3. Let X^* denote the numeraire for \mathcal{P} , and consider a larger null hypothesis $\mathcal{P}' \supset \mathcal{P}$. If X^* is still an e-variable for \mathcal{P}' , then it is also the numeraire for \mathcal{P}' .

Proof. If X is an e-variable for \mathcal{P}' , it is also an e-variable for \mathcal{P} and one has $\mathbb{E}_{\mathbb{Q}}[X/X^*] \leq 1$ by assumption. This proves the lemma. \square

A numeraire X^* is also *log-optimal* in the sense that

$$\mathbb{E}_{\mathbb{Q}} \left[\log \frac{X}{X^*} \right] \leq 0 \text{ for every e-variable } X, \quad (2.1)$$

where the left-hand side may be $-\infty$. This follows directly from Jensen's inequality and the numeraire property. The converse is also true, and we record the equivalence in the following proposition; see also [Cover and Thomas \(2006, Theorem 16.2.2\)](#). The proof shows that the numeraire property is really the first-order condition for log-optimality. Note also that a numeraire X^* is the \mathbb{Q} -almost surely unique log-optimal e-variable in the sense of [\(2.1\)](#) even if $\mathbb{E}_{\mathbb{Q}}[\log X^*]$ happens to be infinite.

Proposition 2.4. *Let X^* be a \mathbb{Q} -almost surely strictly positive e-variable. Then X^* is a numeraire if and only if it is log-optimal. In particular, a log-optimal e-variable is unique up to \mathbb{Q} -nullsets.*

Proof. The forward direction was argued above. To prove the converse we assume X^* is log-optimal. For any e-variable X and $t \in (0, 1)$, $X(t) = (1 - t)X^* + tX$ is an e-variable. Thus by log-optimality, $\mathbb{E}_{\mathbb{Q}}[t^{-1} \log(X(t)/X^*)] \leq 0$. The expression inside the expectation equals $t^{-1} \log(1 - t + tX/X^*)$, which converges to $X/X^* - 1$ as t tends to zero and is bounded below by $t^{-1} \log(1 - t)$, hence by $-2 \log 2$ for $t \in (0, 1/2)$. Fatou's lemma thus yields $\mathbb{E}_{\mathbb{Q}}[X/X^* - 1] \leq 0$, showing that X^* is a numeraire. Finally, the uniqueness statement follows from the equivalence just proved together with uniqueness of the numeraire up to \mathbb{Q} -nullsets. \square

Remark 2.5. In some situations, for instance if $\log X^*$ has infinite \mathbb{Q} -expectation, it may be of interest measure evidence relative to a *benchmark*. By a benchmark we simply mean any \mathbb{Q} -almost surely strictly positive random variable X_b such that $\log(X^*/X_b)$ has finite \mathbb{Q} -expectation. Note that the numeraire itself is a possible benchmark. For any $X \in \mathcal{E}$, the log-optimality property [\(2.1\)](#) of X^* yields

$$\mathbb{E}_{\mathbb{Q}} \left[\log \frac{X}{X_b} \right] = \mathbb{E}_{\mathbb{Q}} \left[\log \frac{X}{X^*} \right] + \mathbb{E}_{\mathbb{Q}} \left[\log \frac{X^*}{X_b} \right] \leq \mathbb{E}_{\mathbb{Q}} \left[\log \frac{X^*}{X_b} \right].$$

Thus the numeraire maximizes the 'benchmarked' quantity $\mathbb{E}_{\mathbb{Q}}[\log(X/X_b)]$ over \mathcal{E} , and the optimal value is finite. This is true regardless of which benchmark is used, so the numeraire is *benchmark-invariant*. We thank an anonymous referee for pointing this out.

The following theorem shows that a numeraire always exists without any assumptions whatsoever on \mathcal{P} or \mathbb{Q} . The theorem is related to an important family of results in the mathematical finance literature, known collectively as the Fundamental Theorem of Asset Pricing, which are versions of the statement that a numeraire portfolio with finite wealth exists if and only if unbounded profits must come with unbounded downside risk. The following theorem contains the corresponding statement in the context of e-variables.

Theorem 2.6. *A numeraire exists. Moreover, the following conditions are equivalent:*

- (i) *The numeraire is \mathbb{Q} -almost surely finite.*
- (ii) *Every e-variable is \mathbb{Q} -almost surely finite.*
- (iii) *\mathbb{Q} is absolutely continuous with respect to \mathcal{P} .*
- (iv) *The set of e-variables \mathcal{E} is bounded in probability under \mathbb{Q} .*

Proof. The existence of a numeraire is Lemma 2.10 below, which relies on the preceding Lemmas 2.8 and 2.9. Here we focus on the equivalence of the four conditions. Clearly (ii) implies (i). To see that (i) implies (ii), let X^* be a \mathbb{Q} -almost surely finite numeraire and note that the numeraire property implies that any e-variable X must be \mathbb{Q} -almost surely finite too. Next, we claim that (ii) is equivalent to (iii). Suppose (iii) fails. Then there is an event A with $\mathbb{Q}(A) > 0$ and $\mathbb{P}(A) = 0$ for all $\mathbb{P} \in \mathcal{P}$, in which case $X = \infty \mathbf{1}_A$ is an e-variable that is not \mathbb{Q} -almost surely finite. Thus (ii) fails also. The converse follows by noting that every e-variable is finite \mathbb{P} -almost surely for all $\mathbb{P} \in \mathcal{P}$. Further, (iv) implies (ii) because under the former condition we have for any e-variable X that $\mathbb{Q}(X = \infty) = \lim_{a \rightarrow \infty} \mathbb{Q}(X > a) = 0$. It remains to prove that (ii) implies (iv). This more involved implication is the content of Lemma 2.7 below. \square

Conditions (ii) and (iv) of the theorem describe the existence of a finite numeraire in financial terms. Condition (ii) simply states that infinite profits are impossible. Condition (iv) is a ‘pre-limiting’ expression of the same fact, in the sense that it excludes sequences of bets that achieve unbounded profits with a fixed positive probability. Condition (iii) has statistical rather than financial meaning. If it were violated, the alternative would assign positive probability to some event that has zero probability under every element of the null; if such an event occurs, the null can be immediately rejected at any level, meaning that any e-variable could be set to equal infinity on that event, violating the other statements.

The rest of this section contains the remaining parts of the proof of Theorem 2.6. This material can be skipped without compromising the understanding of Sections 3, 4, or 5.

Lemma 2.7. *If every e-variable is \mathbb{Q} -almost surely finite then \mathcal{E} is bounded in probability under \mathbb{Q} .*

Proof. Consider the space $L_+^0(\mathbb{Q})$ consisting of all nonnegative real-valued random variables modulo \mathbb{Q} -almost sure equivalence. Let $C \subset L_+^0(\mathbb{Q})$ be the set of (equivalence classes of) random variables that are \mathbb{Q} -almost surely equal to a finite e-variable. The set C is convex, solid, and closed in probability. Suppose now that \mathcal{E} is not bounded in probability under \mathbb{Q} . Then C is not bounded in probability. The decomposition result in Lemma 2.3 of Brannath and Schachermayer (1999) then shows that there is an event A with $\mathbb{Q}(A) > 0$ such that $n\mathbf{1}_A$ belongs to C for all $n \in \mathbb{N}$. By definition of C , for each n there is a finite e-variable X_n that is \mathbb{Q} -almost surely equal to $n\mathbf{1}_A$, say on an event B_n with $\mathbb{Q}(B_n) = 1$. By Fatou’s lemma, $X = \liminf_{n \rightarrow \infty} X_n$ is again an e-variable. Moreover, $X = \infty$ on $A \cap (\bigcap_{n \in \mathbb{N}} B_n)$, which coincides with A up to a \mathbb{Q} -nullset and thus has positive probability under \mathbb{Q} . In other words, X is an e-variable that is not \mathbb{Q} -almost surely finite. \square

The following result, usually referred to as a ‘Komlós-type’ lemma, is the basis of the proof that a numeraire exists. For a proof, see Lemma A1.1 and the subsequent Remark 1 in Delbaen and Schachermayer (1994).

Lemma 2.8. *Let $(X_n)_{n \in \mathbb{N}}$ be any sequence of $[0, \infty]$ -valued random variables. One can choose $\tilde{X}_n \in \text{conv}(X_n, X_{n+1}, \dots)$ such that \tilde{X}_n converges \mathbb{Q} -almost surely to a $[0, \infty]$ -valued random variable.*

Next, we consider a certain concave maximization problem. This involves a *utility function*, by which we mean a function $U: (0, \infty) \rightarrow \mathbb{R}$ that is continuous, nondecreasing, concave,

and differentiable. By convention we extend the domain of U to $[0, \infty]$ by setting $U(0) = \lim_{x \rightarrow 0} U(x)$ and $U(\infty) = \lim_{x \rightarrow \infty} U(x)$, and we define $U'(0)$ and $U'(\infty)$ similarly. These limits exist by monotonicity and concavity of U , and could be $\pm\infty$. The following result assumes in addition that U is bounded above, which forces $U'(\infty) = 0$.

Lemma 2.9. *Let U be a utility function that is bounded from above. There exists an e-variable X^* that attains $\sup_{X \in \mathcal{E}} \mathbb{E}_{\mathbf{Q}}[U(X)]$ and has the property that every e-variable is \mathbf{Q} -almost surely finite on $\{X^* < \infty\}$. If $\mathbb{E}_{\mathbf{Q}}[U'(X^*)X^*] < \infty$, one has the first-order condition*

$$\mathbb{E}_{\mathbf{Q}} [U'(X^*)(X - X^*)] \leq 0, \quad X \in \mathcal{E}. \quad (2.2)$$

Here $U'(X^*)X^*$ and $U'(X^*)(X - X^*)$ are understood as zero on $\{X^* = \infty\}$.

Proof. Let $X_n \in \mathcal{E}$ form a sequence such that $\mathbb{E}_{\mathbf{Q}}[U(X_n)]$ increases to the supremum. Lemma 2.8 yields a \mathbf{Q} -almost surely convergent sequence of random variables of the form $\tilde{X}_n = \sum_{i=n}^{m_n} \lambda_i^n X_i$ for convex weights $\lambda_n^n, \dots, \lambda_{m_n}^n$. Define $X^* = \liminf_{n \rightarrow \infty} \tilde{X}_n$ and note that while this may be a proper limit inferior for some $\omega \in \Omega$, we have \mathbf{Q} -almost surely that $X^* = \lim_{n \rightarrow \infty} \tilde{X}_n$. Each \tilde{X}_n belongs to \mathcal{E} since this set is convex, and Fatou's lemma then implies that X^* belongs to \mathcal{E} as well. Next, since U is concave and $\mathbb{E}[U(X_n)]$ is an increasing sequence,

$$\mathbb{E}_{\mathbf{Q}}[U(\tilde{X}_n)] \geq \sum_{i=n}^{m_n} \lambda_i^n \mathbb{E}_{\mathbf{Q}}[U(X_i)] \geq \mathbb{E}_{\mathbf{Q}}[U(X_n)]. \quad (2.3)$$

Since U is bounded above and $X^* = \lim_{n \rightarrow \infty} \tilde{X}_n$ \mathbf{Q} -almost surely, we may use Fatou's lemma together with (2.3) to get

$$\mathbb{E}_{\mathbf{Q}}[U(X^*)] \geq \limsup_{n \rightarrow \infty} \mathbb{E}_{\mathbf{Q}}[U(\tilde{X}_n)] \geq \limsup_{n \rightarrow \infty} \mathbb{E}_{\mathbf{Q}}[U(X_n)] = \sup_{X \in \mathcal{E}} \mathbb{E}_{\mathbf{Q}}[U(X)].$$

This shows that X^* attains the supremum.

We next show that X^* can be chosen so that every e-variable is \mathbf{Q} -almost surely finite on $\{X^* < \infty\}$. This uses Lemma A.1, which in particular states that there is an event N that is \mathcal{P} -negligible, meaning that $\mathbf{P}(N) = 0$ for all $\mathbf{P} \in \mathcal{P}$, and satisfies $\mathbf{Q}(N' \setminus N) = 0$ for any other \mathcal{P} -negligible event N' . Define $X^{**} = X^* \mathbf{1}_{N^c} + \infty \mathbf{1}_N$. This is still an e-variable since N is \mathcal{P} -negligible, and it attains the supremum since it is at least as large as X^* . Furthermore, since $N' = \{X = \infty\}$ is \mathcal{P} -negligible for any $X \in \mathcal{E}$, we have $\mathbf{Q}(\{X = \infty\} \setminus \{X^{**} = \infty\}) = 0$. Equivalently, by taking complements, X is \mathbf{Q} -almost surely finite on $\{X^{**} < \infty\}$. Replacing X^* by X^{**} thus ensures the claimed property.

We now verify the first-order condition (2.2). Pick $X \in \mathcal{E}$ such that $X \geq \delta$ for some $\delta > 0$; this restriction will be removed later. Define $X(t) = tX + (1-t)X^*$ for $t \in [0, 1]$ and use the optimality of $X^* = X(0)$ to get

$$\mathbb{E}_{\mathbf{Q}} \left[\frac{U(X(t)) - U(X(0))}{t} \right] \leq 0, \quad t \in (0, 1].$$

By concavity of $t \mapsto U(X(t))$, as t tends to zero the difference quotient increases to $\frac{d}{dt} \Big|_{t=0} U(X(t)) = U'(X^*)(X - X^*)$, understood to be zero on $\{X^* = \infty\}$, and is bounded below by $U(X(1)) - U(X(0)) \geq U(\delta) - \sup_{x>0} U(x)$. Monotone convergence then yields (2.2) for all X bounded

below by δ . We now remove this restriction and consider an arbitrary $X \in \mathcal{E}$. For each $\delta \in (0, 1)$ the random variable $\delta + (1 - \delta)X$ still belongs to \mathcal{E} and is bounded below by δ . By what we just proved,

$$\mathbb{E}_{\mathbb{Q}} [U'(X^*)(\delta + (1 - \delta)X - X^*)] \leq 0.$$

The expression inside the expectation is bounded below by $-U'(X^*)X^*$, understood to be zero on $\{X^* = \infty\}$, and this is integrable by assumption. We may then send δ to zero and use Fatou's lemma to deduce (2.2). \square

Lemma 2.10. *A numeraire exists.*

Proof. If we could have taken $U(x) = \log(x)$ in Lemma 2.9, the first-order condition (2.2) would have been the numeraire property. We cannot do this directly, because the logarithm is not bounded above. Instead we consider bounded approximations of the logarithm, each giving an associated optimal e-variable and first-order condition. We then pass to the limit in the first-order conditions using the Komlós-type lemma. We now turn to the details.

For each $n \in \mathbb{N}$ we define $U_n(x)$ by letting $U_n(1) = 0$ and setting the derivative to

$$U'_n(x) = \begin{cases} x^{-1}, & x \leq n, \\ 2n^{-1} - xn^{-2}, & n < x \leq 2n, \\ 0, & x > 2n. \end{cases}$$

Then U_n satisfies the properties required in Lemma 2.9, and we record for later use that U'_n is convex and that

$$U'_n(x) \leq U'_i(x) \text{ for } n \leq i \text{ and } x \in (0, \infty]. \quad (2.4)$$

Applying Lemma 2.9 with $U = U_n$ gives an optimal e-variable X_n^* for each n . These have the additional property that the sets $\{X_n^* < \infty\}$ all coincide up to \mathbb{Q} -nullsets, say with a fixed set A . Because $U'_n(x)x \leq 1$ for all $x \in (0, \infty]$, we may use the first-order condition (2.2) to deduce

$$\mathbb{E}_{\mathbb{Q}} [U'_n(X_n^*)X\mathbf{1}_A] \leq \mathbb{Q}(A), \quad X \in \mathcal{E}. \quad (2.5)$$

Next, Lemma 2.8 yields a \mathbb{Q} -almost surely convergent sequence of random variables $\tilde{X}_n^* = \sum_{i=n}^{m_n} \lambda_i^n X_i^*$ for some convex weights $\lambda_n^1, \dots, \lambda_n^{m_n}$. We define $X^* = \liminf_{n \rightarrow \infty} \tilde{X}_n^*$, which belongs to \mathcal{E} and is \mathbb{Q} -almost surely the limit of \tilde{X}_n^* . Note that A then coincides \mathbb{Q} -almost surely with the sets where X^* and the \tilde{X}_n^* are finite. Using first (2.5), then (2.4), and finally the convexity of U'_n we get

$$\mathbb{Q}(A) \geq \mathbb{E}_{\mathbb{Q}} \left[\sum_{i=n}^{m_n} \lambda_i^n U'_i(X_i^*)X\mathbf{1}_A \right] \geq \mathbb{E}_{\mathbb{Q}} \left[\sum_{i=n}^{m_n} \lambda_i^n U'_n(X_i^*)X\mathbf{1}_A \right] \geq \mathbb{E}_{\mathbb{Q}} [U'_n(\tilde{X}_n^*)X\mathbf{1}_A].$$

The expression inside the last expectation is nonnegative and converges \mathbb{Q} -almost surely to $X\mathbf{1}_A/X^*$. In particular, X^* is \mathbb{Q} -almost surely strictly positive on A . Fatou's lemma now yields $\mathbb{E}_{\mathbb{Q}}[X\mathbf{1}_A/X^*] \leq \mathbb{Q}(A)$. On A^c , where $X^* = \infty$, we have $X/X^* \leq 1$ thanks to the convention (1.1). Hence $\mathbb{E}_{\mathbb{Q}}[X/X^*] \leq \mathbb{E}_{\mathbb{Q}}[X\mathbf{1}_A/X^*] + \mathbb{Q}(A^c) \leq 1$. This shows that X^* is a numeraire. \square

3 Duality and reverse information projection

We continue to consider an arbitrary nonempty composite null hypothesis \mathcal{P} and a simple alternative \mathbb{Q} on a measurable space (Ω, \mathcal{F}) . There is a duality relationship between \mathcal{P} and its set of e-variables \mathcal{E} . To develop this perspective we make the following crucial definition, where we recall that M_+ denotes the set of all nonnegative measures on Ω .

Definition 3.1. The *effective null hypothesis* is the set

$$\mathcal{P}_{\text{eff}} = \{\mathbb{P} \in M_+ : \mathbb{E}_{\mathbb{P}}[X] \leq 1 \text{ for all } X \in \mathcal{E}\}.$$

The effective null hypothesis is the largest family of distributions whose set of e-variables is exactly \mathcal{E} . Thus it consists of those distributions against which e-variables in \mathcal{E} do not provide any evidence. In particular, a ‘nontrivial’ e-variable X , in the sense that $\mathbb{E}_{\mathbb{Q}}[X] > 1$, exists if and only if $\mathbb{Q} \notin \mathcal{P}_{\text{eff}}$. Because the constant $X = 1$ is always an e-variable, elements of \mathcal{P}_{eff} have total mass at most one, and sometimes strictly less than one. For example, if \mathbb{P} belongs to \mathcal{P}_{eff} , then so does any \mathbb{P}' in M_+ that is set-wise dominated by \mathbb{P} , meaning that $\mathbb{P}'(A) \leq \mathbb{P}(A)$ for all A . In Section 4 we discuss more explicit descriptions of the effective null; see Example 4.4, Lemma 4.5, and Corollary 4.6. See also Example 3.12 below.

Remark 3.2. There is a convex analytic interpretation of \mathcal{E} and \mathcal{P}_{eff} as the so-called *polar* and *bipolar* of \mathcal{P} , usually denoted by \mathcal{P}° and $\mathcal{P}^{\circ\circ}$.¹ For some of our results we make use of this connection by invoking a suitable *bipolar theorem*; see the proof of Lemma 4.5.

The following observation is simple, yet key to the theory developed in this paper.

Lemma 3.3. *The numeraire e-variable X^* gives rise to an element $\mathbb{P}^* \in \mathcal{P}_{\text{eff}}$ defined by $d\mathbb{P}^*/d\mathbb{Q} = 1/X^*$, understood as zero on $\{X^* = \infty\}$. Note that $\mathbb{P}^* \ll \mathbb{Q}$ by definition.*

Proof. Since X^* is \mathbb{Q} -almost surely unique and strictly positive, \mathbb{P}^* is well-defined. It belongs to \mathcal{P}_{eff} because $\mathbb{E}_{\mathbb{P}^*}[X] = \mathbb{E}_{\mathbb{Q}}[\mathbf{1}_{\{X^* < \infty\}} X/X^*] \leq \mathbb{E}_{\mathbb{Q}}[X/X^*] \leq 1$ for all $X \in \mathcal{E} = \mathcal{P}^\circ$ thanks to the numeraire property. \square

Our results show that \mathbb{P}^* can be interpreted as a *reverse information projection*, as we explain after Theorem 3.6 below. This concept first appeared in Csiszár and Tusnády (1984) and later in the PhD thesis of Li (1999, Chapter 4) as well as in Csiszar and Matus (2003). It was used centrally in the context of e-variables in Grünwald et al. (2024), and was developed further in a recent preprint by Lardy et al. (2024). This earlier work assumes the existence of a reference measure, an assumption we are able to dispense with but nonetheless examine further in Section 4.

Definition 3.4. We refer to the measure \mathbb{P}^* in Lemma 3.3 as the *reverse information projection* (RIPr) of \mathbb{Q} on \mathcal{P}_{eff} .

¹This terminology is slightly nonstandard in that F_+ , which appears in the definition of \mathcal{E} (see the notation paragraph in the introduction), is not a subset of a locally convex topological vector space, whereas the standard theory of dual pairs works with such spaces. There are however exceptions, for example the bipolar theorems of Brannath and Schachermayer (1999) and Bartl and Kupper (2019).

Our use of the RIPr terminology for \mathbf{P}^* is justified by what follows. Recall that the entropy of \mathbf{Q} relative to \mathbf{P} is defined by

$$H(\mathbf{Q} \mid \mathbf{P}) = \begin{cases} \mathbb{E}_{\mathbf{Q}} \left[\log \frac{d\mathbf{Q}}{d\mathbf{P}} \right] & \text{if } \mathbf{Q} \ll \mathbf{P}, \\ +\infty & \text{otherwise.} \end{cases}$$

We will show that \mathbf{P}^* minimizes the relative entropy $H(\mathbf{Q} \mid \mathbf{P})$ over \mathcal{P}_{eff} . For this it will be convenient to use a slightly different expression for the relative entropy. For any $\mathbf{P} \in M_+$ we let \mathbf{P}^a denote its absolutely continuous part with respect to \mathbf{Q} . One then has

$$H(\mathbf{Q} \mid \mathbf{P}) = \mathbb{E}_{\mathbf{Q}} \left[-\log \frac{d\mathbf{P}^a}{d\mathbf{Q}} \right]. \quad (3.1)$$

To see this, suppose first $d\mathbf{P}^a/d\mathbf{Q}$ is strictly positive \mathbf{Q} -almost surely. Then $\mathbf{Q} \ll \mathbf{P}^a \ll \mathbf{P}$ and $d\mathbf{Q}/d\mathbf{P} = d\mathbf{Q}/d\mathbf{P}^a = 1/(d\mathbf{P}^a/d\mathbf{Q})$ \mathbf{Q} -almost surely, so (3.1) holds. If instead $d\mathbf{P}^a/d\mathbf{Q}$ is zero with positive \mathbf{Q} -probability, then the right-hand side of (3.1) is infinite (the expectation is always well-defined by Jensen's inequality), and so is the relative entropy because $\mathbf{Q} \not\ll \mathbf{P}$. Thus (3.1) holds in this case, too.

Remark 3.5. In the above, \mathbf{P} can be any element of M_+ . We do not adjust the definition of relative entropy for the possibility that its total mass is not one.

We now make the trivial observation that

$$\mathbb{E}_{\mathbf{Q}} \left[X \frac{d\mathbf{P}^a}{d\mathbf{Q}} \right] \leq 1 \text{ for all } X \in \mathcal{E} \text{ and } \mathbf{P} \in \mathcal{P}_{\text{eff}}. \quad (3.2)$$

Indeed, the left-hand side equals $\mathbb{E}_{\mathbf{P}^a}[X]$ which is dominated by $\mathbb{E}_{\mathbf{P}}[X]$, hence by one. This leads to the following theorem, which expresses the duality between the null hypothesis and its set of e-variables and justifies calling \mathbf{P}^* the RIPr. We first treat the notationally simpler and statistically more relevant case where $\mathbf{Q} \ll \mathcal{P}$ and, hence, the numeraire X^* is finite. The general case is in Theorem 3.13.

Theorem 3.6. *Assume $\mathbf{Q} \ll \mathcal{P}$ and let \mathbf{P}^* be an element of \mathcal{P}_{eff} equivalent to \mathbf{Q} . The following statements are equivalent, where as above we write \mathbf{P}^a for the absolutely continuous part of \mathbf{P} with respect to \mathbf{Q} (and hence \mathbf{P}^*).*

- (i) \mathbf{P}^* is the RIPr.
- (ii) $\mathbb{E}_{\mathbf{Q}} \left[\frac{d\mathbf{P}^a}{d\mathbf{P}^*} \right] \leq 1$ for all $\mathbf{P} \in \mathcal{P}_{\text{eff}}$.
- (iii) $\mathbb{E}_{\mathbf{Q}} \left[\log \frac{d\mathbf{P}^a}{d\mathbf{P}^*} \right] \leq 0$ for all $\mathbf{P} \in \mathcal{P}_{\text{eff}}$.

Letting X^* be the numeraire and \mathbf{P}^* the RIPr, one has the strong duality relation

$$\mathbb{E}_{\mathbf{Q}}[\log X^*] = \sup_{X \in \mathcal{E}} \mathbb{E}_{\mathbf{Q}}[\log X] = \inf_{\mathbf{P} \in \mathcal{P}_{\text{eff}}} H(\mathbf{Q} \mid \mathbf{P}) = H(\mathbf{Q} \mid \mathbf{P}^*), \quad (3.3)$$

where some, and then all, of these quantities may be $+\infty$. Here $\mathbb{E}_{\mathbf{Q}}[\log X]$ is understood as $-\infty$ whenever $\mathbb{E}_{\mathbf{Q}}[(\log X)^-] = \infty$.

Property (ii) satisfied by the RIPr can be viewed as dual to the numeraire property and generalizes Property 2 in Li (1999, Theorem 4.3). The equivalence of (ii) and (iii) is analogous to the equivalence between the numeraire property and log-optimality of the numeraire established in Proposition 2.4; we elaborate on this analogy in Remark 3.7 below. Finally, (3.3) states that strong duality holds between the problem of maximizing the expected logarithm over e-variables and the problem of minimizing relative entropy over the effective null. Note that $\mathbb{E}_Q[\log X^*]$ is well-defined and nonnegative, possibly infinite, thanks to the numeraire property and the fact that $\log X^* \geq -1/X^* + 1$. In the spirit of Remark 2.5, we discuss a ‘benchmarked’ version of (3.3) in Remark 3.8 below.

The minimum relative entropy property of P^* justifies calling it the RIPr of Q on \mathcal{P}_{eff} as in Definition 3.4. In the terminology of Lardy et al. (2024), property (iii) states that the *description gain* of switching from P^* to any other element of \mathcal{P}_{eff} is nonpositive, and this provides another justification for calling P^* the RIPr. We elaborate on the connection to their work in Remark 3.9 below. Existing literature on the RIPr assumes the existence of a reference measure along with much stronger versions of the condition that $Q \ll P$. As we show in Section 4, in such settings our P^* coincides with existing notions of the RIPr of Q on \mathcal{P} ; see Theorem 4.7 and its corollaries. However, our measure P^* exists in complete generality and satisfies properties normally associated with the RIPr; see Theorem 3.13 below for the fully general case. This leads us to Definition 3.4, which we argue is the natural definition of the RIPr in the absence of any assumptions.

Proof of Theorem 3.6. (i) \Leftrightarrow (ii): Suppose P^* is the RIPr. Because P^* and Q are equivalent we have $dP^a/dP^* = (dP^a/dQ)/(dP^*/dQ) = X^*dP^a/dQ$ up to Q -nullsets for any $P \in \mathcal{P}_{\text{eff}}$. Thus we obtain (ii) by taking $X = X^*$ in (3.2). This proves the forward direction. To prove the converse, we show that the RIPr is the unique element of \mathcal{P}_{eff} that is equivalent to Q and satisfies (ii). This is argued exactly as uniqueness of the numeraire: if $P_1^*, P_2^* \in \mathcal{P}_{\text{eff}}$ are both equivalent to Q and satisfy $\mathbb{E}_Q[dP^a/dP_i^*] \leq 1$ for all $P \in \mathcal{P}_{\text{eff}}$, then with $Y = dP_2^*/dP_1^*$ we have $1 \leq 1/\mathbb{E}_Q[Y] \leq \mathbb{E}_Q[1/Y] \leq 1$. Thus Y is Q -almost surely equal to one and $P_1^* = P_2^*$.

(ii) \Leftrightarrow (iii): The forward implication follows from Jensen’s inequality, and the converse uses the same argument as the proof of Proposition 2.4. Specifically, $P(t) = (1-t)P^* + tP$ is in \mathcal{P}_{eff} for all $t \in (0, 1)$, so $\mathbb{E}_Q[t^{-1} \log(dP(t)^a/dP^*)] \leq 0$. Sending t to zero yields (ii).

Finally, we argue (3.3) assuming that P^* is the RIPr. Jensen’s inequality and (3.2) give

$$\mathbb{E}_Q \left[\log \left(X \frac{dP^a}{dQ} \right) \right] \leq 0 \tag{3.4}$$

for all $X \in \mathcal{E}$ and $P \in \mathcal{P}_{\text{eff}}$. This implies the ‘weak duality’ inequality

$$\mathbb{E}_Q [\log X] \leq \mathbb{E}_Q \left[-\log \frac{dP^a}{dQ} \right] = H(Q | P). \tag{3.5}$$

Indeed, if $(\log X)^-$ has infinite expectation, (3.5) is trivially true. Otherwise we first replace X by $X \wedge n$, in which case (3.4) can be rearranged to (3.5) with $X \wedge n$ in place of X . Sending n to infinity and using monotone convergence then gives (3.5). Since equality is achieved by $X = X^*$ and $P = P^*$, we deduce (3.3). \square

Remark 3.7. The equivalence of (ii) and (iii) in Theorem 3.6, and the equivalence stated in Proposition 2.4, have commonalities that can be understood by considering two optimization

problems. The first is to maximize $f(X) = \mathbb{E}_Q[\log(X/X^*)]$ over $X \in \mathcal{E}$. The second is to minimize $g(P) = \mathbb{E}_Q[-\log(dP^a/dP^*)]$ over $P \in \mathcal{P}_{\text{eff}}$. The numeraire property can be interpreted as the first-order condition that the Gateaux (i.e. directional) derivative $f'(X^*; X - X^*) = \lim_{t \downarrow 0} t^{-1}(f(X^* + t(X - X^*)) - f(X^*))$ in any feasible direction is nonpositive at X^* . Thanks to concavity, this is necessary and sufficient for optimality. Analogously, property (ii) of Theorem 3.6 states that the Gateaux derivative $g'(P^*; P - P^*) = \lim_{t \downarrow 0} t^{-1}(g(P^* + t(P - P^*)) - g(P^*))$ in any feasible direction is nonnegative at P^* . This is again necessary and sufficient for optimality, expressed by (iii). The first-order condition for f states that $dP^* = (1/X^*)dQ$ yields an element of \mathcal{P}_{eff} ; and the first-order condition for $g(P)$ states that $X^* = dQ/dP^*$ is an element of \mathcal{E} . Furthermore, the two objective functions are related by $f(X) \leq 0 \leq g(P)$, with equality precisely at the optimal solutions X^* and P^* . We refer to Ekeland and Témam (1999) for more information on convex optimization in infinite dimension; see in particular Chapter II therein.

Remark 3.8. Following Remark 2.5, let us consider a benchmark X_b . In the presence of a benchmark, the strong duality relation (3.3) takes the form

$$\mathbb{E}_Q \left[\log \frac{X^*}{X_b} \right] = \sup_{X \in \mathcal{E}} \mathbb{E}_Q \left[\log \frac{X}{X_b} \right] = \inf_{P \in \mathcal{P}_{\text{eff}}} \mathbb{E}_Q \left[-\log \left(\frac{dP^a}{dQ} X_b \right) \right] = \mathbb{E}_Q \left[-\log \left(\frac{dP^*}{dQ} X_b \right) \right],$$

where now all four quantities are finite. The optimizers X^* and P^* do not depend on the choice of benchmark. This ‘benchmarked’ version of strong duality is proved by replacing X by X/X_b and dP^a/dQ by $X_b(dP^a/dQ)$ in (3.4) and the subsequent manipulations.

Remark 3.9. In the spirit of Lardy et al. (2024), one can define the *maximum description gain* of switching from $P \in \mathcal{P}_{\text{eff}}$ to some other element of \mathcal{P}_{eff} by

$$H(Q | P \rightsquigarrow \mathcal{P}_{\text{eff}}) = \sup_{P' \in \mathcal{P}_{\text{eff}}} \mathbb{E}_Q \left[\log \frac{dP'^a}{dP} \right]$$

if Q is absolutely continuous with respect to P , and $+\infty$ otherwise. Thanks to condition (iii) of Theorem 3.6, assuming that $Q \ll \mathcal{P}$, the RIPr is characterized as the unique element of \mathcal{P}_{eff} equivalent to Q such that $H(Q | P^* \rightsquigarrow \mathcal{P}_{\text{eff}}) = 0$, and this is the smallest possible value that the maximum description gain can take.

Remark 3.10. The strong duality in (3.3) implies that there exists an e-variable with positive expected logarithm under Q if and only if $Q \notin \mathcal{P}_{\text{eff}}$. Therefore, thanks to Corollary 4.6 below, Theorem 3.6 extends a result by Zhang et al. (2024) showing that if \mathcal{P} is finite, then such an e-variable exists if and only if Q is not in the convex hull of \mathcal{P} .

Remark 3.11. The RIPr P^* is not a probability measure in general. Equivalently, it may happen that $\mathbb{E}_Q[1/X^*] < 1$. We will see this phenomenon in Sections 5.1 and 5.4, and other examples are given by Lardy et al. (2024). It is also worth noting that there may exist elements $\tilde{P} \in \mathcal{P}_{\text{eff}}$ that are different from P^* , but whose absolutely continuous part \tilde{P}^a is equal to P^* . We will see this in Section 5.4 with \tilde{P} being the symmetrization of Q . Other examples are obtained from Example 3.12 below.

Example 3.12 (Example 2.2 continued). Consider a simple null hypothesis $\mathcal{P} = \{P_0\}$ with $Q \ll P_0$ so that the numeraire X^* is finite. In this case $P^* = P_0^a$ and $X^* = 1/(dP^*/dQ) =$

dQ/dP_0 up to Q -nullsets. Here one can easily check that \mathcal{P}_{eff} consists of all sub-probability measures P such that $P(A) \leq P_0(A)$ for all A . Indeed, any such P must belong to \mathcal{P}_{eff} since $\mathbb{E}_P[X] \leq \mathbb{E}_{P_0}[X] \leq 1$ for all e -variables X . Conversely, if P is not of this form, then $P(A) > P_0(A)$ for some A , and $X = \mathbf{1}_A/P_0(A)$ is an e -variable such that $\mathbb{E}_P[X] > 1$, showing that $P \notin \mathcal{P}_{\text{eff}}$. If $P_0(A) = 0$ we set $X = \infty \mathbf{1}_A$ in the previous sentence.

We end this section with the general version of Theorem 3.6. The proof reduces the general case to the one treated in Theorem 3.6 by introducing the conditional probability measure $Q^*(A) = Q(A \mid X^* < \infty)$ if the quantity $\lambda^* = Q(X^* < \infty)$ is positive, where X^* is the numeraire. The fully degenerate case $\lambda^* = 0$ is treated separately. It is evident from Definition 3.4 that the RIPr P^* is absolutely continuous with respect to Q , and we know from Lemma 3.3 that it belongs to \mathcal{P}_{eff} . As shown below, one actually has more precise information, namely that

$$P^* \text{ is equivalent to } Q \text{ on } \{X^* < \infty\} \text{ and belongs to } \lambda^* \cdot \mathcal{P}_{\text{eff}}, \quad (3.6)$$

where we write $\lambda^* \cdot \mathcal{P}_{\text{eff}}$ for the set $\{\lambda^* P : P \in \mathcal{P}_{\text{eff}}\} \subset \mathcal{P}_{\text{eff}}$. We continue to write P^a for the absolutely continuous part of P with respect to Q . When applying this theorem it is useful to know that the event $\{X^* < \infty\}$ can be determined without first computing the numeraire; see Remark 3.14 below.

Theorem 3.13. *Let X^* be the numeraire e -variable and define $\lambda^* = Q(X^* < \infty)$. The RIPr P^* satisfies (3.6). Furthermore, for any sub-probability measure P^* that is absolutely continuous with respect to Q and satisfies (3.6), the following statements are equivalent.*

- (i) P^* is the RIPr.
- (ii) $\mathbb{E}_Q \left[\frac{dP^a}{dP^*} \mathbf{1}_{\{X^* < \infty\}} \right] \leq 1$ for all $P \in \mathcal{P}_{\text{eff}}$.
- (iii) $\mathbb{E}_Q \left[\log \left(\lambda^* \frac{dP^a}{dP^*} \right) \mathbf{1}_{\{X^* < \infty\}} \right] \leq 0$ for all $P \in \mathcal{P}_{\text{eff}}$.

If any of these hold, one has the strong duality relation (3.3). Finally, suppose $\lambda^* \in (0, 1)$, let P^* be the RIPr, and define $Q^* = Q(\cdot \mid X^* < \infty)$. Then the numeraire and RIPr associated with Q^* are X^* and $P^{**} = P^*/\lambda^*$, and one has the ‘conditional’ strong duality

$$\mathbb{E}_{Q^*}[\log X^*] = \sup_{X \in \mathcal{E}} \mathbb{E}_{Q^*}[\log X] = \inf_{P \in \mathcal{P}_{\text{eff}}} H(Q^* \mid P) = H(Q^* \mid P^{**}). \quad (3.7)$$

Proof. We first let P^* be the RIPr and check (3.6). Equivalence with Q on $\{X^* < \infty\}$ is evident from Definition 3.4. To show that P^* belongs to $\lambda^* \cdot \mathcal{P}_{\text{eff}}$ we refine the calculation in the proof of Lemma 3.3. Consider any $X \in \mathcal{E}$ and set $X' = X \mathbf{1}_{\{X^* < \infty\}} + \infty \mathbf{1}_{\{X^* = \infty\}}$. This still belongs to \mathcal{E} since it is P -almost surely equal to X for all $P \in \mathcal{P}$. Thanks to the convention (1.1) we have $X'/X^* = \mathbf{1}_{\{X^* < \infty\}} X'/X^* + \mathbf{1}_{\{X^* = \infty\}}$ and hence

$$\mathbb{E}_{P^*}[X] = \mathbb{E}_{P^*}[X'] = \mathbb{E}_Q \left[\mathbf{1}_{\{X^* < \infty\}} \frac{X'}{X^*} \right] = \mathbb{E}_Q \left[\frac{X'}{X^*} \right] - Q(X^* = \infty) \leq Q(X^* < \infty),$$

using the numeraire property in the last step. It follows that $P^* \in \lambda^* \cdot \mathcal{P}_{\text{eff}}$.

We now let \mathbb{P}^* be any sub-probability measure that is absolutely continuous with respect to \mathbb{Q} and satisfies (3.6), and prove the claimed equivalences and the strong duality statement. First, in the fully degenerate case $\lambda^* = 0$, (3.6) just states that $\mathbb{P}^* = 0$. This is also the RIPr, so (i) is true. Moreover, (ii) and (iii) reduce to the statements that $0 \leq 1$ and $0 \leq 0$. Thus all three statements are true and hence trivially equivalent. The strong duality (3.3) holds because all four terms are infinite. Second, in the ‘fully non-degenerate’ case $\lambda^* = 1$, the theorem simply reduces to Theorem 3.6, which we have already proved.

Consider now the case $\lambda^* \in (0, 1)$ and define $\mathbb{Q}^* = \mathbb{Q}(\cdot \mid X^* < \infty)$. Taking \mathbb{Q}^* as our new alternative hypothesis, X^* is still the numeraire. Indeed, it is a \mathbb{Q}^* -almost surely strictly positive e-variable, and for any other e-variable X we define X' as above and get

$$\mathbb{E}_{\mathbb{Q}^*} \left[\frac{X}{X^*} \right] = \frac{1}{\lambda^*} \mathbb{E}_{\mathbb{Q}} \left[\mathbf{1}_{\{X^* < \infty\}} \frac{X}{X^*} \right] = \frac{1}{\lambda^*} \left(\mathbb{E}_{\mathbb{Q}} \left[\frac{X'}{X^*} \right] - \mathbb{Q}(X^* = \infty) \right) \leq 1.$$

Since $\mathbb{Q}^*(X^* < \infty) = 1$ by construction, we have from Theorem 2.6 that $\mathbb{Q}^* \ll \mathcal{P}$. We also define $\mathbb{P}^{**} = \mathbb{P}^*/\lambda^*$ which, thanks to (3.6), is an element of \mathcal{P}_{eff} equivalent to \mathbb{Q}^* . Now, the statement (i) that \mathbb{P}^* is the RIPr of \mathbb{Q} is equivalent to \mathbb{P}^{**} being the RIPr of \mathbb{Q}^* . Furthermore, (ii) and (iii) can be equivalently expressed in terms of \mathbb{Q}^* and \mathbb{P}^{**} as

$$\mathbb{E}_{\mathbb{Q}^*} \left[\frac{d\mathbb{P}^a}{d\mathbb{P}^{**}} \right] \leq 1 \text{ for all } \mathbb{P} \in \mathcal{P}_{\text{eff}} \quad \text{and} \quad \mathbb{E}_{\mathbb{Q}^*} \left[\log \frac{d\mathbb{P}^a}{d\mathbb{P}^{**}} \right] \leq 0 \text{ for all } \mathbb{P} \in \mathcal{P}_{\text{eff}},$$

respectively. We may thus simply apply Theorem 3.6 with \mathbb{Q}^* and \mathbb{P}^{**} in place of \mathbb{Q} and \mathbb{P}^* to deduce that the three statements are indeed equivalent and that the ‘conditional’ strong duality (3.7) holds. The ‘unconditional’ strong duality (3.3) holds because all four terms are infinite. \square

Remark 3.14. The event $\{X^* < \infty\}$ can be determined without first computing the numeraire. Let $\mathbb{Q} = \mathbb{Q}^r + \mathbb{Q}^s$ be the generalized Lebesgue decomposition of \mathbb{Q} with respect to \mathcal{P} given by Lemma A.1, and let N be the event given in the lemma. Since N is \mathcal{P} -negligible, $X = \infty \mathbf{1}_N$ is an e-variable, so the numeraire must be infinite on N , \mathbb{Q} -almost surely. On the other hand, because the event $N' = \{X^* = \infty\}$ is \mathcal{P} -negligible, the lemma also yields $\mathbb{Q}(N' \setminus N) = 0$. Thus N and $\{X^* = \infty\}$ are both \mathcal{P} -negligible and coincide up to \mathbb{Q} -nullsets, so the former can be used in place of the latter in Theorem 3.13.

Following up on Example 3.12, the following corollary records the RIPr and numeraire for a point null, generalizing a well known result that the likelihood ratio is the optimal e-variable if the two distributions are equivalent; see for example Shafer (2021).

Corollary 3.15. *Consider an arbitrary point null $\mathcal{P} = \{\mathbb{P}_0\}$. The RIPr is the absolutely continuous part $\mathbb{P}^* = \mathbb{P}_0^a$ with respect to \mathbb{Q} , and the numeraire is*

$$X^* = \infty \mathbf{1}_N + \frac{d\mathbb{Q}^*}{d\mathbb{P}_0} \mathbf{1}_{N^c}, \quad (3.8)$$

where N is a set such that $\mathbb{P}_0(N) = 0$ and \mathbb{Q} restricted to N^c is absolutely continuous with respect to \mathbb{P}_0 , and $\mathbb{Q}^* = \mathbb{Q}(\cdot \mid N^c)$ whenever $\mathbb{Q}(N^c) > 0$.

Proof. Remark 3.14 implies that N equals the event where the numeraire is infinite up to \mathbb{Q} -nullsets, so \mathbb{Q}^* defined in the corollary coincides with \mathbb{Q}^* defined in Theorem 3.13. Moreover, thanks to Example 3.12, the numeraire for \mathbb{Q}^* is $d\mathbb{Q}^*/d\mathbb{P}_0$ up to \mathbb{Q}^* -nullsets. From these facts and Theorem 3.13, we conclude that the numeraire for \mathbb{Q} is given by (3.8). Finally, since $1/X^* = d\mathbb{P}_0^a/d\mathbb{Q}$ up to \mathbb{Q} -nullsets, the RIPr is $\mathbb{P}^* = \mathbb{P}_0^a$ as claimed. \square

4 Finding the numeraire and RIPr

So far we have studied the abstract existence, uniqueness, and duality theory of the numeraire and RIPr, but we have not said much about how to actually compute them. In this section we develop tools for doing so, which we then apply to several examples in Section 5.

We focus on the statistically most relevant situation where $\mathbb{Q} \ll \mathcal{P}$. In this case the numeraire X^* , and indeed every e-variable, is \mathbb{Q} -almost surely finite and the RIPr \mathbb{P}^* is equivalent to \mathbb{Q} . The general case can be reduced to this case by replacing \mathbb{Q} with the conditional measure \mathbb{Q}^* in Theorem 3.13 (except in the fully degenerate situation where the numeraire is infinity and the RIPr zero.)

Proposition 2.4 suggests that X^* can be found by solving an optimization problem. Unfortunately this may not be straightforward because the optimization is over the set of *all* e-variables, which can be difficult to characterize. For the same reason, even if a candidate X^* has been found it may not be clear how to check that it does, in fact, satisfy the numeraire property. Alternatively, the strong duality of Theorem 3.6 suggests minimizing relative entropy to find \mathbb{P}^* . However, the optimization is again over a set that can be difficult to describe explicitly, in this case the effective null \mathcal{P}_{eff} . Moreover, checking that a candidate \mathbb{P}^* is optimal involves comparing it to *all* elements of \mathcal{P}_{eff} .

We first present a verification theorem which simplifies the task of checking that candidates X^* and \mathbb{P}^* are in fact optimal. This result strengthens Corollary 2 of Grünwald et al. (2024). Despite the simple proof, the result turns out to be remarkably useful.

Theorem 4.1. *Assume that $\mathbb{Q} \ll \mathcal{P}$. Let X^* be a \mathbb{Q} -almost surely strictly positive e-variable and let \mathbb{P}^* be the measure given by $d\mathbb{P}^*/d\mathbb{Q} = 1/X^*$. Then X^* is a numeraire if and only if \mathbb{P}^* belongs to the effective null \mathcal{P}_{eff} . Thus the numeraire is the only e-variable that is also a likelihood ratio between \mathbb{Q} and some equivalent element of \mathcal{P}_{eff} . Finally, $\mathbb{E}_{\mathbb{P}^*}[X^*] = 1$ and \mathbb{P}^* is maximal in the sense that no other $\mathbb{P} \in \mathcal{P}_{\text{eff}}$ equivalent to \mathbb{Q} can satisfy $\mathbb{P}(A) \geq \mathbb{P}^*(A)$ for all A with strict inequality for some A .*

Proof. Recall that all e-variables are \mathbb{Q} -almost surely finite thanks to Theorem 2.6 and the assumption that $\mathbb{Q} \ll \mathcal{P}$; in particular, \mathbb{P}^* is equivalent to \mathbb{Q} . Now, if X^* is a numeraire, then Lemma 3.3 yields that $\mathbb{P}^* \in \mathcal{P}_{\text{eff}}$. For the converse, the definition of \mathbb{P}^* and the fact that it belongs to \mathcal{P}_{eff} yield $\mathbb{E}_{\mathbb{Q}}[X/X^*] = \mathbb{E}_{\mathbb{P}^*}[X] \leq 1$ for any e-variable X . This is the numeraire property. Finally, the definition of \mathbb{P}^* immediately gives $\mathbb{E}_{\mathbb{P}^*}[X^*] = 1$, and \mathbb{P}^* must be maximal because if a ‘dominating’ equivalent $\mathbb{P} \in \mathcal{P}_{\text{eff}}$ existed we would get the contradiction $1 = \mathbb{E}_{\mathbb{P}^*}[X^*] < \mathbb{E}_{\mathbb{P}}[X^*] \leq 1$. \square

4.1 Generated null hypotheses

The following corollary of Theorem 4.1 applies in settings where the null hypothesis is ‘generated’, in the sense of (4.1) below, by a subset \mathcal{E}_0 of e-variables. Null hypotheses of this kind

were studied by [Casgrain et al. \(2024\)](#) in a sequential setting. The corollary shows that once a candidate X^* has been found, it is enough to check the numeraire property for the generating family \mathcal{E}_0 , provided that it holds with equality for the e-variable which is identically one. The latter condition means that the RIPr \mathcal{P}^* is a probability measure.

Corollary 4.2. *Assume that $\mathcal{Q} \ll \mathcal{P}$. Let \mathcal{E}_0 be a family of $[0, \infty]$ -valued random variables that generates the null hypothesis in the sense that*

$$\mathcal{P} = \{\mathbb{P} \in M_1: \mathbb{E}_{\mathbb{P}}[X] \leq 1 \text{ for all } X \in \mathcal{E}_0\}. \quad (4.1)$$

Let X^* be a \mathcal{Q} -almost surely strictly positive e-variable such that

$$\mathbb{E}_{\mathcal{Q}} \left[\frac{1}{X^*} \right] = 1 \text{ and } \mathbb{E}_{\mathcal{Q}} \left[\frac{X}{X^*} \right] \leq 1 \text{ for all } X \in \mathcal{E}_0. \quad (4.2)$$

Then X^* is the numeraire and the RIPr belongs to \mathcal{P} .

Proof. Let \mathbb{P}^* be given by $d\mathbb{P}^*/d\mathcal{Q} = 1/X^*$. Then (4.2) says that \mathbb{P}^* is a probability measure such that $\mathbb{E}_{\mathbb{P}^*}[X] \leq 1$ for all $X \in \mathcal{E}_0$. Thus by (4.1), \mathbb{P}^* belongs to \mathcal{P} , hence to \mathcal{P}_{eff} . The result now follows from Theorem 4.1. \square

Thanks to Lemma 2.3, Corollary 4.2 also holds if one has ‘ \supset ’ instead of ‘ $=$ ’ in (4.1). This is intuitive: if one enlarges the generating set \mathcal{E}_0 beyond what is necessary to specify \mathcal{P} , then the right hand side of (4.1) can become smaller than \mathcal{P} . In that case, (4.2) only gets harder to satisfy, and if we find an e-variable X^* satisfying it, it surely must still be the numeraire.

One may wonder if the equality in (4.2) could be replaced by an inequality; this would amount to a significant strengthening of the corollary. The following example shows that this is unfortunately not possible.

Example 4.3. *Consider the single coin toss space $\Omega = \{0, 1\}$ and the random variable X_0 given by $X_0(0) = 2$ and $X_0(1) = 1/2$. Let the singleton set $\mathcal{E}_0 = \{X_0\}$ generate the null hypothesis \mathcal{P} , which then consists of all probability measures \mathbb{P} with $\mathbb{P}(0) \leq 1/3$. Next, let the alternative \mathcal{Q} be given by $\mathcal{Q}(0) = 1$, $\mathcal{Q}(1) = 0$. We then have $\mathbb{E}_{\mathcal{Q}}[1/X_0] = 1/2 < 1$ and, trivially, $\mathbb{E}_{\mathcal{Q}}[X/X_0] \leq 1$ for $X \in \mathcal{E}_0$. If Corollary 4.2 were to remain true with an inequality instead of an equality in (4.2), we would conclude that X_0 is the numeraire. However, this is not so. Indeed, for the e-variable X^* given by $X^*(0) = 3$, $X^*(1) = 0$, we have $\mathbb{E}_{\mathcal{Q}}[X^*/X_0] = 3/2 > 1$. As suggested by the notation, X^* is in fact the numeraire in this example. To see this, just note that every e-variable X must satisfy $X(0) \leq 3$ since \mathcal{P} contains the probability measure \mathbb{P} with $\mathbb{P}(0) = 1/3$, and hence $\mathbb{E}_{\mathcal{Q}}[X/X^*] = X(0)/3 \leq 1$. It is also worth noting that $\mathbb{E}_{\mathcal{Q}}[1/X^*] = 1/3 < 1$, which means that we cannot use Corollary 4.2 to detect the numeraire property of X^* . The corollary is not a universal tool, but nonetheless a useful one as we will see in Section 5.*

4.2 When a reference measure exists

We now consider the dual approach of finding \mathbb{P}^* directly, either by minimizing relative entropy or, if a candidate is available, checking that (ii) or (iii) of Theorem 3.6 are satisfied. The difficulty is that the domain is all of \mathcal{P}_{eff} , which can be hard to describe explicitly. Matters would become simpler if one could instead work directly with \mathcal{P} , but this is not possible in general. Here is one example of what can go wrong.

Example 4.4. *This example shows that Theorem 3.6(ii) may no longer characterize the RIPr \mathcal{P}^* uniquely if \mathcal{P}_{eff} is replaced by \mathcal{P} . Let $\Omega = [0, 1]$ with its Borel σ -algebra, and let $\mathcal{P} = \{\delta_\omega : \omega \in [0, 1]\}$ be the set of all Dirac point masses. In this case the set of e-variables is trivial in the sense that a random variable X is an e-variable if and only if $0 \leq X(\omega) \leq 1$ for all $\omega \in [0, 1]$. It follows that \mathcal{P}_{eff} contains all sub-probabilities on Ω . Now, let the alternative hypothesis \mathcal{Q} be the uniform distribution on the unit interval. We then have the extreme situation that \mathcal{P} does not contain any distribution absolutely continuous with respect to \mathcal{Q} , whereas \mathcal{P}_{eff} contains all such distributions. In particular, Theorem 3.6(ii) would be vacuously true for any candidate \mathcal{P}^* if \mathcal{P}_{eff} were replaced by \mathcal{P} . In contrast, as currently stated, Theorem 3.6(ii) will only hold for $\mathcal{P}^* = \mathcal{Q}$, which is the correct answer.*

The situation improves in the presence of a reference measure μ with respect to which \mathcal{Q} and all $\mathbb{P} \in \mathcal{P}$ are absolutely continuous. In this setting, if \mathbb{P}_n and \mathbb{P} are measures with densities $p_n = d\mathbb{P}_n/d\mu$ and $p = d\mathbb{P}/d\mu$ we say that \mathbb{P}_n μ -converges to \mathbb{P} if p_n converges to p in probability under μ . A set C of absolutely continuous measures is called μ -closed if it is closed with respect to μ -convergence. It is called *solid* if it contains every (nonnegative) measure \mathbb{P}' that is dominated by some $\mathbb{P} \in C$ in the sense that $\mathbb{P}'(A) \leq \mathbb{P}(A)$ for all A . The following result characterizes \mathcal{P}_{eff} in terms of these concepts.

Lemma 4.5. *Suppose all $\mathbb{P} \in \mathcal{P}$ are absolutely continuous with respect to a probability measure μ . Then all elements of \mathcal{P}_{eff} are also absolutely continuous with respect to μ , and \mathcal{P}_{eff} is the smallest μ -closed convex solid set that contains \mathcal{P} .*

Proof. We identify \mathcal{P} with its set of densities $C = \{p \in L_+^0(\mu) : p = d\mathbb{P}/d\mu, \mathbb{P} \in \mathcal{P}\}$. The polar of C in $L_+^0(\mu)$ is $C^\circ = \{f \in L_+^0(\mu) : \mathbb{E}_\mu[pf] \leq 1 \text{ for all } p \in C\}$, which is the set of (equivalence classes of) random variables that are μ -almost surely equal to a finite e-variable. We claim that all elements of \mathcal{P}_{eff} are absolutely continuous with respect to μ , and that its set of densities is precisely the bipolar $C^{\circ\circ} = \{p \in L_+^0(\mu) : \mathbb{E}_\mu[pf] \leq 1 \text{ for all } f \in C^\circ\}$. To check absolute continuity, let $\mathbb{P} \in \mathcal{P}_{\text{eff}}$ and consider any event A such that $\mu(A) = 0$. Then $\infty \mathbf{1}_A$ is an e-variable, hence $\mathbb{E}_\mathbb{P}[\infty \mathbf{1}_A] \leq 1$, and thus $\mathbb{P}(A) = 0$. This shows absolute continuity and lets us identify \mathcal{P}_{eff} with its set of densities, namely, the set of $p \in L_+^0(\mu)$ such that $\mathbb{E}_\mu[pX] \leq 1$ for all finite e-variables X . Here we may restrict to finite e-variables because $\mathbb{E}_\mu[pX] \leq 1$ if and only if $\mathbb{E}_\mu[p(X \wedge n)] \leq 1$ for all $n \in \mathbb{N}$. The finite e-variables are, up to μ -nullsets, precisely the elements of C° , so we deduce that \mathcal{P}_{eff} is indeed identified with $C^{\circ\circ}$. Finally, the bipolar theorem of Brannath and Schachermayer (1999, Theorem 1.3) states that $C^{\circ\circ}$ is the smallest μ -closed convex solid set that contains C . This completes the proof. \square

Corollary 4.6. *If \mathcal{P} is finite, then \mathcal{P}_{eff} consists of all sub-probabilities that are set-wise dominated by an element of $\text{conv}(\mathcal{P})$. In particular, $\mathcal{P}_{\text{eff}} \cap M_1 = \text{conv}(\mathcal{P})$.*

Proof. For finite $\mathcal{P} = \{\mathbb{P}_1, \dots, \mathbb{P}_n\}$ we may take $\mu = (\mathbb{P}_1 + \dots + \mathbb{P}_n)/n$ as reference measure. In this case $\text{conv}(\mathcal{P})$ is already μ -closed and the form of \mathcal{P}_{eff} follows from Lemma 4.5. \square

Lemma 4.5 is powerful because it characterizes *all* elements of \mathcal{P}_{eff} . The following theorem makes use of this to simplify the conditions of Theorem 3.6. Lemma 4.5 also provides a way of checking that a particular measure belongs to \mathcal{P}_{eff} . For example, it shows that limits in probability of (densities of) measures in the convex hull of \mathcal{P} must belong to \mathcal{P}_{eff} . The latter

is however also easy to check directly from the definition of \mathcal{P}_{eff} and Fatou's lemma, and does not require the bipolar theorem used in the proof of Lemma 4.5.

If a reference measure μ is given, we identify any $P \in \mathcal{P}$ with its density $p = dP/d\mu$. Thanks to Lemma 4.5, we may similarly identify \mathcal{P}_{eff} with its set of densities. For the sake of brevity we abuse notation and write, for example, $p \in \mathcal{P}$ even though, strictly speaking, \mathcal{P} is a set of measures, not densities. Next, given Q with density q , the absolutely continuous part of any $p \in \mathcal{P}_{\text{eff}}$ has density $p^a = p\mathbf{1}_{\{q>0\}}$, and this is Q -almost surely equal to p itself. In view of these observations, the conditions (ii) and (iii) of Theorem 3.6 characterizing the RIPr state, respectively, that $\mathbb{E}_Q[p/p^*] \leq 1$ and $\mathbb{E}_Q[\log(p/p^*)] \leq 0$, for all $p \in \mathcal{P}_{\text{eff}}$. Here $p^* = dP^*/d\mu$. The following theorem simplifies these conditions by showing that, essentially, it is enough to check them for elements of \mathcal{P} rather than \mathcal{P}_{eff} .

Theorem 4.7. *Assume that $Q \ll \mathcal{P}$ and that Q and all elements of \mathcal{P} are absolutely continuous with respect to a probability measure μ . Let $p^* \in \mathcal{P}_{\text{eff}}$ be equivalent to Q . Then each of the conditions (i), (ii), and (iii) of Theorem 3.6 is equivalent to*

$$\mathbb{E}_Q \left[\frac{p}{p^*} \right] \leq 1 \text{ for all } p \in \mathcal{P}, \quad (4.3)$$

as well as to

$$\mathbb{E}_Q \left[\log \frac{(1-t)p^* + tp}{p^*} \right] \leq 0 \text{ for all } t \in (0,1) \text{ and } p \in \mathcal{P}. \quad (4.4)$$

In either case, p^* is (the density of) the RIPr.

Remark 4.8. The reason for the asymmetry between Theorem 3.6(iii) and (4.4) is that the latter only involves p from \mathcal{P} , not \mathcal{P}_{eff} . Since \mathcal{P} may not contain p^* , nor be convex, we need to explicitly include convex combinations with p^* in the numerator on the left-hand side of (4.4). Let us also emphasize that $t = 1$ is excluded in (4.4). This sometimes makes the condition easier to check, for instance in the proof of Corollary 4.10 below.

Proof of Theorem 4.7. We know that (i), (ii), and (iii) of Theorem 3.6 are equivalent for elements of \mathcal{P}_{eff} that are equivalent to Q . To prove the theorem it is therefore enough to show that that Theorem 3.6(ii), (4.3), and (4.4) are equivalent.

We first show that Theorem 3.6(ii) is equivalent to (4.3). The former is clearly a stronger condition because it states that (4.3) holds with the larger set \mathcal{P}_{eff} in place of \mathcal{P} . For the converse, suppose (4.3) holds. Let \mathcal{P}' be the set of all sub-probability densities dominated by some element of the convex hull of \mathcal{P} , and let \mathcal{P}'' be the closure in μ -probability of \mathcal{P}' . It is clear that (4.3) holds for all $p \in \mathcal{P}'$ and then, by Fatou's lemma, also for all $p \in \mathcal{P}''$. We claim that \mathcal{P}'' is a closed convex solid set that contains \mathcal{P} ; it must then also contain \mathcal{P}_{eff} since this is the *smallest* such set by Lemma 4.5. From this we conclude that Theorem 3.6(ii) holds. To see that \mathcal{P}'' is closed, convex, and solid, note that by definition, every $p'' \in \mathcal{P}''$ is the limit in probability of some $p'_n \in \mathcal{P}'$, meaning that $p'_n \leq p_n$ for some p_n in the convex hull of \mathcal{P} . From this description convexity follows directly, closedness is obtained from a diagonal argument, and solidity follows by noting that any $\tilde{p} \leq p''$ is the limit of $\tilde{p} \wedge p'_n$, which is still dominated by p_n and thus belongs to \mathcal{P}' .

The equivalence of (4.3) and (4.4) is argued as the equivalence of (ii) and (iii) of Theorem 3.6, or the equivalence of the numeraire property and log-optimality in Proposition 2.4.

Specifically, the forward implication uses Jensen's inequality and (4.3) to get $\mathbb{E}_{\mathbf{Q}}[\log(1 - t + tp/p^*)] \leq \log(1 - t + t\mathbb{E}_{\mathbf{Q}}[p/p^*]) \leq 0$, whereas for the reverse implication one divides (4.4) by t , which is then sent to zero to yield (4.3). \square

Even with a reference measure, it is possible that the relative entropy $H(\mathbf{Q} \mid \mathbf{P})$ is infinite for all $\mathbf{P} \in \mathcal{P}$ but finite for the RIPr \mathbf{P}^* , similarly to what happens in Example 4.4. The following example illustrates this.

Example 4.9. *Let $\Omega = \mathbb{R}$ with its Borel σ -algebra, let \mathcal{P} be the set of all finite mixtures of Gaussians, and let \mathbf{Q} be the standard Cauchy distribution. We may take $\mu = \mathbf{Q}$ as the reference measure, for example. With this setup, $H(\mathbf{Q} \mid \mathbf{P}) = \infty$ for all $\mathbf{P} \in \mathcal{P}$. On the other hand, the Cauchy density can be obtained as the limit in μ -probability (even in the uniform norm) of elements of \mathcal{P} , so it follows from Lemma 4.5 that $\mathbf{Q} \in \mathcal{P}_{\text{eff}}$. Thus the RIPr is $\mathbf{P}^* = \mathbf{Q}$ and $H(\mathbf{Q} \mid \mathbf{P}^*) = 0$.*

We continue to identify measures with their μ -densities. Theorem 4.7 can be used to recover the following result of Lardy et al. (2024), who define the *maximum description gain* of switching from a particular $p \in \mathcal{P}$ to some other element of \mathcal{P} by

$$H(\mathbf{Q} \mid p \rightsquigarrow \mathcal{P}) = \sup_{p' \in \mathcal{P}} \mathbb{E}_{\mathbf{Q}} \left[\log \frac{p'}{p} \right]$$

if \mathbf{Q} is absolutely continuous with respect to p , and $+\infty$ otherwise. Note that this differs from the quantity $H(\mathbf{Q} \mid \mathbf{P} \rightsquigarrow \mathcal{P}_{\text{eff}})$ in Remark 3.9, which involves \mathcal{P}_{eff} rather than \mathcal{P} .

Corollary 4.10. *Assume that $\mathbf{Q} \ll \mathcal{P}$ and that all elements of \mathcal{P} are absolutely continuous with respect to a probability measure μ . If \mathcal{P} is convex and $(p_n)_{n \in \mathbb{N}}$ is a sequence in \mathcal{P} with $H(\mathbf{Q} \mid p_n \rightsquigarrow \mathcal{P}) \rightarrow 0$, then $\log p_n$ converges to $\log p^*$ in $L^1(\mathbf{Q})$, where p^* is the RIPr.*

Proof. The key observation of Lardy et al. (2024) is that the description gain controls the $L^1(\mathbf{Q})$ distance between the log-densities $\log p_n$, showing that they form a Cauchy sequence. We give a self-contained version of their argument. We claim that for any $t > 0$ one has

$$\frac{1}{9} (|\log t| \wedge |\log t|^2) \leq \log \left(\frac{2+t+t^{-1}}{4} \right) = \log \left(\frac{1+t^{-1}}{2} \right) + \log \left(\frac{t+1}{2} \right). \quad (4.5)$$

To verify the inequality, note that it is equivalent to $f(x) \geq 0$ for $x > 0$, where $f(x) = \log((2 + e^x + e^{-x})/4) - (x \wedge x^2)/9$. Since $f''(x) = 2/(2 + e^x + e^{-x}) - (2/9)\mathbf{1}_{(0,1)}(x) > 0$ for $x \neq 1$, and since $f(0) = f'(0) = 0$, we deduce that $f(x) > 0$ for $x > 0$.

Setting $t = p_m/p_n$ in (4.5), taking expectation under \mathbf{Q} , and using that \mathcal{P} is convex and hence contains $(p_m + p_n)/2$, we get

$$\frac{1}{9} \mathbb{E}_{\mathbf{Q}} \left[\left| \log \frac{p_m}{p_n} \right| \wedge \left| \log \frac{p_m}{p_n} \right|^2 \right] \leq \mathbb{E}_{\mathbf{Q}} \left[\log \left(\frac{p_m + p_n}{2p_m} \right) \right] + \mathbb{E}_{\mathbf{Q}} \left[\log \left(\frac{p_m + p_n}{2p_n} \right) \right] \leq H_m + H_n,$$

where we use the shorthand notation $H_n = H(\mathbf{Q} \mid p_n \rightsquigarrow \mathcal{P})$. From the general inequality $\mathbb{E}[Y] \leq \mathbb{E}[Y \wedge Y^2] + \sqrt{\mathbb{E}[Y \wedge Y^2]}$ for nonnegative random variables Y we then deduce

$$\mathbb{E}_{\mathbf{Q}} [|\log p_m - \log p_n|] \leq 9(H_m + H_n) + 3\sqrt{H_m + H_n}.$$

This shows that $\log p_n$, $n \in \mathbb{N}$, is a Cauchy sequence in $L^1(\mathbb{Q})$ and converges to some limit which we may write as $\log p^*$ for some \mathbb{Q} -almost surely positive random variable p^* . To ensure that p^* is the density of a measure that is equivalent to \mathbb{Q} we choose it to be zero on the \mathbb{Q} -nullset $\{q = 0\}$, where $q = d\mathbb{Q}/d\mu$. This is easily achieved by replacing p^* with $p^* \mathbf{1}_{\{q > 0\}}$ if necessary. Now, convergence in L^1 implies convergence in probability, which is preserved under continuous transformations, so p_n converges to p^* in probability under \mathbb{Q} . Furthermore, for any $t \in (0, 1)$ and $p \in \mathcal{P}$ we have $\log \frac{(1-t)p_n + tp}{p_n} \geq \log(1-t) > -\infty$, so Fatou's lemma yields

$$\mathbb{E}_{\mathbb{Q}} \left[\log \frac{(1-t)p^* + tp}{p^*} \right] \leq \liminf_{n \rightarrow \infty} \mathbb{E}_{\mathbb{Q}} \left[\log \frac{(1-t)p_n + tp}{p_n} \right] \leq \lim_{n \rightarrow \infty} H_n = 0 \text{ for all } p \in \mathcal{P}.$$

Theorem 4.7 now shows that p^* is the RIPr, provided we can argue that p^* belongs to \mathcal{P}_{eff} . This will use that $p^* = 0$ on $\{q = 0\}$. Specifically, note that \mathcal{P}_{eff} contains p_n and hence, being solid, also the smaller random variable $p_n \mathbf{1}_{\{q > 0\}}$. Convergence in \mathbb{Q} -probability of p_n to p^* implies convergence in μ -probability of $p_n \mathbf{1}_{\{q > 0\}}$ to $p^* \mathbf{1}_{\{q > 0\}} = p^*$. Since \mathcal{P}_{eff} is closed in μ -probability, it must contain p^* . This completes the proof. \square

The maximum description gain is useful in situations where the relative entropy $H(\mathbb{Q} | p)$ is infinite for all $p \in \mathcal{P}$. On the other hand, in many cases the relative entropies are finite and one recovers the original result of Li (1999) on existence of the reverse information projection, as was also observed by Lardy et al. (2024).

Corollary 4.11. *Assume that $\mathbb{Q} \ll \mathcal{P}$ and that all elements of \mathcal{P} are absolutely continuous with respect to a probability measure μ . If \mathcal{P} is convex, $H(\mathbb{Q} | p) < \infty$ for some $p \in \mathcal{P}$, and $(p_n)_{n \in \mathbb{N}}$ is an entropy minimizing sequence in \mathcal{P} , meaning that $H(\mathbb{Q} | p_n) \rightarrow \inf_{p \in \mathcal{P}} H(\mathbb{Q} | p)$, then $\log p_n$ converges to $\log p^*$ in $L^1(\mathbb{Q})$, where p^* is the RIPr. In particular, if some $p' \in \mathcal{P}$ achieves the infimum, then p' is the RIPr. Finally, $H(\mathbb{Q} | p^*) = \inf_{p \in \mathcal{P}} H(\mathbb{Q} | p)$.*

Proof. Since the quantities involved are finite for n large enough, $\lim_{n \rightarrow \infty} H(\mathbb{Q} | p_n \rightsquigarrow \mathcal{P}) = \lim_{n \rightarrow \infty} H(\mathbb{Q} | p_n) - \inf_{p \in \mathcal{P}} H(\mathbb{Q} | p) = 0$, and the first claim follows from Corollary 4.10. The second claim follows by simply taking $p_n = p'$ for all n . The final claim follows because $H(\mathbb{Q} | p_n) \rightarrow H(\mathbb{Q} | p^*)$ thanks to the $L^1(\mathbb{Q})$ convergence of the log-densities. \square

Corollaries 4.10 and 4.11 confirm that the definition of the RIPr used in this paper, Definition 3.4, is consistent with existing definitions in the literature. Specifically, Li (1999) defines the RIPr as the density p^* with the property that whenever $(p_n)_{n \in \mathbb{N}}$ is a sequence in \mathcal{P} such that $H(\mathbb{Q} | p_n) \rightarrow \inf_{p \in \mathcal{P}} H(\mathbb{Q} | p)$, we have $\log p_n \rightarrow \log p^*$ in $L^1(\mathbb{Q})$. Corollary 4.11 states that this indeed yields the RIPr of Definition 3.4. Similarly, Lardy et al. (2024) define the RIPr as the density p^* with the property that whenever $(p_n)_{n \in \mathbb{N}}$ is a sequence in \mathcal{P} such that $H(\mathbb{Q} | p_n \rightsquigarrow \mathcal{P}) \rightarrow \inf_{p \in \mathcal{P}} H(\mathbb{Q} | p \rightsquigarrow \mathcal{P})$ (which must be zero whenever it is finite), p_n converges to p^* in a certain metric which is equivalent to $L^1(\mathbb{Q})$ -convergence of the log-densities. Thanks to Corollary 4.10, this again yields the RIPr of Definition 3.4.

5 Examples

We now turn to examples. In these examples we consider a single random observation Z in a measurable space \mathcal{Z} . For concreteness we take (Ω, \mathcal{F}) to be \mathcal{Z} with its σ -algebra, and

Z the canonical random variable. We will always have $\mathbb{Q} \ll \mathcal{P}$. We start with a simple parametric example, before moving on to more sophisticated nonparametric examples where there exists no reference measure, in order to display the power and generality of our theory. We end with another parametric example. Except for the first and the last example, past work like [Wasserman et al. \(2020\)](#); [Grünwald et al. \(2024\)](#); [Lardy et al. \(2024\)](#) cannot handle the situations discussed here.

5.1 A parametric null

We start with a simple parametric example from [Lardy et al. \(2024\)](#) involving a real-valued observation Z in $\mathcal{Z} = \mathbb{R}$, and then generalize our observations at the end of this subsection. The null hypothesis \mathcal{P} consists of two unit-variance Gaussians \mathbb{P}_1 and \mathbb{P}_2 with mean $+1$ and -1 , and the alternative hypothesis \mathbb{Q} is the standard Cauchy distribution. This example was considered by the aforementioned authors because $H(\mathbb{Q} | \mathbb{P}_j) = \infty$ for $j \in \{1, 2\}$.

We claim that the RIPr is $\mathbb{P}^* = (\mathbb{P}_1 + \mathbb{P}_2)/2$. This probability measure is equivalent to \mathbb{Q} and belongs to \mathcal{P}_{eff} , which always contains the convex hull of \mathcal{P} . Furthermore, writing p_1, p_2, p^*, q for the Lebesgue densities, we have $p_2(z) = p_1(-z)$ and $q(z) = q(-z)$ and hence

$$\mathbb{E}_{\mathbb{Q}} \left[\frac{p_1(Z)}{p^*(Z)} \right] = \int_{-\infty}^{\infty} \frac{2p_1(z)q(z)}{p_1(z) + p_2(z)} dz = \mathbb{E}_{\mathbb{Q}} \left[\frac{p_2(Z)}{p^*(Z)} \right].$$

Since $p_1 + p_2 = 2p^*$ we also have $\mathbb{E}_{\mathbb{Q}}[p_1(Z)/p^*(Z)] + \mathbb{E}_{\mathbb{Q}}[p_2(Z)/p^*(Z)] = 2$, and conclude that

$$\mathbb{E}_{\mathbb{Q}} \left[\frac{p_i(Z)}{p^*(Z)} \right] = 1, \quad i = 1, 2.$$

Thus (4.3) of Theorem 4.7 is satisfied, and \mathbb{P}^* is the RIPr. The numeraire is

$$X^* = \frac{2q(Z)}{p_1(Z) + p_2(Z)}. \tag{5.1}$$

It is worth remarking that the universal inference e-variable ([Wasserman et al., 2020](#)) would be given by $q(Z)/\max\{p_1(Z), p_2(Z)\}$, which is clearly smaller than X^* . Note also that we could have taken \mathbb{Q} to be any symmetric distribution with a strictly positive density. More generally, we could let \mathbb{Q} be symmetric with a density q that is not necessarily strictly positive. In this case the RIPr would be a sub-probability measure with density $p^* = \frac{1}{2}(p_1 + p_2)\mathbf{1}_{\{q>0\}}$, and the numeraire would still be given by (5.1).

5.2 Bounded mean

This example is a variant of one studied in [Waudby-Smith and Ramdas \(2023\)](#), with the corresponding duality theory derived explicitly in [Honda and Takemura \(2010\)](#), and generalized to other classes of distributions in [Agrawal et al. \(2020\)](#). Let $\mathcal{Z} = [0, 1]$ so that the random variable Z takes values in the unit interval. Fix $\mu \in (0, \frac{1}{2})$ and consider the null hypothesis that the mean of Z is at most μ ,

$$\mathcal{P} = \{\mathbb{P} \in M_1 : \mathbb{E}_{\mathbb{P}}[Z] \leq \mu\}.$$

There is no single dominating measure for \mathcal{P} since it contains the uncountable non-dominated family $\{\delta_z : z \in [0, \mu]\}$. However, \mathcal{P} is generated by $\mathcal{E}_0 = \{Z/\mu\}$ in the sense of Corollary 4.2,

so our strategy will be to locate a candidate numeraire and then apply the corollary to verify that the candidate is, in fact, the numeraire. The alternative hypothesis \mathbf{Q} is the uniform distribution on $[0, 1]$, and we have $\mathbf{Q} \ll \mathcal{P}$ for the simple reason that $\mathbf{P}(A) = 0$ for all $\mathbf{P} \in \mathcal{P}$ implies that A must actually be empty.

To find a candidate numeraire, we observe that two natural e-variables are Z/μ and the constant one. All convex combinations of these are also e-variables; equivalently, $1 + \lambda(Z - \mu)$ is an e-variable for each $\lambda \in [0, \mu^{-1}]$. We now look for a log-optimal e-variable in this class by directly maximizing $f(\lambda) = \mathbb{E}_{\mathbf{Q}}[\log(1 + \lambda(Z - \mu))]$ over $\lambda \in [0, \mu^{-1}]$. This is a strictly concave function whose derivative $f'(\lambda) = \mathbb{E}_{\mathbf{Q}}[(Z - \mu)/(1 + \lambda(Z - \mu))]$ satisfies $f'(0) = \mathbb{E}_{\mathbf{Q}}[Z] - \mu > 0$ and $f'(\mu^{-1}) = \mathbb{E}_{\mathbf{Q}}[\mu - \mu^2/Z] = -\infty$. Thus there is a unique interior maximizer $\lambda^* \in (0, \mu^{-1})$, which is characterized by the first-order condition

$$\mathbb{E}_{\mathbf{Q}} \left[\frac{Z - \mu}{1 + \lambda^*(Z - \mu)} \right] = 0. \quad (5.2)$$

Since \mathbf{Q} is the standard uniform distribution we can be more explicit. Nothing changes if we first multiply both sides by λ^* , and then the left-hand side becomes

$$\int_0^1 \frac{\lambda^*(z - \mu)}{1 + \lambda^*(z - \mu)} dz = 1 - \int_0^1 \frac{1}{1 + \lambda^*(z - \mu)} dz = 1 - \frac{1}{\lambda^*} \log \left(\frac{1 + \lambda^*(1 - \mu)}{1 - \lambda^*\mu} \right).$$

Thus (5.2) for $\lambda^* \in (0, \mu^{-1})$ is equivalent to

$$\frac{1 + \lambda^*(1 - \mu)}{1 - \lambda^*\mu} = e^{\lambda^*},$$

which is easily solved numerically. This leads us to the candidate numeraire

$$X^* = 1 + \lambda^*(Z - \mu),$$

which is strictly positive and finite. To verify that this is indeed the numeraire, we use (5.2) to get, for any e-variable of the form $X = 1 + \lambda(Z - \mu) = X^* + (\lambda - \lambda^*)(Z - \mu)$, that

$$\mathbb{E}_{\mathbf{Q}} \left[\frac{X}{X^*} \right] = 1 + (\lambda - \lambda^*) \mathbb{E}_{\mathbf{Q}} \left[\frac{Z - \mu}{1 + \lambda^*(Z - \mu)} \right] = 1. \quad (5.3)$$

Taking $\lambda = 0$ and $\lambda = 1/\mu$ we see that (4.2) of Corollary 4.2 is satisfied and, hence, that X^* is the numeraire and the RIPr \mathbf{P}^* belongs to \mathcal{P} .

Remark 5.1. For the null hypothesis considered here, we conjecture that the set of all e-variables is given by

$$\mathcal{E} = \{X \in F_+ : X \leq 1 + \lambda(Z - \mu) \text{ for some } \lambda \in [0, \mu^{-1}]\}.$$

It is clear that every such X is an e-variable, but it is less obvious that every e-variable is of this form. Nonetheless, equipped with this knowledge the calculation in (5.3) would directly yield the numeraire property of X^* . The usefulness of Corollary 4.2 is that a complete description of \mathcal{E} is not required; it is enough to check (5.3) for $X = 1$ and for X in the generating set \mathcal{E}_0 .

5.3 Sub-Gaussian with nonpositive mean

Now take $\mathcal{Z} = \mathbb{R}$ and consider the null hypothesis that the observation Z has a 1-sub-Gaussian distribution with nonpositive mean. That is, we set

$$\mathcal{P} = \left\{ \mathbb{P} \in M_1 : \mathbb{E}_{\mathbb{P}}[e^{\lambda Z - \lambda^2/2}] \leq 1 \text{ for all } \lambda \in [0, \infty) \right\}.$$

The ‘one-sided’ restriction $\lambda \in [0, \infty)$ implies that Z has nonpositive (potentially nonzero and even infinite) mean under any $\mathbb{P} \in \mathcal{P}$. Indeed, monotone convergence and the definition of \mathcal{P} yield $\mathbb{E}_{\mathbb{P}}[Z] = \lim_{\lambda \downarrow 0} \mathbb{E}_{\mathbb{P}}[(e^{\lambda Z} - 1)/\lambda] \leq \lim_{\lambda \downarrow 0} (e^{\lambda^2/2} - 1)/\lambda = 0$. As in the previous example, \mathcal{P} does not admit any dominating measure. It is generated by the family $\mathcal{E}_0 = \{e^{\lambda Z - \lambda^2/2} : \lambda \in [0, \infty)\}$, so we will again look for a candidate numeraire and verify it using Corollary 4.2. We let the alternative hypothesis \mathbb{Q} be normal with mean $\mu > 0$ and unit variance. As in the previous example, and for the same reason, we have $\mathbb{Q} \ll \mathcal{P}$.

To find a candidate numeraire we maximize $\mathbb{E}_{\mathbb{Q}}[\log X]$ over $X \in \mathcal{E}_0$. That is, we maximize $\mathbb{E}_{\mathbb{Q}}[\lambda Z - \lambda^2/2] = \lambda\mu - \lambda^2/2$ over $\lambda \in [0, \infty)$. The maximizer is $\lambda^* = \mu$, which yields the candidate

$$X^* = e^{\mu Z - \mu^2/2}.$$

This is finite and strictly positive. Moreover, (4.2) of Corollary 4.2 is satisfied because

$$\mathbb{E}_{\mathbb{Q}} \left[\frac{e^{\lambda Z - \lambda^2/2}}{e^{\mu Z - \mu^2/2}} \right] = \mathbb{E}_{\mathbb{Q}} \left[e^{(\lambda - \mu)(Z - \mu) - (\lambda - \mu)^2/2} \right] = 1. \quad (5.4)$$

We conclude that X^* is the numeraire and, consequently, that the RIPr \mathbb{P}^* is the standard normal distribution. This example can be easily generalized to σ -sub-Gaussian distributions for $\sigma \neq 1$, but we omit this for brevity.

Remark 5.2. We conjecture that the set of all e-variables is given by

$$\mathcal{E} = \left\{ X \in F_+ : X \leq \int_{[0, \infty)} e^{\lambda Z - \lambda^2/2} \pi(d\lambda) \text{ for some probability measure } \pi \text{ on } [0, \infty) \right\}.$$

Although it is clear that every such X is an e-variable, showing that this really includes all e-variables seems to require sophisticated functional analytic methods that are beyond the scope of this paper. Fortunately, Corollary 4.2 sidesteps this issue. Nonetheless, taking the description of \mathcal{E} for granted the numeraire property can be checked directly. Indeed, for a general e-variable $X \leq \int_{[0, \infty)} e^{\lambda Z - \lambda^2/2} \pi(d\lambda)$, where π is a probability measure on $[0, \infty)$, Tonelli’s theorem and (5.4) give

$$\mathbb{E}_{\mathbb{Q}} \left[\frac{X}{X^*} \right] \leq \int_{[0, \infty)} \mathbb{E}_{\mathbb{Q}} \left[\frac{e^{\lambda Z - \lambda^2/2}}{e^{\mu Z - \mu^2/2}} \right] \pi(d\lambda) = 1.$$

Finally, let us mention that Ramdas et al. (2020) prove that $e^{\lambda Z - \lambda^2/2}$ is not just an e-variable, but an admissible one for every $\lambda \geq 0$.

5.4 Symmetric distributions

Continue to take $\mathcal{Z} = \mathbb{R}$. We now consider the null hypothesis that Z is symmetric,

$$\mathcal{P} = \{P \in M_1 : Z \text{ and } -Z \text{ have the same distribution under } P\},$$

which is again a non-dominated family. We also fix an alternative hypothesis Q that admits a Lebesgue density q . It is natural to conjecture that the RIPr is given by the symmetrization \tilde{P} of Q , whose density is $\tilde{p}(z) = (q(z) + q(-z))/2$. However, this cannot quite be true in general because \tilde{P} need not be equivalent to Q . Instead, we claim that the RIPr is the measure P^* with density

$$p^*(z) = \frac{1}{2} (q(z) + q(-z)) \mathbf{1}_{\{q(z) > 0\}}.$$

This is the absolutely continuous part of \tilde{P} with respect to Q . It is a probability measure if Q has symmetric support, and otherwise a proper sub-probability measure. Note that \tilde{P} belongs to \mathcal{P} , which in particular shows that $Q \ll \mathcal{P}$ since $Q \ll \tilde{P}$. (Alternatively, we again have that $P(A) = 0$ for all $P \in \mathcal{P}$ implies that A is empty, which also yields $Q \ll \mathcal{P}$.)

To check that P^* is the RIPr we will show that the implied candidate numeraire is, in fact, the numeraire. It is given by

$$X^* = \frac{dQ}{dP^*} = \frac{2q(Z)}{q(Z) + q(-Z)}.$$

This expression also appears in [Koning \(2023, Section 5–6\)](#). We claim that the set of all e-variables is

$$\mathcal{E} = \{X \in F_+ : X \leq 1 + \phi(Z) \text{ for some odd function } \phi\},$$

where we recall that a function ϕ is odd if $\phi(z) + \phi(-z) = 0$ for all $z \in \mathbb{R}$. Indeed, any X of this form satisfies $\mathbb{E}_P[X] \leq 1 + \mathbb{E}_P[\phi(Z)] = 1$, where the symmetry of Z and the oddness of ϕ were used to get $\mathbb{E}_P[\phi(Z)] = \mathbb{E}_P[\phi(-Z)] = -\mathbb{E}_P[\phi(Z)]$ and hence $\mathbb{E}_P[\phi(Z)] = 0$. Conversely, for any e-variable $X = f(Z)$ we may write $X = \frac{1}{2}(f(Z) + f(-Z)) + \phi(Z)$ where $\phi(z) = \frac{1}{2}(f(z) - f(-z))$ is the odd part of $f(z)$. For any $z \in \mathbb{R}$ we use the symmetric distribution $P = \frac{1}{2}(\delta_z + \delta_{-z})$ and the fact that X is an e-variable to get $\frac{1}{2}(f(z) + f(-z)) = \mathbb{E}_P[f(Z)] \leq 1$. Hence $X \leq 1 + \phi(Z)$ as required. We could also have reached this conclusion by using that every *admissible* e-variable is of the form $1 + \phi(Z)$ for some odd function ϕ as shown by [Ramdas et al. \(2020\)](#).

We can now verify the numeraire property. First, X^* is Q -almost surely strictly positive and finite, and it is an e-variable because it can be written as $X^* = 1 + \phi^*(Z)$ where

$$\phi^*(z) = \frac{q(z) - q(-z)}{q(z) + q(-z)}$$

is odd. Next, for any e-variable $X \leq 1 + \phi(Z)$ where ϕ is odd we get

$$\mathbb{E}_Q \left[\frac{X}{X^*} \right] \leq \mathbb{E}_Q \left[(1 + \phi(Z)) \frac{q(Z) + q(-Z)}{2q(Z)} \right] = \mathbb{E}_{\tilde{P}} \left[(1 + \phi(Z)) \mathbf{1}_{\{q(Z) > 0\}} \right] \leq \mathbb{E}_{\tilde{P}}[1 + \phi(Z)] = 1.$$

This confirms the numeraire property. Finally, let us remark that it is not really necessary that Q admit a density. The symmetrization of Q is still well-defined as $\tilde{P} = \frac{1}{2}(Q + \tilde{Q})$ where \tilde{Q} is the distribution of $-Z$ under Q , or equivalently, the pushforward of Q under the reflection map $z \mapsto -z$. The RIPr is then the absolutely continuous part $P^* = \tilde{P}^a$, and the numeraire is any nonnegative version of the Radon–Nikodym derivative dQ/dP^* such that $dQ/dP^* - 1$ is odd.

5.5 Exponential family with one-dimensional sufficient statistic

Here \mathcal{Z} can be general and is equipped with a reference measure μ . The following example shows how to use the results in Section 4.2 to obtain the result in (Grünwald et al., 2024, Example 4). Consider an exponential family of densities

$$p_\theta(z) = e^{\theta T(z) - A(\theta)}$$

with respect to μ , where A is convex and differentiable, the sufficient statistic $T(z)$ is one-dimensional, and the natural parameter θ ranges in some interval $\Theta \subset \mathbb{R}$. The null hypothesis is $\mathcal{P} = \{p_\theta : \theta \in \Theta_0\}$ for some closed subset $\Theta_0 \subset \Theta$, and the alternative is $q = p_{\theta_1}$ for some $\theta_1 \in \Theta$. We suppose that Θ_0 has a smallest element θ^* and that $\theta_1 < \theta^*$. It is natural to conjecture that p_{θ^*} is the RIPr. To confirm this, it suffices to verify (4.3). Using the standard formula for the moment generating function of the sufficient statistic we get

$$\mathbb{E}_{\mathbb{Q}} \left[\frac{p_\theta}{p_{\theta^*}} \right] = \mathbb{E}_{\mathbb{Q}} \left[e^{(\theta - \theta^*)T(Z) - A(\theta) + A(\theta^*)} \right] = e^{A(\theta_1 + \theta - \theta^*) - A(\theta_1) - A(\theta) + A(\theta^*)}$$

for all $\theta \in \Theta_0$. Since A is convex, its derivative A' is increasing. Together with the fundamental theorem of calculus, as well as $\theta_1 < \theta^*$ and $\theta - \theta^* \geq 0$ for $\theta \in \Theta_0$, this gives

$$\begin{aligned} A(\theta_1 + \theta - \theta^*) - A(\theta_1) &= \int_0^1 (\theta - \theta^*) A'(\theta_1 + t(\theta - \theta^*)) dt \\ &\leq \int_0^1 (\theta - \theta^*) A'(\theta^* + t(\theta - \theta^*)) dt \\ &= A(\theta) - A(\theta^*). \end{aligned}$$

We conclude that $\mathbb{E}_{\mathbb{Q}}[p_\theta/p_{\theta^*}] \leq 1$ for all $\theta \in \Theta_0$ so that (4.3) holds and p_{θ^*} is indeed the RIPr. As a result the numeraire is the likelihood ratio

$$X^* = \frac{p_{\theta_1}(Z)}{p_{\theta^*}(Z)} = e^{(\theta_1 - \theta^*)T(Z) - A(\theta_1) + A(\theta^*)}.$$

We refer to Csizsar and Matus (2003) for an extensive study of information projections involving exponential families.

Remark 5.3. When \mathcal{P} is a parametric family of distributions it is often easier to first find the RIPr and then use it to find the numeraire. Here we used Theorem 4.7 to confirm that the candidate p_{θ^*} was indeed the RIPr. We could also have minimized relative entropy and used Corollary 4.11. Since the corollary assumes that \mathcal{P} is convex, which fails here, we would have had to apply it with the convex hull of \mathcal{P} and perform the minimization over this set.

6 Beyond the logarithm: reverse Rényi projection

The duality between log-utility and relative entropy can be extended to other utilities and divergences. One case where this can be done rigorously using the results presented so far is the power utility

$$U(x) = \frac{x^{1-\gamma}}{1-\gamma}$$

for $\gamma > 1$. In this section we outline the resulting theory using a well-known method from mathematical finance pioneered by [Karatzas et al. \(1991\)](#) and [Kramkov and Schachermayer \(1999\)](#). However, in contrast to that literature we work as in [Sections 2 and 3](#) with completely arbitrary \mathcal{P} and \mathbb{Q} .

The utility function U is continuous, increasing, concave, differentiable, and bounded from above by zero. [Lemma 2.9](#) is therefore applicable and yields an optimal e-variable for the maximization problem

$$\sup_{X \in \mathcal{E}} \mathbb{E}_{\mathbb{Q}}[U(X)].$$

We denote the maximizer by X_{γ}^* and call it the *U-optimal e-variable*. Observe that

$$\mathbb{E}_{\mathbb{Q}}[U'(X_{\gamma}^*)X_{\gamma}^*] = \mathbb{E}_{\mathbb{Q}}[(X_{\gamma}^*)^{1-\gamma}] = (1-\gamma)\mathbb{E}_{\mathbb{Q}}[U(X_{\gamma}^*)] \leq (1-\gamma)U(1) = 1,$$

where we use the convention $U'(\infty)\infty = 0$ and where the inequality relies on the fact that the constant e-variable $X = 1$ is suboptimal and that $\gamma > 1$. Thus [Lemma 2.9](#) further implies the first-order condition

$$\mathbb{E}_{\mathbb{Q}}[(X_{\gamma}^*)^{-\gamma}X] \leq \mathbb{E}_{\mathbb{Q}}[(X_{\gamma}^*)^{1-\gamma}], \quad X \in \mathcal{E}, \quad (6.1)$$

where $(X_{\gamma}^*)^{-\gamma}X$ is understood as zero on $\{X_{\gamma}^* = \infty\}$. This will allow us to identify an analog of the RIPr which we denote by \mathbb{P}_{γ}^* . In the fully degenerate case where $X_{\gamma}^* = \infty$, \mathbb{Q} -almost surely, we simply set $\mathbb{P}_{\gamma}^* = 0$. In the more interesting case where X_{γ}^* is finite with positive \mathbb{Q} -probability, we let \mathbb{P}_{γ}^* be the measure whose density with respect to \mathbb{Q} is

$$\frac{d\mathbb{P}_{\gamma}^*}{d\mathbb{Q}} = \frac{(X_{\gamma}^*)^{-\gamma}}{\mathbb{E}_{\mathbb{Q}}[(X_{\gamma}^*)^{1-\gamma}]}. \quad (6.2)$$

Thanks to [\(6.1\)](#), \mathbb{P}_{γ}^* belongs to \mathcal{P}_{eff} . In analogy with [\(3.3\)](#) we expect \mathbb{P}_{γ}^* to minimize a suitable dual objective function. To discover its form, we look for a weak duality inequality and introduce the Legendre transform $V(y) = \sup_{x>0}\{U(x) - xy\}$. Evaluating the supremum shows that

$$V(y) = -\frac{y^{1-1/\gamma}}{1-1/\gamma}.$$

Now, for any e-variable X , any $\mathbb{P} \in \mathcal{P}_{\text{eff}}$, and any $y \geq 0$, we have

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}[U(X)] &\leq \mathbb{E}_{\mathbb{Q}}[U(X)] + y(1 - \mathbb{E}_{\mathbb{P}}[X]) \\ &\leq \mathbb{E}_{\mathbb{Q}}\left[U(X) - y\frac{d\mathbb{P}^a}{d\mathbb{Q}}X\right] + y \\ &\leq \mathbb{E}_{\mathbb{Q}}\left[V\left(y\frac{d\mathbb{P}^a}{d\mathbb{Q}}\right)\right] + y \\ &= -\frac{y^{1-1/\gamma}}{1-1/\gamma}\mathbb{E}_{\mathbb{Q}}\left[\left(\frac{d\mathbb{P}^a}{d\mathbb{Q}}\right)^{1-1/\gamma}\right] + y. \end{aligned}$$

To obtain a tight bound we minimize the right-hand side over y . A computation shows that the minimizer is $y = \mathbb{E}_{\mathbb{Q}}[(d\mathbb{P}^a/d\mathbb{Q})^{1-1/\gamma}]^{\gamma}$, which is finite, possibly zero, thanks to Jensen's

inequality. With this choice we get, for any e-variable X and any $P \in \mathcal{P}_{\text{eff}}$, the weak duality inequality

$$\mathbb{E}_{\mathbf{Q}}[U(X)] \leq \frac{1}{1-\gamma} \mathbb{E}_{\mathbf{Q}} \left[\left(\frac{dP^a}{d\mathbf{Q}} \right)^{1-1/\gamma} \right]^\gamma.$$

Equality is achieved by taking $X = X_\gamma^*$ and $P = P_\gamma^*$, so we conclude that P_γ^* solves the minimization problem

$$\inf_{P \in \mathcal{P}_{\text{eff}}} \frac{1}{1-\gamma} \mathbb{E}_{\mathbf{Q}} \left[\left(\frac{dP^a}{d\mathbf{Q}} \right)^{1-1/\gamma} \right]^\gamma$$

with optimal value equal to $\mathbb{E}_{\mathbf{Q}}[U(X_\gamma^*)]$. Observe that the quantity being minimized is an increasing function of the Rényi divergence of order $1/\gamma$ of \mathbf{Q} from P ,

$$R_{1/\gamma}(\mathbf{Q} \mid P) = \frac{1}{1/\gamma - 1} \log \mathbb{E}_{\mathbf{Q}} \left[\left(\frac{dP^a}{d\mathbf{Q}} \right)^{1-1/\gamma} \right].$$

Thus P_γ^* also minimizes the Rényi divergence, and as such it is known as the *reverse Rényi projection* of order $1/\gamma$ of \mathbf{Q} on \mathcal{P}_{eff} . In particular, we see that the reverse Rényi projection always exists if the order $1/\gamma$ is in $(0, 1)$. We summarize the above observations, along with a few others, in the following theorem.

Theorem 6.1. *Let $\gamma > 1$ and $U(x) = x^{1-\gamma}/(1-\gamma)$. Then*

$$\sup_{X \in \mathcal{E}} \mathbb{E}_{\mathbf{Q}}[U(X)] = \frac{1}{1-\gamma} \exp \left((1-\gamma) \inf_{P \in \mathcal{P}_{\text{eff}}} R_{1/\gamma}(\mathbf{Q} \mid P) \right), \quad (6.3)$$

and the supremum and infimum are attained by $X_\gamma^* \in \mathcal{E}$ and $P_\gamma^* \in \mathcal{P}_{\text{eff}}$ related by (6.2). Furthermore, X_γ^* is the \mathbf{Q} -almost surely unique maximizer, and P_γ^* is the unique minimizer among the elements of \mathcal{P}_{eff} that are absolutely continuous with respect to \mathbf{Q} . Finally, X_γ^* can be expressed in terms of P_γ^* by

$$X_\gamma^* = \left(\frac{dP_\gamma^*}{d\mathbf{Q}} \right)^{-1/\gamma} / \mathbb{E}_{\mathbf{Q}} \left[\left(\frac{dP_\gamma^*}{d\mathbf{Q}} \right)^{1-1/\gamma} \right], \quad (6.4)$$

which is understood as $+\infty$ on the set where $dP_\gamma^*/d\mathbf{Q} = 0$.

Proof. The equality in (6.3) and the form of the solutions were shown above. The \mathbf{Q} -almost sure uniqueness of X_γ^* follows from strict concavity of the function $f(X) = \mathbb{E}_{\mathbf{Q}}[U(X)]$ on \mathcal{E} . Since $\gamma > 1$, the Rényi divergence is a decreasing transformation of the function $g(P) = \mathbb{E}_{\mathbf{Q}}[(dP^a/d\mathbf{Q})^{1-1/\gamma}]$, which is thus maximized by P_γ^* . The function $g(P)$ is strictly concave on the set of elements of \mathcal{P}_{eff} that are absolutely continuous with respect to \mathbf{Q} . Thus P_γ^* is the unique maximizer of $g(P)$, and the unique minimizer of the Rényi divergence. Finally, (6.4) is obtained by noting that

$$\mathbb{E}_{\mathbf{Q}} [(X_\gamma^*)^{1-\gamma}] = \mathbb{E}_{\mathbf{Q}} \left[\left(\frac{dP_\gamma^*}{d\mathbf{Q}} \right)^{1-1/\gamma} \right]^\gamma$$

and then inverting (6.2). □

Remark 6.2. Lardy et al. (2024, Section V) mention the e-variable $(dP_\gamma^*/dQ)^{-1/\gamma}$ constructed using the reverse Rényi projection. This e-variable is strictly smaller than X_γ^* , because the denominator in (6.4) is strictly smaller than one by Jensen’s inequality.

Remark 6.3. The method described above can be applied to more general utility functions, resulting in projections associated with more general divergences. An interesting question is whether all f -divergences can be covered, at least under suitable technical assumptions.

7 Universal inference is inadmissible

We assume that there exists a common reference measure μ , and follow the notational conventions introduced before Theorem 4.7 to identify distributions with their respective densities, written in lowercase. For a point alternative q over the data Z , the method of universal inference (Wasserman et al., 2020) boils down to constructing the e-variable $X^{\text{UI}} = q(Z)/p_{\max}(Z)$ where $p_{\max}(Z) = \sup_{p \in \mathcal{P}} p(Z)$ is the maximum likelihood. (Strictly speaking, the supremum here should be understood as an essential supremum under the reference measure.) To check that this is indeed an e-variable, simply note that for any $\mathbf{P} \in \mathcal{P}$ with density p we have

$$\mathbb{E}_{\mathbf{P}} \left[\frac{q(Z)}{p_{\max}(Z)} \right] \leq \mathbb{E}_{\mathbf{P}} \left[\frac{q(Z)}{p(Z)} \right] = 1.$$

To compare the universal inference e-variable with the numeraire $X^* = q(Z)/p^*(Z)$, we first claim that the RIPr satisfies $p^*(Z) \leq p_{\max}(Z)$ up to μ -nullsets. This is obvious if the RIPr belongs to \mathcal{P} , but in general it only belongs to \mathcal{P}_{eff} . In this case the claim follows by applying the following lemma with $f = p_{\max}$.

Lemma 7.1. *Let f be a random variable such that $f \geq p$, μ -almost surely, for all $p \in \mathcal{P}$. Then $f \geq p$, μ -almost surely, for all $p \in \mathcal{P}_{\text{eff}}$.*

Proof. It is clear that $f \geq p$, μ -almost surely, for all sub-probability densities p dominated by some element of the convex hull of \mathcal{P} , and then also for all p in the μ -probability closure of this set. This was called \mathcal{P}'' in the proof of Theorem 4.7, where it was shown that \mathcal{P}'' is closed, convex, and solid. By Lemma 4.5, \mathcal{P}'' contains \mathcal{P}_{eff} , completing the proof. \square

We thus see that $p^*(Z) \leq p_{\max}(Z)$, and hence $X^* \geq X^{\text{UI}}$, up to μ -nullsets. In nondegenerate situations involving a composite null hypothesis, the inequality will be strict with positive μ -probability. Thus, in such cases, the relatively general method of universal inference is in fact inadmissible. We end by noting that our numeraire imposes weaker assumptions than the universal inference work, which needs a reference measure to define likelihoods.

8 Summary

We established that under no assumptions on the composite null \mathcal{P} and point alternative Q , there exists a unique special e-variable X^* , the numeraire, which satisfies $\mathbb{E}_Q[X/X^*] \leq 1$ for every other e-variable X , and is also log-optimal. Further, X^* can be written as a likelihood ratio dQ/dP^* for a unique (sub-)probability distribution P^* , absolutely continuous with respect to Q and belonging to the effective null hypothesis \mathcal{P}_{eff} associated with \mathcal{P} .

The numeraire is the only e-variable of this form. A strong duality theory establishes that $\mathbb{E}_Q[\log X^*]$ equals the minimum relative entropy of Q to \mathcal{P}_{eff} , achieved by P^* . This fully generalizes (to composite nulls) the fact that the likelihood ratio is the log-optimal e-variable for point nulls. We gave several sufficient conditions to identify and certify the numeraire, and showed that these were easy to verify in several nonparametric examples without a reference measure that were out of the reach of earlier methods. We also showed how to generalize our theory beyond log-optimality.

One natural direction for future work is to extend our theory to composite alternatives, using the *method of mixtures* or the *plug-in method*. A second direction is to extend our theory to the sequential setting using nonnegative supermartingales and *e-processes* (sequential generalizations of e-variables). See Ramdas et al. (2023) for an elaboration of the italicized terms and existing work. The second topic introduces additional subtleties through dealing with filtrations and stopping times, a direction we are currently pursuing.

Acknowledgments

We thank the referees for insightful comments and helpful suggestions.

References

- Shubhada Agrawal, Sandeep Juneja, and Peter Glynn. Optimal δ -correct best-arm selection for heavy-tailed distributions. In *Algorithmic Learning Theory*, pages 61–110. PMLR, 2020. [21](#)
- Daniel Bartl and Michael Kupper. A pointwise bipolar theorem. *Proc. Amer. Math. Soc.*, 147(4):1483–1495, 2019. ISSN 0002-9939,1088-6826. doi: 10.1090/proc/14231. URL <https://doi.org/10.1090/proc/14231>. [9](#)
- Dirk Becherer. The numeraire portfolio for unbounded semimartingales. *Finance Stoch.*, 5(3):327–341, 2001. ISSN 0949-2984,1432-1122. doi: 10.1007/PL00013535. URL <https://doi.org/10.1007/PL00013535>. [2](#), [4](#)
- W. Brannath and W. Schachermayer. A bipolar theorem for $L_+^0(\Omega, \mathcal{F}, \mathbb{P})$. In *Séminaire de Probabilités, XXXIII*, volume 1709 of *Lecture Notes in Math.*, pages 349–354. Springer, Berlin, 1999. ISBN 3-540-66342-8. doi: 10.1007/BFb0096525. URL <https://doi.org/10.1007/BFb0096525>. [6](#), [9](#), [17](#)
- Philippe Casgrain, Martin Larsson, and Johanna Ziegel. Sequential testing for elicitable functionals via supermartingales. *Bernoulli*, 30(2):1347–1374, 2024. ISSN 1350-7265,1573-9759. doi: 10.3150/23-bej1634. URL <https://doi.org/10.3150/23-bej1634>. [16](#)
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, second edition, 2006. ISBN 978-0-471-24195-9; 0-471-24195-4. [5](#)
- I. Csiszar and F. Matus. Information projections revisited. *IEEE Transactions on Information Theory*, 49(6):1474–1490, 2003. doi: 10.1109/TIT.2003.810633. [9](#), [25](#)

- I. Csiszár and G. Tusnády. Information geometry and alternating minimization procedures. *Statist. Decisions*, pages 205–237, 1984. ISSN 0721-2631. Recent Results in Estimation Theory and Related Topics. [2](#), [9](#)
- Jakša Cvitanić and Ioannis Karatzas. Generalized Neyman-Pearson lemma via convex duality. *Bernoulli*, 7(1):79–97, 2001. ISSN 1350-7265,1573-9759. doi: 10.2307/3318603. URL <https://doi.org/10.2307/3318603>. [2](#)
- Freddy Delbaen and Walter Schachermayer. A general version of the fundamental theorem of asset pricing. *Math. Ann.*, 300(3):463–520, 1994. ISSN 0025-5831,1432-1807. doi: 10.1007/BF01450498. URL <https://doi.org/10.1007/BF01450498>. [6](#)
- Ivar Ekeland and Roger Témam. *Convex Analysis and Variational Problems*, volume 28 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, english edition, 1999. ISBN 0-89871-450-8. doi: 10.1137/1.9781611971088. URL <https://doi.org/10.1137/1.9781611971088>. Translated from the French. [12](#)
- Peter Grünwald, Rianne de Heide, and Wouter M Koolen. Safe testing. *Journal of the Royal Statistical Society, Series B (Methodology), with discussion*, 2024. [2](#), [4](#), [9](#), [15](#), [21](#), [25](#)
- Paul R Halmos and Leonard J Savage. Application of the Radon-Nikodym theorem to the theory of sufficient statistics. *The Annals of Mathematical Statistics*, 20(2):225–241, 1949. [2](#)
- Junya Honda and Akimichi Takemura. An asymptotically optimal bandit algorithm for bounded support models. In *Conference on Learning Theory*, pages 67–79. Citeseer, 2010. [21](#)
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 2020. [2](#)
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 2021. [2](#)
- 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). [2](#)
- Ioannis Karatzas and Constantinos Kardaras. The numéraire portfolio in semimartingale financial models. *Finance Stoch.*, 11(4):447–493, 2007. ISSN 0949-2984,1432-1122. doi: 10.1007/s00780-007-0047-3. URL <https://doi.org/10.1007/s00780-007-0047-3>. [2](#), [4](#)
- Ioannis Karatzas and Constantinos Kardaras. *Portfolio Theory and Arbitrage—A Course in Mathematical Finance*, volume 214 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2021. ISBN 978-1-4704-6014-3. doi: 10.1090/gsm/214. URL <https://doi.org/10.1090/gsm/214>. [2](#), [4](#)
- Ioannis Karatzas, John P. Lehoczky, Steven E. Shreve, and Gan-Lin Xu. Martingale and duality methods for utility maximization in an incomplete market. *SIAM J. Control Optim.*,

- 29(3):702–730, 1991. ISSN 0363-0129. doi: 10.1137/0329039. URL <https://doi.org/10.1137/0329039>. 26
- Nick W Koning. Post-hoc and anytime valid permutation and group invariance testing. *arXiv preprint arXiv:2310.01153*, 2023. 24
- D. Kramkov and W. Schachermayer. The asymptotic elasticity of utility functions and optimal investment in incomplete markets. *Ann. Appl. Probab.*, 9(3):904–950, 1999. ISSN 1050-5164,2168-8737. doi: 10.1214/aoap/1029962818. URL <https://doi.org/10.1214/aoap/1029962818>. 26
- Tyron Lardy, Peter Grünwald, and Peter Harremoës. Reverse information projections and optimal e-statistics. *IEEE Transactions on Information Theory (accepted)*, 2024. 2, 4, 9, 11, 12, 19, 20, 21, 28
- Qiang Jonathan Li. *Estimation of Mixture Models*. PhD thesis, Yale University, 1999. 2, 9, 11, 20
- John B. Long. The numeraire portfolio. *Journal of Financial Economics*, 26(1):29–69, 1990. ISSN 0304-405X. doi: [https://doi.org/10.1016/0304-405X\(90\)90012-O](https://doi.org/10.1016/0304-405X(90)90012-O). URL <https://www.sciencedirect.com/science/article/pii/0304405X90900120>. 2, 4
- Eckhard Platen. A benchmark approach to finance. *Math. Finance*, 16(1):131–151, 2006. ISSN 0960-1627,1467-9965. doi: 10.1111/j.1467-9965.2006.00265.x. URL <https://doi.org/10.1111/j.1467-9965.2006.00265.x>. 2, 4
- Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter Koolen. Admissible anytime-valid sequential inference must rely on nonnegative martingales. *arXiv:2009.03167*, 2020. 23, 24
- Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023. 2, 29
- Herbert Robbins. Statistical methods related to the law of the iterated logarithm. *The Annals of Mathematical Statistics*, 41(5):1397–1409, 1970. 2
- Glenn Shafer. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society, with discussion*, 184(2):407–431, 2021. 2, 3, 4, 14
- Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3):1736–1754, 2021. 2
- Larry Wasserman, Aaditya Ramdas, and Sivaraman Balakrishnan. Universal inference. *Proceedings of the National Academy of Sciences*, 117(29):16880–16890, 2020. 2, 21, 28
- Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society, Series B (Methodology), with discussion*, 2023. 2, 21
- Zhenyuan Zhang, Aaditya Ramdas, and Ruodu Wang. On the existence of powerful p-values and e-values for composite hypotheses. *Annals of Statistics (to appear)*, 2024. 12

A A generalized Lebesgue decomposition

The following generalized Lebesgue decomposition into a ‘regular’ and ‘singular’ part is used in the proof of Lemma 2.9. As always, we work on a given measurable space (Ω, \mathcal{F}) . Given a nonempty subset $\mathcal{P} \subset M_+$, an event $N \in \mathcal{F}$ is called \mathcal{P} -negligible if $\mathbf{P}(N) = 0$ for all $\mathbf{P} \in \mathcal{P}$.

Lemma A.1. *For any $\mathbf{Q} \in M_+$ and nonempty $\mathcal{P} \subset M_+$ there is a unique decomposition $\mathbf{Q} = \mathbf{Q}^r + \mathbf{Q}^s$ with $\mathbf{Q}^r, \mathbf{Q}^s \in M_+$ such that $\mathbf{Q}^r \ll \mathcal{P}$ and $\mathbf{Q}^r, \mathbf{Q}^s$ are singular, meaning that there exists a \mathcal{P} -negligible event $N \in \mathcal{F}$ such that $\mathbf{Q}^s(N^c) = 0$. The event N is essentially unique, i.e., unique up to \mathbf{Q} - and \mathbf{P} -nullsets for all $\mathbf{P} \in \mathcal{P}$. In particular, any other \mathcal{P} -negligible event N' satisfies $\mathbf{Q}(N' \setminus N) = 0$.*

Proof. Define $\alpha = \sup_{N \in \mathcal{N}} \mathbf{Q}(N)$ where $\mathcal{N} = \{N \in \mathcal{F} : \mathbf{P}(N) = 0 \text{ for all } \mathbf{P} \in \mathcal{P}\}$ is the family of all \mathcal{P} -negligible sets. Note that \mathcal{N} is closed under countable unions. Pick a maximizing sequence N_n and set $N = \bigcup_n N_n$. Then $\alpha \geq \mathbf{Q}(N) \geq \mathbf{Q}(N_n) \rightarrow \alpha$, so $\mathbf{Q}(N) = \alpha$. Now define $\mathbf{Q}^s(A) = \mathbf{Q}(A \cap N)$ and $\mathbf{Q}^r(A) = \mathbf{Q}(A \cap N^c)$. To see that $\mathbf{Q}^r \ll \mathcal{P}$, suppose this is not the case. Then there is some $N' \in \mathcal{N}$ with $\mathbf{Q}^r(N') > 0$, and hence $\mathbf{Q}(N \cup N') \geq \mathbf{Q}^s(N) + \mathbf{Q}^r(N') > \alpha$, a contradiction. To show uniqueness, suppose we have two tuples $\mathbf{Q}_i^r, \mathbf{Q}_i^s, N_i$ for $i = 1, 2$ with the stated properties. Set $N = N_1 \cup N_2 \in \mathcal{N}$ and note that $\mathbf{Q}_i^s(N^c) \leq \mathbf{Q}_i^s(N_i^c) = 0$. Then for any A , we have $\mathbf{Q}(A \cap N) = \mathbf{Q}_i^s(A \cap N) = \mathbf{Q}_i^s(A)$. Thus $\mathbf{Q}_1^s = \mathbf{Q}_2^s$, and then also $\mathbf{Q}_1^r = \mathbf{Q}_2^r$. Furthermore, the symmetric difference $N_1 \Delta N_2$ is a \mathbf{Q} -nullset because $\mathbf{Q}(N_1 \cap N_2^c) = \mathbf{Q}_2^r(N_1) = 0$ and similarly $\mathbf{Q}(N_2 \cap N_1^c) = 0$. Since also $N_1 \Delta N_2$ belongs to \mathcal{N} , the essential uniqueness statement follows. Finally, if $N' \in \mathcal{N}$ then $\mathbf{Q}(N' \setminus N) = \mathbf{Q}^r(N' \setminus N) + \mathbf{Q}^s(N' \setminus N) \leq \mathbf{Q}^r(N') + \mathbf{Q}^s(N^c) = 0$. \square

B Overview of finiteness properties

The table below gives an overview of the relation between absolute continuity properties and finiteness of the numeraire and relative entropies. Each cell shows the possibilities for the property or quantity indicated in the column header, under the assumption indicated for that row. In the table, X^* is the numeraire, P^* is the RPr, the relative entropy $H(Q | P^*)$ is in (3.1), and the maximal description gain $H(Q | P^* \rightsquigarrow \mathcal{P}_{\text{eff}})$ is defined in Remark 3.9. We specifically highlight the possibility that $H(Q | P)$ may be infinite for all $P \in \mathcal{P}$, but finite or even zero for the RPr $P^* \in \mathcal{P}_{\text{eff}}$; this happens, for instance, in Examples 4.4 and 4.9. Furthermore, although $H(Q | P^* \rightsquigarrow \mathcal{P}_{\text{eff}}) = \infty$ if $Q \not\ll \mathcal{P}$, the modified quantity in Theorem 3.13(iii) will be equal to zero for $P = P^*$.

	$Q(X^* < \infty) = 1$	$P^* \ll Q$	$Q \ll P^*$	$H(Q P^*)$	$H(Q P^* \rightsquigarrow \mathcal{P}_{\text{eff}})$
$Q \not\ll \mathcal{P}$	no	yes	no	$= \infty$	$= \infty$
$Q \ll \mathcal{P}$	yes	yes	yes	$\leq \infty$	$= 0$
$Q \ll \mathcal{P}$ $H(Q P) < \infty$ for some $P \in \mathcal{P}$	yes	yes	yes	$< \infty$	$= 0$
$Q \ll \mathcal{P}$ $H(Q P) = \infty$ for all $P \in \mathcal{P}$	yes	yes	yes	$\leq \infty$	$= 0$
$Q \ll \mathcal{P} \ll \mu$ for some μ $H(Q P) = \infty$ for all $P \in \mathcal{P}$ \mathcal{P} is convex	yes	yes	yes	$\leq \infty$	$= 0$