

Universal Sampling Rate Distortion

Vinay Praneeth Boda^{1b} and Prakash Narayan^{2b}

Abstract—We examine the coordinated and universal rate-efficient sampling of a subset of correlated discrete memoryless sources followed by lossy compression of the sampled sources. The goal is to reconstruct a predesignated subset of sources within a specified level of distortion. The combined sampling mechanism and rate distortion code are universal in that they are devised to perform robustly without exact knowledge of the underlying joint probability distribution of the sources. In Bayesian as well as nonBayesian settings, single-letter characterizations are provided for the universal sampling rate distortion function for fixed-set sampling, independent random sampling, and memoryless random sampling. It is illustrated how these sampling mechanisms are successively better. Our achievability proofs bring forth new schemes for joint source distribution-learning and lossy compression.

Index Terms—Discrete memoryless multiple source, fixed-set sampling, independent random sampling, joint distribution-learning, memoryless random sampling, sampling rate distortion function, universal rate distortion, universal sampling rate distortion function.

I. INTRODUCTION

CONSIDER a set \mathcal{M} of m discrete memoryless sources with joint probability mass function (pmf) known only to belong to a given family of pmfs. At time instants $t = 1, \dots, n$, possibly different subsets A_t of size $k \leq m$ are sampled “spatially” and compressed jointly by a (block) source code, with the objective of reconstructing a predesignated subset $B \subseteq \mathcal{M}$ of sources from the compressed representations within a specified level of distortion. In forming an efficient rate distortion code that yields the best compression rate for a given distortion level, what are the tradeoffs – under optimal processing – among *causal* sampling procedure, inferential methods for approximating the underlying joint pmf of the memoryless sources, compression rate and distortion level? “Universality” requires that the combined sampling mechanism and lossy compression code be fashioned in the face of imprecise knowledge of the underlying pmf. This paper is a progression of our work in [4] on sampling rate distortion for multiple sources with *known* joint pmf.

Manuscript received June 21, 2017; accepted May 22, 2018. Date of publication July 20, 2018; date of current version November 20, 2018. This work was supported by the U.S. National Science Foundation under Grants CCF-0917057 and CCF-1319799. This paper was presented in part at the 2017 International Symposium on Information Theory.

V. P. Boda was with the Department of Electrical and Computer Engineering and the Institute for Systems Research, University of Maryland at College Park, College Park, MD 20742 USA. He is now with LinkedIn Corporation, Sunnyvale, CA 94085 USA (e-mail: praneeth@umd.edu).

P. Narayan is with the Department of Electrical and Computer Engineering and the Institute for Systems Research, University of Maryland at College Park, College Park, MD 20742 USA (e-mail: prakash@umd.edu).

Communicated by D. L. Neuhoff, Associate Editor for Source Coding.
Digital Object Identifier 10.1109/TIT.2018.2857829

Motivating applications include in-network computation [10] and dynamic thermal management in multicore processor chips [35].

The study of problems of combined sampling and compression has a classical and distinguished history. Recent relevant works include the lossless compression of analog sources in an information theoretic setting [33]; compressed sensing with an allowed detection error rate or quantization distortion [25]; sub-Nyquist temporal sampling followed by lossy reconstruction [14]; and rate distortion function for multiple sources with time-shared sampling [20]. See also [12], [29]. Closer to our approach that entails *spatial* sampling, the rate distortion function has been characterized when multiple Gaussian signals from a random field are sampled and quantized (centralized or distributed) in [23]. In a setting of distributed acoustic sensing and reconstruction, centralized as well as distributed coding schemes and sampling lattices are studied in [15]. In [13], a Gaussian random field on the interval $[0, 1]$ and i.i.d. in time, is reconstructed from compressed versions of k sampled sequences under a mean-squared error distortion criterion. All the sampling problems above assume a knowledge of the underlying pmf.

In the realm of rate distortion theory where a complete knowledge of the signal statistics is unknown, there is a rich literature that considers various formulations of universal coding; only a sampling is listed here. Directions include classical Bayesian and nonBayesian methods [36], [22], [24], [26]; “individual sequences” studies [31], [32], [37]; redundancy in quantization rate or distortion [16]–[18]; and lossy compression of noisy or remote signals [8], [19], [30]. These works propose a variety of distortion measures to investigate universal reconstruction performance.

Interesting variations of our models, especially for the “remote estimation” of a subset of multiple correlated sources, have been considered recently in the control literature [11], [21], [27], [1]. While redolent of our approach in spirit, there are marked differences. For instance, in [11], [21], and [27], sampling is temporal rather than spatial. On the other hand, [1] treats a “scalable” model in which a subset of correlated sources are compressed first in a “coarse” manner and, based on receiver feedback, another possibly different subset of sources are compressed in a refined manner. Connections between these works and ours below remain intriguing and unexplored as yet.

Our work differs materially from the approaches above. Sampling is spatial rather than temporal, unlike in most of the settings above. Furthermore, we introduce new forms of randomized sampling that can depend on the observed source realizations, and which yield a clear gain in performance.

We restrict ourselves to universality that involves a lack of specific knowledge of source pmf within a *finite* family of pmfs. Accordingly, in Bayesian and nonBayesian settings, we consider average and peak distortion criteria, respectively, with an emphasis on the former. Extensions to an infinite family of pmfs are currently under study. It is important to make clear that our models entail a centralized processing of the sources. Note that decentralized encoding would subsume the long-standing general open problem of distributed lossy source coding (cf. e.g., [2], [28]).

Our technical contributions are as follows. In Bayesian and nonBayesian settings, we consider a new formulation involving an *universal sampling rate distortion function* (USRDF), with the objective of capturing the interplay and characterizing inherent tradeoffs among sampling mechanism, approximation of underlying (unknown) pmf, lossy compression rate and distortion level. Our results build on the concept of sampling rate distortion function [4], which in turn uses as an ingredient the rate distortion function for a “remote” source-receiver model with known pmf [2], [3], [9], [34]. We begin with fixed-set sampling where the encoder observes the same set of k sampled sources at every time instant. Recognizing that only the k -marginal pmf of the sources – pertaining to the sampling set – can be learned by the encoder, the corresponding USRDf is characterized. In general, allowing randomization in sampling affords two distinct advantages over fixed-set sampling: better approximation of the underlying joint pmf and improved compression performance enabled by sampling different subsets of sources in apposite proportions. An *independent* random sampler chooses different k -subsets of the sources independently of source realizations and independently in time, and can learn all k -marginals of the joint pmf. This reduction in pmf uncertainty (vis-à-vis fixed-set sampling) aids in improving USRDf. Interestingly, our achievability proof shows how this USRDf can be attained without informing the decoder *explicitly* of the sampling sequence. Lastly, we consider a more powerful sampler, namely the memoryless random sampler, whose choice of sampling sets can depend on instantaneous source realizations. Surprisingly, this latitude allows the encoder to learn the entire joint pmf, and that, too, only from the sampling sequence without recourse to the sampled source realizations. Furthermore, we show how USRDf can be attained by means of a sampling sequence that depends *deterministically* on source realizations, thereby reducing code complexity. Thus, all our achievability proofs bring out new ideas for joint source pmf-learning and lossy compression.

Our model is described in Section II. The main results, illustrated by examples, are stated in Section III. In Section IV, we present the achievability proofs in the increasing order of sampler complexity, with an emphasis on the Bayesian setting; a unified converse proof is presented thereafter.

II. PRELIMINARIES

Denote $\mathcal{M} = \{1, \dots, m\}$, and let $X_{\mathcal{M}} = (X_1, \dots, X_m)$ be a $\mathcal{X}_{\mathcal{M}} = \times_{i=1}^m \mathcal{X}_i$ -valued random variable (rv) where each \mathcal{X}_i

is a finite alphabet. For a (nonempty) set $A \subseteq \mathcal{M}$, we denote by X_A the rv $(X_i, i \in A)$ with values in $\times_{i \in A} \mathcal{X}_i$, and denote n repetitions of X_A by $X_A^n = (X_i^n, i \in A)$ with values in $\mathcal{X}_A^n = \times_{i \in A} \mathcal{X}_i^n$, where $X_i^n = (X_{i1}, \dots, X_{in})$ takes values in the n -fold product space $\mathcal{X}_i^n = \mathcal{X}_i \times \dots \times \mathcal{X}_i$. For $1 \leq k \leq m$, let $\mathcal{A}_k = \{A : A \subseteq \mathcal{M}, |A| = k\}$ be the set of all k -sized subsets of \mathcal{M} and let $A^c = \mathcal{M} \setminus A$. Let $\mathcal{Y}_{\mathcal{M}} = \times_{i=1}^m \mathcal{Y}_i$, where \mathcal{Y}_i is a finite reproduction alphabet for X_i . All logarithms and exponentiations are with respect to the base 2.

Let Θ be a finite set (of parameters) and θ a Θ -valued rv with pmf μ_{θ} of assumed full support. We consider a discrete memoryless multiple source (DMMS) $\{X_{\mathcal{M}t}\}_{t=1}^{\infty}$ consisting of i.i.d. repetitions of the rv $X_{\mathcal{M}}$ with pmf known only to the extent of belonging to a finite family of pmfs $\mathcal{P} = \{P_{X_{\mathcal{M}}|\theta=\tau}, \tau \in \Theta\}$ of assumed full support. Two settings are studied: in a Bayesian formulation, the pmf μ_{θ} is taken to be *known* while in a nonBayesian formulation θ is an *unknown constant* in Θ .

Definition 1: For each $n \geq 1$, in the Bayesian setting, a k -random sampler (k -RS), $1 \leq k \leq m$, collects causally[†] at each $t = 1, \dots, n$, random samples $X_{S_t} \triangleq X_{S_t t}$ from $X_{\mathcal{M}t}$, where S_t is a rv with values in \mathcal{A}_k with (conditional) pmf $P_{S_t|X_{\mathcal{M}}^{t-1}, S^{t-1}}$, with $X_{\mathcal{M}}^t = (X_{\mathcal{M}1}, \dots, X_{\mathcal{M}t})$ and $S^{t-1} = (S_1, \dots, S_{t-1})$. Such a k -RS is specified by a (conditional) pmf $P_{S^n|X_{\mathcal{M}}^n, \theta}$ with the requirement

$$P_{S^n|X_{\mathcal{M}}^n, \theta} = P_{S^n|X_{\mathcal{M}}^n} = \prod_{t=1}^n P_{S_t|X_{\mathcal{M}}^{t-1}, S^{t-1}}. \quad (1)$$

In the nonBayesian setting, the first equality above is redundant. In both settings, a k -RS is unaware of the underlying pmf of the DMMS.

The output of a k -RS is (S^n, X_S^n) where $X_S^n = (X_{S_1}, \dots, X_{S_n})$. Successively restrictive choices of a k -RS in (1) corresponding to

$$P_{S_t|X_{\mathcal{M}}^{t-1}, S^{t-1}} = P_{S_t|X_{\mathcal{M}t}}, \quad t = 1, \dots, n, \quad (2)$$

$$P_{S_t|X_{\mathcal{M}}^{t-1}, S^{t-1}} = P_{S_t}, \quad t = 1, \dots, n, \quad (3)$$

and, for a given $A \subseteq \mathcal{M}$,

$$P_{S_t|X_{\mathcal{M}}^{t-1}, S^{t-1}} = \mathbb{1}(S_t = A), \quad t = 1, \dots, n \quad (4)$$

will be termed the k -memoryless random sampler, k -independent random sampler and the k -fixed-set sampler abbreviated as k -MRS, k -IRS and k -FS, respectively.

Remark: Another formulation of interest allows k to vary with t , subject to a time-average constraint. This is not treated here.

Our objective is to reconstruct a subset of DMMS components with indices in an arbitrary but fixed *recovery set* $B \subseteq \mathcal{M}$, namely X_B^n , from a compressed representation of the k -RS output (S^n, X_S^n) , under a suitable distortion criterion.

[†]By causality, we mean

$$S_t \text{ --- } X_{\mathcal{M}}^t, S^{t-1} \text{ --- } X_{\mathcal{M}, t+1}^n, \quad t = 1, \dots, n$$

as implicit in the second equality in (1).

Definition 2: An n -length block code with k -RS for a DMMS $\{X_{\mathcal{M}t}\}_{t=1}^{\infty}$ with alphabet $\mathcal{X}_{\mathcal{M}}$ and reproduction alphabet \mathcal{Y}_B is a triple $(P_{S^n|X_{\mathcal{M}}^n}, f_n, \varphi_n)$ where $P_{S^n|X_{\mathcal{M}}^n}$ is a k -RS as in (1), and (f_n, φ_n) are a pair of mappings where the encoder f_n maps the k -RS output $(S^n, X_{\mathcal{M}}^n)$ into some finite set $\mathcal{J} = \{1, \dots, J\}$ and the decoder φ_n , with access to S^n and the encoder output, maps $\mathcal{A}_k^n \times \mathcal{J}$ into \mathcal{Y}_B^n . We shall use the compact notation $(P_{S|X_{\mathcal{M}}}, f, \varphi)$, suppressing n . The rate of the code with k -RS $(P_{S|X_{\mathcal{M}}}, f, \varphi)$ is $\frac{1}{n} \log \|f\| = \frac{1}{n} \log J$. (An encoder that operates by forming first an explicit estimate of θ from $(S^n, X_{\mathcal{M}}^n)$ is subsumed by this definition.)

Remark: We note that the decoder φ above is taken to be informed of the sequence of sampling sets S^n . Clearly, this always holds for a k -FS, while the assumption is meaningful for a k -IRS and k -MRS. For a k -IRS, it will be shown to be not needed.

For a given (single-letter) finite-valued distortion measure $d : \mathcal{X}_B \times \mathcal{Y}_B \rightarrow \mathbb{R}^+ \cup \{0\}$, an n -length block code with k -RS $(P_{S|X_{\mathcal{M}}}, f, \varphi)$ will be required to satisfy one of the following distortion criteria (d, Δ) depending on the setting.

(i) Bayesian: The *expected* distortion criterion is

$$\begin{aligned} & \mathbb{E} \left[d \left(X_B^n, \varphi \left(S^n, f \left(S^n, X_{\mathcal{M}}^n \right) \right) \right) \right] \\ & \triangleq \mathbb{E} \left[\frac{1}{n} \sum_{t=1}^n d \left(X_{Bt}, \left(\varphi \left(S^n, f \left(S^n, X_{\mathcal{M}}^n \right) \right) \right)_t \right) \right] \\ & = \sum_{\tau \in \Theta} \mu_{\theta}(\tau) \mathbb{E} \left[\frac{1}{n} \sum_{t=1}^n d \left(X_{Bt}, \left(\varphi \left(S^n, f \left(S^n, X_{\mathcal{M}}^n \right) \right) \right)_t \right) \middle| \theta = \tau \right] \\ & \leq \Delta. \end{aligned} \quad (5)$$

(ii) NonBayesian: The *peak* distortion criterion is

$$\max_{\tau \in \Theta} \mathbb{E} \left[d \left(X_B^n, \varphi \left(S^n, f \left(S^n, X_{\mathcal{M}}^n \right) \right) \right) \middle| \theta = \tau \right] \leq \Delta, \quad (6)$$

where the “conditional” expectation denotes, in fact, $\mathbb{E}_{P_{X_{\mathcal{M}}^n}^{S^n|\theta=\tau}} = \mathbb{E}_{P_{X_{\mathcal{M}}^n|\theta=\tau} P_{S^n|X_{\mathcal{M}}^n}}$.

Definition 3: A number $R \geq 0$ is an achievable universal k -RS coding rate at distortion level Δ if for every $\epsilon > 0$ and sufficiently large n , there exist n -length block codes with k -RS of rate less than $R + \epsilon$ and satisfying the distortion criterion $(d, \Delta + \epsilon)$ in (5) or (6) above; and (R, Δ) will be termed an achievable universal k -RS rate distortion pair under the expected or peak distortion criterion. The infimum of such achievable rates is denoted by $R_A(\Delta)$, $R_i(\Delta)$ and $R_m(\Delta)$ for a k -FS, k -IRS and k -MRS, respectively. We shall refer to $R_A(\Delta)$, $R_i(\Delta)$ as well as $R_m(\Delta)$ as the *universal sampling rate distortion function* (USRdf), suppressing the dependence on k .

Remark: Clearly, the USRdf under (5) will be no larger than that under (6).

III. MAIN RESULTS

We make the following main contributions. First, a (single-letter) characterization is provided of the USRdf for fixed-set sampling, i.e., k -FS, in the Bayesian and nonBayesian settings. Second, building on this, a characterization of the USRdf is obtained for a k -IRS in these settings, and it is shown that




Fig. 1. Ambiguity atoms.

randomized sampling can outperform strictly the “best” fixed-set sampler. Indeed, this USRdf can be attained even upon dispensing with the a priori assumption that the decoder is informed of the sequence of sampling sets. Finally, the USRdf for a k -MRS is characterized and shown to be achievable by a sampler that is determined by the instantaneous realizations of the DMMS at each time instant. We note that the USRdfs for a k -FS and k -IRS can be deduced from that of a k -MRS. Nevertheless, for the sake of expository convenience, we develop the three sampling models in succession; this will also facilitate the presentation of the achievability proofs.

Throughout this paper, a salient theme that recurs is this: An encoder without prior knowledge of θ and with access to only k instantaneously sampled components of the DMMS $\{X_{\mathcal{M}t}\}_{t=1}^{\infty}$ can form only a limited estimate of θ . The quality of said estimate improves steadily from k -FS to k -IRS to k -MRS.

Consider first fixed-set sampling with $A \subseteq \mathcal{M}$ in (4). An encoder f with access to X_A^n cannot distinguish among pmfs in \mathcal{P} (indexed by τ) that have the same $P_{X_A|\theta=\tau}$. Accordingly, let Θ_1 be a partition of Θ comprising “ambiguity” atoms, with each such atom consisting of τ s with *identical marginal pmfs* $P_{X_A|\theta=\tau}$. Indexing the elements of Θ_1 by τ_1 , let θ_1 be a Θ_1 -valued rv with pmf μ_{θ_1} induced by μ_{θ} . For each $\tau_1 \in \Theta_1$, let $\Lambda(\tau_1)$ be the collection of τ s in the atom of Θ_1 indexed by τ_1 . In the Bayesian setting,

$$P_{X_A|\theta_1=\tau_1} = P_{X_A|\theta=\tau}, \quad \tau \in \Lambda(\tau_1).$$

In the nonBayesian setting, in order to retain the same notation, we choose $P_{X_A|\theta_1=\tau_1}$ to be the right-side above.

When the pmf of the DMMS $\{X_{\mathcal{M}t}\}_{t=1}^{\infty}$ is *known*, say $P_{X_{\mathcal{M}}}$ – corresponding to $|\Theta| = 1$ – we recall from [4] that the (U)SRdf for fixed $A \subseteq \mathcal{M}$ is

$$R_A(\Delta) = \min_{\substack{X_{\mathcal{M}} \rightarrow X_A \rightarrow Y_B \\ \mathbb{E}[d(X_B, Y_B)] \leq \Delta} } I(X_A \wedge Y_B), \quad \Delta_{\min} \leq \Delta \leq \Delta_{\max}, \quad (7)$$

with

$$\begin{aligned} \Delta_{\min} &= \mathbb{E} \left[\min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d(X_B, y_B) | X_A] \right], \\ \Delta_{\max} &= \min_{y_B \in \mathcal{Y}_B} \left[\mathbb{E}[d(X_B, y_B) | X_A] \right], \end{aligned}$$

which can be interpreted as the (standard) rate distortion function for the DMMS $\{X_{A_t}\}_{t=1}^{\infty}$ using a modified distortion

measure \tilde{d} defined by

$$\tilde{d}(x_A, y_B) = \mathbb{E}[d(X_B, y_B)|X_A = x_A].$$

This fact will serve as a stepping stone to our analysis of USRDf for a k -random sampler. In the Bayesian setting, we consider a modified distortion measure d_{τ_1} , $\tau_1 \in \Theta_1$, given by

$$d_{\tau_1}(x_A, y_B) \triangleq \mathbb{E}[d(X_B, y_B)|X_A = x_A, \theta_1 = \tau_1]; \quad (8)$$

the set of (constrained) pmfs

$$\kappa_A^{\mathcal{B}}(\delta, \tau_1) \triangleq \{P_{\theta X_{\mathcal{M}} Y_B} : \theta, X_{\mathcal{M}} \text{ --- } \theta_1, X_A \text{ --- } Y_B, \mathbb{E}[d_{\tau_1}(X_A, Y_B)|\theta_1 = \tau_1] \leq \delta\}, \quad (9)$$

and the (minimized) conditional mutual information

$$\rho_A^{\mathcal{B}}(\delta, \tau_1) \triangleq \min_{\kappa_A^{\mathcal{B}}(\delta, \tau_1)} I(X_A \wedge Y_B|\theta_1 = \tau_1) \quad (10)$$

which is akin to (7) and will play a basal role. In the nonBayesian setting, the counterparts of (9) and (10) are

$$\kappa_A^{n\mathcal{B}}(\delta, \tau_1) \triangleq \{P_{X_{\mathcal{M}} Y_B|\theta=\tau} = P_{X_{\mathcal{M}}|\theta=\tau} P_{Y_B|X_A, \theta_1=\tau_1} : \mathbb{E}[d(X_B, Y_B)|\theta = \tau] \leq \delta, \tau \in \Lambda(\tau_1)\} \quad (11)$$

and

$$\rho_A^{n\mathcal{B}}(\delta, \tau_1) \triangleq \min_{\kappa_A^{n\mathcal{B}}(\delta, \tau_1)} I(X_A \wedge Y_B|\theta_1 = \tau_1). \quad (12)$$

Remarks: (i) The minima in (10) and (12) exist as those of convex functions over convex, compact sets.

(ii) It is seen in a standard manner that $\rho_A^{\mathcal{B}}(\delta, \tau_1)$ in (10) and $\rho_A^{n\mathcal{B}}(\delta, \tau_1)$ in (12) are convex and continuous in δ .

(iii) Clearly, the minimum in (12) under pmf-wise constraints (11) can be no smaller than that in (10) under pmf-averaged constraints (9).

Our first main result states that the USRDf at distortion level Δ for fixed-set sampling in the Bayesian setting is a minmax of quantities in (10), where the maximum is over ambiguity atoms τ_1 in Θ_1 , while the minimum is over distortion thresholds $\delta = \Delta_{\tau_1}$, $\tau_1 \in \Theta_1$ whose mean does not exceed Δ . On the other hand, in the nonBayesian setting, the USRDf at distortion level Δ is a maximum over ambiguity atoms of quantities in (12) with $\delta = \Delta$, and hence is no smaller than its Bayesian counterpart.

Theorem 1: The Bayesian USRDf for fixed $A \subseteq \mathcal{M}$ is

$$R_A(\Delta) = \min_{\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\}} \max_{\tau_1 \in \Theta_1} \rho_A^{\mathcal{B}}(\Delta_{\tau_1}, \tau_1) \quad (13)$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, where

$$\Delta_{\min} = \mathbb{E}\left[\mathbb{E}\left[\min_{y_B \in \mathcal{Y}_B} d_{\theta_1}(X_A, y_B)|\theta_1\right]\right] = \mathbb{E}\left[\min_{y_B \in \mathcal{Y}_B} d_{\theta_1}(X_A, y_B)\right],$$

$$\Delta_{\max} = \mathbb{E}\left[\min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d_{\theta_1}(X_A, y_B)|\theta_1]\right].$$

The nonBayesian USRDf is

$$R_A(\Delta) = \max_{\tau_1 \in \Theta_1} \rho_A^{n\mathcal{B}}(\Delta, \tau_1), \quad \Delta_{\min} \leq \Delta \leq \Delta_{\max} \quad (14)$$

where

$$\begin{aligned} \Delta_{\min} &= \max_{\tau_1 \in \Theta_1} \\ &\min_{P_{Y_B|X_A, \theta_1=\tau_1} = P_{Y_B|X_{\mathcal{M}}, \theta=\tau}} \max_{\tau \in \Lambda(\tau_1)} \mathbb{E}[d(X_B, Y_B)|\theta = \tau] \end{aligned}$$

and

$$\Delta_{\max} = \max_{\tau_1 \in \Theta_1} \min_{y_B \in \mathcal{Y}_B} \max_{\tau \in \Lambda(\tau_1)} \mathbb{E}[d(X_B, y_B)|\theta = \tau].$$

Remarks: (i) In fact, the minimizing pmf $P_{Y_B|X_A, \theta_1}$ in Δ_{\min} is a conditional point-mass.

(ii) We note that for a given distortion level Δ , the set $\{\Delta_{\tau_1}, \tau_1 \in \Theta_1 : \sum_{\tau_1 \in \Theta_1} \mu_{\theta_1}(\tau_1) \Delta_{\tau_1} \leq \Delta\}$ is a convex, compact set in $\mathbb{R}^{|\Theta_1|}$. Next, observing that

$$\max_{\tau_1 \in \Theta_1} \rho_A^{\mathcal{B}}(\Delta_{\tau_1}, \tau_1)$$

is a convex function of $\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\}$, the minimum in (13) exists as that of a convex function over a convex, compact set.

(iii) The minimizing $\{\Delta_{\tau_1}^*, \tau_1 \in \Theta_1\}$ in (13) is characterized by the following special property: For a given $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, for each $\tau_1 \in \Theta_1$, either

$$\rho_A^{\mathcal{B}}(\Delta_{\tau_1}^*, \tau_1) \equiv \max_{\tilde{\tau}_1 \in \Theta_1} \rho_A^{\mathcal{B}}(\Delta_{\tilde{\tau}_1}^*, \tilde{\tau}_1) \quad (15)$$

where the right-side does not depend on τ_1 , or

$$\Delta_{\tau_1}^* = \mathbb{E}\left[\min_{y_B \in \mathcal{Y}_B} d_{\tau_1}(X_A, y_B)|\theta_1 = \tau_1\right].$$

By a standard argument in convex optimization, if $\{\Delta_{\tau_1}^*, \tau_1 \in \Theta_1\}$ does not satisfy the property above, then a small perturbation decreases the maximum in (15) leading to a contradiction.

(iv) The Δ_{\min} and Δ_{\max} for the Bayesian and the non-Bayesian settings can be different.

Example 1: For the probability of error distortion measure

$$d(x_B, y_B) = \mathbb{1}(x_B \neq y_B) = 1 - \prod_{i \in B} \mathbb{1}(x_i = y_i),$$

$$x_B, y_B \in \mathcal{X}_B = \mathcal{Y}_B$$

the Bayesian USRDf for fixed-set sampling with $A \subseteq B$ in (13) simplifies with (10) becoming

$$\rho_A^{\mathcal{B}}(\Delta_{\tau_1}, \tau_1) = \min_{\substack{\mathbb{E}[\alpha_{\tau_1}(X_A) \mathbb{1}(X_A \neq Y_A)|\theta_1 = \tau_1] \\ \leq \Delta_{\tau_1} - (1 - \mathbb{E}[\alpha_{\tau_1}(X_A)|\theta_1 = \tau_1])}} I(X_A \wedge Y_A|\theta_1 = \tau_1) \quad (16)$$

where

$$\alpha_{\tau_1}(x_A) = \max_{\tilde{x} \in \mathcal{X}_B} P_{X_B|X_A, \theta_1}(\tilde{x}|x_A, \tau_1) \quad (17)$$

is the *maximum a posteriori* (MAP) estimate of X_B on the basis of $X_A = x_A$ under pmf $P_{X_{\mathcal{M}}|\theta_1=\tau_1}$.

The proof of (16), (17) is along the lines of that of [4, Proposition 1] under the pmf $P_{X_{\mathcal{M}}|\theta_1=\tau_1}$ (rather than $P_{X_{\mathcal{M}}}$ as in [4]), and so is not repeated here. Furthermore,

$$\Delta_{\min} = 1 - \mathbb{E}[\alpha_{\theta_1}(X_A)] \quad \text{and}$$

$$\Delta_{\max} = 1 - \mathbb{E}\left[\max_{x_B \in \mathcal{X}_B} P_{X_B|\theta_1}(x_B|\theta_1)\right].$$

The form of the Bayesian USRDf in (16) suggests a simple achievability scheme comprising two steps. Using a *maximum a posteriori* (MAP) or *maximum likelihood* (ML) estimate $\hat{\tau}_1$ of θ_1 on the basis of $X_A^n = x_A^n$, the first step entails a lossy reconstruction of x_A^n by its codeword y_A^n , under pmf $P_{X_{\mathcal{M}}|\theta_1=\hat{\tau}_1}$ and for a modified distortion measure

$$\tilde{d}_{\hat{\tau}_1}(x_A, y_A) \triangleq \alpha_{\hat{\tau}_1}(x_A) \mathbb{1}(x_A \neq y_A)$$

with a corresponding reduced threshold

$$\Delta_{\hat{\tau}_1} = (1 - \mathbb{E}[\alpha_{\hat{\tau}_1}(X_A)|\theta_1 = \hat{\tau}_1]).$$

This is followed by a second step of reconstructing x_B^n from the output y_A^n of the previous step as a MAP estimate

$$y_B^n = \arg \max_{y_B^n \in \mathcal{Y}_B^n} P_{X_B|X_A\theta_1}(y_B^n|y_A^n, \hat{\tau}_1);$$

the corresponding probability of estimation error coincides with the mentioned reduction $1 - \mathbb{E}[\alpha_{\hat{\tau}_1}(X_A)|\theta_1 = \hat{\tau}_1]$ in the threshold.

In the nonBayesian setting, the USRDf in (14), (12) simplifies with

$$\begin{aligned} \rho_A^{nB}(\Delta, \tau_1) &= \min_{\substack{P_{Y_A|X_A, \theta_1=\tau_1} P_{Y_B|X_A, \theta_1=\tau_1} P_{Y_B|X_{\mathcal{M}}, \theta=\tau} \\ \mathbb{E}[\mathbb{1}(X_B \neq Y_B)|\theta=\tau] \leq \Delta, \tau \in \Lambda(\tau_1)}} I(X_A \wedge Y_A|\theta_1 = \tau_1), \end{aligned} \quad (18)$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, where

$$\Delta_{\min} = \max_{\tau_1 \in \Theta_1} \min_{P_{Y_B|X_A, \theta_1=\tau_1}} \max_{\tau \in \Lambda(\tau_1)} (1 - P(X_B = Y_B|\theta = \tau))$$

and

$$\Delta_{\max} = \max_{\tau_1 \in \Theta_1} \min_{Y_B \in \mathcal{Y}_B} \max_{\tau \in \Lambda(\tau_1)} (1 - P_{X_B|\theta}(Y_B|\tau)).$$

This leads to the following achievability scheme. With $\hat{\tau}_1$ as the ML estimate of θ_1 formed from $X_A^n = x_A^n$, first x_A^n is reconstructed as y_A^n according to $P_{Y_A|X_A, \theta_1=\hat{\tau}_1}$ resulting from the minimization in (18). This is followed by the reconstruction of x_B^n from y_A^n by means of the estimate

$$y_B^n = \arg \max_{y_B^n \in \mathcal{Y}_B^n} P_{Y_B|Y_A\theta_1}(y_B^n|y_A^n, \hat{\tau}_1)$$

under pmf $P_{Y_B|Y_A\theta_1}$ which, too, is obtained from the minimization in (18). \square

Example 2: Let $\mathcal{M} = \{1, 2\}$ and $\mathcal{X}_1 = \mathcal{X}_2 = \{0, 1\}$, consider a DMMS with $P_{X_1 X_2|\theta=\tau}$ represented by a virtual binary symmetric channel (BSC) shown in Figure III, where $p_\tau, q_\tau \leq 0.5$, $\tau \in \Theta$, where Θ is a given finite set. For $A = \{1\}$, $B = \{1, 2\}$, and the probability of error distortion measure of Example 1, the Bayesian USRDf reduces to

$$R_{\{1\}}(\Delta) = \min_{\substack{\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\} \\ \mathbb{E}[\Delta_{\theta_1}] \leq \Delta}} \max_{\tau_1 \in \Theta_1} \left(h(p_{\tau_1}) - h\left(\frac{\Delta_{\tau_1} - q_{\tau_1}}{1 - q_{\tau_1}}\right) \right),$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, where

$$\Delta_{\min} = \mathbb{E}[q_{\theta_1}], \quad \Delta_{\max} = \mathbb{E}[p_{\theta_1} + q_{\theta_1} - p_{\theta_1}q_{\theta_1}];$$




Fig. 2. Virtual BSC (q).

and $q_{\tau_1} = P_{X_2|X_1\theta_1}(0|1, \tau_1)$, $\tau_1 \in \Theta_1$; and the nonBayesian USRDf is

$$R_{\{1\}}(\Delta) = \max_{\tau_1 \in \Theta_1} \left(h(p_{\tau_1}) - \min_{\tau \in \Lambda(\tau_1)} h\left(\frac{\Delta - q_\tau}{1 - q_\tau}\right) \right)$$

with

$$\Delta_{\min} = \max_{\tau \in \Theta} q_\tau \quad \text{and} \quad \Delta_{\max} = \max_{\tau \in \Theta} (p_\tau + q_\tau - p_\tau q_\tau).$$

First, it is easy to see that Δ_{\min} and Δ_{\max} in the Bayesian setting are smaller than their counterparts in the nonBayesian setting. Next, note that the Bayesian USRDf, with an additional outer minimization over thresholds Δ_{τ_1} is smaller than the nonBayesian USRDf which involves a maximum over members in each ambiguity atom. \square

Example 3: This example, albeit concocted, shows that for fixed-set sampling with A and recovery set B , a choice of A outside B can be best. Let $\mathcal{M} = \{1, 2, 3\}$, $B = \{1, 2\}$ and $\mathcal{X}_i = \mathcal{Y}_j = \{0, 1\}$, $i = 1, 2, 3$; $j = 1, 2$. Consider a DMMS with $P_{X_1 X_2|\theta=\tau}$ as in Figure III and $X_3 = X_1 \oplus X_2$ where \oplus denotes addition modulo 2. Here, $p_\tau = 0.5$, $q_\tau \leq 0.5$, $\tau \in \Theta$, with the q_τ s, $\tau \in \Theta$, being distinct. For distortion measure $d(x_B, y_B) \triangleq \mathbb{1}((x_1 \oplus x_2) \neq (y_1 \oplus y_2))$, the Bayesian USRDf for fixed-set sampling is

$$\begin{aligned} R_{\{1\}}(\Delta) &= h(0.5) - h\left(\frac{\Delta - \tilde{q}}{1 - 2\tilde{q}}\right), \\ \Delta_{\min} = \tilde{q} \leq \Delta \leq \Delta_{\max} &= 0.5, \end{aligned} \quad (19)$$

where $\tilde{q} = \sum_{\tau \in \Theta} \mu_\theta(\tau) q_\tau$. Since $P_{X_1|\theta=\tau}$ is the same for all $\tau \in \Theta$, note that $|\Theta_1| = 1$. The nonBayesian USRDf is

$$\begin{aligned} R_{\{1\}}(\Delta) &= h(0.5) - \min_{\tau \in \Theta} h\left(\frac{\Delta - q_\tau}{1 - 2q_\tau}\right), \\ \Delta_{\min} = \max_{\tau \in \Theta} q_\tau \leq \Delta \leq \Delta_{\max} &= 0.5. \end{aligned} \quad (20)$$

Also, $R_{\{1\}}(\Delta) = R_{\{2\}}(\Delta)$. For sampling set $A = \{3\}$, $\Theta_1 = \Theta$ and the Bayesian USRDf is

$$\begin{aligned} R_{\{3\}}(\Delta) &= \min_{\{\Delta_\tau, \tau \in \Theta\}: \mathbb{E}[\Delta_\theta] \leq \Delta} \max_{\tau \in \Theta} h(q_\tau) - h(\Delta_\tau), \\ \Delta_{\min} = 0 \leq \Delta \leq \Delta_{\max} &= \tilde{q}, \end{aligned} \quad (21)$$

and the nonBayesian USRDf is

$$\begin{aligned} R_{\{3\}}(\Delta) &= \max_{\tau \in \Theta} h(q_\tau) - h(\Delta), \\ \Delta_{\min} = 0 \leq \Delta \leq \Delta_{\max} &= \max_{\tau \in \Theta} q_\tau. \end{aligned} \quad (22)$$

Since Δ_{\min} in (19) equals Δ_{\max} in (21), in the Bayesian setting $R_3(\Delta)$ is strictly smaller than $R_1(\Delta)$. A similar

observation holds in the nonBayesian setting upon examining (20) and (22). \square

Turning to a k -IRS in (3), the freedom now given to the sampler to rove over all k -sized subsets in \mathcal{A}_k engenders a partition Θ_2 of Θ_1 (and hence a finer partition of Θ) with smaller ambiguity atoms. Let $A_1, \dots, A_{|\mathcal{A}_k|}$, where $|\mathcal{A}_k| = \binom{m}{k}$, be any fixed ordering of \mathcal{A}_k . Let Θ_2 be a partition of Θ consisting of ambiguity atoms, with each atom formed by τ s with *identical (ordered) collections of marginal pmfs* $(P_{X_{A_i}|\theta=\tau}, i = 1, \dots, |\mathcal{A}_k|)$.

Note that Θ_2 is a refinement of Θ_1 (for any A_i). Indexing the elements of Θ_2 by τ_2 , let θ_2 be a Θ_2 -valued rv with pmf μ_{θ_2} derived from μ_θ . For each τ_2 in Θ_2 , let $\Lambda(\tau_2)$ be the collection of τ s in the atom indexed by τ_2 . In analogy with (10) and (12), we define counterparts in the Bayesian and nonBayesian settings as

$$\rho_i^{\mathcal{B}}(\delta, P_S, \tau_2) \triangleq \min_{\kappa_i^{\mathcal{B}}(\delta, P_S, \tau_2)} I(X_S \wedge Y_B | S, \theta_2 = \tau_2); \quad (23)$$

$$\rho_i^{n\mathcal{B}}(\delta, P_S, \tau_2) \triangleq \min_{\kappa_i^{n\mathcal{B}}(\delta, P_S, \tau_2)} I(X_S \wedge Y_B | S, \theta_2 = \tau_2), \quad (24)$$

where d_{τ_2} is defined as in (8) with $\theta_2 = \tau_2$ replacing $\theta_1 = \tau_1$, and

$$\begin{aligned} \kappa_i^{\mathcal{B}}(\delta, P_S, \tau_2) &\triangleq \{P_{\theta X_{\mathcal{M}} S Y_B} = \mu_\theta P_{X_{\mathcal{M}}|\theta} P_S P_{Y_B|S X_S \theta_2} \\ &: \sum_{A \in \mathcal{A}_k} P_S(A) \mathbb{E}[d_{\tau_2}(X_A, Y_B) | S = A, \theta_2 = \tau_2] \leq \delta\}, \end{aligned}$$

$$\begin{aligned} \kappa_i^{n\mathcal{B}}(\delta, P_S, \tau_2) &\triangleq \{P_{X_{\mathcal{M}} S Y_B | \theta = \tau} = P_{X_{\mathcal{M}}|\theta = \tau} P_S P_{Y_B|S X_S, \theta_2 = \tau_2} \\ &: \sum_{A \in \mathcal{A}_k} P_S(A) \mathbb{E}[d(X_B, Y_B) | S = A, \theta = \tau] \leq \delta, \tau \in \Lambda(\tau_2)\}. \end{aligned}$$

Theorem 2: The Bayesian USRDf for a k -IRS is

$$R_i(\Delta) = \min_{P_S, \{\Delta_{\tau_2}, \tau_2 \in \Theta_2\}} \max_{\substack{\tau_2 \in \Theta_2 \\ \mathbb{E}[\Delta_{\theta_2}] \leq \Delta}} \rho_i^{\mathcal{B}}(\Delta_{\tau_2}, P_S, \tau_2), \quad \Delta_{\min} \leq \Delta \leq \Delta_{\max}, \quad (25)$$

where

$$\Delta_{\min} = \min_{A \in \mathcal{A}_k} \mathbb{E}[\mathbb{E}[\min_{y_B \in \mathcal{Y}_B} d_{\theta_2}(X_A, y_B) | \theta_2]]$$

and

$$\Delta_{\max} = \min_{A \in \mathcal{A}_k} \mathbb{E}[\min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d_{\theta_2}(X_A, y_B) | \theta_2]].$$

The nonBayesian USRDf is

$$R_i(\Delta) = \min_{P_S} \max_{\tau_2 \in \Theta_2} \rho_i^{n\mathcal{B}}(\Delta, P_S, \tau_2), \quad \Delta_{\min} \leq \Delta \leq \Delta_{\max}, \quad (26)$$

for

$$\Delta_{\min} = \min_{P_S} \max_{\tau_2 \in \Theta_2} \sum_{A \in \mathcal{A}_k} P_S(A) \min_{P_{Y_B|S X_S, \theta_2 = \tau_2} = P_{Y_B|S X_{\mathcal{M}}, \theta = \tau}} \max_{\tau \in \Lambda(\tau_2)} \mathbb{E}[d(X_B, Y_B) | S = A, \theta = \tau]$$

and

$$\Delta_{\max} = \max_{\tau_2 \in \Theta_2} \min_{y_B \in \mathcal{Y}_B} \max_{\tau \in \Lambda(\tau_2)} \mathbb{E}[d(X_B, y_B) | \theta = \tau].$$

Corollary 1: The USRDfs in the Bayesian and nonBayesian settings remain unchanged upon a restriction to n -length block codes (f, φ) with uninformed decoder, i.e., with $\varphi = \varphi(f(S^n, X_S^n))$.

Remarks: (i) For a k -IRS we restrict ourselves to the interesting case of $k < |B|$, for otherwise it would suffice to choose $S_t = B$, $t = 1, \dots, n$.

(ii) Akin to a k -FS, the optimizing P_S , $\{\Delta_{\tau_2}^*, \tau_2 \in \Theta_2\}$ in (25) has the following special property: For a given $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, for each $\tau_2 \in \Theta_2$, either

$$\rho_i^{\mathcal{B}}(\Delta_{\tau_2}^*, P_S, \tau_2) = \max_{\tilde{\tau}_2 \in \Theta_2} \rho_i^{\mathcal{B}}(\Delta_{\tilde{\tau}_2}^*, P_S, \tilde{\tau}_2)$$

or

$$\Delta_{\tau_2}^* = \sum_{A \in \mathcal{A}_k} P_S(A) \mathbb{E}[\min_{y_B \in \mathcal{Y}_B} d_{\tau_2}(X_A, y_B) | \theta_2 = \tau_2].$$

(iii) In general, a k -IRS will outperform a k -FS in two ways. First, the former enables a better approximation of θ in the form of θ_2 whereas the latter estimates $\theta_1 = \theta_1(\theta_2)$. Second, random sampling enables a ‘‘time-sharing’’ over various fixed-set samplers, that can outperform strictly the best fixed-set choice. Both these advantages of a k -IRS over fixed-set sampling are illustrated in Examples 4 and 5.

Example 4: This example illustrates that a k -IRS can perform strictly better than the best k -FS. For $\mathcal{M} = B = \{1, 2\}$, and $\mathcal{X}_i = \mathcal{Y}_i = \{0, 1\}$, $i = 1, 2$, consider a DMMS with $P_{X_1 X_2 | \theta = \tau} = P_{X_1 | \theta = \tau} P_{X_2 | \theta = \tau}$ where

$$P_{X_1 | \theta}(0 | \tau) = 1 - p_\tau, \quad P_{X_2 | \theta}(0 | \tau) = 1 - q_\tau, \quad \tau \in \Theta,$$

and $0 < p_\tau, q_\tau < 0.5$. Under the distortion measure $d(x_B, y_B) = \mathbb{1}(x_1 \neq y_1) + \mathbb{1}(x_2 \neq y_2)$, for a k -FS, with $k = 1$, the Bayesian USRDf for sampling set $A = \{1\}$ is

$$R_{\{1\}}(\Delta) = \min_{\substack{\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\} \\ \mathbb{E}[\Delta_{\theta_1}] \leq \Delta}} \max_{\tau_1 \in \Theta_1} (h(p_{\tau_1}) - h(\Delta_{\tau_1} - q_{\tau_1})),$$

for $\mathbb{E}[q_\theta] \leq \Delta \leq \mathbb{E}[p_\theta + q_\theta]$, where $q_{\tau_1} = \mathbb{E}[q_\theta | \theta_1 = \tau_1]$, and the nonBayesian USRDf is

$$R_{\{1\}}(\Delta) = \max_{\tau_1 \in \Theta_1} \left(h(p_{\tau_1}) - \min_{\tau \in \Lambda(\tau_1)} h(\Delta - q_\tau) \right)$$

for $\max_{\tau \in \Theta} q_\tau \leq \Delta \leq \max_{\tau \in \Theta} (p_\tau + q_\tau)$.

Turning to a k -IRS with $k = 1$, clearly, $\Theta_2 = \Theta$. For a k -IRS the Bayesian USRDf is

$$R_i(\Delta) = \min_{P_S, \{\Delta_\tau, \tau \in \Theta\}} \max_{\substack{\tau \in \Theta \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \min_{\substack{\Delta_{1\tau}, \Delta_{2\tau} \\ P_S(\{1\})\Delta_{1\tau} + P_S(\{2\})\Delta_{2\tau} \leq \Delta}} I,$$

for $\min\{\mathbb{E}[p_\theta], \mathbb{E}[q_\theta]\} \leq \Delta \leq \mathbb{E}[p_\theta + q_\theta]$ and the nonBayesian USRDf is

$$R_i(\Delta) = \min_{P_S} \max_{\tau \in \Theta} \min_{\substack{\Delta_{1\tau}, \Delta_{2\tau} \\ P_S(\{1\})\Delta_{1\tau} + P_S(\{2\})\Delta_{2\tau} \leq \Delta}} I, \quad (27)$$

for $\min_{0 \leq \alpha \leq 1} \max_{\tau \in \Theta} (\alpha p_\tau + (1 - \alpha) q_\tau) \leq \Delta \leq \max_{\tau \in \Theta} (p_\tau + q_\tau)$ where I equals

$$P_S(\{1\})(h(p_\tau) - h(\Delta_{1\tau} - q_\tau)) + P_S(\{2\})(h(q_\tau) - h(\Delta_{2\tau} - p_\tau)).$$

An analytical comparison of the USRDf shows the strict superiority of the k -IRS over the k -FS, as seen – for instance – by the lower values of Δ_{\min} for the former. \square

Example 5: In Example 4, assume that

$$p_\tau \geq q_\tau, \quad \tau \in \Theta.$$

For a k -FS with $k = 1$, the nonBayesian USRDf is

$$\begin{aligned} R_{\{1\}}(\Delta) &= \max_{\tau_1 \in \Theta_1} \left(h(p_{\tau_1}) - \min_{\tau \in \Lambda(\tau_1)} h(\Delta - q_\tau) \right), \\ R_{\{2\}}(\Delta) &= \max_{\tau_1 \in \Theta_1} \left(h(q_{\tau_1}) - \min_{\tau \in \Lambda(\tau_1)} h(\Delta - p_\tau) \right). \end{aligned} \quad (28)$$

Now, observe that for each $\tau \in \Theta$

$$h(p_\tau) - h(\delta - q_\tau) \leq h(q_\tau) - h(\delta - p_\tau)$$

holds for $p_\tau \leq \delta \leq p_\tau + q_\tau$. Thus, for a k -IRS with $k = 1$, the nonBayesian USRDf in (27) simplifies to

$$R_i(\Delta) = \max_{\tau \in \Theta} h(p_\tau) - h(\Delta - q_\tau)$$

which is strictly smaller than the USRDf for the better k -FS in (28). The superior performance of the k -IRS is enabled by its ability to estimate simultaneously both $P_{X_1|\theta}$ and $P_{X_2|\theta}$ (and thereby $P_{X_1 X_2|\theta}$); a k -FS can estimate only one of $P_{X_1|\theta}$ or $P_{X_2|\theta}$. \square

Lastly, for a k -MRS in (2), the ability of the sampler to depend instantaneously on the current realization of the DMMS enables an encoder with access to the sampler output to distinguish among all the pmfs in \mathcal{P} . Accordingly, for a k -MRS, Θ itself serves as the counterpart of the partitions Θ_1 (for a k -FS) and Θ_2 for a k -IRS. For a rv U with fixed pmf P_U on some finite set \mathcal{U} , and for fixed $P_{S|X_{\mathcal{M}}U}$, we define the counterparts of (23) and (24) as

$$\begin{aligned} \rho_m^{\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau) \\ \triangleq \min_{\kappa_m^{\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau)} I(X_S \wedge Y_B | S, U, \theta = \tau), \end{aligned} \quad (29)$$

and

$$\begin{aligned} \rho_m^{n\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau) \\ \triangleq \min_{\kappa_m^{n\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau)} I(X_S \wedge Y_B | S, U, \theta = \tau), \end{aligned} \quad (30)$$

where the minimization in (29) and (30), in effect, is with respect to $P_{Y_B|S X_S U \theta}$ and the sets of (constrained) pmfs are

$$\begin{aligned} \kappa_m^{\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau) \\ \triangleq \{P_{\theta U X_{\mathcal{M}} S Y_B} = \mu_\theta P_U P_{X_{\mathcal{M}}|\theta} P_{S|X_{\mathcal{M}}U} P_{Y_B|S X_S U \theta} \\ : \mathbb{E}[d(X_B, Y_B)|\theta = \tau] \leq \delta\}, \end{aligned}$$

and

$$\begin{aligned} \kappa_m^{n\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau) \\ \triangleq \{P_{U X_{\mathcal{M}} S Y_B|\theta=\tau} = P_U P_{X_{\mathcal{M}}|\theta=\tau} P_{S|X_{\mathcal{M}}U} P_{Y_B|S X_S U, \theta=\tau} \\ : \mathbb{E}[d(X_B, Y_B)|\theta = \tau] \leq \delta\}. \end{aligned}$$

Here, U plays the role of a “time-sharing” rv, as will be seen below.

Theorem 3: For a k -MRS, the Bayesian USRDf is

$$R_m(\Delta) = \min_{P_U, P_{S|X_{\mathcal{M}}U}, \{\Delta_\tau, \tau \in \Theta\}} \max_{\substack{\tau \in \Theta \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \rho_m^{\mathcal{B}}(\Delta_\tau, P_U, P_{S|X_{\mathcal{M}}U}, \tau), \quad (31)$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, where

$$\Delta_{\min} = \min_{P_{S|X_{\mathcal{M}}}} \mathbb{E} \left[\min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d(X_B, y_B) | S, X_S, \theta] \right]$$

and

$$\Delta_{\max} = \min_{P_{S|X_{\mathcal{M}}}} \mathbb{E} \left[\min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d(X_B, y_B) | S, \theta] \right]. \quad (32)$$

The nonBayesian USRDf is

$$R_m(\Delta) = \min_{P_U, P_{S|X_{\mathcal{M}}U}} \max_{\tau \in \Theta} \rho_m^{n\mathcal{B}}(\Delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau), \quad (33)$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, where

$$\begin{aligned} \Delta_{\min} &= \min_{P_{S|X_{\mathcal{M}}}} \left(\max_{\tau \in \Theta} \mathbb{E} \left[\min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d(X_B, y_B) | S, X_S, \theta = \tau] | \theta = \tau \right] \right) \end{aligned} \quad (34)$$

and

$$\begin{aligned} \Delta_{\max} &= \min_{P_{S|X_{\mathcal{M}}}} \max_{\tau \in \Theta} \sum_{A_i \in \mathcal{A}_k} P_{S|\theta}(A_i | \tau) \\ &\quad \times \min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d(X_B, y_B) | S = A_i, \theta = \tau]. \end{aligned} \quad (35)$$

It suffices to take $|\mathcal{U}| \leq 2|\Theta| + 1$.

In (32) and (34), (35), it is readily seen that conditionally deterministic samplers (defined below) attain the minima in Δ_{\min} and Δ_{\max} . In fact, such samplers will be seen to be optimal for every $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$.

For a mapping $w : \mathcal{X}_{\mathcal{M}} \times \mathcal{U} \rightarrow \mathcal{A}_k$, a *deterministic sampler* is specified in terms of a conditional point-mass pmf

$$\begin{aligned} P_{S|X_{\mathcal{M}}U}(s | x_{\mathcal{M}}, u) \\ = \delta_{w(x_{\mathcal{M}}, u)}(s) \\ \triangleq \begin{cases} 1, & s = w(x_{\mathcal{M}}, u) \\ 0, & \text{otherwise, } (x_{\mathcal{M}}, u) \in \mathcal{X}_{\mathcal{M}} \times \mathcal{U}, s \in \mathcal{A}_k. \end{cases} \end{aligned} \quad (36)$$

Theorem 3 is equivalent to

Proposition 1: For a k -MRS, the Bayesian USRDf is

$$R_m(\Delta) = \min_{P_U, \{\delta_w, \{\Delta_\tau, \tau \in \Theta\}\}} \max_{\substack{\tau \in \Theta \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \rho_m^{\mathcal{B}}(\Delta_\tau, P_U, \delta_w, \tau), \quad (37)$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$ with Δ_{\min} and Δ_{\max} as in (32), and the nonBayesian USRDf is

$$R_m(\Delta) = \min_{P_U, \delta_w} \max_{\tau \in \Theta} \rho_m^{n\mathcal{B}}(\Delta, P_U, \delta_w, \tau), \quad (38)$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, with Δ_{\min} and Δ_{\max} as in (34) and (35), respectively. It suffices if $|\mathcal{U}| \leq 2|\Theta| + 1$.

Proof: See Appendix B.

The achievability proof of Theorem 3, by dint of Proposition 1, will use a deterministic sampler based on the minimizing w from (37) or (38).

Example 6: This example compares the USRDfs for a k -MRS and a k -IRS and is an adaptation of Example 2 above (and also of ([4], Example 2)). Consider Example 2 with $q_\tau = 0.5$ for every $\tau \in \Theta$, whereby $P_{X_1 X_2 | \theta = \tau} = P_{X_1 | \theta = \tau} P_{X_2 | \theta = \tau}$. Clearly, $\Theta_2 = \Theta$. For a k -IRS, the Bayesian USRDf is

$$\begin{aligned} R_i(\Delta) &= \min_{\substack{\{\Delta_\tau, \tau \in \Theta\} \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \max_{\tau \in \Theta} \left(h(0.5) - h\left(\frac{\Delta_\tau - p_\tau}{1 - p_\tau}\right) \right) \\ &= h(0.5) - h\left(\frac{\Delta - p}{1 - p}\right) \end{aligned}$$

for $0 \leq \Delta \leq p$, where $p = \mathbb{E}[p_\theta]$, and the nonBayesian USRDf is

$$R_i(\Delta) = h(0.5) - \min_{\tau \in \Theta} h\left(\frac{\Delta - p_\tau}{1 - p_\tau}\right), \quad 0 \leq \Delta \leq \max_{\tau \in \Theta} p_\tau.$$

For a k -MRS, in $\rho_m^{\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau)$ as well as $\rho_m^{n\mathcal{B}}(\delta, P_U, P_{S|X_{\mathcal{M}}U}, \tau)$, $P_U = \text{a point-mass}$ and

$$\begin{aligned} P_{S|X_{\mathcal{M}}U}(s|x_{\mathcal{M}}, u) &= P_{S|X_{\mathcal{M}}}(s|x_{\mathcal{M}}) \\ &= \begin{cases} 1, & s = 1, x_{\mathcal{M}} = 00 \text{ or } 11 \\ 1, & s = 2, x_{\mathcal{M}} = 01 \text{ or } 10 \\ 0, & \text{otherwise} \end{cases} \end{aligned}$$

are uniformly optimal for all $0 \leq \delta \leq p_\tau$ and for all $\tau \in \Theta$. Then, the Bayesian USRDf is

$$R_m(\Delta) = \min_{\substack{\{\Delta_\tau, \tau \in \Theta\} \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \max_{\tau \in \Theta} (h(p_\tau) - h(\Delta_\tau)), \quad 0 \leq \Delta \leq p,$$

and the nonBayesian USRDf is

$$R_m(\Delta) = \max_{\tau \in \Theta} h(p_\tau) - h(\Delta), \quad 0 \leq \Delta \leq \max_{\tau \in \Theta} p_\tau.$$

Clearly, in both the Bayesian and nonBayesian settings $R_m(\Delta) < R_i(\Delta)$. \square

In closing this section, standard properties of the USRDf for the fixed-set sampler, k -IRS and k -MRS in the Bayesian and nonBayesian settings are summarized below, with the proof provided in Appendix C.

Lemma 1: The right-sides of (13), (14), (25), (26), (31) and (33) are finite-valued, decreasing, convex, continuous functions of $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$.

IV. PROOFS

A. Achievability Proofs

Our achievability proofs emphasize the Bayesian setting. Counterpart proofs in the nonBayesian setting use similar sets of ideas, and so we limit ourselves to pointing out only the distinctions between these and their Bayesian brethren. In the Bayesian setting, the achievability proofs successively build upon each other according to increasing complexity of the sampler, and are presented in the order: fixed-set sampler, k -IRS and k -MRS.

A common theme in the achievability proofs for a k -FS, a k -IRS and a k -MRS involves forming estimates $\hat{\tau}_1$ of the underlying τ_1 in Θ_1 , $\hat{\tau}_2$ of τ_2 in Θ_2 and $\hat{\tau}$ of τ in Θ , respectively. The assumed finiteness of Θ enables $\hat{\tau}_1$ or $\hat{\tau}_2$ to be conveyed rate-free to the decoder. Codes for achieving

USRDf at a prescribed distortion level Δ are chosen from among fixed-set sampling rate distortion codes for τ_1 s in Θ_1 or from among IRS codes for τ_2 s in Θ_2 or from among MRS codes for τ s in Θ . Such codes, in the Bayesian setting, correspond to appropriate distortion thresholds that, in effect, average to yield a distortion level Δ ; in the nonBayesian setting, a suitable ‘‘worst-case’’ distortion must not exceed Δ . A chosen code corresponds to an estimate $\hat{\tau}_1$, $\hat{\tau}_2$ or $\hat{\tau}$.

A mainstay of our achievability proofs is the existence of sampling rate distortion codes with fixed-set sampling for a DMMS with known pmf Q .

Lemma 2: Consider a DMMS $\{X_{\mathcal{M}t}\}_{t=1}^\infty$ with known pmf $Q = Q_{X_{\mathcal{M}}}$. Let $A, B \subseteq \mathcal{M}$ be fixed sampling and recovery sets, respectively, and define

$$d_A(x_A, y_B) \triangleq \mathbb{E}[d(X_B, y_B) | X_A = x_A].$$

For every $\epsilon > 0$ and $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, there exists a sampling rate distortion code (f, φ) of rate

$$\frac{1}{n} \log \|f\| \leq \min_{\mathbb{E}_Q[d_A(X_A, Y_B)] \leq \Delta} I_Q(X_A \wedge Y_B) + \epsilon$$

and expected distortion

$$\begin{aligned} \mathbb{E}_Q[d(X_B^n, \varphi(f(X_A^n)))] &= \mathbb{E}_Q[d_A(X_A^n, \varphi(f(X_A^n)))] \\ &\leq \Delta + \epsilon \end{aligned}$$

for all n large enough. Here,

$$\Delta_{\min} = \mathbb{E}[\min_{y_B \in \mathcal{Y}_B} d_A(X_A, y_B)]$$

and

$$\Delta_{\max} = \min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d_A(X_A, y_B)].$$

Proof: The proof of the lemma follows from the achievability proof of [4, Proposition 1] upon replacing the recovery set \mathcal{M} therein by B . \square

Theorem 1: Considering first the Bayesian setting, observe that

$$\begin{aligned} \Delta_{\min} &= \min_{\theta, X_{\mathcal{M}}} \min_{\theta_1, X_A} \min_{Y_B} \mathbb{E}[d(X_B, Y_B)] \\ &= \min_{\theta, X_{\mathcal{M}}} \min_{\theta_1, X_A} \min_{Y_B} \mathbb{E}[\mathbb{E}[d(X_B, Y_B) | X_A, \theta_1]] \\ &= \min_{\theta, X_{\mathcal{M}}} \min_{\theta_1, X_A} \min_{Y_B} \mathbb{E}[d_{\theta_1}(X_A, Y_B)] \quad \text{by (8)} \\ &= \mathbb{E}[\mathbb{E}[\min_{y_B \in \mathcal{Y}_B} d_{\theta_1}(X_A, y_B) | \theta_1]] \end{aligned}$$

and

$$\begin{aligned} \Delta_{\max} &= \min_{\theta, X_{\mathcal{M}}} \min_{\theta_1, X_A} \min_{Y_B} \mathbb{E}[d(X_B, Y_B)] \\ &= \mathbb{E}[\min_{P_{X_A Y_B | \theta_1} = P_{X_A | \theta_1} P_{Y_B | \theta_1}} \min_{P_{X_A Y_B | \theta_1} = P_{X_A | \theta_1} P_{Y_B | \theta_1}} \mathbb{E}[d_{\theta_1}(X_A, Y_B) | \theta_1]] \\ &= \mathbb{E}[\min_{y_B \in \mathcal{Y}_B} \mathbb{E}[d_{\theta_1}(X_A, y_B) | \theta_1]]. \end{aligned}$$

Now, consider a partition Θ_1 of Θ as in Section III. Based on the sampler output X_A^n , the encoder forms an ML estimate of θ_1 as

$$\hat{\tau}_{1,n} = \hat{\tau}_{1,n}(X_A^n) \triangleq \arg \max_{\tau_1 \in \Theta_1} P_{X_A^n | \theta_1}(X_A^n | \tau_1).$$

For each τ_1 in Θ_1 , observe that $\{X_{At}\}_{t=1}^\infty$ is a DMMS with pmf $P_{\tau_1} \triangleq P_{X_A|\theta_1=\tau_1}$. As shown in Appendix D, the sequence of ML estimates

$$\{\widehat{\tau}_{1,n}\}_n \text{ converges to } \tau_1 \text{ in } P_{\tau_1}\text{-probability, } \tau_1 \in \Theta_1, \quad (39)$$

so that for every $\epsilon > 0$ and τ_1 in Θ_1 , there exists an $N_1(\epsilon, \tau_1)$ such that

$$P_{\tau_1}(\widehat{\tau}_{1,n} \neq \tau_1) = P_{\tau_1}(\widehat{\tau}_{1,n}(X_A^n) \neq \tau_1) \leq \frac{\epsilon}{2 d_{\max}}, \quad n \geq N_1(\epsilon, \tau_1),$$

where $d_{\max} = \max_{x_B \in \mathcal{X}_B, y_B \in \mathcal{Y}_B} d(x_B, y_B)$. By the finiteness of Θ_1 , there exists an $N(\epsilon)$ such that simultaneously for all $\tau_1 \in \Theta_1$,

$$P_{\tau_1}(\widehat{\tau}_{1,n} \neq \tau_1) \leq \frac{\epsilon}{2 d_{\max}}, \quad n \geq N(\epsilon)$$

and consequently

$$\begin{aligned} P(\widehat{\tau}_{1,n} \neq \theta_1) &= \sum_{\tau_1 \in \Theta_1} \mu_{\theta_1}(\tau_1) P_{\tau_1}(\widehat{\tau}_{1,n} \neq \tau_1) \\ &\leq \frac{\epsilon}{2 d_{\max}}, \quad n \geq N(\epsilon). \end{aligned} \quad (40)$$

For a fixed $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, let $\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\}$ yield the minimum in (13). For each τ_1 in Θ_1 , for the DMMS $\{X_{\mathcal{M}t}\}_{t=1}^\infty$ with pmf $P_{X_{\mathcal{M}t}|\theta_1=\tau_1}$ and distortion measure d_{τ_1} , there exists by Lemma 2 – with $Q = P_{X_{\mathcal{M}t}|\theta_1=\tau_1}$ and $d_A = d_{\tau_1}$ – a fixed-set sampling rate distortion code $(f_{\tau_1}, \varphi_{\tau_1})$, $f_{\tau_1} : \mathcal{X}_A^n \rightarrow \{1, \dots, J\}$ and $\varphi_{\tau_1} : \{1, \dots, J\} \rightarrow \mathcal{Y}_B^n$ of rate $\frac{1}{n} \log J \leq \max_{\tau_1 \in \Theta_1} \rho_A^B(\Delta_{\tau_1}, \tau_1) + \frac{\epsilon}{2} = R_A(\Delta) + \frac{\epsilon}{2}$ and with expected distortion

$$\mathbb{E}[d_{\tau_1}(X_A^n, \varphi_{\tau_1}(f_{\tau_1}(X_A^n)))|\theta_1 = \tau_1] \leq \Delta_{\tau_1} + \frac{\epsilon}{2}$$

for all $n \geq N_2(\epsilon, \tau_1)$.

A code (f, φ) , with f taking values in $\mathcal{J} \triangleq \{1, \dots, |\Theta_1|\} \times \{1, \dots, J\}$ is constructed as follows. Order (in any manner) the elements of Θ_1 . The encoder f , dictated by the estimate $\widehat{\tau}_{1,n}$, is

$$f(x_A^n) \triangleq (\widehat{\tau}_{1,n}(x_A^n), f_{\widehat{\tau}_{1,n}}(x_A^n)), \quad x_A^n \in \mathcal{X}_A^n.$$

The decoder is

$$\varphi(\widehat{\tau}_{1,n}, j) \triangleq \varphi_{\widehat{\tau}_{1,n}}(j), \quad (\widehat{\tau}_{1,n}, j) \in \mathcal{J}.$$

The rate of the code is

$$\frac{1}{n} \log |\mathcal{J}| = \frac{1}{n} \log |\Theta_1| + \frac{1}{n} \log J \leq R_A(\Delta) + \epsilon, \quad (41)$$

for all n large enough, by the finiteness of Θ_1 .

The code (f, φ) is seen to satisfy

$$\begin{aligned} &\mathbb{E}[d(X_B^n, \varphi(f(X_A^n)))] \\ &\leq \mathbb{E}[\mathbb{1}(\widehat{\tau}_{1,n} = \theta_1) d(X_B^n, \varphi_{\widehat{\tau}_{1,n}}(f_{\widehat{\tau}_{1,n}}(X_A^n)))] \\ &\quad + P(\widehat{\tau}_{1,n} \neq \theta_1) d_{\max} \\ &= \mathbb{E}[\mathbb{1}(\widehat{\tau}_{1,n} = \theta_1) d(X_B^n, \varphi_{\theta_1}(f_{\theta_1}(X_A^n)))] \\ &\quad + P(\widehat{\tau}_{1,n} \neq \theta_1) d_{\max} \\ &\leq \mathbb{E}[d(X_B^n, \varphi_{\theta_1}(f_{\theta_1}(X_A^n)))] + P(\widehat{\tau}_{1,n} \neq \theta_1) d_{\max}. \end{aligned} \quad (42)$$

The first term on the right-side of (42) is

$$\begin{aligned} &\mathbb{E}\left[\frac{1}{n} \sum_{t=1}^n d(X_{Bt}, (\varphi_{\theta_1}(f_{\theta_1}(X_A^n)))_t)\right] \\ &= \mathbb{E}\left[\frac{1}{n} \sum_{t=1}^n \mathbb{E}[d(X_{Bt}, (\varphi_{\theta_1}(f_{\theta_1}(X_A^n)))_t) | X_A^n, \theta_1]\right] \\ &= \mathbb{E}\left[\frac{1}{n} \sum_{t=1}^n \mathbb{E}[d(X_{Bt}, (\varphi_{\theta_1}(f_{\theta_1}(X_A^n)))_t) | X_{At}, \theta_1]\right], \\ &\quad \text{since } P_{X_{\mathcal{M}t}|\theta} = \prod_{t=1}^n P_{X_{\mathcal{M}t}|\theta} \\ &= \mathbb{E}\left[\frac{1}{n} \sum_{t=1}^n \mathbb{E}[d(X_{Bt}, (\varphi_{\theta_1}(f_{\theta_1}(X_A^n)))_t) | X_{At}, \theta_1]\right], \\ &\quad \text{since } \theta \text{ --- } \theta_1 \text{ --- } X_A^n \\ &= \mathbb{E}\left[\frac{1}{n} \sum_{t=1}^n d_{\theta_1}(X_{At}, (\varphi_{\theta_1}(f_{\theta_1}(X_A^n)))_t)\right], \text{ by (8)} \\ &= \mathbb{E}[d_{\theta_1}(X_A^n, \varphi_{\theta_1}(f_{\theta_1}(X_A^n)))] \end{aligned} \quad (43)$$

Combining (42) and (43),

$$\begin{aligned} &\mathbb{E}[d(X_B^n, \varphi(f(X_A^n)))] \\ &\leq \mathbb{E}[d_{\theta_1}(X_A^n, \varphi_{\theta_1}(f_{\theta_1}(X_A^n)))] + P(\widehat{\tau}_{1,n} \neq \theta_1) d_{\max} \\ &\leq \mathbb{E}[\Delta_{\theta_1}] + \epsilon \leq \Delta + \epsilon, \end{aligned} \quad (44)$$

by (40) for all n large enough. Finally, we note that (41) and (44) hold simultaneously for all n large enough.

In the nonBayesian setting, the achievability proof follows by adapting the steps above with the following differences. For each τ_1 in Θ_1 , a fixed-set sampling rate distortion code $(f_{\tau_1}, \varphi_{\tau_1})$ is chosen now with expected distortion $\mathbb{E}[d(X_B^n, \varphi_{\tau_1}(f_{\tau_1}(X_A^n)))|\theta = \tau_1] \leq \Delta + \frac{\epsilon}{2}$ for every τ_1 in $\Lambda(\tau_1)$ and of rate $\frac{1}{n} \log |f_{\tau_1}| \leq R_A(\Delta) + \frac{\epsilon}{2}$, where $R_A(\Delta)$ is the nonBayesian USRDf for a fixed-set sampler. \square

Theorem 2: In the Bayesian setting, for a given $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, consider the $P_S, \{\Delta_{\tau_2}, \tau_2 \in \Theta_2\}$ that attain the (outer) minimum in (25). For the corresponding minimizing $P_{Y_B|S X_S \theta_2}$ in (25) (by way of (23))

$$\begin{aligned} &\max_{\tau_2 \in \Theta_2} \rho_i^B(\Delta_{\tau_2}, P_S, \tau_2) \\ &= \max_{\tau_2 \in \Theta_2} \sum_{A_i \in \mathcal{A}_k} P_S(A_i) I(X_{A_i} \wedge Y_B | S = A_i, \theta_2 = \tau_2) \end{aligned} \quad (45)$$

and let

$$\Delta_{A_i, \tau_2} \triangleq \mathbb{E}[d(X_B, Y_B) | S = A_i, \theta_2 = \tau_2], \quad A_i \in \mathcal{A}_k, \tau_2 \in \Theta_2.$$

The second expression in (45) suggests an achievability scheme using an IRS code (see [4]) governed by θ_2 . Our achievability proof comprises two phases. In the first phase an estimate $\widehat{\tau}_2$ of θ_2 is formed based on the output of a k -IRS that chooses each A_i in \mathcal{A}_k repeatedly for N time instants. The second phase, of length n , entails choosing each $S_t = A_i$ repeatedly for $\approx n P_S(A_i)$ time instants and an IRS code governed by $\widehat{\tau}_2$ of expected distortion

$$\sum_i P_S(A_i) \Delta_{A_i, \widehat{\tau}_2}$$

is applied to the output of the sampler. This predetermined selection of sampling sets obviates the need for the decoder to be additionally informed.

Denote $|\mathcal{A}_k|$ by $M_k = \binom{m}{k}$. Fix $\epsilon > 0$ and $0 < \epsilon' < \epsilon$. In the first phase, a k -IRS is chosen to sample each $A_i \in \mathcal{A}_k$ over disjoint time-sets μ_i of length N . The union of the time-sets μ_i , $i \in \mathcal{M}_k \triangleq \{1, \dots, M_k\}$ is denoted by $\mu \triangleq \{1, \dots, M_k N\}$. Based on the sampler output, an ML estimate $\widehat{\tau}_{2,N} = \widehat{\tau}_{2,N}(S^\mu, X_S^\mu)$ of θ_2 is formed with

$$P(\widehat{\tau}_{2,N} \neq \theta_2) \leq \frac{\epsilon'}{2d_{\max}}, \quad (46)$$

for $N \geq N_{\epsilon'}$, say.

In the second phase, we denote the next set of n time instants, i.e., $\{M_k N + 1, \dots, M_k N + n\}$ simply by $\nu \triangleq \{1, \dots, n\}$. Further, for each i in \mathcal{M}_k , define the time-sets $\nu_{A_i} \subset \nu$, made up of consecutive time instants, as

$$\nu_{A_i} = \left\{ t : \left[n \sum_{j=1}^{i-1} P_S(A_j) \right] + 1 \leq t \leq \left[n \sum_{j=1}^i P_S(A_j) \right] \right\},$$

and note that the union of ν_{A_i} s is ν , and

$$\left| \frac{|\nu_{A_i}|}{n} - P_S(A_i) \right| \leq \frac{1}{n}, \quad i \in \mathcal{M}_k.$$

In this phase, the k -IRS is now chosen (deterministically) as follows:

$$S_t = s_t = A_i, \quad t \in \nu_{A_i}, \quad i \in \mathcal{M}_k.$$

For each DMMS $\{X_{\mathcal{M}t}\}_{t=1}^\infty$ with pmf $P_{X_{\mathcal{M}t}|\theta_2=\tau_2}$, $\tau_2 \in \Theta_2$, and for each A_i in \mathcal{A}_k and its corresponding distortion measure d_{τ_2} , there exists by Lemma 2 – with $Q = P_{X_{\mathcal{M}t}|\theta_2=\tau_2}$ and $d_A = d_{\tau_2}$ – a fixed-set sampling rate distortion code $(f_{A_i}^{\tau_2}, \varphi_{A_i}^{\tau_2})$, $f_{A_i}^{\tau_2} : \mathcal{X}_{A_i}^{\nu_{A_i}} \rightarrow \{1, \dots, J_{A_i}^{\tau_2}\}$ and $\varphi_{A_i}^{\tau_2} : \{1, \dots, J_{A_i}^{\tau_2}\} \rightarrow \mathcal{Y}_B^{\nu_{A_i}}$ of rate $\frac{1}{|\nu_{A_i}|} \log J_{A_i}^{\tau_2} \leq I(X_{A_i} \wedge Y_B | S = A_i, \theta_2 = \tau_2) + \frac{\epsilon'}{4}$ (cf. (45)) and with

$$\mathbb{E} \left[d_{\tau_2}(X_{A_i}^{\nu_{A_i}}, \varphi_{A_i}^{\tau_2}(f_{A_i}^{\tau_2}(X_{A_i}^{\nu_{A_i}}))) | \theta_2 = \tau_2 \right] \leq \Delta_{A_i, \tau_2} + \frac{\epsilon'}{2},$$

for all $|\nu_{A_i}| \geq N_{A_i}(\epsilon', \tau_2)$. Note that

$$\sum_{\tau_2 \in \Theta_2} \mu_{\theta_2}(\tau_2) \sum_{i=1}^{M_k} P_S(A_i) \Delta_{A_i, \tau_2} \leq \Delta$$

and

$$\sum_{i=1}^{M_k} P_S(A_i) I(X_{A_i} \wedge Y_B | S = A_i, \theta_2 = \tau_2) \leq R_i(\Delta)$$

for every τ_2 in Θ_2 .

Consider a (composite) code (f, φ) as follows. Denote $n' \triangleq |\mu| + |\nu| = M_k N + n$, and the encoder f consisting of a concatenation of encoders is defined by

$$f(s^{n'}, x^{n'}) \triangleq (\widehat{\tau}_{2,N}, f_{A_1}^{\widehat{\tau}_{2,N}}(x_{A_1}^{\nu_{A_1}}), \dots, f_{A_{M_k}}^{\widehat{\tau}_{2,N}}(x_{A_{M_k}}^{\nu_{A_{M_k}}})) .$$

The decoder φ , which is aware of the predetermined sequence of sampling sets, is defined by

$$\begin{aligned} \varphi(s^{n'}, \widehat{\tau}_{2,N}, j_1, \dots, j_{M_k}) \\ &= \varphi(\widehat{\tau}_{2,N}, j_1, \dots, j_{M_k}) \\ &\triangleq \underbrace{(y_1, \dots, y_1)}_{\text{first phase}}, \underbrace{(\varphi_{A_1}^{\widehat{\tau}_{2,N}}(j_1), \dots, \varphi_{A_{M_k}}^{\widehat{\tau}_{2,N}}(j_{M_k}))}_{\text{second phase}}, \end{aligned}$$

for each encoder output $(\widehat{\tau}_{2,N}, j_1, \dots, j_{M_k})$, where $y_1 \in \mathcal{Y}_{\mathcal{M}}$ is an arbitrary symbol. Clearly, $|\Theta_2| \times \max_{\tau_2 \in \Theta_2} \prod_{i=1}^{M_k} J_{A_i}^{\tau_2}$ indices would suffice to describe all possible encoder outputs.

The rate of the code is

$$\begin{aligned} &\frac{1}{n'} \log |\Theta_2| + \max_{\tau_2 \in \Theta_2} \frac{1}{n'} \sum_{i=1}^{M_k} \log J_{A_i}^{\tau_2} \\ &\leq \max_{\tau_2 \in \Theta_2} \sum_{i=1}^{M_k} \frac{|\nu_{A_i}|}{n} \frac{1}{|\nu_{A_i}|} \log J_{A_i}^{\tau_2} + \frac{1}{n'} \log |\Theta_2| \\ &\leq \max_{\tau_2 \in \Theta_2} \sum_{i=1}^{M_k} \left(P_S(A_i) + \frac{1}{n} \right) \\ &\quad \times \left(I(X_{A_i} \wedge Y_B | S = A_i, \theta_2 = \tau_2) + \frac{\epsilon'}{4} \right) \\ &\quad + \frac{1}{n'} \log |\Theta_2| \\ &\leq \max_{\tau_2 \in \Theta_2} \sum_{i=1}^{M_k} P_S(A_i) I(X_{A_i} \wedge Y_B | S = A_i, \theta_2 = \tau_2) + \epsilon' \\ &< R_i(\Delta) + \epsilon, \end{aligned} \quad (47)$$

where the previous inequality holds for all n large enough.

Denoting the output of the decoder by $Y_B^{n'} \triangleq \varphi(f(s^{n'}, X_S^{n'}))$

$$\begin{aligned} &\mathbb{E}[d(X_B^{n'}, Y_B^{n'})] \\ &= \frac{1}{n'} \mathbb{E} \left[\sum_{t \in \mu} d(X_{Bt}, Y_{Bt}) + \sum_{t \in \nu} (\mathbb{1}(\widehat{\tau}_{2,N} \neq \theta_2) d(X_{Bt}, Y_{Bt}) \right. \\ &\quad \left. + \mathbb{1}(\widehat{\tau}_{2,N} = \theta_2) d(X_{Bt}, Y_{Bt})) \right]. \end{aligned} \quad (48)$$

The first two terms on the right-side of (48) are

$$\begin{aligned} &\mathbb{E} \left[\frac{1}{n'} \sum_{t \in \mu} d(X_{Bt}, Y_{Bt}) + \frac{\mathbb{1}(\widehat{\tau}_{2,N} \neq \theta_2)}{n'} \sum_{t \in \nu} d(X_{Bt}, Y_{Bt}) \right] \\ &\leq \frac{M_k N d_{\max}}{n'} + \frac{\epsilon'}{2}, \end{aligned} \quad (49)$$

by (46) for N large enough, and the last term on the right-side of (48) is

$$\begin{aligned} &\mathbb{E} \left[\frac{\mathbb{1}(\widehat{\tau}_{2,N} = \theta_2)}{n'} \sum_{t \in \nu} d(X_{Bt}, Y_{Bt}) \right] \\ &\leq \sum_{i=1}^{M_k} \frac{|\nu_{A_i}|}{n} \mathbb{E} \left[\mathbb{1}(\widehat{\tau}_{2,N} = \theta_2) d(X_B^{\nu_{A_i}}, \varphi_{A_i}^{\widehat{\tau}_{2,N}}(f_{A_i}^{\widehat{\tau}_{2,N}}(X_{A_i}^{\nu_{A_i}}))) \right] \\ &\leq \sum_{i=1}^{M_k} \frac{|\nu_{A_i}|}{n} \mathbb{E} [d(X_B^{\nu_{A_i}}, \varphi_{A_i}^{\theta_2}(f_{A_i}^{\theta_2}(X_{A_i}^{\nu_{A_i}})))] \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^{M_k} \frac{|v_{A_i}|}{n} \mathbb{E}[d_{\theta_2}(X_A^{v_{A_i}}, \varphi_{A_i}^{\theta_2}(f_{A_i}^{\theta_2}(X_{A_i}^{v_{A_i}})))] \\
&\leq \sum_{i=1}^{M_k} \left(P_S(A_i) + \frac{1}{n} \right) \mathbb{E}[\Delta_{A_i, \theta_2} + \frac{\epsilon'}{2}] \\
&\leq \Delta + \frac{\epsilon'}{2} + \frac{1}{n} \sum_{i=1}^{M_k} \mathbb{E}[\Delta_{A_i, \theta_2}] + \frac{M_k \epsilon'}{2}. \tag{50}
\end{aligned}$$

From (48)-(50), we have

$$\mathbb{E}[d(X_B^n, Y_B^n)] \leq \Delta + \epsilon, \tag{51}$$

for n and N large enough. Finally, we note that (47) and (51) hold simultaneously for all n and N large enough.

The Corollary is immediate by the choice of codes with “uninformed” decoder in the proof above.

For the nonBayesian setting, achievability follows by adapting the proof above in a manner similar to that for a k -FS in Theorem 1. \square

Theorem 3: The achievability proof relies on the deterministic sampler justified by Proposition 1. In the Bayesian setting, for a given $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$, let P_U , $P_{S|X_{\mathcal{M}}U} = \delta_w$, $\{\Delta_\tau, \tau \in \Theta\}$ attain the minimum in (37). For the corresponding minimizing $P_{Y_B|S X_S U \theta}$ in (29), the right-side of (37) is

$$\begin{aligned}
&\max_{\tau \in \Theta} \rho_m^{\mathcal{B}}(\Delta_\tau, P_U, \delta_w, \tau) \\
&= \max_{\tau \in \Theta} \sum_{u \in \mathcal{U}} P_U(u) I(X_S \wedge Y_B | S, U = u, \theta = \tau) \tag{52}
\end{aligned}$$

and we set

$$\begin{aligned}
\Delta_{A_i, u, \tau} &\triangleq \mathbb{E}[d(X_B, Y_B) | S = A_i, U = u, \theta = \tau], \\
&A_i \in \mathcal{A}_k, \quad \tau \in \Theta, \quad u \in \mathcal{U}.
\end{aligned}$$

Our achievability proof uses a k -MRS in two distinct modes. First, a deterministic k -MRS is chosen so as to form an estimate $\hat{\tau}$ of θ from the sampler output. Next, for each $U = u$, a suitable deterministic k -MRS is chosen in accordance with $w(x_{\mathcal{M}}, u)$, and an MRS code (see [4]) governed by $\hat{\tau}$ of expected distortion

$$\tilde{\leq} \sum_{A_i} P_{S|U\theta}(A_i | u, \hat{\tau}) \Delta_{A_i, u, \hat{\tau}}$$

is applied to the sampler output. Concatenation of such codes corresponding to various $u \in \mathcal{U}$ yields, in effect, time-sharing that serves to achieve (52). To simplify the notation, the conditioning on $U = u$ will be suppressed except when needed.

Fix $\epsilon > 0$ and $0 < \epsilon' < \epsilon$.

(i) We devise a deterministic k -MRS on a time-set μ , based on whose output an estimate $\hat{\tau}_N = \hat{\tau}_N(S^\mu, X_S^\mu) = \hat{\tau}_N(S^\mu)$ of θ is formed with

$$P(\hat{\tau}_N \neq \theta) \leq \frac{\epsilon'}{4d_{\max}}, \tag{53}$$

for $N \geq N_{\epsilon'}$. The estimate $\hat{\tau}_N$ is formed from *only the sampling sequence* S^μ and thus is available to the encoder

as well as the decoder. The k -MRS is chosen on the time-set μ , to signal the occurrences of each $x \in \mathcal{X}_{\mathcal{M}}$ to the encoder and decoder through S^μ above; for each $x \in \mathcal{X}_{\mathcal{M}}$, a distinct $A \in \mathcal{A}_k$ is chosen. If $|\mathcal{A}_k| \geq |\mathcal{X}_{\mathcal{M}}|$, a trivial one-to-one mapping from $\mathcal{X}_{\mathcal{M}}$ to \mathcal{A}_k enables S^μ to determine $X_{\mathcal{M}}^\mu$, where S^μ is of length N , say. Then $\hat{\tau}_N$ is taken to be the ML estimate of θ based on $X_{\mathcal{M}}^\mu$, which satisfies (53).

When $|\mathcal{A}_k| < |\mathcal{X}_{\mathcal{M}}|$, a k -MRS is chosen attuned variously to disjoint subsets of $\mathcal{X}_{\mathcal{M}}$, of size $|\mathcal{A}_k| - 1$, on corresponding disjoint time-sets μ_l of length N , $l = 1, \dots, \lceil \frac{|\mathcal{X}_{\mathcal{M}}|}{|\mathcal{A}_k| - 1} \rceil$, as follows. In each μ_l , the k -MRS signals the occurrence (or not) of $X_{\mathcal{M}l} = x$ in the l^{th} -subset of $\mathcal{X}_{\mathcal{M}}$ in a (deterministic) manner by choosing $|\mathcal{A}_k| - 1$ distinct sampling sets in \mathcal{A}_k ; the nonoccurrence of symbols from this l^{th} -subset of $\mathcal{X}_{\mathcal{M}}$ is indicated by the remaining (dummy) sampling set in \mathcal{A}_k . We denote $\bigcup_l \mu_l$ by μ . Finally, $\hat{\tau}_N$ is taken as the ML estimate of θ based on the sampling sequence S^μ of length $\lceil \frac{|\mathcal{X}_{\mathcal{M}}|}{|\mathcal{A}_k| - 1} \rceil N = N'$, say.

(ii) Next, for each $U = u$, a k -MRS is chosen according to $P_{S|X_{\mathcal{M}}, U=u} = \delta_w(\cdot, u)$ for n time instants. Then, for a DMMS $\{X_{\mathcal{M}t}\}_{t=1}^\infty$ with pmf $P_{X_{\mathcal{M}t}|\theta=\hat{\tau}_N}$ an MRS code comprising a concatenation of fixed-set sampling rate distortion codes corresponding to the A_i s in \mathcal{A}_k is applied to the sampler output.

Denote the set of n time instants $\{N'+1, \dots, N'+n\}$ simply by $\gamma \triangleq \{1, \dots, n\}$. Define time-sets $\gamma_{S^n}(A_i) \triangleq \{t : 1 \leq t \leq n, S_t = A_i\}$, $i \in \mathcal{M}_k$, and note that $\gamma_{S^n}(A_i)$ s cover γ , i.e.,

$$\gamma = \bigcup_{A_i \in \mathcal{A}_k} \gamma_{S^n}(A_i).$$

Denote the set of the first $\max\{\lceil (nP_{S|\theta}(A_i|\hat{\tau}_N)) - \epsilon' \rceil, 0\}$ time instants in each $\gamma_{S^n}(A_i)$ by v_{A_i} (suppressing the dependence on $\hat{\tau}_N$). Defining the (typical) set for each τ in Θ

$$\mathcal{T}^{(n)}(\epsilon', \tau) \triangleq \left\{ s^n \in \mathcal{A}_k^n : \left| \frac{|\gamma_{S^n}(A_i)|}{n} - P_{S|\theta}(A_i|\tau) \right| \leq \epsilon', \right. \\ \left. i \in \mathcal{M}_k \right\},$$

we have that

$$\begin{aligned}
P(S^\gamma \notin \mathcal{T}^{(n)}(\epsilon', \hat{\tau}_N)) &= P(S^\gamma \notin \mathcal{T}^{(n)}(\epsilon', \hat{\tau}_N), \hat{\tau}_N = \theta) \\
&\quad + P(S^\gamma \notin \mathcal{T}^{(n)}(\epsilon', \hat{\tau}_N), \hat{\tau}_N \neq \theta) \\
&\leq \frac{\epsilon'}{2d_{\max}} \tag{54}
\end{aligned}$$

for all n large enough.

By Lemma 2, for *each* DMMS $\{X_{\mathcal{M}t}\}_{t=1}^\infty$ with pmf $P_{X_{\mathcal{M}t}|S=A_i, \theta=\tau}$, $i \in \mathcal{M}_k$, $\tau \in \Theta$, there exists a code $(f_{A_i}^\tau, \varphi_{A_i}^\tau)$, $f_{A_i}^\tau : \mathcal{X}_{A_i}^{v_{A_i}} \rightarrow \{1, \dots, J_{A_i}^\tau\}$ and $\varphi_{A_i}^\tau : \{1, \dots, J_{A_i}^\tau\} \rightarrow \mathcal{Y}_B^{v_{A_i}}$ of rate

$$\frac{1}{|v_{A_i}|} \log J_{A_i}^\tau \leq I(X_{A_i} \wedge Y_B | S = A_i, \theta = \tau) + \frac{\epsilon'}{2} \tag{55}$$

and with

$$\begin{aligned}
&\mathbb{E}\left[d(X_B^{v_{A_i}}, \varphi_{A_i}^\tau(f_{A_i}^\tau(X_{A_i}^{v_{A_i}}))) | S^{v_{A_i}} = A_i^{v_{A_i}}, \theta = \tau \right] \\
&\leq \Delta_{A_i, \tau} + \frac{\epsilon'}{4} \tag{56}
\end{aligned}$$

for all $|v_{A_i}| \geq N_{A_i}(\epsilon', \tau)$. Such codes are considered for each $U = u$.

Consider a (composite) code (f, φ) as follows. Denoting $N' + n$ by n' , an encoder f consisting of a concatenation of encoders is defined as

$$f(s^{n'}, x_s^{n'}) \triangleq \begin{cases} (f_{A_1}^{\widehat{\tau}_N}(x_{A_1}^{v_{A_1}}), \dots, f_{A_{M_k}}^{\widehat{\tau}_N}(x_{A_{M_k}}^{v_{A_{M_k}}})) & s^{n'} \in \mathcal{T}^{(n)}(\epsilon', \widehat{\tau}_N) \\ (1, \dots, 1) & s^{n'} \notin \mathcal{T}^{(n)}(\epsilon', \widehat{\tau}_N). \end{cases}$$

For $t = 1, \dots, n'$, and each encoder output (j_1, \dots, j_{M_k}) , the decoder φ , which can recover the estimate $\widehat{\tau}_N$ from its knowledge of the sampling sequence $S^{n'} = s^{n'}$, is given by

$$\left(\varphi(s^{n'}, j_1, \dots, j_{M_k}) \right)_t \triangleq \begin{cases} \left(\varphi_{A_i}^{\widehat{\tau}_N}(j_i) \right)_t & s^{n'} \in \mathcal{T}^{(n)}(\epsilon', \widehat{\tau}_N) \text{ and } t \in v_{A_i}, i \in \mathcal{M}_k \\ y_1 & \text{otherwise,} \end{cases}$$

where y_1 is a fixed but arbitrary symbol in $\mathcal{Y}_{\mathcal{M}}$.

Finally, for N and n large enough, the codes (f, φ) corresponding to each $U = u$ are concatenated so as to effect the time-sharing prescribed by P_U , in a standard manner. It is shown in Appendix A that the rate of the resulting code is

$$\begin{aligned} & \lesssim \left(\max_{\tau \in \Theta} \sum_{u \in \mathcal{U}} P_U(u) \sum_{A_i \in \mathcal{A}_k} P_{S|U\theta}(A_i|u, \tau) \right. \\ & \quad \left. I(X_{A_i} \wedge Y_B | S = A_i, U = u, \theta = \tau) \right) + \epsilon' \\ & \lesssim R_m(\Delta) + \epsilon, \end{aligned} \quad (57)$$

using (55) and the expected distortion is

$$\begin{aligned} & \lesssim \mathbb{E}[\Delta_{S,U,\theta}] + \epsilon \\ & \lesssim \Delta + \epsilon, \end{aligned} \quad (58)$$

from (53), (54), (56) and the definition of $\Delta_{A_i,u,\tau}$. \square

B. Converse Proof

In contrast with the achievability proofs, we present a unified converse proof for Theorems 3, 2 and 1 according to successive weakening of the sampler, viz. k -MRS, k -IRS and fixed-set sampler. We begin with the technical Lemma 3 that is used subsequently in the converse proof.

Lemma 3: Let finite-valued rvs C, D^n, E^n, F^n , be such that (D_t, E_t) , $t = 1, \dots, n$, are conditionally mutually independent given C , i.e.,

$$P_{D^n E^n | C} = \prod_{t=1}^n P_{D_t E_t | C} \quad (59)$$

and satisfy

$$C, D^n \text{ --- } E^n \text{ --- } F^n. \quad (60)$$

For any function $g(C)$ of C , such that

$$E^n \text{ --- } g(C) \text{ --- } C \text{ and } P_{E^n | g(C)} = \prod_{t=1}^n P_{E_t | g(C)}, \quad (61)$$

it holds that

$$C, D_t \text{ --- } g(C), E_t \text{ --- } F_t, \quad t = 1, \dots, n. \quad (62)$$

Proof: First, from (60), we have

$$\begin{aligned} 0 &= I(C, D^n \wedge F^n | E^n) \\ &= I(C \wedge F^n | E^n) + I(D^n \wedge F^n | E^n, C) \\ &= I(C, g(C) \wedge F^n | E^n) + I(D^n \wedge F^n | E^n, C) \\ &\geq I(C \wedge F^n | E^n, g(C)) + I(D^n \wedge F^n | E^n, C). \end{aligned} \quad (63)$$

Now, the second term on the right-side of (63) is

$$\begin{aligned} 0 &= I(D^n \wedge F^n | E^n, C) \\ &= H(D^n | E^n, C) - H(D^n | E^n, F^n, C) \\ &= \sum_{t=1}^n (H(D_t | E_t, C) - H(D_t | D^{t-1}, E^n, F^n, C)), \text{ by (59)} \\ &\geq \sum_{t=1}^n (H(D_t | E_t, C) - H(D_t | E_t, F_t, C)) \\ &= \sum_{t=1}^n I(D_t \wedge F_t | E_t, C). \end{aligned} \quad (64)$$

Next, the first part of (61) along with (63) implies that

$$\begin{aligned} 0 &= I(C \wedge E^n | g(C)) + I(C \wedge F^n | E^n, g(C)) \\ &= I(C \wedge E^n, F^n | g(C)), \end{aligned}$$

and hence

$$I(C \wedge E_t, F_t | g(C)) = 0, \quad t = 1, \dots, n. \quad (65)$$

Now, by (64) and (65), for $t = 1, \dots, n$,

$$\begin{aligned} & I(C, D_t \wedge F_t | E_t, g(C)) \\ &= I(C \wedge F_t | E_t, g(C)) + I(D_t \wedge F_t | E_t, C) = 0, \end{aligned}$$

which is the claim (62). \square

Converse: In the Bayesian setting, we provide first a converse proof for Theorem 3, which is then refashioned to give converse proofs for Theorems 2 and 1.

Let $(\{P_{S_t | X_{\mathcal{M}t}} = P_{S_t | X_{\mathcal{M}t}}\}_{t=1}^{\infty}, f, \varphi)$ be an n -length k -MRS block code of rate R and with decoder output $Y_B^n = \varphi(S^n, f(S^n, X_S^n))$ satisfying $\mathbb{E}[d(X_B^n, Y_B^n)] \leq \Delta$. The hypothesis of Lemma 3 is met with $C = \theta$, $D^n = X_{\mathcal{M}}^n$, $E^n = (S^n, X_S^n)$, $F^n = Y_B^n$ and $g(\theta) = \theta$, since

$$\begin{aligned} P_{X_{\mathcal{M}}^n S^n | \theta} &= P_{X_{\mathcal{M}}^n | \theta} P_{S^n | X_{\mathcal{M}}^n} \\ &= \prod_{t=1}^n P_{X_{\mathcal{M}t} | \theta} P_{S_t | X_{\mathcal{M}t}} = \prod_{t=1}^n P_{X_{\mathcal{M}t} S_t | \theta}, \end{aligned} \quad (66)$$

while

$$\theta, X_{\mathcal{M}}^n \text{ --- } S^n, X_S^n \text{ --- } Y_B^n$$

holds by code construction. Also, (66) implies, upon summing over all realizations of $X_{S^c}^n$, that

$$P_{S^n X_S^n | \theta} = \prod_{t=1}^n P_{S_t X_{S_t} | \theta}. \quad (67)$$

Then the claim of the lemma implies that

$$\theta, X_{\mathcal{M}t} \text{ --- } \theta, S_t, X_{S_t} \text{ --- } Y_{Bt}, \quad t = 1, \dots, n. \quad (68)$$

Let Δ_τ denote $\mathbb{E}[d(X_B^n, Y_B^n)|\theta = \tau] = \frac{1}{n} \sum_{t=1}^n \mathbb{E}[d(X_{Bt}, Y_{Bt})|\theta = \tau]$ for each τ in Θ and note that $\mathbb{E}[\Delta_\theta] \leq \Delta$. For every τ in Θ , the following holds:

$$\begin{aligned} R &= \frac{1}{n} \log \|f\| \geq \frac{1}{n} H(f(S^n, X_S^n)|\theta = \tau) \\ &\geq \frac{1}{n} H(f(S^n, X_S^n)|S^n, \theta = \tau) \\ &\geq \frac{1}{n} H(\varphi(S^n, f(S^n, X_S^n))|S^n, \theta = \tau) \\ &= \frac{1}{n} H(Y_B^n|S^n, \theta = \tau) \\ &= \frac{1}{n} I(X_S^n \wedge Y_B^n|S^n, \theta = \tau) \\ &= \frac{1}{n} \sum_{t=1}^n \left(H(X_{S_t}|S^n, X_S^{t-1}, \theta = \tau) \right. \\ &\quad \left. - H(X_{S_t}|S^n, X_S^{t-1}, Y_B^n, \theta = \tau) \right) \\ &\geq \frac{1}{n} \sum_{t=1}^n \left(H(X_{S_t}|S^n, X_S^{t-1}, \theta = \tau) \right. \\ &\quad \left. - H(X_{S_t}|S_t, Y_{Bt}, \theta = \tau) \right) \\ &= \frac{1}{n} \sum_{t=1}^n \left(H(X_{S_t}|S_t, \theta = \tau) - H(X_{S_t}|S_t, Y_{Bt}, \theta = \tau) \right), \\ &\hspace{15em} \text{by (67)} \\ &= \frac{1}{n} \sum_{t=1}^n I(X_{S_t} \wedge Y_{Bt}|S_t, \theta = \tau). \end{aligned} \quad (69)$$

By (68),

$$\left(\left(\frac{1}{n} \sum_{t=1}^n \mathbb{E}[d(X_{Bt}, Y_{Bt})|\theta = \tau], \right. \right. \\ \left. \left. \frac{1}{n} \sum_{t=1}^n I(X_{S_t} \wedge Y_{Bt}|S_t, \theta = \tau) \right), \tau \in \Theta \right)$$

lies in the convex hull of

$$\mathcal{C} \triangleq \left\{ \left(\mathbb{E}[d(X_B, Y_B)|\theta = \tau], I(X_S \wedge Y_B|S, \theta = \tau) \right), \tau \in \Theta \right\} \\ : P_{\theta X_{\mathcal{M}} S Y_B} = \mu_\theta P_{X_{\mathcal{M}}|\theta} P_{S|X_{\mathcal{M}}} P_{Y_B|S X_S \theta} \} \subset \mathbb{R}^{2|\Theta|}.$$

By the Carathéodory Theorem [5], every point in the convex hull of \mathcal{C} can be represented as a convex combination of at most $2|\Theta| + 1$ elements in \mathcal{C} . The corresponding pmfs are indexed by the values of a rv U with

$$P_{U\theta X_{\mathcal{M}} S Y_B} = P_U \mu_\theta P_{X_{\mathcal{M}}|\theta} P_{S|X_{\mathcal{M}} U} P_{Y_B|S X_S \theta U}, \quad (70)$$

where the pmf of U has support of size $\leq 2|\Theta| + 1$. Then, in a standard manner, (69) leads to

$$R \geq \min_{\substack{P_{Y_B|S X_S U, \theta = \tau} \\ \mathbb{E}[d(X_B, Y_B)|\theta = \tau] \leq \Delta_\tau}} I(X_S \wedge Y_B|S, U, \theta = \tau) \quad (71)$$

$$= \rho_m^{\mathcal{B}}(\Delta_\tau, P_U, P_{S|X_{\mathcal{M}} U}, \tau). \quad (72)$$

Now, (72) holds for every $\tau \in \Theta$, and hence

$$\begin{aligned} R &\geq \max_{\tau \in \Theta} \rho_m^{\mathcal{B}}(\Delta_\tau, P_U, P_{S|X_{\mathcal{M}} U}, \tau) \\ &\geq \min_{\substack{P_U, P_{S|X_{\mathcal{M}} U, \{\Delta_\tau, \tau \in \Theta\}} \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \max_{\tau \in \Theta} \rho_m^{\mathcal{B}}(\Delta_\tau, P_U, P_{S|X_{\mathcal{M}} U}, \tau) \\ &= R_m(\Delta) \end{aligned} \quad (73)$$

for $\Delta \geq \Delta_{\min}$.

Turning next to Theorems 2 and 1, an n -length k -IRS code or a fixed-set sampling block code can be viewed as restrictions of a k -MRS code. Specifically, in Theorem 2, for a k -IRS code of rate R with P_{S_t} , $g(\theta) = \theta_2$ instead of $P_{S_t|X_{\mathcal{M}t}}$, $g(\theta) = \theta$ (for a k -MRS), the hypothesis of Lemma 3 holds. Denote $\mathbb{E}[d(X_B^n, Y_B^n)|\theta_2 = \tau_2]$ by Δ_{τ_2} , $\tau_2 \in \Theta_2$. Then, the pmfs in (70) satisfy

$$P_{U\theta X_{\mathcal{M}} S Y_B} = P_U \mu_\theta P_{X_{\mathcal{M}}|\theta} P_{S|U} P_{Y_B|S X_S \theta U}. \quad (74)$$

The counterpart of (71) is

$$\begin{aligned} R &\geq \min_{\substack{P_{Y_B|S X_S U, \theta_2 = \tau_2} \\ \mathbb{E}[d(X_B, Y_B)|\theta_2 = \tau_2] \leq \Delta_{\tau_2}}} I(X_S \wedge Y_B|S, U, \theta_2 = \tau_2) \\ &= \min_{\substack{P_{Y_B|S X_S U, \theta_2 = \tau_2} \\ \mathbb{E}[d(X_B, Y_B)|\theta_2 = \tau_2] \leq \Delta_{\tau_2}}} \sum_{A, u} P_S(A) P_U|S(u|A) \\ &\quad \times I(X_A \wedge Y_B|S = A, U = u, \theta_2 = \tau_2), \end{aligned}$$

noting from (74) that $P_U|S, \theta_2 = P_U|S$. Using the convexity of the mutual information terms above with respect to $P_{Y_B|S X_S U \theta_2}$, we get

$$\begin{aligned} R &\geq \min_{\substack{P_{Y_B|S X_S U, \theta_2 = \tau_2} \\ \mathbb{E}[d(X_B, Y_B)|\theta_2 = \tau_2] \leq \Delta_{\tau_2}}} \sum_A P_S(A) I(X_A \wedge Y_B|S = A, \theta_2 = \tau_2) \\ &= \rho_1^{\mathcal{B}}(\Delta_{\tau_2}, P_S, \tau_2). \end{aligned} \quad (75)$$

Since (75) holds for every $\tau_2 \in \Theta_2$

$$\begin{aligned} R &\geq \max_{\tau_2 \in \Theta_2} \rho_1^{\mathcal{B}}(\Delta_{\tau_2}, P_S, \tau_2) \\ &\geq \min_{\substack{P_S, \{\Delta_{\tau_2}, \tau_2 \in \Theta_2\} \\ \mathbb{E}[\Delta_{\theta_2}] \leq \Delta}} \max_{\tau_2 \in \Theta_2} \rho_1^{\mathcal{B}}(\Delta_{\tau_2}, P_S, \tau_2) \\ &= R_1(\Delta), \end{aligned}$$

i.e., $R \geq R_1(\Delta)$, $\Delta \geq \Delta_{\min}$, completing the converse proof of Theorem 2.

In a manner analogous to a k -IRS, in Theorem 1 for a fixed-set sampler the hypothesis of Lemma 3 holds with $P_{S_t} = \mathbb{1}(S_t = A)$, $g(\theta) = \theta_1$. Defining $\Delta_{\tau_1} \triangleq \mathbb{E}[d(X_B^n, Y_B^n)|\theta_1 = \tau_1]$, $\tau_1 \in \Theta_1$, the counterpart of the right-side of (73) reduces to $\max_{\tau_1 \in \Theta_1} \rho_A^{\mathcal{B}}(\Delta_{\tau_1}, \tau_1)$. It then follows that

$$R \geq \min_{\substack{\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\} \\ \mathbb{E}[\Delta_{\theta_1}] \leq \Delta}} \max_{\tau_1 \in \Theta_1} \rho_A^{\mathcal{B}}(\Delta_{\tau_1}, \tau_1), \quad \Delta \geq \Delta_{\min}$$

providing the converse proof for Theorem 1.

In the nonBayesian setting, the analog of Lemma 3 is obtained similarly with $C = c$, $g(C) = g(c)$, and (59)–(62) expressed in terms of appropriate conditional pmfs. The converse proofs for a k -MRS, k -IRS and k -FS are obtained as above but by excluding the outer minimizations over $\{\Delta_\tau, \tau \in \Theta\}$, $\{\Delta_{\tau_2}, \tau_2 \in \Theta_2\}$ and $\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\}$, respectively. \square

V. DISCUSSION

Our formulation of universality requires optimum sampling rate distortion performance when the “true” underlying pmf of the DMMS belongs to a finite family $\mathcal{P} = \{P_{X_{\mathcal{M}}|\theta=\tau}, \tau \in \Theta\}$. The assumed finiteness of Θ affords two benefits in addition to mathematical ease: (i) simple proofs of estimator consistency uniformly over Θ_1 , Θ_2 or Θ ; and (ii) rate-free conveyance of corresponding estimates $\widehat{\tau}_1$, $\widehat{\tau}_2$ or $\widehat{\tau}$ to the decoder. General extensions to the case when Θ is an infinite set (countable or uncountable) remain open.

Unlike for a k -IRS, the assumption in a k -MRS that the decoder is informed of the sampling sequence S^n plays an important role. Specifically, embedded information regarding $X_{\mathcal{M}}^n$ is conveyed implicitly to the decoder through S^n . Also, as a side-benefit, the decoder can replicate the estimate of θ formed by the encoder based on S^n alone, obviating the need for explicitly transmitting it. However, if the decoder were denied a knowledge of S^n , what is the USRDf? This question, too, remains unanswered.

Underlying our achievability proofs of Theorems 2 and 3 for a k -IRS and k -MRS, are schemes for distribution-estimation based on $(S^n, X_{\mathcal{S}}^n)$. A distinguishing feature from classical estimation settings is the additional degree of (spatial) freedom in the choice of the sampling sequence S^n . This motivates questions of the following genre: How should S^n , consisting of (possibly different) k -sized subsets, be chosen to form “best” estimates of the underlying joint pmf? How does the degree of the allowed dependence of S^n on $X_{\mathcal{M}}^n$ affect estimator performance? For instance, our choice of sampling sequence and estimation procedure in the achievability proof of Theorem 3 is a simple starting point. How must we devise *efficient* sampling mechanisms to exploit an implicit embedding of DMMS realization in the sampler output? These questions are of independent interest in statistical learning theory.

APPENDIX

A. Proof of (57) and (58)

For the code formed by concatenating (f, φ) for each $u \in \mathcal{U}$, the rate is

$$\begin{aligned} & \lesssim \max_{\tau \in \Theta} \sum_{u \in \mathcal{U}} P_U(u) \frac{1}{n'} \sum_{i=1}^{M_k} \log J_{A_i}^{u, \tau} \\ & \leq \max_{\tau \in \Theta} \sum_{u \in \mathcal{U}} P_U(u) \left(\sum_{i=1}^{M_k} \frac{|v_{A_i}^{u, \tau}|}{n} \frac{1}{|v_{A_i}^{u, \tau}|} \log J_{A_i}^{u, \tau} \right) \\ & \leq \max_{\tau \in \Theta} \sum_{u \in \mathcal{U}} P_U(u) \left(\sum_{i=1}^{M_k} P_{S|U\theta}(A_i|u, \tau) \right. \\ & \quad \left. \times \left(I(X_{A_i} \wedge Y_B | S = A_i, U = u, \theta = \tau) + \frac{\epsilon'}{2} \right) \right), \text{ by (55)} \\ & \leq \max_{\tau \in \Theta} \sum_{u \in \mathcal{U}} P_U(u) I(X_S \wedge Y_B | S, U = u, \theta = \tau) + \epsilon' \\ & \leq R_m(\Delta) + \epsilon, \end{aligned}$$

for all n large enough.

For each $U = u$, let $\Delta_u \triangleq \sum_{\tau \in \Theta, A_i \in \mathcal{A}_k} \mu_{\theta}(\tau) P_{S|U\theta}(A_i|u, \tau)$

$$\begin{aligned} & \Delta_{A_i, u, \tau}. \text{ Denoting the output of the decoder by } Y_B^{n'}, \text{ we get} \\ & \mathbb{E}[d(X_B^{n'}, Y_B^{n'})] \\ & \leq P(\widehat{\tau}_N \neq \theta) d_{\max} + \mathbb{E}[\mathbb{1}(\widehat{\tau}_N = \theta) d(X_B^{n'}, Y_B^{n'})] \\ & \leq P(\widehat{\tau}_N \neq \theta) d_{\max} + P(S' \notin \mathcal{T}^{(n)}(\epsilon', \widehat{\tau}_N)) d_{\max} \\ & \quad + \mathbb{E}[\mathbb{E}[\mathbb{1}(\widehat{\tau}_N = \theta, S' \in \mathcal{T}^{(n)}(\epsilon', \widehat{\tau}_N)) d(X_B^{n'}, Y_B^{n'}) | S', \theta]] \\ & \leq \mathbb{E}[\Delta_{S, U, \theta} | U = u] + \epsilon \\ & = \Delta_u + \epsilon \end{aligned} \tag{76}$$

for all n, N large enough, where the previous inequality is shown below. Then, expected distortion for the code formed by concatenating (f, φ) for each $u \in \mathcal{U}$, is

$$\lesssim \mathbb{E}[\Delta_U] + \epsilon \leq \Delta + \epsilon.$$

It remains to show (77). Now, (77) follows from the following: In (76), for each $\tau \in \Theta$ and $s^n \in \mathcal{T}^{(n)}(\epsilon', \widehat{\tau}_N)$,

$$\begin{aligned} & \mathbb{E}[\mathbb{1}(\widehat{\tau}_N = \theta) d(X_B^{n'}, Y_B^{n'}) | S' = s^n, \theta = \tau] \\ & = \mathbb{E} \left[\frac{\mathbb{1}(\widehat{\tau}_N = \theta)}{n'} \sum_{t \in \mu} d(X_{B_t}, Y_{B_t}) \right. \\ & \quad \left. + \frac{\mathbb{1}(\widehat{\tau}_N = \theta)}{n'} \sum_{t \in \gamma} d(X_{B_t}, Y_{B_t}) | S' = s^n, \theta = \tau \right] \\ & \leq \frac{N'}{n'} d_{\max} \\ & \quad + \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^{M_k} \sum_{t \in \gamma, s^n(A_i) \setminus v_{A_i}} d(X_{B_t}, Y_{B_t}) | S' = s^n, \theta = \tau \right] \\ & \quad + \sum_{i=1}^{M_k} \mathbb{E} \left[\frac{|v_{A_i}|}{n} \mathbb{1}(\widehat{\tau}_N = \theta) \right. \\ & \quad \left. d(X_B^{v_{A_i}}, \varphi_{A_i}^{\theta}(f_{A_i}^{\theta}(X_{A_i}^{v_{A_i}}))) | S^{v_{A_i}} = A_i^{v_{A_i}}, \theta = \tau \right] \\ & \leq d \frac{N'}{n'} d_{\max} + M_k \epsilon' d_{\max} \\ & \quad + \sum_{i=1}^{M_k} P_{S|U\theta}(A_i|u, \tau) \left(\Delta_{A_i, u, \tau} + \frac{\epsilon'}{4} \right), \text{ by (56)} \\ & \leq \mathbb{E}[\Delta_{S, U, \theta} | U = u, \theta = \tau] + M_k \epsilon' d_{\max} + \frac{N'}{n'} d_{\max} + \frac{\epsilon'}{4} \\ & \leq \mathbb{E}[\Delta_{S, U, \theta} | U = u, \theta = \tau] + \epsilon, \end{aligned}$$

for all n large enough and ϵ' chosen appropriately. \square

B. Proof of Proposition 1

First, for the Bayesian setting, by Theorem 3, the claim entails showing that

$$\min_{\{\Delta_{\tau}, \tau \in \Theta\}, \mathbb{E}[\Delta_{\theta}] \leq \Delta} \max_{\tau \in \Theta} \min_{P_{Y_B | S X_S U, \theta = \tau}, \mathbb{E}[d(X_B, Y_B) | \theta = \tau] \leq \Delta_{\tau}} I(X_S \wedge Y_B | S, U, \theta = \tau) \tag{78}$$

$$= \min_{\substack{P_U, \delta_w, \{\Delta_{\tau}, \tau \in \Theta\} \\ \mathbb{E}[\Delta_{\theta}] \leq \Delta}} \max_{\tau \in \Theta} \min_{\substack{P_{Y_B | S X_S U, \theta = \tau} \\ \mathbb{E}[d(X_B, Y_B) | \theta = \tau] \leq \Delta_{\tau}}} I(X_S \wedge Y_B | S, U, \theta = \tau), \tag{79}$$

for $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$. Denote the expressions in (78) and (79) by $q(\Delta)$ and $r(\Delta)$, respectively. Now, from the conditional version of Topsøe's identity [6, Lemma 8.5], observe that $q(\Delta)$ equals

$$\min_{\substack{P_U, P_{S|X_{\mathcal{M}}U}, \{\Delta_\tau, \tau \in \Theta\} \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \max_{\tau \in \Theta} \min_{\substack{P_{Y_B|SX_SU, \theta=\tau} \\ \mathbb{E}[d(X_B, Y_B)|\theta=\tau] \leq \Delta_\tau}} \left(D(P_{Y_B|SX_SU, \theta=\tau} \| Q_{Y_B|SU, \theta=\tau} | P_{SX_SU|\theta=\tau}) \right). \quad (80)$$

Note that the inner max and min can be interchanged in (80). Denoting $D(P_{Y_B|SX_SU, \theta=\tau} \| Q_{Y_B|SU, \theta=\tau} | P_{SX_SU|\theta=\tau})$ by D_τ , $\tau \in \Theta$, we write (80) as

$$\begin{aligned} & \min_{\substack{P_U, P_{S|X_{\mathcal{M}}U}, \{\Delta_\tau, \tau \in \Theta\} \\ \mathbb{E}[\Delta_\theta] \leq \Delta}} \min_{\substack{P_{Y_B|SX_SU, \theta=\tau} \\ \mathbb{E}[d(X_B, Y_B)|\theta=\tau] \leq \Delta_\tau, \tau \in \Theta}} \max_{\tau \in \Theta} D_\tau \\ &= \min_{\substack{P_U, P_{S|X_{\mathcal{M}}U}, P_{Y_B|SX_SU, \theta=\tau} \\ \mathbb{E}[d(X_B, Y_B)] \leq \Delta}} \max_{\tau \in \Theta} D_\tau \\ &= \min_{\substack{t, P_U, P_{S|X_{\mathcal{M}}U}, P_{Y_B|SX_SU, \theta=\tau} \\ D_\tau \leq t, \tau \in \Theta \\ \mathbb{E}[d(X_B, Y_B)] \leq \Delta}} t, \end{aligned} \quad (81)$$

which is the epigraph form. Also, $r(\Delta)$ can be expressed in a similar manner. Based on (81), we define $G_q(\alpha, \{\lambda_\tau, \tau \in \Theta\})$ and $G_r(\alpha, \{\lambda_\tau, \tau \in \Theta\})$ in terms of the Lagrangians of $q(\Delta)$ and $r(\Delta)$, respectively, in a standard way. Specifically,

$$\begin{aligned} & G_q(\alpha, \{\lambda_\tau, \tau \in \Theta\}) \\ &= \min_{\substack{t, P_U, P_{S|X_{\mathcal{M}}U} \\ P_{Y_B|SX_SU, \theta=\tau}}} t + \sum_{\tau \in \Theta} \lambda_\tau (D_\tau - t) + \alpha \mathbb{E}[d(X_B, Y_B)] \\ &= \min_{\substack{t, P_U, P_{S|X_{\mathcal{M}}U} \\ P_{Y_B|SX_SU, \theta=\tau}}} t \left(1 - \sum_{\tau \in \Theta} \lambda_\tau \right) + \sum_{\tau \in \Theta} \lambda_\tau D_\tau + \alpha \mathbb{E}[d(X_B, Y_B)] \\ &= \begin{cases} \min_{\substack{P_U, P_{S|X_{\mathcal{M}}U} \\ P_{Y_B|SX_SU, \theta=\tau}}} \sum_{\tau \in \Theta} \lambda_\tau D_\tau + \alpha \mathbb{E}[d(X_B, Y_B)], \\ \text{if } \sum_{\tau \in \Theta} \lambda_\tau = 1, \\ -\infty, \text{ otherwise.} \end{cases} \end{aligned} \quad (82)$$

Let $P_\tau \triangleq P_{X_{\mathcal{M}}|\theta=\tau}$. When $\sum_{\tau \in \Theta} \lambda_\tau = 1$, from (82), $G_q(\alpha, \{\lambda_\tau, \tau \in \Theta\})$ equals

$$\begin{aligned} & \min_{\substack{P_U, Q_{Y_B|SU, \theta} \\ P_{Y_B|SX_SU, \theta}}} \sum_{u, x_{\mathcal{M}}} P_U(u) \min_{P_{S|X_{\mathcal{M}}U}} \sum_{s \in \mathcal{A}_k} P_{S|X_{\mathcal{M}}U}(s|x_{\mathcal{M}}, u) \\ & \times \left(\mathbb{E} \left[\sum_{\tau \in \Theta} \lambda_\tau P_\tau(x_{\mathcal{M}}) \log \frac{P_{Y_B|SX_SU, \theta}(Y_B|s, x_s, u, \tau)}{Q_{Y_B|SU, \theta}(Y_B|s, u, \tau)} \right. \right. \\ & \left. \left. + \alpha \sum_{\tau \in \Theta} \mu_\theta(\tau) P_\tau(x_{\mathcal{M}}) d(x_B, Y_B) \right] \right) \\ & \left. S = s, X_S = x_s, U = u, \theta = \tau \right), \end{aligned}$$

where the expectation above is with respect to $P_{Y_B|S=s, X_S=x_s, U=u, \theta=\tau}$. Noting that the term (\dots) above is

a function of $s, x_{\mathcal{M}}, u$, we get

$$\begin{aligned} & G_q(\alpha, \{\lambda_\tau, \tau \in \Theta\}) \\ &= \min_{\substack{P_U, Q_{Y_B|SU, \theta} \\ P_{Y_B|SX_SU, \theta}}} \sum_{u, x_{\mathcal{M}}} P_U(u) \\ & \times \min_{s \in \mathcal{A}_k} \left(\mathbb{E} \left[\sum_{\tau \in \Theta} \lambda_\tau P_\tau(x_{\mathcal{M}}) \log \frac{P_{Y_B|SX_SU, \theta}(Y_B|s, x_s, u, \tau)}{Q_{Y_B|SU, \theta}(Y_B|s, u, \tau)} \right. \right. \\ & \left. \left. + \alpha \sum_{\tau \in \Theta} \mu_\theta(\tau) P_\tau(x_{\mathcal{M}}) d(x_B, Y_B) \right] \right) \\ & \left. S = s, X_S = x_s, U = u, \theta = \tau \right) \\ &= \min_{\substack{P_U, Q_{Y_B|SU, \theta} \\ P_{Y_B|SX_SU, \theta}}} \sum_{u, x_{\mathcal{M}}} P_U(u) \min_{\delta_w(\cdot, \cdot)} \sum_{s \in \mathcal{A}_k} \delta_w(x_{\mathcal{M}}, u)(s) \\ & \times \left(\mathbb{E} \left[\sum_{\tau \in \Theta} \lambda_\tau P_\tau(x_{\mathcal{M}}) \log \frac{P_{Y_B|SX_SU, \theta}(Y_B|s, x_s, u, \tau)}{Q_{Y_B|SU, \theta}(Y_B|s, u, \tau)} \right. \right. \\ & \left. \left. + \alpha \sum_{\tau \in \Theta} \mu_\theta(\tau) P_\tau(x_{\mathcal{M}}) d(x_B, Y_B) \right] \right) \\ & \left. S = s, X_S = x_s, U = u, \theta = \tau \right) \\ &= \min_{\substack{P_U, Q_{Y_B|SU, \theta} \\ P_{Y_B|SX_SU, \theta}, \delta_w}} \left[\left(\sum_{\tau \in \Theta} \lambda_\tau \right. \right. \\ & \left. \left. \times D(P_{Y_B|SX_SU, \theta=\tau} \| Q_{Y_B|SU, \theta=\tau} | P_{SX_SU|\theta=\tau}) \right. \right. \\ & \left. \left. + \alpha \mathbb{E}[d(X_B, Y_B)] \right) \right] \\ &= G_r(\alpha, \{\lambda_\tau, \tau \in \Theta\}). \end{aligned}$$

Since $q(\Delta)$ and $r(\Delta)$ are convex in Δ , they can be expressed in terms of their respective Lagrangians as

$$q(\Delta) = \max_{\alpha \geq 0, \{\lambda_\tau \geq 0, \tau \in \Theta\}} G_q(\alpha, \{\lambda_\tau, \tau \in \Theta\}) - \alpha \Delta$$

and

$$r(\Delta) = \max_{\alpha \geq 0, \{\lambda_\tau \geq 0, \tau \in \Theta\}} G_r(\alpha, \{\lambda_\tau, \tau \in \Theta\}) - \alpha \Delta. \quad (83)$$

Thus,

$$\begin{aligned} & q(\Delta) = \max_{\alpha \geq 0, \{\lambda_\tau \geq 0, \tau \in \Theta\}} G_q(\alpha, \{\lambda_\tau, \tau \in \Theta\}) - \alpha \Delta \\ &= \max_{\substack{\alpha \geq 0, \{\lambda_\tau \geq 0, \tau \in \Theta\} \\ \sum_{\tau \in \Theta} \lambda_\tau = 1}} G_q(\alpha, \{\lambda_\tau, \tau \in \Theta\}) - \alpha \Delta \\ &= \max_{\substack{\alpha \geq 0, \{\lambda_\tau \geq 0, \tau \in \Theta\} \\ \sum_{\tau \in \Theta} \lambda_\tau = 1}} G_r(\alpha, \{\lambda_\tau, \tau \in \Theta\}) - \alpha \Delta \\ &= r(\Delta), \end{aligned}$$

upon observing that the maxima in (83) are attained when $\sum_{\tau \in \Theta} \lambda_\tau = 1$.

The proof for the nonBayesian setting is similar. \square

C. Proof of Lemma 1

Clearly, for each $\tau_1 \in \Theta_1$, $\rho_A^B(\delta, \tau_1)$ and $\rho_A^{nB}(\delta, \tau_1)$ are finite-valued and, hence, so are the right-sides of (13) and (14). Also, they are also nonincreasing in Δ . The convexity of the right-sides of (13) and (14) follows from the convexity of

$\rho_A^B(\delta, \tau_1)$ and $\rho_A^{nB}(\delta, \tau_1)$ in δ along with a standard argument shown below; continuity for $\Delta > \Delta_{\min}$ is a consequence. Continuity at Δ_{\min} holds, for instance, as in [6, Lemma 7.2]. The claimed properties of the right-sides of (25), (26), (31) and (33) follow in a similar manner.

The convexity of the right-side of (13) can be shown explicitly as follows. Let $\tau_1(1)$ and $\tau_1(2)$ attain the maximum in (13) at $\Delta = \Delta_1$ and $\Delta = \Delta_2$, respectively, where $\Delta_1 < \Delta_2$. The corresponding minimizing $\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\}$ are denoted by $\{\Delta_{\tau_1}^1, \tau_1 \in \Theta_1\}$ and $\{\Delta_{\tau_1}^2, \tau_1 \in \Theta_1\}$, respectively. For any $0 < \alpha < 1$, for $i = 1, \dots, |\Theta_1|$

$$\begin{aligned} & \alpha R_A(\Delta_1) + (1 - \alpha)R_A(\Delta_2) \\ &= \alpha \rho_A^B(\Delta_{\tau_1(1)}^1, \tau_1(1)) + (1 - \alpha)\rho_A^B(\Delta_{\tau_1(2)}^2, \tau_1(2)) \\ &\geq \alpha \rho_A^B(\Delta_{\tau_1(i)}^1, \tau_1(i)) + (1 - \alpha)\rho_A^B(\Delta_{\tau_1(i)}^2, \tau_1(i)) \\ &\geq \rho_A^B(\alpha \Delta_{\tau_1(i)}^1 + (1 - \alpha)\Delta_{\tau_1(i)}^2, \tau_1(i)), \end{aligned} \quad (84)$$

where the inequality above follows by Remark (iii) preceding Theorem 1 in Section III. Now, (84) holds for every $i = 1, \dots, |\Theta_1|$, hence

$$\begin{aligned} & \alpha R_A(\Delta_1) + (1 - \alpha)R_A(\Delta_2) \\ &\geq \max_i \rho_A^B(\alpha \Delta_{\tau_1(i)}^1 + (1 - \alpha)\Delta_{\tau_1(i)}^2, \tau_1(i)) \\ &\geq \min_{\{\Delta_{\tau_1}, \tau_1 \in \Theta_1\}} \max_{\mathbb{E}[\Delta_{\theta_1}] \leq \alpha \Delta_1 + (1 - \alpha)\Delta_2} \rho_A^B(\Delta_{\tau_1}, \tau_1) \\ &= R_A(\alpha \Delta_1 + (1 - \alpha)\Delta_2). \end{aligned}$$

□

D. Proof of (39)

Since for each τ_1 in Θ_1 and x_A^n in \mathcal{X}_A^n ,

$$P_{X_A^n | \theta_1}(x_A^n | \tau_1) = \exp\left[-n(D(Q_n(x_A^n) || P_{\tau_1}) + H(Q_n(x_A^n)))\right]$$

where $Q_n(x_A^n)$ is the type of x_A^n (cf. e.g., [6, Lemma 2.6]), we have

$$\begin{aligned} \widehat{\tau}_{1,n}(X_A^n) &= \arg \max_{\tau_1 \in \Theta_1} P_{X_A^n | \theta_1}(X_A^n | \tau_1) \\ &= \arg \min_{\tau_1 \in \Theta_1} D(Q_n(X_A^n) || P_{\tau_1}). \end{aligned} \quad (85)$$

Next, for each τ_1' in Θ_1 , we know by the strong law of large numbers that $\lim_n Q_n(X_A^n) = P_{\tau_1'}$, $P_{\tau_1'}$ -a.s. (cf. e.g., [7, Corollary 2.1]). Since each P_{τ_1} , τ_1 in Θ_1 , has (full) support \mathcal{X}_A , $D(\cdot || P_{\tau_1})$ is continuous so that

$$\lim_n D(Q_n(X_A^n) || P_{\tau_1}) = D(P_{\tau_1'} || P_{\tau_1}) \quad P_{\tau_1'}\text{-a.s.} \quad (86)$$

It follows from (85) and (86) that

$$\lim_n \widehat{\tau}_{1,n}(X_A^n) = \tau_1' \quad P_{\tau_1'}\text{-a.s.},$$

which implies (39). □

REFERENCES

- [1] E. Akyol, U. Mitra, E. Tuncel, and K. Rose, "Source coding in the presence of exploration-exploitation tradeoff," *Proc. Int. Symp. Inf. Theory*, vol. 4, Jun. 2014, pp. 2057–2061.
- [2] T. Berger, *Rate Distortion Theory: A Mathematical Basis for Data Compression*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1971.
- [3] T. Berger, "Multiterminal source coding," in *The Information Theory Approach to Communications* (CISM Courses and Lectures), vol. 229, G. Longo, Ed. New York, NY, USA: Springer-Verlag, 1978, pp. 171–231.
- [4] V. P. Boda and P. Narayan, "Sampling rate distortion," *IEEE Trans. Inf. Theory*, vol. 63, no. 1, pp. 563–574, Jan. 2017.
- [5] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2012.
- [6] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [7] I. Csiszár and P. C. Shields, "Information theory and statistics: A tutorial," *Found. Trends Commun. Inf. Theory*, vol. 1, no. 4, pp. 417–528, 2004.
- [8] A. Dembo and T. Weissman, "The minimax distortion redundancy in noisy source coding," *IEEE Trans. Inf. Theory*, vol. 49, no. 11, pp. 3020–3030, Nov. 2003.
- [9] R. Dobrushin and B. Tsybakov, "Information transmission with additional noise," *IRE Trans. Inf. Theory*, vol. IT-8, no. 5, pp. 293–304, Sep. 1962.
- [10] A. Giridhar and P. R. Kumar, "Computing and communicating functions over sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 23, pp. 755–764, Apr. 2005.
- [11] O. C. Imer and T. Basar, "Optimal estimation with limited measurements," *Int. J. Syst., Control Commun.*, vol. 2, nos 1–3, pp. 5–29, Jan. 2010.
- [12] A. Kumar, P. Ishwar, and K. Ramachandran, "High-resolution distributed sampling of bandlimited fields with low-precision sensors," *IEEE Trans. Inf. Theory*, vol. 57, no. 1, pp. 476–492, Jan. 2011.
- [13] A. Kashyap, L. A. Lastras-Montano, C. Xia, and Z. Liu, "Distributed source coding in dense sensor networks," in *Proc. Data Compress. Conf.*, Mar. 2005, pp. 13–22.
- [14] A. Kipnis, A. Goldsmith, Y. C. Eldar, and T. Weissman, "Distortion rate function of sub-Nyquist sampled Gaussian sources," *IEEE Trans. Inf. Theory*, vol. 62, no. 1, pp. 401–429, Jan. 2016.
- [15] R. L. Konsbruck, E. Telatar, and M. Vetterli, "On sampling and coding for distributed acoustic sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3198–3214, May 2012.
- [16] T. Linder, "On the training distortion of vector quantizers," *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1617–1623, Jul. 2000.
- [17] T. Linder, "Learning-theoretic methods in vector quantization," in *Principles of Nonparametric Learning*. Vienna, Austria: Springer, 2002, pp. 163–210.
- [18] T. Linder, G. Lugosi, and K. Zeger, "Fixed-rate universal lossy source coding and rates of convergence for memoryless sources," *IEEE Trans. Inf. Theory*, vol. 41, no. 3, pp. 665–676, May 1995.
- [19] T. Linder, G. Lugosi, and K. Zeger, "Empirical quantizer design in the presence of source noise or channel noise," *IEEE Trans. Inf. Theory*, vol. 43, no. 2, pp. 612–623, Mar. 1997.
- [20] X. Liu, O. Simeone, and E. Erkip, "Lossy computing of correlated sources with fractional sampling," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Sep. 2012, pp. 232–236.
- [21] K. Nar and T. Basar, "Sampling multidimensional Wiener processes," in *Proc. IEEE 53rd Annu. Conf. Decis. Control (CDC)*, Dec. 2014, pp. 3426–3431.
- [22] D. Neuhoff, R. Gray, and L. Davison, "Fixed rate universal block source coding with a fidelity criterion," *IEEE Trans. Inf. Theory*, vol. IT-21, no. 5, pp. 511–523, Sep. 1975.
- [23] D. L. Neuhoff and S. S. Pradhan, "Information rates of densely sampled data: Distributed vector quantization and scalar quantization with transforms for Gaussian sources," *IEEE Trans. Inf. Theory*, vol. 59, no. 9, pp. 5641–5664, Sep. 2013.
- [24] D. Neuhoff and P. Shields, "Fixed-rate universal codes for Markov sources," *IEEE Trans. Inf. Theory*, vol. IT-24, no. 3, pp. 360–367, May 1978.
- [25] G. Reeves and M. Gastpar, "The sampling rate-distortion tradeoff for sparsity pattern recovery in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3065–3092, May 2012.
- [26] J. Rissanen, "Universal coding, information, prediction, and estimation," *IEEE Trans. Inf. Theory*, vol. IT-30, no. 4, pp. 629–636, Jul. 1984.
- [27] Y. Sun, Y. Polyanskiy, and E. Uysal-Biyikoglu. (2017). "Optimal sampling and remote estimation of the Wiener process over a channel with random delay." [Online]. Available: <https://arxiv.org/abs/1707.02531>

- [28] S.-Y. Tung, "Multiterminal source coding," Ph.D. dissertation, School Elect. Eng., Cornell Univ., Ithaca, NY, USA, May 1978.
- [29] C. Weidmann and M. Vetterli, "Rate distortion behavior of sparse sources," *IEEE Trans. Inf. Theory*, vol. 58, no. 8, pp. 4969–4992, Aug. 2012.
- [30] T. Weissman, "Universally attainable error exponents for rate-distortion coding of noisy sources," *IEEE Trans. Inf. Theory*, vol. 50, no. 6, pp. 1229–1246, Jun. 2001.
- [31] T. Weissman and N. Merhav, "Universal prediction of individual binary sequences in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 47, no. 6, pp. 2151–2173, Sep. 2001.
- [32] T. Weissman and N. Merhav, "On limited-delay lossy coding and filtering of individual sequences," *IEEE Trans. Inf. Theory*, vol. 48, no. 3, pp. 721–733, Mar. 2002.
- [33] Y. Wu and S. Verdú, "Rényi information dimension: Fundamental limits of almost lossless analog compression," *IEEE Trans. Inf. Theory*, vol. 56, no. 8, pp. 3721–3748, Aug. 2010.
- [34] H. Yamamoto and K. Itoh, "Source coding theory for multiterminal communication systems with a remote source," *IEICE Trans.*, vol. E63, no. 10, pp. 700–706, Oct. 1980.
- [35] Y. Zhang, A. Srivastava, and M. Zahran, "On-chip sensor-driven efficient thermal profile estimation algorithms," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 15, no. 3, p. 25, May 2010.
- [36] J. Ziv, "Coding of sources with unknown statistics—II: Distortion relative to a fidelity criterion," *IEEE Trans. Inf. Theory*, vol. IT-18, no. 3, pp. 389–394, May 1972.
- [37] J. Ziv, "Distortion-rate theory for individual sequences," *IEEE Trans. Inf. Theory*, vol. IT-26, no. 2, pp. 137–143, Mar. 1980.

Vinay Praneeth Boda received the Bachelor of Technology degree in Electrical Engineering and the Master of Technology degree in Communication Systems, both from the Indian Institute of Technology, Madras, India in 2011. He received his Ph.D degree in Electrical and Computer Engineering from University of Maryland, College Park and is currently working at LinkedIn Corp.

Prakash Narayan received the Bachelor of Technology degree in Electrical Engineering from the Indian Institute of Technology, Madras in 1976. He received the M.S. degree in Systems Science and Mathematics in 1978, and the D.Sc. degree in Electrical Engineering, both from Washington University, St. Louis, MO.

He is Professor of Electrical and Computer Engineering at the University of Maryland, College Park, with a joint appointment at the Institute for Systems Research. His research interests are in multiuser information theory, communication theory, communication networks, cryptography, and information theory and statistics.