

Minimization Problems Based on Relative α -Entropy I: Forward Projection

M. Ashok Kumar and Rajesh Sundaresan

Abstract

Minimization problems with respect to a one-parameter family of generalized relative entropies are studied. These relative entropies, which we term relative α -entropies (denoted \mathcal{J}_α), arise as redundancies under mismatched compression when cumulants of compressed lengths are considered instead of expected compressed lengths. These parametric relative entropies are a generalization of the usual relative entropy (Kullback-Leibler divergence). Just like relative entropy, these relative α -entropies behave like squared Euclidean distance and satisfy the Pythagorean property. Minimizers of these relative α -entropies on closed and convex sets are shown to exist. Such minimizations generalize the maximum Rényi or Tsallis entropy principle. The minimizing probability distribution (termed forward \mathcal{J}_α -projection) for a linear family is shown to obey a power-law. Other results in connection with statistical inference, namely subspace transitivity and iterated projections, are also established. In a companion paper, a related minimization problem of interest in robust statistics that leads to a reverse \mathcal{J}_α -projection is studied.

Index Terms

Best approximant; exponential family; information geometry; Kullback-Leibler divergence; linear family; power-law family; projection; Pythagorean property; relative entropy; Rényi entropy; Tsallis entropy.

I. INTRODUCTION

Relative entropy¹ or Kullback-Leibler divergence $\mathcal{J}(P||Q)$ between two probability measures is a fundamental quantity that arises in a variety of situations in probability theory, statistics, and information theory. In probability theory, it arises as the rate function for estimating the probability of a large deviation for the empirical measure of independent samplings. In statistics, for example, it arises as the best error exponent in deciding between two hypothetical distributions for observed data. In Shannon theory, it is the penalty in expected compressed length, namely the gap from Shannon entropy $H(P)$, when the compressor assumes (for a finite-alphabet source) a mismatched probability measure Q instead of the true probability measure P .

Relative entropy also brings statistics and probability theory together to provide a foundation for the well-known maximum entropy principle for decision making under uncertainty. This is an idea that goes back to L. Boltzmann, was popularized by E. T. Jaynes [3], and has its foundation in the theory of large deviation. Suppose that an ensemble average measurement (say sample mean, sample second moment, or any other similar linear statistic) is made on the realization of a sequence of independent and identically distributed (i.i.d.) random variables. The realization must then have an empirical measure that obeys the constraint placed by the measurement – the empirical measure must belong to an appropriate convex set, say \mathbb{E} . Large deviation theory tells us that a special member of \mathbb{E} , denoted P^* , is overwhelmingly more likely than the others. If the alphabet \mathbb{X} is finite (with cardinality $|\mathbb{X}|$), and the prior probability (before measurement) is the uniform measure U on \mathbb{X} , then P^* is the one that minimizes the relative entropy

$$\mathcal{J}(P||U) = \log |\mathbb{X}| - H(P),$$

which is the same as the one that maximizes (Shannon) entropy, subject to $P \in \mathbb{E}$. This explains why the principle is called maximum entropy principle. In Jaynes' words, "... it is maximally noncommittal to the missing information" [3].

M. Ashok Kumar was supported by a Council for Scientific and Industrial Research (CSIR) fellowship and by the Department of Science and Technology. R. Sundaresan was supported in part by the University Grants Commission by Grant Part (2B) UGC-CAS-(Ph.IV) and in part by the Department of Science and Technology. A part of the material in this paper was presented at the IEEE International Symposium on Information Theory (ISIT 2011), St. Petersburg, Russia, August 2011 [1]. A part of the material in the Introduction has overlap with a conference article [2] presented at the National Conference on Communication (NCC 2015), Mumbai, India, held during February 2015.

M. Ashok Kumar and R. Sundaresan are with the ECE Department, Indian Institute of Science, Bangalore 560012, India.

¹The relative entropy of P with respect to Q is defined as

$$\mathcal{J}(P||Q) := \sum_{x \in \mathbb{X}} P(x) \log \frac{P(x)}{Q(x)}$$

and the Shannon entropy of P is defined as

$$H(P) := - \sum_{x \in \mathbb{X}} P(x) \log P(x).$$

The usual convention is $p \log \frac{p}{q} = 0$ if $p = 0$ and $+\infty$ if $p > q = 0$.

As a physical example, let us tag a particular molecule in the atmosphere. Let X denote the height of the molecule in the atmosphere. Then the potential energy of the molecule is mgX . Let us suppose that the average potential energy is held constant, that is, $E[mgX] = c$, a constant. Then the probability distribution of the height of the molecule is taken to be the exponential distribution $\lambda \exp(-\lambda x)$, where $\lambda = mg/c$. This is also the maximum entropy probability distribution subject to first moment constraint [4].

More generally, if the prior probability (before measurement) is Q , then P^* minimizes $\mathcal{I}(P||Q)$ subject to $P \in \mathbb{E}$. Something more specific can be said: P^* is the limiting conditional distribution of a “tagged” particle under the conditioning imposed by the measurement. This is called *the conditional limit theorem* or the *Gibbs conditioning principle*; see for example Campenhouet and Cover [5] or Csiszár [6] for a more general result.

It is well-known that $\mathcal{I}(P||Q)$ behaves like “squared Euclidean distance” and has the “Pythagorean property” (Csiszár [7]). In view of this and since P^* minimizes $\mathcal{I}(P||Q)$ subject to $P \in \mathbb{E}$, one says that P^* is “closest” to Q in the relative entropy sense amongst the measures in \mathbb{E} , or in other words, “ P^* is the forward \mathcal{I} -projection of Q on \mathbb{E} ”. Motivated by the above maximum entropy and Gibbs conditioning principles, \mathcal{I} -projection was extensively studied by Csiszár [6], [7], Csiszár and Matúš [8], Csiszár and Shields [9], and Csiszár and Tusnády [10]. More recently, minimizations of general entropy functionals with convex integrands were studied by Csiszár and Matúš [11]. These include Bregman’s divergences and Csiszár’s f -divergences. \mathcal{I} -minimization also arises in the contraction principle in large deviation theory (see for example Dembo and Zeitouni’s [12, p.126]).

This paper is on projections or minimization problems associated with a parametric generalization of relative entropy. To see how this parametric generalization arises, we return to our remark on how relative entropy arises in Shannon theory. For this, we must first recall how Rényi entropies are a parametric generalization of the Shannon entropy.

Rényi entropies $H_\alpha(P)$ for $\alpha \in (0, 1)$ play the role of Shannon entropy when the *normalized cumulant* of compression length is considered instead of expected compression length. Campbell [13] showed that

$$\min \frac{1}{n\rho} \log E[\exp\{\rho L_n(X^n)\}] \rightarrow H_\alpha(P) \text{ (as } n \rightarrow \infty)$$

for an i.i.d. source with marginal P . The minimum is over all compression strategies L_n that satisfy the Kraft inequality², $\alpha = 1/(1 + \rho)$, and $\rho > 0$ is the cumulant parameter. We also have $\lim_{\alpha \rightarrow 1} H_\alpha(P) = H(P)$, so that Rényi entropy may be viewed as a generalization of Shannon entropy.

If the compressor assumed that the true probability measure is Q , instead of P , then the gap in the normalized cumulant’s optimal value is an analogous parametric divergence quantity³, which we shall denote $\mathcal{S}_\alpha(P, Q)$ [15]. The same quantity⁴ also arises when we study the gap from optimality of mismatched guessing exponents. See Arikan [16] and Hanawal and Sundaresan [17] for general results on guessing, and see Sundaresan [18], [15] on how $\mathcal{S}_\alpha(P, Q)$ arises in the context of mismatched guessing. Recently, Bunte and Lapidoth [19] have shown that the $\mathcal{S}_\alpha(P, Q)$ also arises as redundancy in a mismatched version of the problem of coding for tasks.

As one might expect, it is known that (see for example, Sundaresan [15, Sec. V-5]) or Johnson and Vignat [20, A.1]) $\lim_{\alpha \rightarrow 1} \mathcal{S}_\alpha(P, Q) = \mathcal{I}(P||Q)$, so that we may think of relative entropy as $\mathcal{S}_1(P, Q)$. Thus \mathcal{S}_α is a generalization of relative entropy, i.e., a *relative α -entropy*⁵.

Not surprisingly, the maximum Rényi entropy principle has been considered as a natural alternative to the maximum entropy principle of decision making under uncertainty. This principle is equivalent to another principle of maximizing the so-called Tsallis entropy which happens to be a monotone function of the Rényi entropy. Rényi entropy maximizers under moment constraints are distributions with a power-law decay (when $\alpha < 1$). See Costa et al. [22] or Johnson and Vignat [20]. Many statistical physicists have studied this principle in the hope that it may “explain” the emergence of power-laws in many naturally occurring physical and socio-economic systems, beginning with Tsallis [23]. Based on our explorations of the vast literature on this topic, we feel that our understanding, particularly one that ought to involve a modeling of the *dynamics* of such systems with the observed power-law profiles as equilibria in the asymptotics of large time, is not yet as mature as our understanding of the classical Boltzmann-Gibbs setting. But, by noting that $\mathcal{S}_\alpha(P, U) = \log |\mathbb{X}| - H_\alpha(P)$, we see that both the maximum Rényi entropy principle and the maximum Tsallis entropy principle are particular instances of a “minimum relative α -entropy principle”:

$$\text{minimize } \mathcal{S}_\alpha(P, Q) \text{ over } P \in \mathbb{E}.$$

We shall call the minimizing P^* as the forward \mathcal{S}_α -projection of Q on \mathbb{E} .

The main aim of this paper is to study forward \mathcal{S}_α -projections in general measure spaces. Our main contributions are on existence, uniqueness, and structure of these projections. We have several motivations to publish our work.

²A compression strategy $L_n: \mathbb{X}^n \rightarrow \{0, 1, 2, \dots\}$ assigns a target codeword length $L_n(x^n)$ to each string $x^n \in \mathbb{X}^n$.

³Blumer and McEliece [14], in their attempt to find better upper and lower bounds on the redundancy of generalized Huffman coding, were indirectly bounding this parameterized divergence.

⁴We suggest the pronunciation “I-alpha” for \mathcal{S}_α .

⁵This terminology is from Lutwak, et al. [21].

- We provide a rather general sufficient condition on the constraint set under which a forward \mathcal{I}_α -projection exists and is unique. This can enable statistical physicists to speak of the Rényi entropy maximizer and explore its properties even if the maximizer is not known explicitly. While the existence and uniqueness of \mathcal{I}_α -projection for closed convex sets \mathbb{E} was shown for the finite alphabet case by Sundaresan [15], here we study more general measure spaces (for example \mathbb{R}^n).
- Unlike relative entropy, its generalization relative α -entropy does not, in general, satisfy the well-known data processing inequality, nor is it in general convex in either of its arguments. Nevertheless, there is a remarkable parallelism between relative entropy and relative α -entropy. In particular, they share the “Pythagorean property” and behave like squared Euclidean distance. This too was explored by Sundaresan [15] for the finite alphabet case, and we wish to extend the parallels to more general alphabet spaces.
- We provide information on the structure of the Rényi entropy maximizer, under linear statistical constraints, whenever the maximizer exists. This can provide statistical physicists a quick means to check if their empirical observations in a particular physical setting conform to the maximum Rényi entropy principle. It also provides a means to estimate the appropriate α for a particular physical setting. Interestingly, the Rényi entropy maximizers belong to a “power-law family” of distributions that are the natural parametric generalizations of the Shannon entropy maximizers, namely the exponential family of distributions.
- In a companion paper, we shall show that a robust parameter estimation problem is a “reverse \mathcal{I}_α -projection” problem, where the minimization is with respect to the *second* argument of \mathcal{I}_α . If this reverse projection is on a power-law family, then one may turn the reverse projection into a forward projection of a specific distribution on an appropriate linear family. In that paper we shall also explore the geometric relationship between the power-law and the linear families.
- One may think of the maximum entropy principle or the minimization of relative entropy as a “projection rule”; see Section VI for projection rules with some desired properties. Three of these properties are “regularity”, “locality”, and “subspace-transitivity”. It turns out that the \mathcal{I}_α -based projection rule is regular, subspace-transitive when $\alpha < 1$, but “nonlocal”. Any regular, subspace-transitive, and local projection rule is generated by Bregman’s divergences of the sum-form [24]. In our, as yet not very successful, attempt to characterize all regular, subspace-transitive, but possibly nonlocal projection rules, we wished to understand as much as we could about a particular nonlocal projection rule. The understanding we have gained may be of use to the wider community interested in axiomatic approaches to abstract inference problems.

It is known (see for example [15]) that $\mathcal{I}_\alpha(P, Q)$ is the more commonly studied Rényi divergence of order $1/\alpha$, not of the original measures P and Q , but of their escort measures P' and Q' , where $P'(x) = P(x)^\alpha/Z(P)$, and $Z(P)$ is the normalization that makes P' a probability measure. Q' is similarly defined. While the Rényi divergences arise naturally in hypothesis testing problems (see for example Csiszár [25]), \mathcal{I}_α arises more naturally as a redundancy for mismatched compression, as discussed earlier. Moreover, $\mathcal{I}_\alpha(P, Q)$ is a certain monotone function of Csiszár’s f -divergence between P' and Q' . As a consequence of the appearance of the escort measures, the data-processing property satisfied by the f -divergences does not hold for the \mathcal{I}_α -divergences. It is therefore all the more intriguing that it is neither the f -divergences nor the Rényi divergences but the \mathcal{I}_α -divergences that share the Pythagorean property with relative entropy. However, quite recently, van Erven and Harremoës [26] showed that Rényi divergences have a Pythagorean property when the forward projection is on a so-called α -convex set.

The paper is organized as follows. In Section II, we formally define \mathcal{I}_α and establish some of its basic algebraic and topological properties, those desired of an information divergence. In Section III, we establish the existence of \mathcal{I}_α -projection on closed (in an appropriate topology) and convex sets. The proof for the case $\alpha < 1$ is analogous to that for relative entropy [7, Th. 2.1]. The proof for the case $\alpha > 1$ exploits some functional analytic tools. In Section IV, we present the Pythagorean property in generality and derive some of its immediate consequences in connection with the forward projection. In Section V, we characterize the forward \mathcal{I}_α -projection on a linear family of probability measures, whenever it exists. In Section VI, we establish a desirable subspace transitivity property and further prove the convergence of an iterative method for finding the forward \mathcal{I}_α -projection on linear families. In the concluding Section VII, we highlight some interesting open questions.

The companion paper [27] will explore the orthogonality between the power-law and the linear families, will exploit this orthogonality in a robust parameter estimation problem, and will study the reverse \mathcal{I}_α -projection in detail.

II. THE RELATIVE α - ENTROPY

We begin by defining relative α -entropy on a general measure space for all $\alpha > 0$ except $\alpha = 1$. As $\alpha \rightarrow 1$ our definition will approach the usual relative entropy or Kullback-Leibler divergence.

Let P and Q be two probability measures on a measure space $(\mathbb{X}, \mathcal{X})$. Let $\alpha \in (0, \infty)$ with $\alpha \neq 1$. Let μ be a dominating σ -finite measure on $(\mathbb{X}, \mathcal{X})$ with respect to which P and Q are both absolutely continuous, denoted $P \ll \mu$ and $Q \ll \mu$. Write $p = dP/d\mu$ and $q = dQ/d\mu$ and assume that p and q belong to the complete topological vector space $L^\alpha(\mu)$ with metric

$$d(h, g) = \begin{cases} \left(\int |h - g|^\alpha d\mu \right)^{1/\alpha} & \text{if } \alpha > 1, \\ \int |h - g|^\alpha d\mu & \text{if } \alpha < 1. \end{cases}$$

We shall use the notation

$$\|h\| := \left(\int |h|^\alpha d\mu \right)^{1/\alpha}$$

even though $\|\cdot\|$, as defined, is not a norm for $\alpha < 1$. For convenience we suppress the dependence of $d(\cdot, \cdot)$ and $\|\cdot\|$ on α ; but this dependence should be borne in mind. Throughout we shall restrict attention to probability measures whose densities with respect to μ are in $L^\alpha(\mu)$. The Rényi entropy of P of order α (with respect to μ) is defined to be

$$H_\alpha(P) := \frac{1}{1-\alpha} \log \left(\int_{\mathbb{X}} p^\alpha d\mu \right). \quad (1)$$

Consider the escort measures P' and Q' having densities p' and q' with respect to μ defined by

$$\frac{dP'}{d\mu} = p' := \frac{p^\alpha}{\int p^\alpha d\mu} \quad \text{and} \quad \frac{dQ'}{d\mu} = q' := \frac{q^\alpha}{\int q^\alpha d\mu}. \quad (2)$$

Once again, the dependence of p' and q' on α is suppressed for convenience. By setting $\alpha = \frac{1}{1+\rho}$, we have the re-parametrization in terms of ρ with $-1 < \rho < \infty$, $\rho \neq 0$, and $\rho = \alpha^{-1} - 1$. Define

$$f(u) := \text{sgn}(\rho) \cdot (u^{1+\rho} - 1), \quad u \geq 0.$$

Csiszár's f -divergence [28] between two probability measures P and Q , both absolutely continuous with respect to μ , is given by

$$I_f(P, Q) := \int q f\left(\frac{p}{q}\right) d\mu. \quad (3)$$

In the above definition we use the following conventions:

$$0 \cdot f\left(\frac{0}{0}\right) = 0,$$

and for $a > 0$,

$$0 \cdot f\left(\frac{a}{0}\right) = \begin{cases} \infty & \text{if } \rho > 0, \\ 0 & \text{if } \rho < 0. \end{cases}$$

Since f is strictly convex when $\rho \neq 0$, by Jensen's inequality, $I_f(P, Q) \geq 0$ with equality if and only if $P = Q$.

Definition 1 (Relative α -entropy): The α -entropy of P relative to Q (or relative α -entropy of P with respect to Q , or simply relative α -entropy) is defined as

$$\mathcal{J}_\alpha^\mu(P, Q) := \frac{1}{\rho} \log [\text{sgn}(\rho) \cdot I_f(P', Q') + 1]. \quad (4)$$

\mathcal{J}_α^μ depends on the reference measure μ because the densities p' and q' defined in (2) do. However, for brevity, we omit the superscript μ and ask the reader to bear the dependence on μ in mind. For the information theoretic and statistical physics motivating examples in Section I, μ is the counting measure or the Lebesgue measure depending on whether \mathbb{X} is finite or \mathbb{R}^d .

From the conventions used to define I_f , we have $\mathcal{J}_\alpha(P, Q) = \infty$ when either

- $\alpha < 1$ and $P \not\ll Q$, or
- $\alpha > 1$ and P and Q are mutually singular.

Abusing notation a little, when speaking of densities, we shall some times write $\mathcal{J}_\alpha(p, q)$ for $\mathcal{J}_\alpha(P, Q)$. Let us reemphasize that implicit in our definition of $\mathcal{J}_\alpha(P, Q)$ is the assumption that p and q are both in $L^\alpha(\mu)$.

The following are some alternative expressions of \mathcal{J}_α that are used in this paper:

$$\mathcal{J}_\alpha(P, Q) = \frac{\alpha}{1-\alpha} \log \int \frac{p}{\|p\|} \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu \quad (5)$$

$$= \frac{\alpha}{1-\alpha} \log \int p q^{\alpha-1} d\mu - \frac{1}{1-\alpha} \log \int p^\alpha d\mu + \log \int q^\alpha d\mu. \quad (6)$$

When \mathbb{X} is discrete (with μ being the counting measure on \mathbb{X}), the probability measures may be viewed as finite or countably infinite dimensional vectors. In this case, we may write

$$\mathcal{J}_\alpha(P, Q) = \frac{\alpha}{1-\alpha} \log \left[\sum_x \frac{P(x)}{\|P\|} \left(\frac{Q(x)}{\|Q\|} \right)^{\alpha-1} \right] \quad (7)$$

$$= \frac{\alpha}{1-\alpha} \log \left[\sum_x P(x) Q(x)^{\alpha-1} \right] - \frac{1}{1-\alpha} \log \sum_x P(x)^\alpha + \log \sum_x Q(x)^\alpha. \quad (8)$$

We now summarize some properties of relative α -entropy.

Lemma 2: The following properties hold.

- a) (*Positivity*). $\mathcal{J}_\alpha(P, Q) \geq 0$ with equality if and only if $P = Q$.

- b) (*Generalization of relative entropy*). Let $\mathcal{I}_\alpha(P, Q) < \infty$ for some $\alpha = \alpha_l < 1$ and simultaneously for some $\alpha = \alpha_u > 1$. Then $\mathcal{I}_\alpha(P, Q)$ is well-defined for all $\alpha \in [\alpha_l, \alpha_u] \setminus \{1\}$, and

$$\lim_{\alpha \rightarrow 1} \mathcal{I}_\alpha(P, Q) = \mathcal{I}(P\|Q),$$

where $\mathcal{I}(P\|Q)$ is the relative entropy of P with respect to Q .

- c) (*Relation to Rényi divergence*).

$$\mathcal{I}_\alpha(P, Q) = D_{1/\alpha}(P\|Q),$$

where

$$D_\beta(P\|Q) := \frac{1}{\beta-1} \log \int p^\beta q^{1-\beta} d\mu$$

is the Rényi divergence of order β .

- d) (*Relation to Rényi entropy*). Let $|\mathbb{X}| < \infty$ and let U be the uniform probability measure on \mathbb{X} . Then $\mathcal{I}_\alpha(P, U) = \log |\mathbb{X}| - H_\alpha(P)$.
- e) (*Rényi entropy maximizer under a covariance constraint*). Let $\mathbb{X} = \mathbb{R}^n$ and let μ be the Lebesgue measure on \mathbb{R}^n . For $\alpha > n/(n+2)$ and $\alpha \neq 1$, define the constant $b_\alpha = (1-\alpha)/(2\alpha - n(1-\alpha))$. With C a positive definite covariance matrix, the function

$$\phi_{\alpha, C}(x) = Z_\alpha^{-1} [1 + b_\alpha \cdot x^T C^{-1} x]_+^{\frac{1}{\alpha-1}},$$

with $[a]_+ := \max\{a, 0\}$ and Z_α the normalization constant, is the density function of a probability measure on \mathbb{R}^n whose covariance matrix is C . Furthermore, if g is the density function of any other random vector with covariance matrix C , then

$$\mathcal{I}_\alpha(g, \phi_{\alpha, C}) = H_\alpha(\phi_{\alpha, C}) - H_\alpha(g). \quad (9)$$

Consequently $\phi_{\alpha, C}$ is the density function of the Rényi entropy maximizer among all \mathbb{R}^n -valued random vectors with covariance matrix C .

Proof: See Appendix A. ■

Remark 1: For relative entropy ($\alpha = 1$), the analog of (9) under a covariance constraint is

$$\mathcal{I}(g\|\phi) = H(\phi) - H(g),$$

where H is differential entropy and ϕ is the Gaussian distribution with the same covariance as g [4, Th. 8.6.5]. In Section V we shall study Rényi entropy maximizers under more general linear constraints.

Remark 2: While the numerical value of relative entropy $\mathcal{I}(P\|Q)$ does not depend on the dominating measure μ , recall that $\mathcal{I}_\alpha(P, Q)$ does depend on μ in general.

Analogous to the property that $p \mapsto \mathcal{I}(p\|q)$ is lower semicontinuous in the topology on $L^1(\mu)$ arising from the total variation metric [29, Sec. 2.4, Assertion 5], we have the following.

Proposition 3 (Lower semicontinuity in the first argument): For a fixed q , consider $p \mapsto \mathcal{I}_\alpha(p, q)$ as a function on $L^\alpha(\mu)$. This function is continuous for $\alpha > 1$ and lower semicontinuous for $\alpha < 1$.

Proof: See Appendix B. ■

Remark 3: When $\alpha < 1$, $\mathcal{I}_\alpha(\cdot, Q)$ is lower semicontinuous, but not necessarily continuous. To see this, let \mathbb{X} be finite. Let P_n, P, Q be probability measures on \mathbb{X} such that all P_n 's have full support, i.e., $P_n(x) > 0$ for all $x \in \mathbb{X}$, but $Q(x_0) = 0$ for some $x_0 \in \mathbb{X}$, $P \ll Q$, and finally $P_n \rightarrow P$. Then $\mathcal{I}_\alpha(P_n, Q) = \infty$ for all n , but $\mathcal{I}_\alpha(P, Q) < \infty$.

Remark 4: If however \mathbb{X} is finite and Q has full support, then $\mathcal{I}_\alpha(\cdot, Q)$ is indeed continuous and this can be seen by taking the limit term by term in (7).

We now address the behavior as a function of q .

Proposition 4: Fix $\alpha > 0$, $\alpha \neq 1$. For a fixed p , the mapping $q \mapsto \mathcal{I}_\alpha(p, q)$ is lower semicontinuous in $L^\alpha(\mu)$.

Proof: See Appendix C ■

Remark 5: When \mathbb{X} is finite, with $+\infty$ as a potential limiting value, $\mathcal{I}_\alpha(P, \cdot)$ is continuous for all $\alpha > 0$, $\alpha \neq 1$, as is easily seen by taking term-wise limits in the summation in (7).




Fig. 1: The usual Apollonius theorem would be, with $\lambda = \frac{1}{2}$, $\frac{1}{2}|\overrightarrow{RP_1}|^2 + \frac{1}{2}|\overrightarrow{RP_0}|^2 = \frac{1}{2}|\overrightarrow{R_{1,0}P_1}|^2 + \frac{1}{2}|\overrightarrow{R_{1,0}P_0}|^2 + |\overrightarrow{R_{1,0}R}|^2$. Here, $|\overrightarrow{RP_1}|^2$ is replaced by the asymmetric $I_f(P'_1, R')$, etc., and the equality by an inequality whose direction depends on $\alpha < 1$ or $\alpha > 1$.

We next establish quasi-convexity of \mathcal{S}_α in the first argument, i.e., for every fixed q and real number τ , the lower level sets $\overline{B}(q, \tau) := \{p: \mathcal{S}_\alpha(p, q) \leq \tau\}$ (or “ \mathcal{S}_α -balls”) are convex.

Proposition 5: Fix $\alpha > 0$, $\alpha \neq 1$. For a fixed q , the mapping $p \mapsto \mathcal{S}_\alpha(p, q)$ is quasi-convex in $L^\alpha(\mu)$.

Proof: See Appendix D ■

Remark 6: In general, for both $\alpha < 1$ and $\alpha > 1$, \mathcal{S}_α is not convex in either of its arguments. Moreover, \mathcal{S}_α does not satisfy the data processing inequality while relative entropy and more generally Csiszár’s f -divergences do.

III. EXISTENCE AND UNIQUENESS OF THE FORWARD \mathcal{S}_α -PROJECTION

In this section, we shall introduce the notion of a forward \mathcal{S}_α -projection of a probability measure on a subset of probability measures. We shall also prove a sufficiency result for the existence of the forward \mathcal{S}_α -projection. We begin by first proving a useful inequality relating f -divergences. This is an inequality that turns out to be the analog of the parallelogram identity of [7] for relative entropy ($\alpha = 1$) and the analog of the Apollonius Theorem in plane geometry (see, for e.g., Bhatia [30, p. 85]). While these analogs show an equality, our generalization is at the cost of a weakening of the equality to an inequality.

Proposition 6 (Extension of Apollonius Theorem): Let $\alpha < 1$. Let P_0, P_1, R be probability measures that are absolutely continuous with respect to μ , and let the corresponding Radon-Nikodym derivatives p_0, p_1 , and r be in $L^\alpha(\mu)$. Assume $0 \leq \lambda \leq 1$. We then have

$$\lambda I_f(P'_1, R') + (1 - \lambda) I_f(P'_0, R') - \lambda I_f(P'_1, R'_{1,0}) - (1 - \lambda) I_f(P'_0, R'_{1,0}) \geq I_f(R'_{1,0}, R'), \quad (10)$$

where

$$R_{1,0} = \frac{\frac{\lambda}{\|p_1\|} P_1 + \frac{1-\lambda}{\|p_0\|} P_0}{\frac{\lambda}{\|p_1\|} + \frac{1-\lambda}{\|p_0\|}}. \quad (11)$$

When $\alpha > 1$, the reversed inequality holds in (10).

Proof: See Figure 1 for an interpretation of (10) as an analog of the Apollonius Theorem. We first recognize that

$$I_f(P', Q') = \text{sgn}(\rho) \left[\int \frac{p}{\|p\|} \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu - 1 \right]. \quad (12)$$

Let $r_{1,0} = dR_{1,0}/d\mu$. Using (12), the left-hand side of (10) can be expanded to

$$\begin{aligned} & \text{sgn}(\rho) \int \frac{\lambda p_1}{\|p_1\|} \left[\left(\frac{r}{\|r\|} \right)^{\alpha-1} - \left(\frac{r_{1,0}}{\|r_{1,0}\|} \right)^{\alpha-1} \right] d\mu \\ & + \text{sgn}(\rho) \int \frac{(1-\lambda)p_0}{\|p_0\|} \left[\left(\frac{r}{\|r\|} \right)^{\alpha-1} - \left(\frac{r_{1,0}}{\|r_{1,0}\|} \right)^{\alpha-1} \right] d\mu \\ \stackrel{(a)}{=} & \text{sgn}(\rho) \int \frac{r_{1,0}}{\|r_{1,0}\|} \left[\left(\frac{r}{\|r\|} \right)^{\alpha-1} - \left(\frac{r_{1,0}}{\|r_{1,0}\|} \right)^{\alpha-1} \right] d\mu \\ & \times \left[\frac{\lambda}{\|p_1\|} + \frac{1-\lambda}{\|p_0\|} \right] \|r_{1,0}\| \\ \stackrel{(b)}{=} & \left[\frac{\lambda}{\|p_1\|} + \frac{1-\lambda}{\|p_0\|} \right] \|r_{1,0}\| \cdot I_f(R'_{1,0}, R'), \end{aligned}$$

where (a) follows from (11) and after a multiplication and a division by the scalar $\|r_{1,0}\|$; (b) follows from (12). The lemma would follow if we can show

$$\left(\frac{\lambda}{\|p_1\|} + \frac{1-\lambda}{\|p_0\|} \right) \|r_{1,0}\| \geq 1$$

for $\alpha < 1$, and the reversed inequality for $\alpha > 1$. But these are direct consequences of Minkowski's inequalities for $\alpha < 1$ and $\alpha > 1$ applied to (11). \blacksquare

Let us now formally define what we mean by a forward \mathcal{I}_α -projection.

Definition 7: If \mathbb{E} is a set of probability measures on $(\mathbb{X}, \mathcal{X})$ such that $\mathcal{I}_\alpha(P, R) < \infty$ for some $P \in \mathbb{E}$, a measure $Q \in \mathbb{E}$ satisfying

$$\mathcal{I}_\alpha(Q, R) = \inf_{P \in \mathbb{E}} \mathcal{I}_\alpha(P, R) =: \mathcal{I}_\alpha(\mathbb{E}, R) \quad (13)$$

is called a forward \mathcal{I}_α -projection of R on \mathbb{E} .

For a set \mathbb{E} of probability measures on $(\mathbb{X}, \mathcal{X})$, let

$$\mathcal{E} := \left\{ p = \frac{dP}{d\mu} : P \in \mathbb{E} \right\}$$

be the corresponding set of μ -densities. We shall assume that $\mathcal{E} \subset L^\alpha(\mu)$.

We are now ready to state our first main result on the existence and uniqueness of the forward \mathcal{I}_α -projection.

Theorem 8 (Existence and uniqueness of the forward \mathcal{I}_α -projection): Fix $\alpha > 0$, $\alpha \neq 1$. Let \mathbb{E} be a set of probability measures whose corresponding set of density functions \mathcal{E} is convex and closed in $L^\alpha(\mu)$. Let R be a probability measure (with density r) and suppose that $\mathcal{I}_\alpha(P, R) < \infty$ for some $P \in \mathbb{E}$. Then R has a unique forward \mathcal{I}_α -projection on \mathbb{E} .

Remark 7: This is a generalization of Csiszár's projection result [7, Th. 2.1] for relative entropy ($\alpha = 1$). The analog of " \mathcal{E} is closed in $L^\alpha(\mu)$ " for relative entropy is closure in the topology arising from the total variation metric, one of the hypotheses in [7, Th. 2.1]. The proof ideas are different for the two cases $\alpha < 1$ and $\alpha > 1$. The proof for $\alpha < 1$ is a modification of Csiszár's approach in [7], and is similar to the classical proof of existence and uniqueness of the *best approximant* of a point (in a Hilbert space) from a given closed and convex set of the Hilbert space. (See, for e.g., [30, Ch. 11, Th. 14]). The proof for $\alpha > 1$ exploits the reflexive property of the Banach space $L^\alpha(\mu)$. This alternative approach is required because the inequality in the extension of Apollonius Theorem (Proposition 6) is in a direction that renders the classical approach inapplicable. We are indebted to Pietro Majer for suggesting some key steps for the $\alpha > 1$ case on the mathoverflow.net forum.

Remark 8: In general, when $\alpha \neq 1$, the forward \mathcal{I}_α -projection depends on the reference measure μ . The case $\alpha = 1$ of relative entropy is however special in that the forward \mathcal{I}_1 -projection does not depend on the reference measure μ .

Remark 9: The above result was established by Sundaresan [15, Prop. 23] for finite \mathbb{X} . That proof relied on the compactness of \mathbb{E} . The current proof works for general measure spaces.

Proof: (a) We first consider the case $\alpha < 1$.

Existence of forward projection: Pick a sequence (P_n) in \mathbb{E} such that $I_f(P'_n, R') < \infty$ and

$$I_f(P'_n, R') \rightarrow \inf_{P \in \mathbb{E}} I_f(P', R'). \quad (14)$$

Apply Proposition 6 with $\lambda = \frac{1}{2}$ to get

$$\frac{1}{2}I_f(P'_m, R') + \frac{1}{2}I_f(P'_n, R') - \frac{1}{2}I_f(P'_m, R'_{m,n}) - \frac{1}{2}I_f(P'_n, R'_{m,n}) \geq I_f(R'_{m,n}, R'), \quad (15)$$

where

$$R_{m,n} = \frac{\frac{1}{\|p_m\|}P_m + \frac{1}{\|p_n\|}P_n}{\frac{1}{\|p_m\|} + \frac{1}{\|p_n\|}}.$$

$R_{m,n} \in \mathbb{E}$ on account of the convexity of \mathbb{E} . Using $I_f(\cdot, \cdot) \geq 0$ and then rearranging (15), we get

$$0 \leq \frac{1}{2}I_f(P'_m, R'_{m,n}) + \frac{1}{2}I_f(P'_n, R'_{m,n}) \quad (16)$$

$$\leq \frac{1}{2}I_f(P'_m, R') + \frac{1}{2}I_f(P'_n, R') - I_f(R'_{m,n}, R'). \quad (17)$$

Now let $m, n \rightarrow \infty$. We claim the expression on the right-hand side of (17) must approach 0. Indeed, that the liminf of the right-hand side of (17) is at least 0 is clear from the inequalities (16) and (17). But the limsup is at most 0 because both $I_f(P'_m, R')$ and $I_f(P'_n, R')$ approach the infimum value, and $I_f(R'_{m,n}, R')$ is at least this infimum value for each m and n . This establishes the claim.

Consequently, the right-hand side of (16) converges to 0. Using this and the nonnegativity of $I_f(\cdot, \cdot)$, we get

$$\lim_{m,n \rightarrow \infty} I_f(P'_m, R'_{m,n}) = 0. \quad (18)$$

From [31, Th. 1], a generalization of Pinsker's inequality for f -divergence under $\alpha < 1$, and with $|P - Q|_{TV}$ denoting the total variation distance between probability measures P and Q , we have

$$\lim_{m,n \rightarrow \infty} |P'_m - R'_{m,n}|_{TV} = 0.$$

The triangle inequality for the total variation metric then yields

$$|P'_m - P'_n|_{TV} \leq |P'_m - R'_{m,n}|_{TV} + |P'_m - R'_{m,n}|_{TV} \rightarrow 0$$

as $m, n \rightarrow \infty$, i.e., the sequence (p'_n) is a Cauchy sequence in $L^1(\mu)$. It must therefore converge to some g in $L^1(\mu)$, i.e.,

$$\lim_{n \rightarrow \infty} \int |p'_n - g| d\mu = 0. \quad (19)$$

It follows that $\int p'_n d\mu \rightarrow \int g d\mu$, and since $\int p'_n d\mu = 1$ for all n , we must have $\int g d\mu = 1$.

From the $L^1(\mu)$ convergence in (19), we also have $p'_n \rightarrow g$ in $[\mu]$ -measure.

We will now demonstrate that the probability measure with μ -density proportional to $g^{1/\alpha}$ is in \mathbb{E} and is a forward \mathcal{S}_α -projection, thereby establishing existence.

In view of the convergence in $[\mu]$ -measure and the upper bound

$$\left| (p'_n)^{1/\alpha} - g^{1/\alpha} \right|^\alpha \leq 2^\alpha [p'_n + g],$$

we can apply the generalized version of the dominated convergence theorem ([32, Ch. 2, Ex. 20] or [33, p.139, Problem 19]) to get

$$\frac{p_n}{\|p_n\|} = (p'_n)^{1/\alpha} \rightarrow g^{1/\alpha} \text{ in } L^\alpha(\mu).$$

We next claim that

$$\|p_n\| \text{ is bounded.} \quad (20)$$

Suppose not; then working on a subsequence if needed, we have $\|p_n\| := M_n \rightarrow \infty$. As $\int p_n d\mu = 1$, given any $\epsilon > 0$,

$$\mu(\{p'_n > \epsilon\}) = \mu\left(\left\{p_n > \epsilon^{1/\alpha} M_n\right\}\right) \leq \frac{1}{\epsilon^{1/\alpha} M_n} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

and hence $p'_n \rightarrow 0$ in $[\mu]$ -measure, or $g = 0$ except on a set of $[\mu]$ -measure 0 (i.e., $g = 0$ a.e. $[\mu]$). But this is a contradiction since $\int g d\mu = 1$. Thus (20) holds, and we can pick a subsequence of the sequence $(\|p_n\|)$ that converges to some c . Reindex and work on this subsequence to get $p_n \rightarrow cg^{1/\alpha}$ in $L^\alpha(\mu)$.

It is now that we use the hypothesis that \mathcal{E} is closed in $L^\alpha(\mu)$. We remind the reader that \mathcal{E} is the set of μ -densities of members of \mathbb{E} . The closedness implies that the limiting function $cg^{1/\alpha} = q$ for some $q \in \mathcal{E}$, and so q must be the density of a probability measure, say Q . Since we also have $\int g d\mu = 1$, it follows that $c = \|q\|$ and $g = q^\alpha / \|q\|^\alpha$. As $p_n \rightarrow q$ in $L^\alpha(\mu)$, lower semicontinuity of $\mathcal{S}_\alpha(\cdot, r)$ (Proposition 3) implies

$$\mathcal{S}_\alpha(Q, R) \leq \liminf_{n \rightarrow \infty} \mathcal{S}_\alpha(P_n, R) = \mathcal{S}_\alpha(\mathbb{E}, R). \quad (21)$$

Since $Q \in \mathbb{E}$, $\mathcal{S}_\alpha(Q, R) \geq \mathcal{S}_\alpha(\mathbb{E}, R)$, and therefore equality must hold in (21), and Q is a forward \mathcal{S}_α -projection of R on \mathbb{E} .

Uniqueness: Our proof of uniqueness is analogous to the usual proof of uniqueness of projection in Hilbert spaces [30, p. 86]. A simpler proof, after the 'Pythagorean property' is established, can be found at the end of Section IV.

Write d for the infimum value in the right-hand side of (14) and let Q_1 and Q_0 attain the infimum. Apply Proposition 6 with $\lambda = 1/2$ and with Q_1 and Q_0 in place of P_1 and P_0 to get

$$\frac{1}{2} I_f(Q'_1, R') + \frac{1}{2} I_f(Q'_0, R') - \frac{1}{2} I_f(Q'_1, R'_{1,0}) - \frac{1}{2} I_f(Q'_0, R'_{1,0}) \geq I_f(R'_{1,0}, R'), \quad (22)$$

where

$$R_{1,0} = \frac{\frac{1}{\|q_1\|} Q_1 + \frac{1}{\|q_0\|} Q_0}{\frac{1}{\|q_1\|} + \frac{1}{\|q_0\|}}.$$

Since $R_{1,0} \in \mathbb{E}$ we have $I_f(R'_{1,0}, R') \geq d$. Use this in (22), substitute $I_f(Q'_i, R') = d$, $i = 0, 1$, and we get

$$\frac{d}{2} + \frac{d}{2} - \frac{1}{2} I_f(Q'_1, R'_{1,0}) - \frac{1}{2} I_f(Q'_0, R'_{1,0}) \geq d,$$

and this implies

$$I_f(Q'_1, R'_{1,0}) + I_f(Q'_0, R'_{1,0}) \leq 0.$$

The nonnegativity of each of the terms then implies that each must be zero, and so $Q_1 = R_{1,0} = Q_0$. The forward \mathcal{J}_α -projection is unique.

This completes the proof for the case when $\alpha < 1$.

(b) We now consider the case when $\alpha > 1$.

Existence of forward projection: Equation (13) can be rewritten (using (5)) as

$$\inf_{P \in \mathbb{E}} \mathcal{J}_\alpha(P, R) = \frac{1}{\rho} \log \left[\sup_{p \in \mathcal{E}} \int \frac{p}{\|p\|} \left(\frac{r}{\|r\|} \right)^{\alpha-1} d\mu \right] \quad (23)$$

$$= \frac{1}{\rho} \log \left[\sup_{h \in \hat{\mathcal{E}}} \int hg d\mu \right], \quad (24)$$

where

$$\hat{\mathcal{E}} := \left\{ s \frac{p}{\|p\|} : p \in \mathcal{E}, 0 \leq s \leq 1 \right\},$$

and $g = (r/\|r\|)^{\alpha-1}$, an element of the dual space $(L^\alpha(\mu))^*$. Allowing $s \in [0, 1]$ makes $\hat{\mathcal{E}}$ convex (as we shall soon show), but does not change the supremum.

We now claim that

$$\hat{\mathcal{E}} \text{ is a closed and convex subset of } L^\alpha(\mu). \quad (25)$$

Assume the claim. Since $L^\alpha(\mu)$ is a reflexive Banach space for $\alpha > 1$, the convex and closed set $\hat{\mathcal{E}}$ is also closed in the *weak topology* [34, Ch. 10, Cor. 23]. Using the Banach-Alaoglu theorem and the fact that $L^\alpha(\mu)$ is a reflexive Banach space, we have that the unit ball is compact in the weak topology. Since $\hat{\mathcal{E}}$ is a (weakly) closed subset of a (weakly) compact set, $\hat{\mathcal{E}}$ is (weakly) compact. The linear functional $h \mapsto \int hg d\mu$ is continuous in the weak topology, and hence the supremum over the (weakly) compact set $\hat{\mathcal{E}}$ is attained. Since the linear functional increases with s , the supremum is attained when $s = 1$, i.e., there exists a $p \in \mathcal{E}$ for which the supremum in (23) is attained.

We now proceed to show the claim (25). To see convexity, let $p_1, p_0 \in \mathcal{E}$, let $0 \leq s_1, s_0 \leq 1$, and let $0 \leq \lambda \leq 1$. The convex combination of $s_1 p_1 / \|p_1\|$ and $s_0 p_0 / \|p_0\|$ is

$$\lambda s_1 \frac{p_1}{\|p_1\|} + (1 - \lambda) s_0 \frac{p_0}{\|p_0\|}.$$

If both λs_1 and $(1 - \lambda) s_0$ are zero, then this convex combination is 0 which is trivially in $\hat{\mathcal{E}}$. Otherwise, we can write the convex combination as

$$\lambda s_1 \frac{p_1}{\|p_1\|} + (1 - \lambda) s_0 \frac{p_0}{\|p_0\|} = s_\lambda \frac{p_\lambda}{\|p_\lambda\|}, \quad (26)$$

where

$$p_\lambda := \frac{\frac{\lambda s_1}{\|p_1\|} p_1 + \frac{(1-\lambda)s_0}{\|p_0\|} p_0}{\frac{\lambda s_1}{\|p_1\|} + \frac{(1-\lambda)s_0}{\|p_0\|}}, \quad (27)$$

$$s_\lambda := \left(\frac{\lambda s_1}{\|p_1\|} + \frac{(1-\lambda)s_0}{\|p_0\|} \right) \cdot \|p_\lambda\|. \quad (28)$$

To show that the convex combination is in $\hat{\mathcal{E}}$, it suffices to show that $p_\lambda \in \mathcal{E}$ and $s_\lambda \in [0, 1]$.

The convexity of \mathcal{E} immediately implies that $p_\lambda \in \mathcal{E}$. It is also clear that $s_\lambda \geq 0$. From Minkowski's inequality (for $\alpha > 1$), we have

$$\begin{aligned} s_\lambda &= \left(\frac{\lambda s_1}{\|p_1\|} + \frac{(1-\lambda)s_0}{\|p_0\|} \right) \cdot \|p_\lambda\| \\ &= \left\| \frac{\lambda s_1}{\|p_1\|} p_1 + \frac{(1-\lambda)s_0}{\|p_0\|} p_0 \right\| \\ &\leq \frac{\lambda s_1}{\|p_1\|} \cdot \|p_1\| + \frac{(1-\lambda)s_0}{\|p_0\|} \cdot \|p_0\| \\ &= \lambda s_1 + (1-\lambda)s_0 \\ &\leq 1. \end{aligned} \quad (29)$$

This establishes that $\hat{\mathcal{E}}$ is convex.

To see that $\hat{\mathcal{E}}$ is closed in $L^\alpha(\mu)$, let (g_n) be a sequence in $\hat{\mathcal{E}}$ such that $g_n \rightarrow g$ for some $g \in L^\alpha(\mu)$. We need to show $g \in \hat{\mathcal{E}}$.




Fig. 2: Pythagorean property with inequality




Fig. 3: Pythagorean property with equality

Write $g_n = s_n p_n / \|p_n\|$, where $p_n \in \mathcal{E}$ and $0 \leq s_n \leq 1$. Since $g_n \rightarrow g$ in $L^\alpha(\mu)$, take norms to get $s_n = \|g_n\| \rightarrow \|g\|$, and so $\|g\| \leq 1$.

If $\|g\| = 0$, then $g = 0$ a.e. $[\mu]$, and so g trivially belongs to $\hat{\mathcal{E}}$. We may therefore assume $\|g\| > 0$. It follows that $p_n / \|p_n\| = g_n / \|g_n\| \rightarrow g / \|g\|$ in $L^\alpha(\mu)$.

Again, as in (20), we claim that $\|p_n\|$ is bounded. Suppose not. As in the proof of (20), move to a subsequence if needed and assume $\|p_n\| := M_n \rightarrow \infty$. As $\int p_n d\mu = 1$, we have

$$\mu \left(\left\{ \frac{p_n}{\|p_n\|} > \epsilon \right\} \right) = \mu (\{p_n > \epsilon M_n\}) \leq \frac{1}{\epsilon M_n} \rightarrow 0$$

as $n \rightarrow \infty$, and $p_n / \|p_n\| \rightarrow 0$ in μ -measure, or its limit $g / \|g\| = 0$ a.e. $[\mu]$. But this contradicts the fact that $\int (g / \|g\|)^\alpha d\mu = 1$. Thus $\|p_n\|$ is bounded.

Focusing on a subsequence, if needed, we may assume $\|p_n\| \rightarrow c$ for some $c \geq 0$. Hence $p_n \rightarrow cg / \|g\|$ in $L^\alpha(\mu)$. Since \mathcal{E} is closed, we must have $cg / \|g\| = p$ for some $p \in \mathcal{E}$, whence $c = \|p\|$ and $g = \|g\| \cdot p / \|p\|$. Since we already established that $\|g\| \leq 1$, it follows that $g \in \hat{\mathcal{E}}$.

Uniqueness: We now proceed to show uniqueness.

Let p_0, p_1 attain the supremum in (23). Set $h_0 = s_0 p_0 / \|p_0\|$ and $h_1 = s_1 p_1 / \|p_1\|$ with $s_0 = s_1 = 1$. Clearly h_0 and h_1 attain the supremum in (24). By convexity of $\hat{\mathcal{E}}$, $\frac{1}{2}h_1 + \frac{1}{2}h_0$ belongs to $\hat{\mathcal{E}}$. This and the linearity of the integral in (24) in the h variable imply that $\frac{1}{2}h_1 + \frac{1}{2}h_0$ attains the supremum in (24). Noticing that $\frac{1}{2}h_1 + \frac{1}{2}h_0 = s_{\frac{1}{2}} p_{\frac{1}{2}} / \|p_{\frac{1}{2}}\|$ as in (26), with $p_{\frac{1}{2}}$ and $s_{\frac{1}{2}}$ as in (27) and (28), respectively, we gather that $s_{\frac{1}{2}} = 1$. Consequently, all the inequalities in the chain (29) must be equalities. But then p_1 and p_0 are scalings of each other (which is the condition for equality in Minkowski's inequality). Since p_0 and p_1 are densities of probability measures with respect to μ , we deduce that the scaling factor must be 1, i.e., $p_0 = p_1$. This completes the proof. \blacksquare

IV. PYTHAGOREAN PROPERTY

In this section, we state and prove the Pythagorean property for relative α -entropy. We define the \mathcal{I}_α -ball with center R and radius τ to be $B(R, \tau) := \{P : \mathcal{I}_\alpha(P, R) < \tau\}$, $0 < \tau \leq \infty$. By virtue of quasi-convexity, $B(R, \tau)$ is a convex set.

Theorem 9 (The Pythagorean property): Let $\alpha > 0$ and $\alpha \neq 1$.

- (a) Let $\mathcal{I}_\alpha(P, R)$ and $\mathcal{I}_\alpha(Q, R)$ be finite. The segment joining P and Q does not intersect the \mathcal{I}_α -ball $B(R, \tau)$ with radius $\tau = \mathcal{I}_\alpha(Q, R)$, i.e., $\mathcal{I}_\alpha(P_\lambda, R) \geq \mathcal{I}_\alpha(Q, R)$ for

$$P_\lambda = \lambda P + (1 - \lambda)Q, \quad \lambda \in [0, 1],$$

(see figure 2) if and only if

$$\mathcal{I}_\alpha(P, R) \geq \mathcal{I}_\alpha(P, Q) + \mathcal{I}_\alpha(Q, R). \quad (30)$$

- (b) Let

$$Q = \lambda P + (1 - \lambda)S, \quad \text{for some fixed } \lambda \in (0, 1), \quad (31)$$

and let $\mathcal{I}_\alpha(Q, R)$ be finite. The segment joining P and S does not intersect $B(R, \tau)$ with $\tau = \mathcal{I}_\alpha(Q, R)$ (see figure 3) if and only if the following two equalities hold:

$$\left. \begin{aligned} \mathcal{I}_\alpha(P, R) &= \mathcal{I}_\alpha(P, Q) + \mathcal{I}_\alpha(Q, R) \\ \mathcal{I}_\alpha(S, R) &= \mathcal{I}_\alpha(S, Q) + \mathcal{I}_\alpha(Q, R). \end{aligned} \right\} \quad (32)$$

Proof: Our proof proceeds as in [15], where the above result is proved for the finite alphabet case, with appropriate functional analytic justifications to account for the generality of the alphabet.

(a) We begin with the “only if” part. Assume $\mathcal{I}_\alpha(P, R)$ and $\mathcal{I}_\alpha(Q, R)$ are finite, and that the segment joining P and Q does not intersect the \mathcal{I}_α -ball $B(R, \tau)$ with radius $\tau = \mathcal{I}_\alpha(Q, R)$. To show (30), since

$$\begin{aligned} I_f(P', R') &= \int r' f\left(\frac{p'}{r'}\right) d\mu \\ &= \operatorname{sgn}(\rho) \left[\int (p')^{1+\rho} \cdot (r')^{-\rho} d\mu - 1 \right] \\ &= \operatorname{sgn}(\rho) \left[\int \frac{p}{\|p\|} \cdot (r')^{-\rho} d\mu - 1 \right], \end{aligned}$$

which follow from (2), (3), and $\alpha(1 + \rho)$, it suffices to show that

$$\operatorname{sgn}(\rho) \int p \cdot (r')^{-\rho} d\mu \geq \frac{\operatorname{sgn}(\rho)}{\|q\|} \int p \cdot (q')^{-\rho} d\mu \cdot \int q \cdot (r')^{-\rho} d\mu. \quad (33)$$

We have

$$I_f(P'_\lambda, R') = \operatorname{sgn}(\rho) \left[\int \frac{p_\lambda}{\|p_\lambda\|} \cdot (r')^{-\rho} d\mu - 1 \right]$$

Let

$$\begin{aligned} s(\lambda) &:= \int p_\lambda \cdot (r')^{-\rho} d\mu, \\ t(\lambda) &:= \|p_\lambda\|. \end{aligned}$$

Clearly, $\mathcal{I}_\alpha(P_\lambda, R) \geq \mathcal{I}_\alpha(Q, R)$ for $\lambda \in (0, 1)$ implies that

$$\frac{I_f(P'_\lambda, R') - I_f(P'_0, R')}{\lambda} \geq 0 \quad \text{for } \lambda \in (0, 1). \quad (34)$$

Therefore, by taking the limit as $\lambda \rightarrow 0$, the derivative of $I_f(P'_\lambda, R')$ with respect to λ evaluated at $\lambda = 0$, should be ≥ 0 . Observe that

$$\begin{aligned} \frac{s(\lambda) - s(0)}{\lambda} &= \frac{1}{\lambda} \left[\int p_\lambda \cdot (r')^{-\rho} d\mu - \int q \cdot (r')^{-\rho} d\mu \right] \\ &= \int \left(\frac{p_\lambda - q}{\lambda} \right) \cdot (r')^{-\rho} d\mu \\ &= \int (p - q) \cdot (r')^{-\rho} d\mu \\ &= \left[\int p \cdot (r')^{-\rho} d\mu - \int q \cdot (r')^{-\rho} d\mu \right]. \end{aligned}$$

So $\dot{s}(0) := \lim_{\lambda \downarrow 0} (s(\lambda) - s(0))/\lambda$ exists and equals the above expression.

Let us now identify $\dot{t}(0)$. For $\alpha > 1$ (i.e., $\rho > 0$), we have

$$\left| \frac{\partial}{\partial \lambda} (p_\lambda)^\alpha \right| = \alpha (p_\lambda)^{\alpha-1} |p - q| \leq \alpha (p + q)^\alpha,$$

while for $\alpha < 1$, notice that for any $0 < l < \frac{1}{2}$, we have

$$\begin{aligned} \left| \frac{\partial}{\partial \lambda} (p_\lambda)^\alpha \right| &= \alpha (p_\lambda)^{\alpha-1} |p - q| = \frac{\alpha |p - q|}{[\lambda p + (1 - \lambda)q]^{1-\alpha}} \leq \frac{\alpha(p + q)}{\min\{\lambda, (1 - \lambda)\}^{1-\alpha} (p + q)^{1-\alpha}} \\ &\leq \frac{\alpha(p + q)^\alpha}{l^{1-\alpha}} \quad \forall \lambda \in (l, 1 - l), \end{aligned}$$

and both upper bounds are in $L^1(\mu)$. Therefore by chain rule and [32, Th. 2.27], we get

$$\dot{t}(\lambda) = \left[\int (p_\lambda)^\alpha d\mu \right]^{\frac{1}{\alpha}-1} \cdot \int (p_\lambda)^{\alpha-1} (p - q) d\mu$$

for $\lambda \in (l, 1-l)$. As $\lambda \downarrow 0$ (by moving l closer to 0), we get

$$\begin{aligned}
\dot{t}(0) &= \left(\int q^\alpha d\mu \right)^{\frac{1}{\alpha}-1} \cdot \int q^{\alpha-1}(p-q)d\mu \\
&= \left(\int q^\alpha d\mu \right)^{\frac{1-\alpha}{\alpha}} \cdot \left(\int pq^{\alpha-1}d\mu - \int q^\alpha d\mu \right) \\
&= \int p \left(\frac{q^\alpha}{\int q^\alpha d\mu} \right)^{\frac{\alpha-1}{\alpha}} d\mu - \left(\int q^\alpha d\mu \right)^{\frac{1}{\alpha}} \\
&= \int p \cdot (q')^{-\rho} d\mu - \|q\|.
\end{aligned}$$

Since

$$\frac{1}{\lambda} \left[\frac{s(\lambda)}{t(\lambda)} - \frac{s(0)}{t(0)} \right] = \frac{1}{t(\lambda)t(0)} \left[t(0) \frac{s(\lambda) - s(0)}{\lambda} - s(0) \frac{t(\lambda) - t(0)}{\lambda} \right],$$

it follows that the derivative of $s(\lambda)/t(\lambda)$ exists at $\lambda = 0$ and is given by $(t(0)\dot{s}(0) - s(0)\dot{t}(0))/t^2(0)$. Equation (34) and $t(0) > 0$ imply that

$$\dot{s}(0) - s(0) \cdot \frac{\dot{t}(0)}{t(0)} \geq 0. \quad (35)$$

Consequently, $\dot{t}(0)$ is necessarily finite. Substituting the values of $s(0)$, $\dot{s}(0)$, $t(0)$ and $\dot{t}(0)$ in (35) we get the required inequality (33).

To prove the converse “if” part, let us assume that

$$\mathcal{I}_\alpha(P, R) \geq \mathcal{I}_\alpha(P, Q) + \mathcal{I}_\alpha(Q, R),$$

which is the same as (33). Since $\mathcal{I}_\alpha(P, R)$ and $\mathcal{I}_\alpha(Q, R)$ are finite, it follows that $\mathcal{I}_\alpha(P, Q)$ is also finite. From the trivial statement

$$\mathcal{I}_\alpha(Q, R) = \mathcal{I}_\alpha(Q, Q) + \mathcal{I}_\alpha(Q, R), \quad (36)$$

we get the following analog of (33) but with equality (replace p in (33) with q):

$$\text{sgn}(\rho) \int q \cdot (r')^{-\rho} d\mu = \frac{\text{sgn}(\rho)}{\|q\|} \int q \cdot (q')^{-\rho} d\mu \cdot \int q \cdot (r')^{-\rho} d\mu. \quad (37)$$

The λ and $(1-\lambda)$ weighted linear combination of (33) and (37), respectively, yields,

$$\text{sgn}(\rho) \int p_\lambda \cdot (r')^{-\rho} d\mu \geq \frac{\text{sgn}(\rho)}{\|q\|} \int p_\lambda \cdot (q')^{-\rho} d\mu \cdot \int q \cdot (r')^{-\rho} d\mu,$$

i.e.,

$$\begin{aligned}
\mathcal{I}_\alpha(P_\lambda, R) &\geq \mathcal{I}_\alpha(P_\lambda, Q) + \mathcal{I}_\alpha(Q, R) \\
&\geq \mathcal{I}_\alpha(Q, R).
\end{aligned}$$

This completes the proof of (a).

(b) The “if” part is a trivial consequence of (a). We proceed to prove the “only if” part.

The finiteness of $\mathcal{I}_\alpha(Q, R)$ implies that $\mathcal{I}_\alpha(P, R)$ and $\mathcal{I}_\alpha(S, R)$ are also finite. Indeed, from (31), it is clear that $p \leq \lambda^{-1}q$ and thus $p/r \leq \lambda^{-1}q/r$. As a consequence, we have

$$\begin{aligned}
\left(\frac{p'}{r'} \right)^{\frac{1}{\alpha}} &= \frac{p}{r} \cdot \frac{\|r\|}{\|p\|} \\
&\leq \lambda^{-1} \frac{q}{r} \cdot \frac{\|r\|}{\|p\|} \\
&= \lambda^{-1} \left(\frac{q'}{r'} \right)^{\frac{1}{\alpha}} \cdot \frac{\|q\|}{\|p\|}.
\end{aligned}$$

Integrating with respect to R' , we get

$$\int \left(\frac{p'}{r'} \right)^{\frac{1}{\alpha}} dR' \leq \lambda^{-1} \frac{\|q\|}{\|p\|} \cdot \int \left(\frac{q'}{r'} \right)^{\frac{1}{\alpha}} dR' < \infty.$$




Fig. 4: Subspace-transitivity

From (4), we have

$$\mathcal{I}_\alpha(P, Q) = \frac{1}{\rho} \cdot \log \left[\text{sgn}(\rho) \int \left(\frac{p'}{q'} \right)^{\frac{1}{\alpha}} dQ' - 1 \right].$$

Hence $\mathcal{I}_\alpha(P, R) \leq \mathcal{I}_\alpha(Q, R) + c$ for some constant c , and therefore $\mathcal{I}_\alpha(P, R)$ is finite. Similarly $\mathcal{I}_\alpha(S, R)$ is also finite.

Applying the first part of the theorem, we get

$$\begin{aligned} \mathcal{I}_\alpha(P, R) &\geq \mathcal{I}_\alpha(P, Q) + \mathcal{I}_\alpha(Q, R) \\ \mathcal{I}_\alpha(S, R) &\geq \mathcal{I}_\alpha(S, Q) + \mathcal{I}_\alpha(Q, R). \end{aligned}$$

The first inequality is the same as (33) while the second inequality is the same as (33) with s , the density of S , in place of p . Suppose one of these were a strict inequality. Then the λ and $(1 - \lambda)$ weighted linear combination of these two inequalities, along with $Q = \lambda P + (1 - \lambda)S$, yields (37) with a strict inequality, which is the same as (36) with a strict inequality, a contradiction. So the two inequalities must be equalities. This proves the “only if” part and completes the proof of (b). ■

Once Theorem 9 is established for general measure spaces, the proofs of the following results are exactly as in [15]. We provide them for the benefit of the reader and for ease of reference. Let us first recall that any $Q \in \mathbb{E}$ is said to be an *algebraic inner point* of \mathbb{E} if for every $P \in \mathbb{E}$ there exists $S \in \mathbb{E}$ and $0 < t < 1$ such that $Q = tP + (1 - t)S$.

Theorem 10: The following statements hold.

- (a) (*Projection and the Pythagorean property*): A probability measure $Q \in \mathbb{E} \cap B(R, \infty)$ is a forward \mathcal{I}_α -projection of R on the convex set \mathbb{E} of probability measures if and only if every $P \in \mathbb{E} \cap B(R, \infty)$ satisfies (30). If the forward \mathcal{I}_α -projection is an algebraic inner point of \mathbb{E} then $\mathbb{E} \subset B(R, \infty)$ and (32) holds for every $P \in \mathbb{E}$.
- (b) (*Subspace-transitivity*): Let \mathbb{E} and $\mathbb{E}_1 \subset \mathbb{E}$ be convex sets of probability measures. Let R have the forward \mathcal{I}_α -projection Q on \mathbb{E} and the forward \mathcal{I}_α -projection Q_1 on \mathbb{E}_1 , and suppose that (32) holds for every $P \in \mathbb{E}$. Then Q_1 is the forward \mathcal{I}_α -projection of Q on \mathbb{E}_1 . (See figure 4).

Remark 10: Th. 10 (a) essentially says that the forward projection is the unique point in \mathbb{E} that has the Pythagorean property. The importance of Th. 10 (b) will be made clear in section VI.

Proof: (a) Consider the first part of the statement. The “if” part is trivial from the nonnegativity of \mathcal{I}_α . The “only if” part easily follows from Theorem 9-(a). Indeed, $Q \in \mathbb{E} \cap B(R, \infty)$ is the forward \mathcal{I}_α -projection of R implies that for every $P \in \mathbb{E}$, we have $\mathcal{I}_\alpha(P, R) \geq \mathcal{I}_\alpha(Q, R)$ where $P_\lambda = \lambda P + (1 - \lambda)Q$. Hence by Theorem 9-(a), (30) holds.

If the forward \mathcal{I}_α -projection Q is an algebraic inner point of \mathbb{E} then for every $P \in \mathbb{E}$, there exists $S \in \mathbb{E}$ and $\lambda \in (0, 1)$ such that $Q = \lambda P + (1 - \lambda)S$. Hence by Theorem 9-(b), (32) holds.

(b) Applying Theorem 10-(a) to \mathbb{E}_1 , we have for every $P \in \mathbb{E}_1$

$$\begin{aligned} \mathcal{I}_\alpha(P, R) &\geq \mathcal{I}_\alpha(P, Q_1) + \mathcal{I}_\alpha(Q_1, R) \\ &= \mathcal{I}_\alpha(P, Q_1) + (\mathcal{I}_\alpha(Q_1, Q) + \mathcal{I}_\alpha(Q, R)), \end{aligned}$$

where the second equality follows from the equality hypothesis that (32) holds. Using this same equality hypothesis, we also have

$$\mathcal{I}_\alpha(P, R) = \mathcal{I}_\alpha(P, Q) + \mathcal{I}_\alpha(Q, R).$$

Thus

$$\mathcal{I}_\alpha(P, Q) \geq \mathcal{I}_\alpha(P, Q_1) + \mathcal{I}_\alpha(Q_1, Q)$$

for every $P \in \mathbb{E}_1$. Applying Theorem 10-(a) once again, we conclude that Q_1 is the forward \mathcal{I}_α -projection of Q on \mathbb{E}_1 . ■

Theorem 10-(a) yields a simple proof of the uniqueness of projection on a convex \mathbb{E} , if the projection exists. Indeed, let Q_1 and Q_2 be two projections of a probability measure R on a convex \mathbb{E} . Then $\mathcal{I}_\alpha(Q_1, R) = \mathcal{I}_\alpha(Q_2, R) < \infty$. By Theorem 10-(a),

$$\mathcal{I}_\alpha(Q_2, R) \geq \mathcal{I}_\alpha(Q_2, Q_1) + \mathcal{I}_\alpha(Q_1, R).$$

Canceling $\mathcal{I}_\alpha(Q_2, R)$ and $\mathcal{I}_\alpha(Q_1, R)$, we get $\mathcal{I}_\alpha(Q_2, Q_1) = 0$ which further implies $Q_1 = Q_2$.

V. EXAMPLE: FORWARD \mathcal{I}_α -PROJECTION FOR A LINEAR FAMILY

In this section we provide an explicit characterization of the forward \mathcal{I}_α -projection on a linear family.

Let Γ be an arbitrary index set and let $f_\gamma: \mathbb{X} \rightarrow \mathbb{R}$, for $\gamma \in \Gamma$, be measurable functions. The family of probability measures defined by

$$\mathbb{L} = \left\{ P: \int f_\gamma dP = 0, \gamma \in \Gamma \right\}, \quad (38)$$

if nonempty, is called a *linear family*⁶.

Our next result is that the forward \mathcal{I}_α -projection on a linear family is a member of an associated α -power-law family⁷ just as forward \mathcal{I} -projection on a linear family is a member of an associated exponential family [7, Th. 3.1]. The proof for $\alpha < 1$ is similar with only minor changes. The proof for $\alpha > 1$ involves some additional conditions. We will explore the geometric relationship between the linear family and the α -power-law family in a companion paper [27].

Theorem 11: Let $\alpha > 0$ and $\alpha \neq 1$. Let \mathbb{L} be a linear family of probability measures as in (38). Let R have μ -density r .

(a) If Q is the forward \mathcal{I}_α -projection of R on \mathbb{L} then the μ -density q of Q satisfies

$$q(x)^{\alpha-1} = c \cdot r(x)^{\alpha-1} + g(x), \quad \forall x \notin N \quad (39)$$

$$q(x) = 0, \quad \forall x \in N, \quad (40)$$

where $N \subseteq \mathbb{X}$ is such that, for every $P \in \mathbb{L} \cap B(R, \infty)$,

$$\begin{cases} P(N) = 0, & \text{if } \alpha < 1 \\ c \int_N r^{\alpha-1} dP \leq \int_{\mathbb{X} \setminus N} g dP, & \text{if } \alpha > 1, \end{cases} \quad (41)$$

$$c = \frac{\int q^\alpha d\mu}{\int q r^{\alpha-1} d\mu}, \quad (42)$$

and g belongs to the $L^1(Q)$ -closure of the linear space spanned by $\{f_\gamma\}_{\gamma \in \Gamma}$.

(b) Conversely, if there is a $Q \in \mathbb{L}$ whose μ -density satisfies (39)-(41) for some scalar c and some g in the linear span of $\{f_\gamma\}_{\gamma \in \Gamma}$, then Q is the forward \mathcal{I}_α -projection of R on \mathbb{L} , (30) holds for every $P \in \mathbb{L} \cap B(R, \infty)$, and further, (30) holds with equality when $\alpha < 1$.

Proof: (a) Let Q be the forward \mathcal{I}_α -projection of R on \mathbb{L} with μ -density q . Let $N = \{x \in \mathbb{X}: q(x) = 0\}$. By definition of the forward \mathcal{I}_α -projection, we have $\mathcal{I}_\alpha(Q, R) < \infty$. When $\alpha < 1$, if $P \in \mathbb{L} \cap B(R, \infty)$, then Theorem 10-(a) implies (30), which further implies $\mathcal{I}_\alpha(P, Q) < \infty$, $P \ll Q$, and thus $P(N) = 0$. We will soon define g on $\mathbb{X} \setminus N$ and will show the inequality in (41) for $\alpha > 1$ later in this proof.

From $\mathcal{I}_\alpha(Q, R) < \infty$, using (6), it is also easy to verify that $0 < \int q r^{\alpha-1} d\mu < \infty$. Define

$$\mathbb{L}_1 := \{P \in \mathbb{L}: p(x) \leq 2q(x) \text{ a.e.}[\mu]\}.$$

Obviously, \mathbb{L}_1 is convex and $Q \in \mathbb{L}_1$. For any $P \in \mathbb{L}_1$, define P_1 to have the density $p_1(x) = 2q(x) - p(x)$. We then have $P_1 \in \mathbb{L}_1$ and $Q = \frac{P+P_1}{2}$. Hence Q is an algebraic inner point of \mathbb{L}_1 . By Theorem 10-(a), (30) holds with equality for all $P \in \mathbb{L}_1$. This equality can be simplified, based on (5), to

$$\int p r^{\alpha-1} d\mu = \int p q^{\alpha-1} d\mu \cdot \frac{\int q r^{\alpha-1} d\mu}{\int q^\alpha d\mu} \quad (43)$$

$$= c^{-1} \int p q^{\alpha-1} d\mu, \quad (44)$$

where c is given by (42). This can be rewritten as

$$\int p \cdot (q^{\alpha-1} - c r^{\alpha-1}) d\mu = 0 \quad \forall P \in \mathbb{L}_1, \quad (45)$$

which with $g(x) := q(x)^{\alpha-1} - c r(x)^{\alpha-1}$, $x \in \mathbb{X} \setminus N$, is the same as

$$\int p g d\mu = 0 \quad \forall P \in \mathbb{L}_1. \quad (46)$$

⁶Let us reiterate the standing assumptions: $P \ll \mu$ and the μ -density $p \in L^\alpha(\mu)$ for every $P \in \mathbb{L}$.

⁷A parametric family of probability distributions that are of the form (39).

We have left g undefined for x with $q(x) = 0$, but this is inconsequential because we now show g belongs to the $L^1(Q)$ -closure of the linear span of $\{f_\gamma\}_{\gamma \in \Gamma}$.

From (46), we get

$$\int g \cdot \frac{dP}{dQ} \cdot dQ = 0 \quad \forall P \in \mathbb{L}_1, \quad (47)$$

and by setting $P = Q$ in (47) we get

$$\int g dQ = 0. \quad (48)$$

Combining (47) and (48) yields

$$\int g \left(\frac{dP}{dQ} - 1 \right) dQ = 0 \quad \forall P \in \mathbb{L}_1. \quad (49)$$

If $h: \mathbb{X} \rightarrow \mathbb{R}$ is a measurable function such that $|h| \leq 1$, a.e. $[Q]$, and further

$$\int h dQ = 0, \text{ and } \int f_\gamma h dQ = 0 \text{ for every } \gamma \in \Gamma, \quad (50)$$

then P defined according to $dP = (h + 1) dQ$ belongs to \mathbb{L}_1 , and from (49), it follows that

$$\int gh dQ = 0. \quad (51)$$

It immediately follows after scaling that if $h \in L^\infty(Q)$, the dual of $L^1(Q)$, and (50) holds, then (51) must also hold. In other words, any continuous linear functional $F_h: L^1(Q) \rightarrow \mathbb{R}$ given by $F_h(f) = \int fh dQ$ that vanishes on the linear subspace spanned by 1 and the f_γ 's also vanishes at $f = g$. By the Hahn-Banach theorem [32, Th. 5.8.a], g is in the $L^1(Q)$ -closure of that linear subspace. From (48), it follows that g is in the $L^1(Q)$ -closure of the subspace spanned by the f_γ 's alone.

We now show the inequality in (41) for $\alpha > 1$. For any $P \in \mathbb{L} \cap B(R, \infty)$, where such a P may be outside \mathbb{L}_1 , let us observe that

$$0 \leq \int_{\mathbb{X}} pq^{\alpha-1} d\mu - c \int_{\mathbb{X}} pr^{\alpha-1} d\mu \quad (52)$$

$$= \int_{\mathbb{X} \setminus N} pq^{\alpha-1} d\mu - c \int_{\mathbb{X}} pr^{\alpha-1} d\mu \quad (53)$$

$$= \int_{\mathbb{X} \setminus N} p \cdot (cr^{\alpha-1} + g) d\mu - c \int_{\mathbb{X}} pr^{\alpha-1} d\mu \quad (54)$$

$$= \int_{\mathbb{X} \setminus N} pg d\mu - c \int_N pr^{\alpha-1} d\mu, \quad (55)$$

where (52) follows from the combination of (6), (30), and (42); consequently, (53) follows from the fact that $q(x) = 0$ for $x \in N$, (54) follows from the definition of $g(x)$ on the set $x \in \mathbb{X} \setminus N$, and (55) follows from the cancellation of a portion of the last integral term on the right-hand side of (54). Inequality (41) for $\alpha > 1$ follows from (55). This completes the proof of (a).

(b) Let $Q \in \mathbb{L}$ have μ -density q which satisfies (39)-(41) where c is some scalar and g is a linear combination of the f_γ 's; so $\int g dP = 0$ for all $P \in \mathbb{L}$. Integrating (39)-(40) with respect to Q and using $\int g dQ = 0$, we get

$$\int q^\alpha d\mu = c \int qr^{\alpha-1} d\mu$$

from which the following are clear:

- $0 < \int qr^{\alpha-1} d\mu < \infty$, and so $\mathcal{I}_\alpha(Q, R) < \infty$;
- $c > 0$ and satisfies (42).

Fix any $P \in \mathbb{L}$ with $\mathcal{I}_\alpha(P, R) < \infty$. As claimed at the beginning of the proof of part (a), we then have $0 < \int pr^{\alpha-1} d\mu < \infty$. Integrating (39)-(40) with respect to P , we now get

$$\int pq^{\alpha-1} d\mu \geq c \int pr^{\alpha-1} d\mu,$$

where

- equality holds when $\alpha < 1$ because of the assumption $P(N) = 0$,
- inequality holds when $\alpha > 1$ because of the inequality assumption in (41); indeed, this assumption is the same as saying that the right-hand side of (55) is ≥ 0 , and one proceeds in the reverse direction in that sequence of equalities to obtain the inequality (52) which is the same as the above inequality.

Since c satisfies (42), we have that (30) holds (with equality when $\alpha < 1$). By Theorem 10-(a) (in the “if” direction) Q is the forward \mathcal{I}_α -projection of R on \mathbb{L} . ■

Remark 11: As in the case of relative entropy ($\alpha = 1$), in Theorem 11-(a), it is possible that the inequality in (30) is strict for some P in the linear family, and in Csiszár’s words [7, p.152], “neither the necessary nor the sufficient condition of Theorem 11 is both necessary and sufficient, in general.” Csiszár’s counterexamples in [7, pp.152-153], but with $q^{\alpha-1} = c \cdot r^{\alpha-1} + g$ instead of $q = c \cdot r \cdot \exp\{g\}$, continue to serve as counterexamples for our parametric setting (see Appendix E).

However, under an additional assumption, Theorem 11 can be leveraged to provide a necessary and sufficient condition for a $Q \in \mathbb{L}$ to be the forward \mathcal{I}_α -projection.

Corollary 12: Let $\alpha > 0$ and $\alpha \neq 1$. Let \mathbb{L} be the linear family as defined in (38). Suppose that the linear space spanned by $\{f_\gamma\}_{\gamma \in \Gamma}$ is $L^1(P)$ -closed for every $P \in \mathbb{L}$. Consider a $Q \in \mathbb{L}$. Q is the forward \mathcal{I}_α -projection of R on \mathbb{L} if and only if the μ -density q of Q satisfies (39)-(41) for some scalar c and some g in the span of $\{f_\gamma\}_{\gamma \in \Gamma}$. Moreover, the inequality in (41) for $\alpha > 1$ is equivalent to

$$\int_N (cr^{\alpha-1} + g) dP \leq 0, \quad \alpha > 1. \quad (56)$$

Proof: The forward direction is immediate from the forward direction of Theorem 11 and the hypothesis that the linear space spanned by $\{f_\gamma\}_{\gamma \in \Gamma}$ is $L^1(Q)$ -closed; so g is in the span of $\{f_\gamma\}_{\gamma \in \Gamma}$. The reverse direction is the same as the reverse direction in Theorem 11.

To prove (56), let us observe that because g is in the span of $\{f_\gamma\}_{\gamma \in \Gamma}$, it is well-defined for all $x \in \mathbb{X}$ and consequently satisfies $\int g dP = 0$ for every $P \in \mathbb{L}$. Adding $\int_N g dP$ to both sides of (41) and using $\int g dP = 0$, we get (56). ■

One example where the linear space spanned by $\{f_\gamma\}_{\gamma \in \Gamma}$ is $L^1(P)$ -closed for every $P \in \mathbb{L}$ is when Γ is finite, i.e., $\Gamma = \{1, 2, \dots, k\}$ for some finite k . If Q is the forward \mathcal{I}_α -projection of R on \mathbb{L} , then the expression

$$q(x)^{\alpha-1} = c \cdot r(x)^{\alpha-1} + \sum_{\gamma=1}^k \theta_\gamma f_\gamma(x),$$

where $(\theta_1, \dots, \theta_k) \in \mathbb{R}^k$, holds for all x with $q(x) > 0$. Moreover, (30) holds for all $P \in \mathbb{L} \cap B(R, \infty)$, and it holds with equality when $\alpha < 1$.

For relative entropy, $\alpha = 1$, Csiszár provides another example: the family of probability measures on a product space $\mathbb{X} = \mathbb{X}_1 \times \mathbb{X}_2$ with the associated product σ -algebra, having specified marginals. We leave the question of whether Corollary 12 is applicable or not to this setting as an open question.

Even though Corollary 12 characterizes the forward \mathcal{I}_α -projection to some extent, existence of the projection is not assured, and one appeals to Theorem 8 or other means to guarantee existence. Let us note in passing two instances when the crucial hypothesis of Theorem 8, that the set of μ -densities is $L^\alpha(\mu)$ -closed, holds.

- (a) If $\alpha > 1$, $\mu(\mathbb{X}) < +\infty$, and $f_\gamma \in L^\infty(\mu)$ for $\gamma = 1, \dots, k$, then a simple application of Lyapunov’s inequality⁸ and the dominated convergence theorem suffices to show that \mathcal{L} , the set of μ -densities of probability measures in \mathbb{L} , is $L^\alpha(\mu)$ -closed.
- (b) If \mathbb{X} is finite, point-wise convergence suffices to establish that \mathcal{L} is $L^\alpha(\mu)$ -closed.

Let us now exploit the understanding we have gained to generalize Lemma 2-e) on Rényi entropy maximizers.

Corollary 13: Let $\alpha > 0$ and $\alpha \neq 1$. Let \mathbb{L} be the linear family as defined in (38). If \mathbb{L} has a member Q whose μ -density q satisfies (39)-(41) for some scalar c , some g in the span of $\{f_\gamma\}_{\gamma \in \Gamma}$, and with $r(x) \equiv 1$, then

$$\mathcal{I}_\alpha(P, Q) \leq H_\alpha(Q) - H_\alpha(P) \quad \forall P \in \mathbb{L}, \quad (57)$$

with equality when $\alpha < 1$. Furthermore, Q is the Rényi entropy maximizer in \mathbb{L} .

Proof: It suffices to prove (57). The second statement immediately follows.

Using (6), (1), and after a simple rearrangement, we get

$$\mathcal{I}_\alpha(P, Q) = \frac{\alpha}{1-\alpha} \left[\log \int pq^{\alpha-1} d\mu - \log \int q^\alpha d\mu \right] + H_\alpha(Q) - H_\alpha(P).$$

⁸Lyapunov’s moment inequality states that if $\mu(\mathbb{X}) < \infty$, $0 < r < s \leq \infty$, and $u \in L^s(\mu)$, then $\|u\|_r \leq \|u\|_s \mu(\mathbb{X})^{(1/r)-(1/s)}$, and consequently $L^s(\mu) \subseteq L^r(\mu)$.

Let us note from (39)-(40) that $\int q^\alpha d\mu = c$. So (57) will hold if we can establish

$$\begin{aligned} \int pq^{\alpha-1} d\mu &= c \quad \forall P \in \mathbb{L}, \quad \text{if } \alpha < 1, \\ \int pq^{\alpha-1} d\mu &\geq c \quad \forall P \in \mathbb{L}, \quad \text{if } \alpha > 1. \end{aligned}$$

Both of these are obvious from the hypotheses of the corollary via (39)-(41), the assumption that $r(x) \equiv 1$, and the fact that P and Q are both probability measures belonging to \mathbb{L} . \blacksquare

Remark 12: When $0 < \mu(\mathbb{X}) < \infty$, with $r(x) \equiv 1$, define the probability measure \tilde{R} with μ -density

$$\tilde{r}(x) := \frac{r(x)}{\mu(\mathbb{X})} \equiv \frac{1}{\mu(\mathbb{X})}.$$

We then have from (6) that $\mathcal{S}_\alpha(P, \tilde{R}) = H_\alpha(\tilde{R}) - H_\alpha(P) = \log \mu(\mathbb{X}) - H_\alpha(P)$, and so the Rényi entropy maximizer on \mathbb{L} is the forward \mathcal{S}_α -projection of \tilde{R} on \mathbb{L} . From (5), it is clear that scale factors are irrelevant, and if we allow the second argument of \mathcal{S}_α to be positive measures, not just probability measures, then the Rényi entropy maximizer on \mathbb{L} can be interpreted as the “forward \mathcal{S}_α -projection of μ on \mathbb{L} ”. When $\mu(\mathbb{X})$ is not finite, there is no probability measure on \mathbb{X} with the uniform μ -density. Nevertheless, Corollary 13 shows that the Rényi entropy maximizer is the “forward \mathcal{S}_α -projection of μ on \mathbb{L} ”.

Remark 13: Student-t and Student-r distributions are maximizers of Rényi entropy under a covariance constraint [20]. Since a Student-r distribution has a compact support, it can be shown to be the forward \mathcal{S}_α -projection of the uniform distribution as described above, when $\alpha > 1$. The support of a Student-t distribution is the whole of \mathbb{R}^d . However, it can also be seen as a limit of forward \mathcal{S}_α -projections of uniform distributions on an increasing sequence of compact subsets of \mathbb{R}^d , when $\alpha < 1$.

VI. TRANSITIVITY AND ITERATED PROJECTIONS FOR A LINEAR FAMILY

In this section we assume \mathbb{X} is finite. Let $\mathcal{P}(\mathbb{X})$ be the space of all probability measures on \mathbb{X} . In a remarkable paper [24] on an axiomatic approach to inference, Csiszár explored some natural axioms for *selection* and *projection rules*, and their consequences on linear families.

A *projection rule* is a mapping that (in our context) takes a probability measure R and a linear family \mathbb{L} and maps them to a probability measure $\Pi(\mathbb{L}|R)$ in \mathbb{L} , such that if $R \in \mathbb{L}$ then $\Pi(\mathbb{L}|R) = R$. $\Pi(\mathbb{L}|R)$ is then called the projection of R on \mathbb{L} . A projection rule is said to be generated by a function $F(P|R)$, $P \in \mathcal{P}(\mathbb{X})$, $R \in \mathcal{P}(\mathbb{X})$, if for each R , $\Pi(\mathbb{L}|R)$ is the unique element of \mathbb{L} where $F(P|R)$ is minimized subject to $P \in \mathbb{L}$. A projection rule may be interpreted as follows: a “prior guess” R is updated to $\Pi(\mathbb{L}|R)$ upon information that the “feasible set” is \mathbb{L} .

Clearly, the forward \mathcal{S}_α -projection of R on a linear family \mathbb{L} is an example of a projection rule that is generated by the function $F(P|R) = \mathcal{S}_\alpha(P, R)$. Csiszár [24, Th. 1] showed that any *regular* and *local* projection rule, see [24, Def. 2-3] for the definitions, is generated by a separable function $F(P|R) = \sum_{x \in \mathbb{X}} \phi_x(P(x)|R(x))$, for some component functions $\phi_x(\cdot|\cdot)$, $x \in \mathbb{X}$, with the value 0 at $P = R$.

Another desired property of a projection rule is subspace-transitivity ([24, Def. 6]). A projection rule is *subspace-transitive* if for any $\mathbb{L}' \subset \mathbb{L}$, both of which are linear families, and any probability measure R , we have

$$\Pi(\mathbb{L}'|R) = \Pi(\mathbb{L}'|\Pi(\mathbb{L}|R)).$$

This can be interpreted as follows: if a “prior guess” R is updated to $\Pi(\mathbb{L}|R)$ upon information that the “feasible set” is \mathbb{L} , and further information restricts the possibilities to a smaller feasible set \mathbb{L}' , then updating the “current guess” $\Pi(\mathbb{L}|R)$ on the basis of all available information yields the same outcome as updating the “prior guess” R directly on the basis of all available information. Csiszár showed [24, Th. 3] that any regular, local, and subspace-transitive projection rule is generated by Bregman’s divergence of the sum-form, i.e.,

$$F(P|R) = \Phi(P) - \Phi(R) - \langle \text{grad } \Phi(R), P - R \rangle,$$

where $\Phi(P) = \sum_x \varphi_x(P(x))$. Squared Euclidean distance and relative entropy \mathcal{S}_1 are examples of such divergences.

\mathcal{S}_α is, in general, neither of the sum-form nor a Bregman’s divergence. Yet when $\alpha < 1$, the projection rule generated by $\mathcal{S}_\alpha(P, R)$ is subspace-transitive. The property fails in general when $\alpha > 1$, but holds even in this case in the special circumstance when the projection is an algebraic inner point. The main goal of this section is to establish subspace transitivity. This suggests that if one is willing to forgo the locality axiom of a projection rule, then there is at least one other family of projection rules, those generated by \mathcal{S}_α , that are regular and subspace-transitive.

To formalize the result, we begin with two simple propositions. For a probability measure P write $\text{Supp}(P)$ for the set of x where $P(x) > 0$. For a family of probability measures \mathbb{E} , write $\text{Supp}(\mathbb{E})$ for the union of the supports of all probability measures in \mathbb{E} . We then have the following.

Proposition 14: Let $\alpha < 1$. Let Q be the forward \mathcal{J}_α -projection of R on \mathbb{E} . If \mathbb{E} is convex, then $\text{Supp}(Q) = \text{Supp}(\mathbb{E}) \cap \text{Supp}(R)$.

Proof: We may restrict attention to those $P \in \mathbb{E}$ such that $P \ll R$. For such a P , let $P_t = (1-t)Q + tP$, $0 \leq t \leq 1$. Since \mathbb{E} is convex, $P \in \mathbb{E}$ implies that $P_t \in \mathbb{E}$. By the mean value theorem, for each $t \in (0, 1)$, there exists $\tilde{t} \in (0, t)$ such that

$$0 \leq \frac{1}{t} \left[\mathcal{J}_\alpha(P_t, R) - \mathcal{J}_\alpha(Q, R) \right] = \frac{d}{ds} \mathcal{J}_\alpha(P_s, R) \Big|_{s=\tilde{t}}. \quad (58)$$

The first inequality follows from the fact that Q is the projection. Using (8), we see that

$$\frac{d}{ds} \mathcal{J}_\alpha(P_s, R) = \frac{\alpha}{1-\alpha} \left[\frac{\sum_x (P(x) - Q(x)) R(x)^{\alpha-1}}{\sum_x P_s(x) R(x)^{\alpha-1}} - \frac{\sum_x (P(x) - Q(x)) P_s(x)^{\alpha-1}}{\sum_x P_s(x)^\alpha} \right]. \quad (59)$$

Suppose $Q(x) = 0$ for an $x \in \text{Supp}(P)$. Then $\alpha < 1$ implies that right-hand side of (59) goes to $-\infty$ as $t \downarrow 0$, which contradicts the nonnegativity requirement in (58). Hence $\text{Supp}(P) \subset \text{Supp}(Q)$ for every $P \in \mathbb{L}$. Also, since Q is the \mathcal{J}_α -projection of R , $\mathcal{J}_\alpha(Q, R) < \infty$, and as a consequence, $\text{Supp}(Q) \subset \text{Supp}(R)$. This establishes the proposition. ■

Consider now the linear family of probability measures on \mathbb{X} given by

$$\mathbb{L} = \left\{ P: \sum_x P(x) f_\gamma(x) = 0, \text{ for all } \gamma = 1, \dots, k \right\}. \quad (60)$$

Since \mathbb{X} is finite, we already saw at the end of the previous section that \mathbb{L} is closed in $L^\alpha(\mu)$, with μ being the counting measure. By Theorem 8, any probability measure R with $\mathcal{J}_\alpha(P, R) < \infty$ for some $P \in \mathbb{L}$ has a forward \mathcal{J}_α -projection on \mathbb{L} . Moreover, we have the following.

Proposition 15: Let $\alpha < 1$. Let R have full support. Let \mathbb{L} be as in (60) and let Q be the forward \mathcal{J}_α -projection of R on \mathbb{L} . Then Q is an algebraic inner point of \mathbb{L} .

Proof: By Proposition 14, $\text{Supp}(Q) = \text{Supp}(\mathbb{L})$. Hence for every $P \in \mathbb{L}$, one can find $t < 0$ such that $P_t = (1-t)Q + tP \in \mathbb{L}$. This implies that

$$Q = \frac{1}{1-t} P_t - \frac{t}{1-t} P,$$

and hence Q is an algebraic inner point of \mathbb{L} . ■

We are now ready to state the main result of this section.

Theorem 16 (Subspace-transitivity): Let $\mathbb{L}_1 \subset \mathbb{L}$ be two linear families of probability measures. Let R be a probability measure with full support. Let R have the forward \mathcal{J}_α -projection Q on \mathbb{L} and the forward \mathcal{J}_α -projection Q_1 on \mathbb{L}_1 . If either (a) $\alpha < 1$ or (b) $\alpha > 1$ and Q is an algebraic inner point of \mathbb{L} , then Q_1 is the forward \mathcal{J}_α -projection of Q on \mathbb{L}_1 .

Proof: If $\alpha < 1$, then by Proposition 15, Q is an algebraic inner point of \mathbb{L} . If $\alpha > 1$, by assumption (b), Q is an algebraic inner point of \mathbb{L} . Apply Theorem 10-(a) to get that (32) holds for all $P \in \mathbb{L}$. Now apply Theorem 10-(b) to conclude that subspace-transitivity holds. ■

Remark 14: As can be observed from Theorem 10-(b), and from the proof above, subspace-transitivity follows whenever there is equality in (32). What is special about linear spaces under $\alpha < 1$ is that this equality comes for free, thanks to Proposition 15.

Example 1: The following example shows that subspace-transitivity for the \mathcal{J}_α -projection rule need not hold when $\alpha > 1$. Take $\alpha = 2$ and $\mathbb{X} = \{1, 2, 3, 4\}$. Take $R = (1/4, 1/4, 1/4, 1/4)$. Consider the two linear families on the probability simplex in \mathbb{R}^4 ,

$$\begin{aligned} \mathbb{L} &= \{P \in \mathcal{P}(\mathbb{X}): 8p_1 + 4p_2 + 2p_3 + p_4 = 7\}, \\ \mathbb{L}_1 &= \{P \in \mathcal{P}(\mathbb{X}): 8p_1 + 4p_2 + 2p_3 + p_4 = 7; p_2 = 1/8\}. \end{aligned}$$

Thus

$$\begin{aligned} \mathbb{L} &= \left\{ P \in \mathcal{P}(\mathbb{X}): \sum_x P(x) f_1(x) = 0 \right\}, \\ \mathbb{L}_1 &= \left\{ P \in \mathcal{P}(\mathbb{X}): \sum_x P(x) f_i(x) = 0, i = 1, 2 \right\}, \end{aligned}$$

where $f_1(\cdot) = (1, -3, -5, -6)$ and $f_2(\cdot) = (-1/8, 7/8, -1/8, -1/8)$.

We claim that the forward \mathcal{I}_α -projection of R on \mathbb{L} is $Q = (3/4, 1/4, 0, 0)$. To check this claim, first note that $Q \in \mathbb{L}$. Also, with $c = 5/2$ and $\theta_1 = 1/8$, we can check that

$$\begin{aligned} 0 &< Q(x) = cR(x) + \theta_1 f_1(x), \quad x = 1, 2, \\ 0 &= Q(3) = cR(3) + \theta_1 f_1(3), \\ 0 &= Q(4) > cR(4) + \theta_1 f_1(4). \end{aligned}$$

One can then easily verify that this Q satisfies (56) (which is equivalent to (41) with $\alpha > 1$) for every $P \in \mathbb{L}$. Hence, by Corollary 12, Q is the forward \mathcal{I}_α -projection of R on \mathbb{L} .

Similarly one can show that the forward \mathcal{I}_α -projection of R on \mathbb{L}_1 is $Q_1 = (19/24, 1/8, 1/12, 0)$. Indeed, with $\theta_1 = 17/144$, $\theta_2 = -7/36$ and $c = 187/72$, we have

$$\begin{aligned} 0 &< Q_1(x) = cR(x) + \theta_1 f_1(x) + \theta_2 f_2(x), \quad x = 1, 2, 3, \\ 0 &= Q_1(4) > cR(4) + \theta_1 f_1(4) + \theta_2 f_2(4). \end{aligned}$$

Again, Q_1 satisfies (56) for every $P \in \mathbb{L}_1$ and, by Corollary 12, must be the forward \mathcal{I}_α -projection of R on \mathbb{L}_1 .

Numerical calculations show that $(0.798, 0.125, 0.038, 0.039)$ is in \mathbb{L}_1 and

$$\mathcal{I}_\alpha((0.798, 0.125, 0.038, 0.039), Q) = 0.0323 < \mathcal{I}_\alpha(Q_1, Q) = 0.0382.$$

If \tilde{Q}_1 is the forward \mathcal{I}_α -projection of Q on \mathbb{L}_1 , it must satisfy $\mathcal{I}_\alpha(\tilde{Q}_1, Q) \leq 0.0323$, which Q_1 does not. Thus, the transitive projection of R on \mathbb{L}_1 via Q is different from Q_1 .

The next theorem provides an iterative way of finding the forward \mathcal{I}_α -projection for $\alpha < 1$ when the set \mathbb{L} is an intersection of several linear families. A similar result is known for relative entropy ($\alpha = 1$); see [7, Th. 3.2].

Theorem 17 (Iterated projections): Let $\alpha < 1$. Suppose that $\mathbb{L}_0, \dots, \mathbb{L}_{m-1}$ are linear families of probability measures on a finite set \mathbb{X} and that $\mathbb{L} = \bigcap_{i=0}^{m-1} \mathbb{L}_i \neq \emptyset$. Let R be a probability measure on \mathbb{X} with full support. Let Q be the forward \mathcal{I}_α -projection of R on \mathbb{L} . Write $Q_0 = R$ and write Q_n for the forward \mathcal{I}_α -projection of Q_{n-1} on \mathbb{L}_{n-1} , where for $n > m$, $\mathbb{L}_n = \mathbb{L}_i$, $i = n \pmod{m}$. Then $Q_n \rightarrow Q$.

Proof: The proof largely follows Csizsár's proof of [7, Th. 3.2] with the main changes being the use of the generalization of Pinsker's inequality [31, Th. 1] and some care to address convergence of the escort measures. Details follow.

First let us observe that if $\text{Supp}(\mathbb{L}_n) \not\subseteq \text{Supp}(Q_{n-1})$, in order to find the projection of Q_{n-1} on \mathbb{L}_n , one may restrict attention to members $P \in \mathbb{L}_n$ with $\text{Supp}(P) \subset \text{Supp}(Q_{n-1})$. If not, $\mathcal{I}_\alpha(P, Q_{n-1}) = \infty$. With this restricted \mathbb{L}_n , by Proposition 15, Q_n is an algebraic inner point of the restricted \mathbb{L}_n . Henceforth we call these simply \mathbb{L}_n and denote their intersection by \mathbb{L} .

Fix a natural number N . In view of Proposition 15, applying Theorem 10-(a), we see that for any $P \in \mathbb{L}$ we have

$$\mathcal{I}_\alpha(P, Q_{n-1}) = \mathcal{I}_\alpha(P, Q_n) + \mathcal{I}_\alpha(Q_n, Q_{n-1}), \quad n = 1, \dots, N. \quad (61)$$

Summing all the N equations, we get

$$\mathcal{I}_\alpha(P, R) = \mathcal{I}_\alpha(P, Q_N) + \sum_{n=1}^N \mathcal{I}_\alpha(Q_n, Q_{n-1}) \quad \forall P \in \mathbb{L}.$$

Now let (Q_{N_k}) be a subsequence of (Q_n) converging to, say, \tilde{Q} . Taking limit as $k \rightarrow \infty$ along this subsequence, we get

$$\mathcal{I}_\alpha(P, R) = \mathcal{I}_\alpha(P, \tilde{Q}) + \sum_{n=1}^{\infty} \mathcal{I}_\alpha(Q_n, Q_{n-1}) \quad \forall P \in \mathbb{L}, \quad (62)$$

which implies that the summation term is finite, and so $\mathcal{I}_\alpha(Q_n, Q_{n-1}) \rightarrow 0$, or $I_f(Q'_n, Q'_{n-1}) \rightarrow 0$, as $n \rightarrow \infty$ in view of (4). Hence, by [31, Th. 1], $|Q'_n - Q'_{n-1}|_{TV} \rightarrow 0$ as $n \rightarrow \infty$. Hence all of the sequences $(Q'_{N_k}), (Q'_{N_k+1}), \dots, (Q'_{N_k+m-1})$ converge to same \tilde{Q}' . Now, for any k , $Q'_{N_k}, Q'_{N_k+1}, \dots, Q'_{N_k+m-1}$, are m consecutive members of the sequence (Q'_n) , and by the periodic construction of the Q_n 's, each is in one of $\mathbb{L}'_0, \dots, \mathbb{L}'_{m-1}$, where $\mathbb{L}'_i = \{P' : P \in \mathbb{L}_i\}$ with P' as in (2). Hence \tilde{Q}' is in each of them which implies $\tilde{Q}' \in \mathbb{L}'$ and $\tilde{Q} \in \mathbb{L}$. Putting $P = \tilde{Q}$ in (62), we get

$$\mathcal{I}_\alpha(\tilde{Q}, R) = \sum_{n=1}^{\infty} \mathcal{I}_\alpha(Q_n, Q_{n-1})$$

for this subsequential limit \tilde{Q} . Substituting this back in (62), we see that

$$\mathcal{I}_\alpha(P, R) = \mathcal{I}_\alpha(P, \tilde{Q}) + \mathcal{I}_\alpha(\tilde{Q}, R) \quad \forall P \in \mathbb{L}.$$

By Theorem 10, \tilde{Q} is the forward \mathcal{I}_α -projection Q of R on \mathbb{L} . By uniqueness of the forward \mathcal{I}_α -projection, every subsequential limit equals Q , and so (Q_n) converges to Q . \blacksquare

Remark 15: Again, the above theorem continues to hold for $\alpha > 1$ under the rather restrictive assumption that each of the forward \mathcal{I}_α -projections satisfies the Pythagorean property (61) with equality.

VII. CONCLUDING REMARKS

We end this paper with some concluding remarks.

- 1) The forward \mathcal{I}_α -projection, in general, depends on the reference measure μ . The dependence on μ however disappears as $\alpha \rightarrow 1$, and in this sense \mathcal{I}_1 -projection or \mathcal{I} -projection is special.
- 2) Throughout this paper, motivated by constraints induced by linear statistics, we restricted \mathbb{E} to be a convex set of probability measures. But it is clear that if p and q are two μ -densities of probability measures, and both belong to $L^\alpha(\mu)$, then, for positive constants c_1 and c_2 , we have $\mathcal{I}_\alpha(c_1 p, c_2 q) = \mathcal{I}_\alpha(p, q)$ because \mathcal{I}_α depends only on the associated escort probability densities of the arguments, and scale factors do not affect these escort densities. It would therefore be interesting to extend our theory of the forward \mathcal{I}_α -projection to general convex and closed subspaces of $L^\alpha(\mu)$.
- 3) The above remark on the insignificance of the scaling factors suggests that perhaps the theory ought to be developed from the view point of escort distributions. However, convexity of \mathbb{E} which is a natural consequence of linear statistics, may be lost in the escort domain.
- 4) Is there a “generalized” forward \mathcal{I}_α -projection Q for a convex \mathbb{E} that is not $L^\alpha(\mu)$ -closed? Further, if (P_n) is a sequence in \mathbb{E} such that $\mathcal{I}_\alpha(P_n, R) \rightarrow \inf_{P \in \mathbb{E}} \mathcal{I}_\alpha(P, R)$ as $n \rightarrow \infty$, does P_n converge to this Q ? A careful examination of the proof of Theorem 8 for the case when $\alpha < 1$ shows that while one can extract a *unique* probability measure Q that satisfies

$$\mathcal{I}_\alpha(Q, R) \leq \lim_{k \rightarrow \infty} \mathcal{I}_\alpha(P_{n_k}, R) = \inf_{P \in \mathbb{E}} \mathcal{I}_\alpha(P, R)$$

for any converging subsequence of densities (p_{n_k}) in $L^\alpha(\mu)$, it is not clear if $p_n \rightarrow q$, the μ -density of Q , in $L^\alpha(\mu)$. However, each subsequential limit is always a scaled version of q . Thus Q can serve as *the* generalized forward \mathcal{I}_α -projection. This too suggests the benefit of a theory modulo scale factors.

- 5) In Section V, we considered projection on linear families. Let us highlight an open question raised in that section. Is Corollary 12 applicable to a family of distributions on a product space with specified marginals? While the answer is true for $\alpha = 1$ ([7, Cor. 3.2]), we have not been able to address the general case of $\alpha > 0, \alpha \neq 1$.
- 6) Suppose that we have a nested sequence $\mathbb{L}_1 \supset \mathbb{L}_2 \supset \dots$ of convex sets of probability measures absolutely continuous with respect to a common σ -finite measure μ such that the respective set of densities \mathcal{L}_n is closed in $L^\alpha(\mu)$. Let

$$\mathcal{L} = \bigcap_{n=1}^{\infty} \mathcal{L}_n$$

and assume that \mathcal{L} is nonempty. Questions of interest are whether the forward \mathcal{I}_α -projections of a probability measure R on the sets \mathbb{L}_n converge to the forward \mathcal{I}_α -projection on the limiting set \mathbb{L} and whether the optimal values on these sets converge to that on the limiting set. Questions of this kind have been studied for entropy by Borwein and Lewis [35] and for ϕ -entropies by Teboulle and Vajda [36].

- 7) Can one characterize the set of all regular and subspace-transitive projection rules? We therefore wish to relax the locality axiom for projection rules. This ought to include all projection rules generated by Bregman’s divergences of the sum-form and additionally the projection rule generated by \mathcal{I}_α .

APPENDIX

A. Proof of Lemma 2:

These properties are well-known. We provide the proofs of a) - d) for completeness. For e) we provide a reference.

- a) By Jensen’s inequality,

$$I_f(P', Q') \geq 0,$$

with equality if and only if $P' = Q'$, which holds if and only if $P = Q$. Substituting this in (4), we get $\mathcal{I}_\alpha(P, Q) \geq 1/\rho \cdot \log(1) = 0$ for both positive and negative ρ , with equality if and only if $P = Q$.

- b) Using (6), we get

$$\mathcal{I}_\alpha(P, Q) = -\alpha \log \left(\int q^{\alpha-1} dP \right)^{\frac{1}{\alpha-1}} + \log \left(\int p^{\alpha-1} dP \right)^{\frac{1}{\alpha-1}} + (\alpha - 1) \log \left(\int q^{\alpha-1} dQ \right)^{\frac{1}{\alpha-1}}. \quad (63)$$

By assumption, $\mathcal{I}_{\alpha_u}(P, Q) < \infty$, where $\alpha_u > 1$. From the fact that p and q are in $L^{\alpha_u}(\mu)$, we have that $\int p^{\alpha-1} dP = \int p^\alpha d\mu$ is finite and nonzero for all $\alpha \in (1, \alpha_u]$, and the same holds for $\int q^{\alpha-1} dQ$. Using these facts in (63), we conclude that $\int q^{\alpha-1} dP$ is finite and nonzero, and consequently so is $\int q^{\alpha-1} dP$ for all $\alpha \in (1, \alpha_u]$. We shall now apply a result [32,

Ch. 6, Ex. 8] which states that if $g \in L^{\beta_u}(\nu)$ for some $\beta_u > 0$ and a probability measure ν , then $g \in L^\beta(\nu)$ for $0 < \beta < \beta_u$, and

$$\lim_{\beta \downarrow 0} \left(\int |g|^\beta d\nu \right)^{1/\beta} = \exp \left\{ \int (\log |g|) d\nu \right\}.$$

By setting $\beta = \alpha - 1$, and by letting $\alpha \downarrow 1$, we apply the above result on each of the terms on the right-hand side of (63) and conclude that $\int (\log q) dP$, $\int (\log p) dP$, and $\int (\log q) dQ$ exist, and the right-hand side of (63) goes to

$$- \int (\log q) dP + \int (\log p) dP + 0 = \mathcal{I}(P\|Q).$$

A similar argument shows that when $\mathcal{I}_{\alpha_l}(P, Q) < \infty$ for some $\alpha_l < 1$, we have $\lim_{\alpha \uparrow 1} \mathcal{I}_\alpha(P, Q) = \mathcal{I}(P\|Q)$.

c) and d) follow directly from the definitions.

e) This was proved by Lutwak et al. [21, Th. 2] for the scalar case and by Costa et al. [22] for the vector case.

B. Proof of Proposition 3:

We shall first prove the lower semicontinuity for $\alpha < 1$: if $p_n \rightarrow p$ in $L^\alpha(\mu)$ then

$$\liminf_{n \rightarrow \infty} \mathcal{I}_\alpha(p_n, q) \geq \mathcal{I}_\alpha(p, q). \quad (64)$$

Fix an $\alpha < 1$; this fixes a $\rho > 0$. From (4), we may write

$$\mathcal{I}_\alpha(p, q) = \frac{1}{\rho} \log [I_f(p', q') + 1],$$

where $f(u) = u^{1+\rho} - 1$ for $u \geq 0$.

Let $p_n \rightarrow p$ in $L^\alpha(\mu)$. Then $\|p_n\| \rightarrow \|p\| > 0$ and since $|p_n^\alpha - p^\alpha| \leq |p_n|^\alpha + |p|^\alpha$. The generalized version of the dominated convergence theorem states that (see [32, Ch. 2, Ex. 20] or [33, p.139, Problem 19]), if $\{u_n\}$ is a sequence of measurable functions on a measurable space $(\mathbb{X}, \mathcal{X})$ such that $u_n \rightarrow u$ μ -a.e. and if $v_n, v \in L^1(\mu)$ are such that $|u_n| \leq v_n$ μ -a.e. and $v_n \rightarrow v$ in $L^1(\mu)$, then $u_n \rightarrow u$ in $L^1(\mu)$. By taking, $u_n = p_n^\alpha$, and $v_n = |p_n|^\alpha + |p|^\alpha$, the above theorem yields $p_n^\alpha \rightarrow p^\alpha$ in $L^1(\mu)$. From these, we have

$$(p_n/\|p_n\|)^\alpha \rightarrow (p/\|p\|)^\alpha \text{ in } L^1(\mu),$$

i.e., $p'_n \rightarrow p'$ in $L^1(\mu)$, which implies $p'_n/q' \rightarrow p'/q'$ in $L^1(Q')$. (Observe that the argument thus far does not use the assumption that $\alpha < 1$ and is therefore equally applicable for an $\alpha > 1$).

Teboulle and Vajda showed in [36, Lemma 1] that the mapping $h \mapsto \int f(h) d\nu = \int h^{1+\rho} d\nu$ is lower semicontinuous in $L^1(\nu)$ for a probability measure ν on $(\mathbb{X}, \mathcal{X})$. Put $h_n = p'_n/q'$, $h = p'/q'$, and $\nu = Q'$. Then, we just established in the previous paragraph that $h_n \rightarrow h$ in $L^1(\nu)$. Using (3) and the lower semicontinuity result of Teboulle and Vajda, we have

$$\liminf_{n \rightarrow \infty} I_f(p'_n, q') \geq I_f(p', q') \geq 0. \quad (65)$$

Since $1/\rho \log(\cdot + 1)$ is increasing and continuous in $[0, \infty)$, using the definition in (4), (65) implies (64) which establishes the lower semicontinuity result for $\alpha < 1$.

We now deal with the other case. Fix $\alpha > 1$. Observe that the dual space of the Banach space $L^\alpha(\mu)$ is $L^\alpha(\mu)^* = L^{\frac{\alpha}{\alpha-1}}(\mu)$, and therefore $(q/\|q\|)^{\alpha-1} \in L^\alpha(\mu)^*$. Consequently, the mapping defined by

$$T: L^\alpha(\mu) \ni h \mapsto T(h) = \int h \cdot \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu \in \mathbb{R}$$

is a bounded linear functional and therefore continuous. If $p_n \rightarrow p$ in $L^\alpha(\mu)$, then $\|p_n\| \rightarrow \|p\|$, and therefore $p_n/\|p_n\| \rightarrow p/\|p\|$ in $L^\alpha(\mu)$. By the continuity of T , we have

$$\begin{aligned} \int \left(\frac{p_n}{\|p_n\|} \right) \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu &= T \left(\frac{p_n}{\|p_n\|} \right) \\ &\rightarrow T \left(\frac{p}{\|p\|} \right), \text{ as } n \rightarrow \infty, \\ &= \int \left(\frac{p}{\|p\|} \right) \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu. \end{aligned}$$

Taking $1/\rho \log(\cdot)$ on both sides, and using (5), we see $\mathcal{I}_\alpha(p_n, q) \rightarrow \mathcal{I}_\alpha(p, q)$ where $\mathcal{I}_\alpha(p, q)$ may possibly be $+\infty$.

C. *Proof of Proposition 4:*

From (4), we may write

$$\mathcal{I}_\alpha(p, q) = \frac{1}{\rho} \log[\operatorname{sgn}(\rho) \cdot I_{\tilde{f}}(q', p') + 1], \quad (66)$$

where $\tilde{f}(u) = \operatorname{sgn}(\rho) \cdot (u^{-\rho} - 1)$, $u \geq 0$.

Let $q_n \rightarrow q$ in $L^\alpha(\mu)$. Then, as in the proof of Proposition 3, we have that $q'_n/p' \rightarrow q'/p'$ in $L^1(P')$. Following the argument of Proposition 3, we apply the lower semicontinuity result of Teboulle and Vajda [36, Lemma 1] with \tilde{f} playing the role of f , and we have

$$\liminf_{n \rightarrow \infty} I_{\tilde{f}}(q'_n, p') \geq I_{\tilde{f}}(q', p') \geq 0. \quad (67)$$

If either (a) $\rho < 0$ and $I_{\tilde{f}}(q', p') = 1$, or (b) $\rho > 0$ and $I_{\tilde{f}}(q', p') = \infty$, then using the first inequality in (67) and using (66) one easily verifies the limit

$$\liminf_{n \rightarrow \infty} \mathcal{I}_\alpha(p, q_n) = \mathcal{I}_\alpha(p, q) = +\infty. \quad (68)$$

For all other cases, we recognize that $1/\rho \log[\operatorname{sgn}(\rho) \cdot u + 1]$ is an increasing continuous function for $u \in [0, 1]$ when $\rho < 0$ and for $u \in [0, \infty)$ when $\rho > 0$. Using this, the first inequality in (67), and (4), we have the following analog of (64)

$$\liminf_{n \rightarrow \infty} \mathcal{I}_\alpha(p, q_n) \geq \mathcal{I}_\alpha(p, q). \quad (69)$$

Equations (68) and (69) together establish the lower semicontinuity in the second argument.

D. *Proof of Proposition 5:*

Let $p_0, p_1 \in \overline{B}(q, \tau)$, i.e., using (5),

$$\operatorname{sgn}(\rho) \int \frac{p_\lambda}{\|p_\lambda\|} \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu \leq \operatorname{sgn}(\rho) \cdot t \quad \text{for } \lambda = 0, 1, \quad (70)$$

where $t = \exp\{\tau\rho\}$. Now, let us consider $\lambda \in [0, 1]$, and define

$$p_\lambda := \lambda p_1 + (1 - \lambda)p_0. \quad (71)$$

We then have the following chain of inequalities:

$$\begin{aligned} & \operatorname{sgn}(\rho) \int \frac{p_\lambda}{\|p_\lambda\|} \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu \\ & \stackrel{(a)}{=} \frac{\operatorname{sgn}(\rho)}{\|p_\lambda\|} \left[\lambda \int p_1 \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu + (1 - \lambda) \int p_0 \left(\frac{q}{\|q\|} \right)^{\alpha-1} d\mu \right] \\ & \stackrel{(b)}{\leq} \frac{\operatorname{sgn}(\rho)}{\|p_\lambda\|} [\lambda \|p_1\| t + (1 - \lambda) \|p_0\| t] \\ & = \operatorname{sgn}(\rho) \cdot t \cdot \frac{[\lambda \|p_1\| + (1 - \lambda) \|p_0\|]}{\|p_\lambda\|} \\ & \stackrel{(c)}{\leq} \operatorname{sgn}(\rho) \cdot t \cdot 1, \end{aligned}$$

where (a) follows by plugging in (71), (b) follows from (70), and (c) follows because Minkowski's inequality gives that, for $\alpha > 1$, $\|p_\lambda\| \leq \lambda \|p_1\| + (1 - \lambda) \|p_0\|$ while for $0 < \alpha < 1$, this inequality is reversed.

Using (5) once again, this time to write the above inequality in terms of \mathcal{I}_α , we get $\mathcal{I}_\alpha(p_\lambda, q) \leq \tau$, which implies $p_\lambda \in \overline{B}(q, \tau)$ for $\lambda \in [0, 1]$.

E. *Counterexamples as indicated Remark 11:*

Let $\mathbb{X} = (0, 1)$. Let $\mu = Q$ be the Lebesgue measure on \mathbb{X} . Let

$$\mathbb{L} = \left\{ P: \int f_n dP = 0, n = 1, 2, 3, \dots \right\}, \quad (72)$$

where

$$f_n(x) := \begin{cases} \frac{3+\sqrt{n}}{4} & 0 < x < \frac{1}{4n} \\ \frac{3}{4} & \frac{1}{4n} \leq x < \frac{1}{4} \\ -\frac{1}{4} - \frac{1}{4\sqrt{n}} & \frac{1}{4} \leq x < \frac{1}{2} \\ -\frac{1}{4} & \frac{1}{2} \leq x < 1. \end{cases} \quad (73)$$

Then $Q \in \mathbb{L}$. Clearly $\lim_{n \rightarrow \infty} f_n = \frac{3}{4} \cdot 1_{\{(0, \frac{1}{4})\}} - \frac{1}{4} \cdot 1_{\{(\frac{1}{4}, 1)\}}$. Now $g = \lim_{n \rightarrow \infty} f_n$ is in the closure of the linear span of $\{f_n\}_{n \geq 1}$, but not in the linear span of $\{f_n\}_{n \geq 1}$. Let R be a probability measure whose μ -density r satisfies $q^{\alpha-1} = cr^{\alpha-1} + g$. Then $c = \int q^\alpha d\mu / \int qr^{\alpha-1}$. Notice that the inequality in (30), using (5), is equivalent to

$$\int p(q^{\alpha-1} - cr^{\alpha-1})d\mu \geq 0 \text{ if } \alpha > 1 \quad (74)$$

$$\leq 0 \text{ if } \alpha < 1. \quad (75)$$

$\alpha > 1$: *Necessary condition is not sufficient*: Let P be a probability measure defined by

$$\frac{dP}{dQ}(x) = \begin{cases} \frac{1}{5\sqrt{x}} & 0 < x < \frac{1}{4} \\ 0 & \frac{1}{4} \leq x < \frac{3}{5} \\ 2 & \text{if } \frac{3}{5} \leq x < 1. \end{cases} \quad (76)$$

It is easy to check that $P \in \mathbb{L}$. The left-hand side of (74), for the P defined above, evaluates to $-1/20 \not\geq 0$. Therefore, by Th. 10, Q cannot be the forward \mathcal{I}_α -projection of R on \mathbb{L} .

$\alpha > 1$: *Sufficient condition is not necessary*: Define R by setting $g = -\lim_{n \rightarrow \infty} f_n$. The left-hand side of (74) is

$$\begin{aligned} \int g dP &= - \int (\lim_{n \rightarrow \infty} f_n) dP \\ &\geq - \lim_{n \rightarrow \infty} \int f_n dP \\ &= 0, \end{aligned}$$

where the last inequality follows by Fatou's lemma. Since this holds for every $P \in \mathbb{L}$, by Th. 10, Q is the \mathcal{I}_α -projection of R on \mathbb{L} .

For $\alpha < 1$, define R by setting $g = -\lim_{n \rightarrow \infty} f_n$ and $g = \lim_{n \rightarrow \infty} f_n$, respectively to show that the necessary condition is not sufficient and vice-versa.

ACKNOWLEDGEMENTS

We thank the reviewers whose comments/suggestions helped improve this manuscript enormously.

REFERENCES

- [1] M. Ashok Kumar and R. Sundaresan, "Further results on geometric properties of a family of relative entropies," in *Proc. of the 2011 IEEE International Symposium on Information Theory*, July 2011, pp. 1940–1944.
- [2] —, "Relative α -entropy minimizers subject to linear statistical constraints," in *2015 National Conference on Communication, NCC 2015, IIT Bombay, Mumbai*, February–March 2015.
- [3] E. T. Jaynes, *Papers on Probability, Statistics and Statistical Physics*, R. D. Rosenkrantz, Ed. P.O. Box 17,3300 AA Dordrecht, The Netherlands.: Kluwer Academic Publishers, 1982.
- [4] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York: John Wiley & Sons, 2006.
- [5] J. M. V. Campenhout and T. M. Cover, "Maximum entropy and conditional probability," *Information Theory, IEEE Transactions on*, vol. 27, no. 4, pp. 483–489, July 1981.
- [6] I. Csiszár, "Sanov property, generalized I -projection, and a conditional limit theorem," *Ann. Prob.*, vol. 12, no. 3, pp. 768–793, 1984.
- [7] —, " I -divergence geometry of probability distributions and minimization problems," *Ann. Prob.*, vol. 3, pp. 146–158, 1975.
- [8] I. Csiszár and F. Matúš, "Information projections revisited," *Information Theory, IEEE Transactions on*, vol. 49, no. 6, pp. 1474–1490, June 2003.
- [9] I. Csiszár and P. Shields, *Information Theory and Statistics: A Tutorial*, ser. Foundations and Trends in Communications and Information Theory. Hanover, USA: Now Publishers Inc, 2004, vol. 1, no. 4.
- [10] I. Csiszár and G. Tusnády, "Information geometry and alternating minimization procedures," *Statistics and Decisions*, Supp. 1, pp. 205–237, 1984.
- [11] I. Csiszár and F. Matúš, "Generalized minimizers of convex integral functionals, Bregman distance, Pythagorean identities," *Kybernetika*, vol. 48, no. 4, pp. 637–689, 2012.
- [12] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed., ser. Applications of Mathematics. New York, USA: Springer-Verlag, 1998, vol. 38.
- [13] L. L. Campbell, "A coding theorem and Rényi's entropy," *Information and Control*, vol. 8, pp. 423–429, 1965.
- [14] A. C. Blumer and R. J. McEliece, "The Rényi redundancy of generalized Huffman codes," *Information Theory, IEEE Transactions on*, vol. 34, no. 5, pp. 1242–1249, September 1988.
- [15] R. Sundaresan, "Guessing under source uncertainty," *Information Theory, IEEE Transactions on*, vol. 53, no. 1, pp. 269–287, January 2007.

- [16] E. Arikan, "An inequality on guessing and its application to sequential decoding," *Information Theory, IEEE Transactions on*, vol. 42, no. 1, pp. 99–105, January 1996.
- [17] M. K. Hanawal and R. Sundaresan, "Guessing revisited: A large deviations approach," *Information Theory, IEEE Transactions on*, vol. 57, no. 1, pp. 70–78, January 2011.
- [18] R. Sundaresan, "A measure of discrimination and its geometric properties," in *Proc. of the 2002 IEEE International Symposium on Information Theory*, Lausanne, Switzerland, June 2002, p. 264.
- [19] C. Bunte and A. Lapidoth, "Codes for tasks and Rényi entropy," *Information Theory, IEEE Transactions on*, vol. 60, no. 9, pp. 5065–5076, September 2014.
- [20] O. T. Johnson and C. Vignat, "Some results concerning maximum Rényi entropy distributions," *Annales de l'Institut Henri Poincaré (B)*, vol. 43, no. 3, pp. 339–351, May–June 2007.
- [21] E. Lutwak, D. Yang, and G. Zhang, "Cramer-Rao and moment-entropy inequalities for Rényi entropy and generalized Fisher information," *Information Theory, IEEE Transactions on*, vol. 51, no. 1, pp. 473–478, January 2005.
- [22] J. Costa, A. Hero, and C. Vignat, "On solutions to multivariate maximum-entropy problems," in *EMMCVPR 2003, Lisbon, Portugal*, ser. Lecture Notes in Computer Science, A. Rangarajan, M. Figueiredo, and J. Zerubia, Eds., vol. 2683. Berlin, Germany: Springer-Verlag, July 2003, pp. 211–228.
- [23] C. Tsallis, "Possible generalization of Boltzmann-Gibbs statistics," *Journal of Statistical Physics*, vol. 52, no. 1-2, pp. 479–487, 1988.
- [24] I. Csiszár, "Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems," *The Annals of Statistics*, vol. 19, no. 4, pp. 2032–2066, 1991.
- [25] —, "Generalized cutoff rates and Rényi's information measures," *Information Theory, IEEE Transactions on*, vol. 41, pp. 26–34, January 1995.
- [26] T. van Erven and P. Harremoës, "Rényi divergence and Kullback-Leibler divergence," *Information Theory, IEEE Transactions on*, vol. 60, no. 7, pp. 3797–3820, July 2014.
- [27] M. Ashok Kumar and R. Sundaresan, "Minimization problems based on a parametric family of relative entropies II: Reverse projection," *arXiv:1410.5550*, October 2014.
- [28] I. Csiszár, "Information-type measures of difference of probability distributions and indirect observations," *Studia Sci. Math. Hungar.*, vol. 2, pp. 299–318, 1967.
- [29] M. S. Pinsker, *Information and Information Stability of Random Variables and Processes*, ser. Holden-Day series in time series analysis. Holden-Day, San Francisco, 1964.
- [30] R. Bhatia, *Notes on Functional Analysis*. New Delhi, India: Hindustan Book Agency, 2009.
- [31] I. Csiszár, "On topological properties of f -divergences," *Studia Sci. Math. Hungar.*, no. 2, pp. 329–339, 1967.
- [32] G. B. Folland, *Real Analysis: Modern Techniques and their Applications*, 2nd ed. John Wiley and Sons, Inc., 1999.
- [33] F. Jones, *Lebesgue Integration on Euclidean Spaces*. Jones and Bartlett Mathematics, Revised Edition, 2001.
- [34] H. L. Royden, *Real Analysis*, 3rd ed. Delhi, India: Pearson Education (Singapore) Pte. Ltd., Indian Branch, 1988.
- [35] J. M. Borwein and A. S. Lewis, "Convergence of best entropy estimates," *SIAM J. Optimization*, vol. 1, pp. 191–205, 1991.
- [36] M. Teboulle and I. Vajda, "Convergence of best entropy estimates," *Information Theory, IEEE Transactions on*, vol. 39, no. 1, pp. 297–301, January 1993.