

Multiterminal Source Coding under Logarithmic Loss

Thomas A. Courtade* and Tsachy Weissman†

January 24, 2012

Abstract

We consider the classical multiterminal source coding problem subject to distortion constraints computed using the logarithmic loss distortion measure. We provide a single-letter description of the achievable rate distortion region for arbitrarily correlated sources with finite alphabets. In doing so, we also give the rate distortion region for the m -encoder CEO problem. Several applications and examples are given.

1 Introduction

A complete characterization of the achievable rate distortion region for the two-encoder source coding problem depicted in Figure 1 has remained an open problem for over three decades. Following tradition, we will refer to this two-encoder source coding network as the *multiterminal source coding problem* throughout this paper. Several special cases have been solved for general source alphabets and distortion measures:

- The lossless case where $D_1 = 0, D_2 = 0$. Slepian and Wolf solved this case in their seminal work [1].
- The case where one source is recovered losslessly: i.e., $D_1 = 0, D_2 = D_{\max}$. This case corresponds to the source coding with side information problem of Ahlswede-Körner-Wyner [2], [3].
- The Wyner-Ziv case [4] where Y_2 is available to the decoder as side information and Y_1 should be recovered with distortion at most D_1 .
- The Berger-Yeung case (which subsumes the previous three cases) [5] where D_1 is arbitrary and $D_2 = 0$.

*Thomas Courtade is with the Department of Electrical Engineering, University of California, Los Angeles. Email: tacourta@ee.ucla.edu.

†Tsachy Weissman is with the Department of Electrical Engineering, Stanford University. Email: tsachy@stanford.edu.

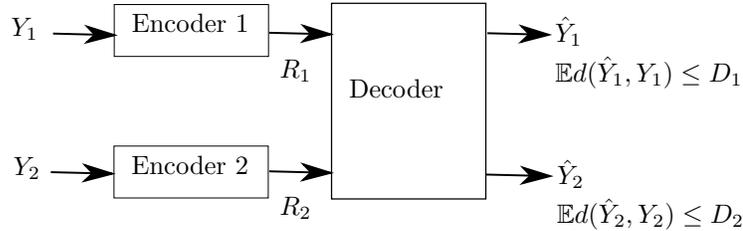


Figure 1: The multiterminal source coding network.

Despite the apparent progress, other seemingly fundamental cases, such as when D_1 is arbitrary and $D_2 = D_{\max}$, remain unsolved except perhaps in very special cases.

Recently, the achievable rate distortion region for the quadratic Gaussian multiterminal source coding problem was given by Wagner, Tavildar, and Viswanath [6]. Until now, this was the only case for which the entire achievable rate distortion region was known. While this is a very important result, it is again a special case from a theoretical point of view: a specific choice of source distribution, and a specific choice of distortion measure.

In the present paper, we determine the achievable rate distortion region of the multiterminal source coding problem for arbitrarily correlated sources with finite alphabets. However, as in [6], we restrict our attention to a specific distortion measure.

At a high level, the roadmap for our argument is similar to that of [6]. In particular, both arguments couple the multiterminal source coding problem to a parametrized family of CEO problems. Then, the parameter in the CEO problem is “tuned” to yield the converse result. Despite this apparent similarity, the proofs in [6] rely heavily on the previously known Gaussian CEO results [7], the Gaussian one-helper results [8], and the calculus performed on the closed-form entropy expressions which arise from the Gaussian source assumption. In our case we do not have this luxury, and our CEO tuning argument essentially relies on an existence lemma to yield the converse result. The success of our approach is largely due to the fact that the distortion measure we consider admits a lower bound in the form of a conditional entropy, much like the quadratic distortion measure for Gaussian sources.

1.1 Our Contributions

In this paper, we give a single-letter characterization of the achievable rate distortion region for the multiterminal source coding problem under logarithmic loss. In the process of accomplishing this, we derive the achievable rate distortion region for the m -encoder CEO problem, also under logarithmic loss. In both settings, we obtain a stronger converse than is standard for rate distortion problems in the sense that augmenting the reproduction alphabet does not

enlarge the rate distortion region. Notably, we make no assumptions on the source distributions, other than that the sources have finite alphabets. In both cases, the Berger-Tung inner bound on the rate distortion region is tight. To our knowledge, this constitutes the first time that the entire achievable rate distortion region has been described for general finite-alphabet sources under nontrivial distortion constraints.

1.2 Organization

This paper is organized as follows. In Section 2 we formally define the logarithmic loss function and the multiterminal source coding problem we consider. In Section 3 we define the CEO problem and give the rate distortion region under logarithmic loss. In Section 4 we return to the multiterminal source coding problem and derive the rate distortion region for the two-encoder setting. Also in Sections 3 and 4, applications to estimation, horse racing, and list decoding are given. In Section 5, we discuss connections between our results and the multiterminal source coding problem with arbitrary distortion measures. Section 6 delivers our concluding remarks and discusses directions for future work.

2 Problem Definition

Throughout this paper, we adopt notational conventions that are standard in the literature. Specifically, random variables are denoted by capital letters (e.g., X) and their corresponding alphabets are denoted by corresponding calligraphic letters (e.g., \mathcal{X}). We abbreviate a sequence (X_1, X_2, \dots, X_n) of n random variables by X^n , and we denote the interval $(X_k, X_{k+1}, \dots, X_j)$ by X_k^j . If the lower index is equal to 1, it will be omitted when there is no ambiguity (e.g., $X^j \triangleq X_1^j$). Frequently, random variables will appear with two subscripts (e.g., $Y_{i,j}$). In this case, we are referring to the j^{th} instance of random variable Y_i . We overload our notation here slightly in that $Y_{i,1}^j$ is often abbreviated as Y_i^j . However, our meaning will always be clear from context.

Let $\{(Y_{1,j}, Y_{2,j})\}_{j=1}^n = (Y_1^n, Y_2^n)$ be a sequence of n independent, identically distributed random variables with finite alphabets \mathcal{Y}_1 and \mathcal{Y}_2 respectively and joint pmf $p(y_1, y_2)$. That is, $(Y_1^n, Y_2^n) \sim \prod_{i=1}^n p(y_{1,j}, y_{2,j})$.

In this paper, we take the reproduction alphabet $\hat{\mathcal{Y}}_i$ to be equal to the set of probability distributions over the source alphabet \mathcal{Y}_i for $i = 1, 2$. Thus, for a vector $\hat{Y}_i^n \in \hat{\mathcal{Y}}_i^n$, we will use the notation $\hat{Y}_{i,j}(y_i)$ to mean the j^{th} coordinate ($1 \leq j \leq n$) of \hat{Y}_i^n (which is a probability distribution on \mathcal{Y}_i) evaluated for the outcome $y_i \in \mathcal{Y}_i$. In other words, the decoder generates ‘soft’ estimates of the source sequences.

We will consider the logarithmic loss distortion measure defined as follows:

$$d(y_i, \hat{y}_i) = \log \left(\frac{1}{\hat{y}_i(y_i)} \right) = D(1_{y_i}(y) \| \hat{y}_i(y)) \text{ for } i = 1, 2.$$

In particular, $d(y_i, \hat{y}_i)$ is the relative entropy (i.e., Kullback-Leibler divergence) between the empirical distribution of the event $\{Y_i = y_i\}$ and the estimate \hat{y}_i . Using this definition for symbol-wise distortion, it is standard to define the distortion between sequences as

$$d(y_i^n, \hat{y}_i^n) = \frac{1}{n} \sum_{j=1}^n d(y_{i,j}, \hat{y}_{i,j}) \text{ for } i = 1, 2.$$

We point out that the logarithmic loss function is a widely used penalty function in the theory of learning and prediction (cf. [9, Chapter 9]). Further, it is a particularly natural loss criterion in settings where the reconstructions are allowed to be ‘soft’, rather than deterministic values. Surprisingly, since distributed learning and estimation problems are some of the most oft-cited applications of lossy multiterminal source coding, it does not appear to have been studied in this context until the recent work [10]. However, we note that this connection has been established previously for the single-encoder case in the study of the information bottleneck problem [11]. Beyond learning and prediction, a similar distortion measure has appeared before in the image processing literature [12]. As we demonstrate through several examples, the logarithmic loss distortion measure has a variety of useful applications in the context of multiterminal source coding.

A rate distortion code (of blocklength n) consists of encoding functions:

$$g_i^{(n)} : \mathcal{Y}_i^n \rightarrow \{1, \dots, M_i^{(n)}\} \text{ for } i = 1, 2$$

and decoding functions

$$\psi_i^{(n)} : \{1, \dots, M_1^{(n)}\} \times \{1, \dots, M_2^{(n)}\} \rightarrow \hat{\mathcal{Y}}_i^n \text{ for } i = 1, 2.$$

A rate distortion vector (R_1, R_2, D_1, D_2) is strict-sense achievable if there exists a blocklength n , encoding functions $g_1^{(n)}, g_2^{(n)}$ and a decoder $(\psi_1^{(n)}, \psi_2^{(n)})$ such that

$$\begin{aligned} R_i &\geq \frac{1}{n} \log M_i^{(n)} \text{ for } i = 1, 2 \\ D_i &\geq \mathbb{E}d(Y_i^n, \hat{Y}_i^n) \text{ for } i = 1, 2. \end{aligned}$$

Where

$$\hat{Y}_i^n = \psi_i^{(n)}(g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n)) \text{ for } i = 1, 2.$$

Definition 1. Let \mathcal{RD}^* denote the set of strict-sense achievable rate distortion vectors and define the set of achievable rate distortion vectors to be its closure, $\overline{\mathcal{RD}^*}$.

Our ultimate goal in the present paper is to give a single-letter characterization of the region $\overline{\mathcal{RD}^*}$. However, in order to do this, we first consider an

associated CEO problem. In this sense, the roadmap for our argument is similar to that of [6]. Specifically, both arguments couple the multiterminal source coding problem to a parametrized family of CEO problems. Then, the parameter in the CEO problem is “tuned” to yield the converse result. Despite this apparent similarity, the proofs are quite different since the results in [6] depend heavily on the peculiarities of the Gaussian distribution.

3 The CEO problem

In order to attack the general multiterminal problem, we begin by studying the CEO problem (See [13] for an introduction.). To this end, let $\{(X_j, Y_{1,j}, Y_{2,j})\}_{j=1}^n = (X^n, Y_1^n, Y_2^n)$ be a sequence of n independent, identically distributed random variables distributed according to the joint pmf $p(x, y_1, y_2) = p(x)p(y_1|x)p(y_2|x)$. That is, $Y_1 \leftrightarrow X \leftrightarrow Y_2$ form a Markov chain, in that order.

In this section, we consider the reproduction alphabet $\hat{\mathcal{X}}$ to be equal to the set of probability distributions over the source alphabet \mathcal{X} . As before, for a vector $\hat{X}^n \in \hat{\mathcal{X}}^n$, we will use the notation $\hat{X}_j(x)$ to mean the j^{th} coordinate of \hat{X}^n (which is a probability distribution on \mathcal{X}) evaluated for the outcome $x \in \mathcal{X}$. As in the rest of this paper, $d(\cdot, \cdot)$ is the logarithmic loss distortion measure.

A rate distortion CEO code (of blocklength n) consists of encoding functions:

$$g_i^{(n)} : \mathcal{Y}_i^n \rightarrow \{1, \dots, M_i^{(n)}\} \text{ for } i = 1, 2$$

and a decoding function

$$\psi^{(n)} : \{1, \dots, M_1^{(n)}\} \times \{1, \dots, M_2^{(n)}\} \rightarrow \hat{\mathcal{X}}^n.$$

A rate distortion vector (R_1, R_2, D) is strict-sense achievable for the CEO problem if there exists a blocklength n , encoding functions $g_1^{(n)}, g_2^{(n)}$ and a decoder $\psi^{(n)}$ such that

$$\begin{aligned} R_i &\geq \frac{1}{n} \log M_i^{(n)} \text{ for } i = 1, 2 \\ D &\geq \mathbb{E}d(X^n, \hat{X}^n). \end{aligned}$$

Where

$$\hat{X}^n = \psi^{(n)}(g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n)).$$

Definition 2. Let \mathcal{RD}_{CEO}^* denote the set of strict-sense achievable rate distortion vectors and define the set of achievable rate distortion vectors to be its closure, $\overline{\mathcal{RD}}_{CEO}^*$.

3.1 Inner Bound

Definition 3. Let $(R_1, R_2, D) \in \mathcal{RD}_{CEO}^i$ if and only if there exists a joint distribution of the form

$$p(x, y_1, y_2)p(u_1|y_1, q)p(u_2|y_2, q)p(q)$$

where $|\mathcal{U}_1| \leq |\mathcal{Y}_1|$, $|\mathcal{U}_2| \leq |\mathcal{Y}_2|$, and $|\mathcal{Q}| \leq 4$, which satisfies

$$\begin{aligned} R_1 &\geq I(Y_1; U_1|U_2, Q) \\ R_2 &\geq I(Y_2; U_2|U_1, Q) \\ R_1 + R_2 &\geq I(U_1, U_2; Y_1, Y_2|Q) \\ D &\geq H(X|U_1, U_2, Q). \end{aligned}$$

Theorem 1. $\mathcal{RD}_{CEO}^i \subseteq \overline{\mathcal{RD}_{CEO}^*}$. That is, all rate distortion vectors $(R_1, R_2, D) \in \mathcal{RD}_{CEO}^i$ are achievable.

Before proceeding with the proof, we cite the following variant of a well-known inner bound:

Proposition 1 (Berger-Tung Inner Bound [14, 15]). *The rate distortion vector (R_1, R_2, D) is achievable if*

$$\begin{aligned} R_1 &\geq I(U_1; Y_1|U_2, Q) \\ R_2 &\geq I(U_2; Y_2|U_1, Q) \\ R_1 + R_2 &\geq I(U_1, U_2; Y_1, Y_2|Q) \\ D &\geq \mathbb{E}[d(X, f(U_1, U_2, Q))] \end{aligned}$$

for a joint distribution

$$p(x)p(y_1|x)p(y_2|x)p(u_1|y_1, q)p(u_2|y_2, q)p(q)$$

and reproduction function

$$f : \mathcal{U}_1 \times \mathcal{U}_2 \times \mathcal{Q} \rightarrow \hat{\mathcal{X}}.$$

The proof of this proposition is a standard exercise in information theory, and is therefore omitted. The interested reader is directed to the text [16] for a modern, detailed treatment. The proposition follows from what is commonly called the Berger-Tung achievability scheme. In this encoding scheme, each encoder quantizes its observation Y_i^n to a codeword U_i^n , such that the empirical distribution of the entries in (Y_i^n, U_i^n) is very close to the true distribution $p(y_i, u_i)$. In order to communicate their respective quantizations to the decoder, the encoders essentially perform Slepian-Wolf coding. For this reason, the Berger-Tung achievability scheme is also referred to as a “quantize-and-bin” coding scheme.

Proof of Theorem 1. Given Proposition 1, the proof of Theorem 1 is immediate. Indeed, if we apply Proposition 1 with the reproduction function $f(U_1, U_2, Q) \triangleq \Pr[X = x|U_1, U_2, Q]$, we note that

$$\mathbb{E}[d(X, f(U_1, U_2, Q))] = H(X|U_1, U_2, Q),$$

which yields the desired result. \square

Thus, from the proof of Theorem 1, we see that our inner bound \mathcal{RD}_{CEO}^i simply corresponds to a specialization of the general Berger-Tung inner bound to the case of logarithmic loss.

3.2 A Matching Outer Bound

A particularly useful property of the logarithmic loss distortion measure is that the expected distortion is lower-bounded by a conditional entropy. A similar property is enjoyed by Gaussian random variables under quadratic distortion. In particular, if G is Gaussian, and \hat{G} is such that $\mathbb{E}(\hat{G} - G)^2 \leq D$, then $\frac{1}{2} \log(2\pi e)D \geq h(G|\hat{G})$. The case for logarithmic loss is similar, and we state it formally in the following lemma which is crucial in the proof of the converse.

Lemma 1. *Let $Z = (g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n))$ be the argument of the reproduction function $\psi^{(n)}$. Then $n\mathbb{E}d(\hat{X}^n, X^n) \geq H(X^n|Z)$.*

Proof. By definition of the reproduction alphabet, we can consider the reproduction \hat{X}^n to be a probability distribution on \mathcal{X}^n conditioned on the argument Z . In particular, if $\hat{x}^n = \psi^{(n)}(z)$, define $s(x^n|z) \triangleq \prod_{j=1}^n \hat{x}_j(x_j)$. It is readily verified that s is a probability measure on \mathcal{X}^n . Then, we obtain the following lower bound on the expected distortion conditioned on $Z = z$:

$$\begin{aligned} \mathbb{E}[d(X^n, \hat{X}^n)|Z = z] &= \frac{1}{n} \sum_{j=1}^n \sum_{x^n \in \mathcal{X}^n} p(x^n|z) \log\left(\frac{1}{\hat{x}_j(x_j)}\right) \\ &= \frac{1}{n} \sum_{x^n \in \mathcal{X}^n} p(x^n|z) \sum_{j=1}^n \log\left(\frac{1}{\hat{x}_j(x_j)}\right) \\ &= \frac{1}{n} \sum_{x^n \in \mathcal{X}^n} p(x^n|z) \log\left(\frac{1}{s(x^n|z)}\right) \\ &= \frac{1}{n} \sum_{x^n \in \mathcal{X}^n} p(x^n|z) \log\left(\frac{p(x^n|z)}{s(x^n|z)}\right) + \frac{1}{n} H(X^n|Z = z) \\ &= \frac{1}{n} D(p(x^n|z)||s(x^n|z)) + \frac{1}{n} H(X^n|Z = z) \\ &\geq \frac{1}{n} H(X^n|Z = z), \end{aligned}$$

where $p(x^n|z) = \Pr(X^n = x^n|Z = z)$ is the true conditional distribution. Averaging both sides over all values of Z , we obtain the desired result. \square

Definition 4. Let $(R_1, R_2, D) \in \mathcal{RD}_{CEO}^o$ if and only if there exists a joint distribution of the form

$$p(x)p(y_1|x)p(y_2|x)p(u_1|y_1, q)p(u_2|y_2, q)p(q),$$

which satisfies

$$\left. \begin{aligned} R_1 &\geq I(Y_1; U_1|X, Q) + H(X|U_2, Q) - D \\ R_2 &\geq I(Y_2; U_2|X, Q) + H(X|U_1, Q) - D \\ R_1 + R_2 &\geq I(U_1; Y_1|X, Q) + I(U_2; Y_2|X, Q) + H(X) - D \\ D &\geq H(X|U_1, U_2, Q). \end{aligned} \right\} \quad (1)$$

Theorem 2. If (R_1, R_2, D) is strict-sense achievable for the CEO problem, then $(R_1, R_2, D) \in \mathcal{RD}_{CEO}^o$.

Proof. Suppose the point (R_1, R_2, D) is strict-sense achievable. Let A be a nonempty subset of $\{1, 2\}$, and let $F_i = g_i^{(n)}(Y_i^n)$ be the message sent by encoder $i \in \{1, 2\}$. Define $U_{i,j} \triangleq (F_i, Y_i^{j-1})$ and $Q_j \triangleq (X^{j-1}, X_{j+1}^n) = X^n \setminus X_j$. To simplify notation, let $Y_A = \cup_{i \in A} Y_i$ (similarly for U_A and F_A).

With these notations established, we have the following string of inequalities:

$$\begin{aligned} n \sum_{i \in A} R_i &\geq \sum_{i \in A} H(F_i) \\ &\geq H(F_A) \\ &\geq I(Y_A^n; F_A | F_{A^c}) \\ &= I(X^n, Y_A^n; F_A | F_{A^c}) \end{aligned} \quad (2)$$

$$= I(X^n; F_A | F_{A^c}) + \sum_{i \in A} I(F_i; Y_i^n | X^n) \quad (3)$$

$$\begin{aligned} &= H(X^n | F_{A^c}) - H(X^n | F_1, F_2) + \sum_{i \in A} \sum_{j=1}^n I(Y_{i,j}; F_i | X^n, Y_i^{j-1}) \\ &\geq H(X^n | F_{A^c}) + \sum_{i \in A} \sum_{j=1}^n I(Y_{i,j}; F_i | X^n, Y_i^{j-1}) - nD \end{aligned} \quad (4)$$

$$\begin{aligned} &= \sum_{j=1}^n H(X_j | F_{A^c}, X^{j-1}) + \sum_{i \in A} \sum_{j=1}^n I(Y_{i,j}; F_i | X^n, Y_i^{j-1}) - nD \\ &= \sum_{j=1}^n H(X_j | F_{A^c}, X^{j-1}) + \sum_{i \in A} \sum_{j=1}^n I(Y_{i,j}; U_{i,j} | X_j, Q_j) - nD \end{aligned} \quad (5)$$

$$\geq \sum_{j=1}^n H(X_j | U_{A^c, j}, Q_j) + \sum_{i \in A} \sum_{j=1}^n I(Y_{i,j}; U_{i,j} | X_j, Q_j) - nD. \quad (6)$$

The nontrivial steps above can be justified as follows:

- (2) follows since F_A is a function of Y_A^n .

- (3) follows since F_i is a function of Y_i^n and $F_1 \leftrightarrow X^n \leftrightarrow F_2$ form a Markov chain (since $Y_1^n \leftrightarrow X^n \leftrightarrow Y_2^n$ form a Markov chain).
- (4) follows since $nD \geq H(X^n|F_1, F_2)$ by Lemma 1.
- (5) follows from the Markov chain $Y_{i,j} \leftrightarrow X^n \leftrightarrow Y_i^{j-1}$, which follows from the i.i.d. nature of the source sequences.
- (6) simply follows from the fact that conditioning reduces entropy.

Therefore, dividing both sides by n , we have:

$$\sum_{i \in A} R_i \geq \frac{1}{n} \sum_{j=1}^n H(X_j|U_{A^c,j}, Q_j) + \sum_{i \in A} \frac{1}{n} \sum_{j=1}^n I(Y_{i,j}; U_{i,j}|X_j, Q_j) - D.$$

Also, using Lemma 1 and the fact that conditioning reduces entropy:

$$D \geq \frac{1}{n} H(X^n|F_1, F_2) \geq \frac{1}{n} \sum_{j=1}^n H(X_j|U_{1,j}, U_{2,j}, Q_j).$$

Observe that Q_j is independent of $(X_j, Y_{1,j}, Y_{2,j})$ and, conditioned on Q_j , we have the long Markov chain $U_{1,j} \leftrightarrow Y_{1,j} \leftrightarrow X_j \leftrightarrow Y_{2,j} \leftrightarrow U_{2,j}$. Finally, by a standard time-sharing argument, we conclude by saying that if (R_1, R_2, D) is strict-sense achievable for the CEO problem, then

$$\begin{aligned} R_1 &\geq I(Y_1; U_1|X, Q) + H(X|U_2, Q) - D \\ R_2 &\geq I(Y_2; U_2|X, Q) + H(X|U_1, Q) - D \\ R_1 + R_2 &\geq I(U_1; Y_1|X, Q) + I(U_2; Y_2|X, Q) + H(X) - D \\ D &\geq H(X|U_1, U_2, Q). \end{aligned}$$

for some joint distribution $p(q)p(x, y_1, y_2)p(u_1|y_1, q)p(u_2|y_2, q)$. \square

Theorem 3. $\mathcal{RD}_{CEO}^o = \mathcal{RD}_{CEO}^i = \overline{\mathcal{RD}_{CEO}^*}$.

Proof. We first remark that the cardinality bounds on the alphabets in the definition of \mathcal{RD}_{CEO}^i can be imposed without any loss of generality. This is a consequence of [17, Lemma 2.2] and is discussed in detail in Appendix A.

Therefore, it will suffice to show $\mathcal{RD}_{CEO}^o \subseteq \mathcal{RD}_{CEO}^i$ without considering the cardinality bounds. To this end, fix $p(q)$, $p(u_1|y_1, q)$, and $p(u_2|y_2, q)$ and

consider the extreme points¹ of polytope defined by the inequalities (1):

$$\begin{aligned}
P_1 &= \left(0, 0, I(Y_1; U_1|X, Q) + I(Y_2; U_2|X, Q) + H(X) \right) \\
P_2 &= \left(I(Y_1; U_1|Q), 0, I(U_2; Y_2|X, Q) + H(X|U_1, Q) \right) \\
P_3 &= \left(0, I(Y_2; U_2|Q), I(U_1; Y_1|X, Q) + H(X|U_2, Q) \right) \\
P_4 &= \left(I(Y_1; U_1|Q), I(Y_2; U_2|U_1, Q), H(X|U_1, U_2, Q) \right) \\
P_5 &= \left(I(Y_1; U_1|U_2, Q), I(Y_2; U_2|Q), H(X|U_1, U_2, Q) \right),
\end{aligned}$$

where the point P_j is a triple $(R_1^{(j)}, R_2^{(j)}, D^{(j)})$. We say a point $(R_1^{(j)}, R_2^{(j)}, D^{(j)})$ is *dominated by* a point in \mathcal{RD}_{CEO}^i if there exists some $(R_1, R_2, D) \in \mathcal{RD}_{CEO}^i$ for which $R_1 \leq R_1^{(j)}$, $R_2 \leq R_2^{(j)}$, and $D \leq D^{(j)}$. Observe that each of the extreme points P_1, \dots, P_5 is dominated by a point in \mathcal{RD}_{CEO}^i :

- First, observe that P_4 and P_5 are both in \mathcal{RD}_{CEO}^i , so these points are not problematic.
- Next, observe that the point $(0, 0, H(X))$ is in \mathcal{RD}_{CEO}^i , which can be seen by setting all auxiliary random variables to be constant. This point dominates P_1 .
- By using auxiliary random variables $(\hat{U}_1, \hat{U}_2, Q) = (U_1, \emptyset, Q)$, the point $(I(Y_1; U_1|Q), 0, H(X|U_1, Q))$ is in \mathcal{RD}_{CEO}^i , and dominates the point P_2 . By a symmetric argument, the point P_3 is also dominated by a point in \mathcal{RD}_{CEO}^i .

Since \mathcal{RD}_{CEO}^o is the convex hull of all such extreme points (i.e., the convex hull of the union of extreme points over all appropriate joint distributions), the theorem is proved. \square

Remark 1. *Theorem 3 can be extended to the general case of m -encoders. Details are provided in Appendix B.*

3.3 A stronger converse result for the CEO problem

As defined, our reproduction sequence \hat{X}^n is restricted to be a product distribution on \mathcal{X}^n . However, for a blocklength n code, we can allow \hat{X}^n to be *any* probability distribution on \mathcal{X}^n and the converse result still holds. In this case, we define the sequence distortion as follows:

$$d(x^n, \hat{x}^n) = \frac{1}{n} \log \left(\frac{1}{\hat{x}^n(x^n)} \right),$$

¹For two encoders, it is easy enough to enumerate the extreme points by inspection. However, this can be formalized by a submodularity argument, which is given in Appendix B.

which is compatible with the original definition when \hat{X}^n is a product distribution. The reader can verify that the result of Lemma 1 is still true for this more general distortion alphabet by setting $s(x^n|z) = \hat{x}^n(x^n)$ in the corresponding proof. Since Lemma 1 is the key tool in the CEO converse result, this implies that the converse holds even if \hat{X}^n is allowed to be any probability distribution on \mathcal{X}^n (rather than being restricted to the set of product distributions).

When this stronger converse result is taken together with the achievability result, we observe that restricting \hat{X}^n to be a product distribution is in fact optimal and can achieve all points in $\overline{\mathcal{RD}}_{CEO}^*$.

3.4 An Example: Distributed compression of a posterior distribution

Suppose two sensors observe sequences Y_1^n and Y_2^n respectively, which are conditionally independent given a hidden sequence X^n . The sensors communicate with a fusion center through rate-limited links of capacity R_1 and R_2 respectively. Given sequences (Y_1^n, Y_2^n) are observed, the sequence X^n cannot be determined in general, so the fusion center would like to estimate the posterior distribution $p(x^n|Y_1^n, Y_2^n)$. Since the communication links are rate-limited, the fusion center cannot necessarily compute $p(x^n|Y_1^n, Y_2^n)$ exactly. In this case, the fusion center would like to generate an estimate $\hat{p}(x^n|g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n))$ that should approximate $p(x^n|Y_1^n, Y_2^n)$ in the sense that, on average:

$$D\left(p(x^n|y_1^n, y_2^n) \parallel \hat{p}(x^n|g_1^{(n)}(y_1^n), g_2^{(n)}(y_2^n))\right) \leq n\varepsilon,$$

where, consistent with standard notation (e.g. [18]), we write $D(p(x^n|y_1^n, y_2^n) \parallel \hat{p}(x^n|g_1^{(n)}(y_1^n), g_2^{(n)}(y_2^n)))$ as shorthand for

$$\sum_{x^n, y_1^n, y_2^n} p(x^n, y_1^n, y_2^n) \log \frac{p(x^n|y_1^n, y_2^n)}{\hat{p}(x^n|g_1^{(n)}(y_1^n), g_2^{(n)}(y_2^n))}.$$

The relevant question here is the following. What is the minimum distortion ε that is attainable given R_1 and R_2 ?

Considering the CEO problem for this setup, we have:

$$\begin{aligned} \mathbb{E}d(\hat{X}^n, X^n) &= \frac{1}{n} \sum_{(x^n, y_1^n, y_2^n)} p(x^n, y_1^n, y_2^n) \log \left(\frac{1}{\hat{x}^n(x^n)} \right) \\ &= \frac{1}{n} D\left(p(x^n|y_1^n, y_2^n) \parallel \hat{x}^n(x^n)\right) + \frac{1}{n} H(X^n|Y_1^n, Y_2^n). \end{aligned}$$

Identifying $\hat{p}(x^n|g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n)) \leftarrow \hat{X}^n(x^n)$, we have:

$$D\left(p(x^n|y_1^n, y_2^n) \parallel \hat{p}(x^n|g_1^{(n)}(y_1^n), g_2^{(n)}(y_2^n))\right) = n\mathbb{E}d(\hat{X}^n, X^n) - nH(X|Y_1, Y_2).$$

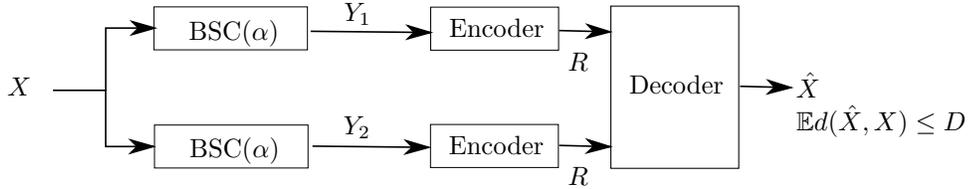


Figure 2: An example CEO problem where $X \sim \text{Bernoulli}(\frac{1}{2})$, $\Pr(Y_i = X) = (1 - \alpha)$, and both encoders are subject to the same rate constraint.

Thus, finding the minimum possible distortion reduces to an optimization problem over $\overline{\mathcal{RD}}_{CEO}^*$. In particular, the minimum attainable distortion ε^* is given by

$$\varepsilon^* = \inf \left\{ D : (R_1, R_2, D) \in \overline{\mathcal{RD}}_{CEO}^* \right\} - H(X|Y_1, Y_2). \quad (7)$$

Moreover, the minimum distortion is obtained by estimating each x_j separately. In other words, there exists an optimal (essentially, for large n) estimate $\hat{p}^*(x^n|\cdot, \cdot)$ (which is itself a function of optimal encoding functions $g_1^{*(n)}(\cdot)$ and $g_2^{*(n)}(\cdot)$) that can be expressed as a product distribution

$$\hat{p}^*(x^n|\cdot, \cdot) = \prod_{j=1}^n \hat{p}_j^* \left(x_j | g_1^{*(n)}(\cdot), g_2^{*(n)}(\cdot) \right).$$

For this choice of $\hat{p}^*(x^n|\cdot, \cdot)$, we have the following relationship:

$$\frac{1}{n} \sum_{j=1}^n D \left(p(x_j | y_{1,j}, y_{2,j}) \left\| \hat{p}_j^* \left(x_j | g_1^{*(n)}(y_1^n), g_2^{*(n)}(y_2^n) \right) \right. \right) = \varepsilon^*.$$

In light of this fact, one can apply Markov's inequality to obtain the following estimate on peak component-wise distortion:

$$\# \left\{ j \mid D \left(p(x_j | y_{1,j}, y_{2,j}) \left\| \hat{p}_j^* \left(x_j | g_1^{*(n)}(y_1^n), g_2^{*(n)}(y_2^n) \right) \right. \right) \geq \zeta \right\} \leq n \frac{\varepsilon^*}{\zeta},$$

where $\#(\cdot)$ is the counting measure.

To make this example more concrete, consider the scenario depicted in Figure 2, where $X \sim \text{Bernoulli}(\frac{1}{2})$ and Y_i is the result of passing X through a binary symmetric channel with crossover probability α for $i = 1, 2$. To simplify things, we constrain the rates of each encoder to be at most R bits per channel use.

By performing a brute-force search over a fine mesh of conditional distributions $\{p(u_i|y_i)\}_{i=1}^2$, we numerically approximate the set of (R, D) pairs such that (R, R, D) is in the achievable region $\overline{\mathcal{RD}}_{CEO}^*$ corresponding to the network in Figure 2. The lower convex envelope of these (R, D) pairs is plotted in Figure 3

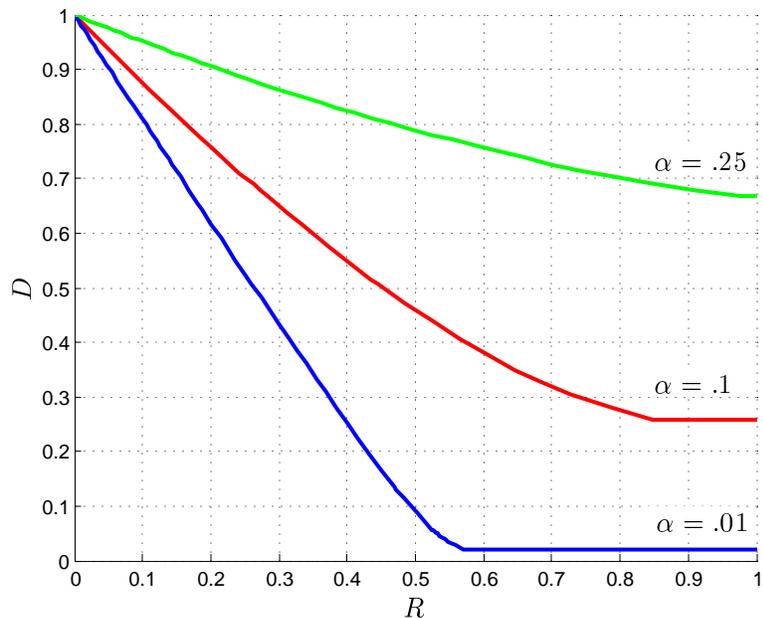


Figure 3: The distortion-rate function of the network in Figure 2 computed for $\alpha \in \{0.01, 0.1, 0.25\}$.

for $\alpha \in \{0.01, 0.1, 0.25\}$. Continuing our example above for this concrete choice of source parameters, we compute the minimum achievable Kullback-Leibler distance ε^* according to (7). The result is given in Figure 4.

These numerical results are intuitively satisfying in the sense that, if Y_1, Y_2 are high-quality estimates of X (e.g., $\alpha = 0.01$), then a small increase in the allowable rate R results in a large relative improvement of $\hat{p}(x|\cdot, \cdot)$, the decoder's estimate of $p(x|Y_1, Y_2)$. On the other hand, if Y_1, Y_2 are poor-quality estimates of X (e.g., $\alpha = 0.25$), then we require a large increase in the allowable rate R in order to obtain an appreciable improvement of $\hat{p}(x|\cdot, \cdot)$.

One field where this example is directly applicable is machine learning. In this case, X_j could represent the class of object j , and $Y_{1,j}, Y_{2,j}$ are observable attributes. In machine learning, one typically estimates the probability that an object belongs to a particular class given a set of observable attributes. For this type of estimation problem, relative entropy is a natural penalty criterion.

Another application is to horse-racing with conditionally independent, rate-limited side informations. In this case, the doubling rate of the gambler's wealth can be expressed in terms of the logarithmic loss distortion measure. This example is consistent with the original interpretation of the CEO problem, where the CEO makes consecutive business decisions (investments) having outcomes X^n , with the objective of maximizing the wealth of the company. We omit the details.

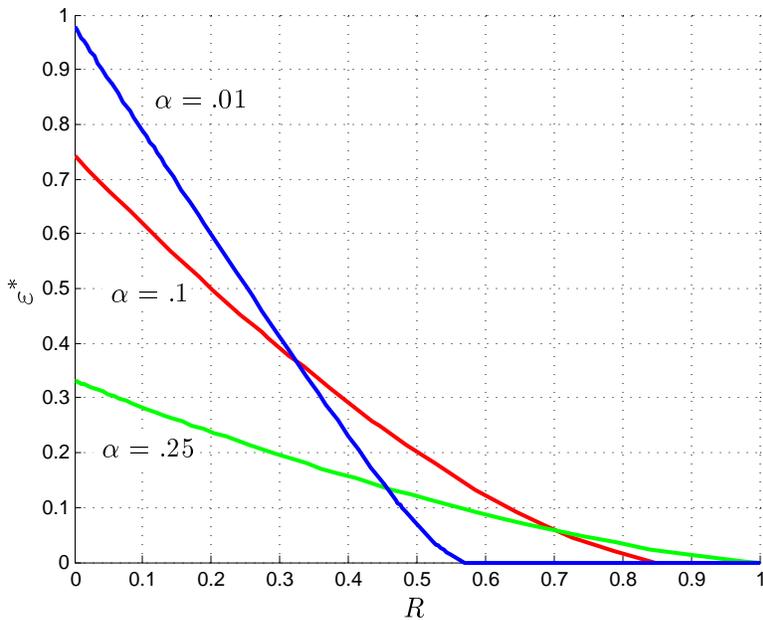


Figure 4: The minimum achievable Kullback-Leibler distance computed according to (7), i.e., the curves here are those of Figure 3, lowered by the constant $H(X|Y_1, Y_2)$.

3.5 An Example: Joint estimation of the encoder observations

Suppose one wishes to estimate the encoder observations (Y_1, Y_2) . In this case, the rate region simplifies considerably. In particular, if we tolerate a distortion D in our estimate of the pair (Y_1, Y_2) , then the achievable rate region is the same as the Slepian-Wolf rate region with each rate constraint relaxed by D bits. Formally:

Theorem 4. *If $X = (Y_1, Y_2)$, then $\overline{\mathcal{RD}}_{CEO}^*$ consists of all vectors (R_1, R_2, D) satisfying*

$$\begin{aligned} R_1 &\geq H(Y_1|Y_2) - D \\ R_2 &\geq H(Y_2|Y_1) - D \\ R_1 + R_2 &\geq H(Y_1, Y_2) - D \\ D &\geq 0. \end{aligned}$$

Proof. First, note that Theorem 3 implies that $\overline{\mathcal{RD}}_{CEO}^*$ is equivalent to the union of (R_1, R_2, D) triples satisfying (1) taken over all joint distributions $p(q)p(x, y_1, y_2)p(u_1|y_1, q)p(u_2|y_2, q)$. Now, since $X = (Y_1, Y_2)$, each of the in-

equalities (1) can be lower bounded as follows:

$$\begin{aligned}
R_1 &\geq I(Y_1; U_1 | Y_1, Y_2, Q) + H(Y_1, Y_2 | U_2, Q) - D \\
&= H(Y_2 | U_2, Q) + H(Y_1 | Y_2) - D \\
&\geq H(Y_1 | Y_2) - D \\
R_2 &\geq I(Y_2; U_2 | Y_1, Y_2, Q) + H(Y_1, Y_2 | U_1, Q) - D \\
&= H(Y_1 | U_1, Q) + H(Y_2 | Y_1) - D \\
&\geq H(Y_2 | Y_1) - D \\
R_1 + R_2 &\geq I(U_1; Y_1 | Y_1, Y_2, Q) + I(U_2; Y_2 | Y_1, Y_2, Q) + H(Y_1, Y_2) - D \\
&= H(Y_1, Y_2) - D \\
D &\geq H(Y_1, Y_2 | U_1, U_2, Q) \\
&\geq 0.
\end{aligned}$$

Finally, observe that by setting $U_i = Y_i$ for $i = 1, 2$, we can achieve any point in this relaxed region (again, a consequence of Theorem 3). \square

We remark that this result was first proved in [10] by Courtade and Wesel using a different method.

4 Multiterminal Source Coding

With Theorem 3 in hand, we are now in a position to characterize the achievable rate distortion region $\overline{\mathcal{RD}}^*$ for the multiterminal source coding problem under logarithmic loss. As before, we prove an inner bound first.

4.1 Inner Bound

Definition 5. Let $(R_1, R_2, D_1, D_2) \in \mathcal{RD}^i$ if and only if there exists a joint distribution of the form

$$p(y_1, y_2)p(u_1 | y_1, q)p(u_2 | y_2, q)p(q)$$

where $|\mathcal{U}_1| \leq |\mathcal{Y}_1|$, $|\mathcal{U}_2| \leq |\mathcal{Y}_2|$, and $|\mathcal{Q}| \leq 5$, which satisfies

$$\begin{aligned}
R_1 &\geq I(Y_1; U_1 | U_2, Q) \\
R_2 &\geq I(Y_2; U_2 | U_1, Q) \\
R_1 + R_2 &\geq I(U_1, U_2; Y_1, Y_2 | Q) \\
D_1 &\geq H(Y_1 | U_1, U_2, Q) \\
D_2 &\geq H(Y_2 | U_1, U_2, Q).
\end{aligned}$$

Theorem 5. $\mathcal{RD}^i \subseteq \overline{\mathcal{RD}}^*$. That is, all rate distortion vectors in \mathcal{RD}^i are achievable.

Again, we require an appropriate version of the Berger-Tung inner bound:

Proposition 2 (Berger-Tung Inner Bound [14, 15]). *The rate distortion vector (R_1, R_2, D_1, D_2) is achievable if*

$$\begin{aligned} R_1 &\geq I(U_1; Y_1 | U_2, Q) \\ R_2 &\geq I(U_2; Y_2 | U_1, Q) \\ R_1 + R_2 &\geq I(U_1, U_2; Y_1, Y_2 | Q) \\ D_1 &\geq \mathbb{E}[d(Y_1, f_1(U_1, U_2, Q))] \\ D_2 &\geq \mathbb{E}[d(Y_2, f_2(U_1, U_2, Q))]. \end{aligned}$$

for a joint distribution

$$p(y_1, y_2)p(u_1|y_1, q)p(u_2|y_2, q)p(q)$$

and reproduction functions

$$f_i : \mathcal{U}_1 \times \mathcal{U}_2 \times \mathcal{Q} \rightarrow \hat{\mathcal{Y}}_i, \text{ for } i = 1, 2.$$

Proof of Theorem 5. To prove the theorem, we simply apply Proposition 2 with the reproduction functions $f_i(U_1, U_2, Q) := \Pr[Y_i = y_i | U_1, U_2, Q]$. \square

Hence, we again see that our inner bound $\mathcal{RD}^i \subseteq \overline{\mathcal{RD}}^*$ is nothing more than the Berger-Tung inner bound specialized to the setting when distortion is measured under logarithmic loss.

4.2 A Matching Outer Bound

The main result of this paper is the following theorem.

Theorem 6. $\mathcal{RD}^i = \overline{\mathcal{RD}}^*$.

Proof. As before, we note that the cardinality bounds on the alphabets in the definition of \mathcal{RD}^i can be imposed without any loss of generality. This is discussed in detail in Appendix A.

Assume (R_1, R_2, D_1, D_2) is strict-sense achievable. Observe that proving that $(R_1, R_2, D_1, D_2) \in \mathcal{RD}^i$ will prove the theorem, since $\mathcal{RD}^i \subseteq \overline{\mathcal{RD}}^*$ and $\overline{\mathcal{RD}}^*$ is closed by definition.

For convenience, define $\mathcal{P}(R_1, R_2)$ to be the set of joint distributions of the form

$$p(y_1, y_2)p(u_1|y_1, q)p(u_2|y_2, q)p(q)$$

with $|\mathcal{U}_1| \leq |\mathcal{Y}_1|$, $|\mathcal{U}_2| \leq |\mathcal{Y}_2|$, and $|\mathcal{Q}| \leq 4$ satisfying

$$\begin{aligned} R_1 &\geq I(U_1; Y_1 | U_2, Q) \\ R_2 &\geq I(U_2; Y_2 | U_1, Q) \\ R_1 + R_2 &\geq I(U_1, U_2; Y_1, Y_2 | Q). \end{aligned}$$

We remark that $\mathcal{P}(R_1, R_2)$ is compact. We also note that it will suffice to show the existence of a joint distribution in $\mathcal{P}(R_1, R_2)$ satisfying $H(Y_1|U_1, U_2, Q) \leq D_1$ and $H(Y_2|U_1, U_2, Q) \leq D_2$ to prove that $(R_1, R_2, D_1, D_2) \in \mathcal{RD}^i$.

With foresight, consider random variable X defined as follows

$$X = \begin{cases} (Y_1, 1) & \text{with probability } t \\ (Y_2, 2) & \text{with probability } 1 - t. \end{cases} \quad (8)$$

In other words, $X = (Y_B, B)$, where B is a Bernoulli random variable independent of Y_1, Y_2 . Observe that $Y_1 \leftrightarrow X \leftrightarrow Y_2$ form a Markov chain, and thus, we are able to apply Theorem 3.

Since (R_1, R_2, D_1, D_2) is strict-sense achievable, the decoder can construct reproductions \hat{Y}_1^n, \hat{Y}_2^n satisfying

$$\frac{1}{n} \sum_{j=1}^n \mathbb{E}d(Y_{i,j}, \hat{Y}_{i,j}) \leq D_i \text{ for } i = 1, 2.$$

Fix the encoding operations and set $\hat{X}_j((y_1, 1)) = t\hat{Y}_{1,j}(y_1)$ and $\hat{X}_j((y_2, 2)) = (1-t)\hat{Y}_{2,j}(y_2)$. Then for the CEO problem defined by (X, Y_1, Y_2) :

$$\begin{aligned} & \frac{1}{n} \sum_{j=1}^n \mathbb{E}d(X_j, \hat{X}_j) \\ &= \frac{t}{n} \sum_{j=1}^n \mathbb{E} \log \left(\frac{1}{t\hat{Y}_{1,j}(Y_{1,j})} \right) + \frac{1-t}{n} \sum_{j=1}^n \mathbb{E} \log \left(\frac{1}{(1-t)\hat{Y}_{2,j}(Y_{2,j})} \right) \\ &= h_2(t) + \frac{t}{n} \sum_{j=1}^n \mathbb{E}d(Y_{1,j}, \hat{Y}_{1,j}) + \frac{1-t}{n} \sum_{j=1}^n \mathbb{E}d(Y_{2,j}, \hat{Y}_{2,j}) \\ &\leq h_2(t) + tD_1 + (1-t)D_2 \end{aligned}$$

where $h_2(t)$ is the binary entropy function. Hence, for this CEO problem, distortion $h_2(t) + tD_1 + (1-t)D_2$ is achievable and Theorem 3 yields a joint distribution² $P_t \in \mathcal{P}(R_1, R_2)$ satisfying

$$\begin{aligned} h_2(t) + tD_1 + (1-t)D_2 &\geq H(X|U_1^{(t)}, U_2^{(t)}, Q^{(t)}) \\ &= h_2(t) + tH(Y_1|U_1^{(t)}, U_2^{(t)}, Q^{(t)}) \\ &\quad + (1-t)H(Y_2|U_1^{(t)}, U_2^{(t)}, Q^{(t)}), \end{aligned}$$

where the second equality follows by definition of X in (8). For convenience, define $H_1(P_t) \triangleq H(Y_1|U_1^{(t)}, U_2^{(t)}, Q^{(t)})$ and $H_2(P_t) \triangleq H(Y_2|U_1^{(t)}, U_2^{(t)}, Q^{(t)})$. Note the following two facts:

1. By continuity of entropy, the functions $H_1(\cdot)$ and $H_2(\cdot)$ are continuous on the compact domain $\mathcal{P}(R_1, R_2)$.

²Henceforth, we use the superscript (t) to explicitly denote the dependence of the auxiliary random variables on the distribution parametrized by t .

2. The above argument proves the existence of a function $\varphi : [0, 1] \rightarrow \mathcal{P}(R_1, R_2)$ which satisfies

$$tH_1(\varphi(t)) + (1-t)H_2(\varphi(t)) \leq tD_1 + (1-t)D_2 \text{ for all } t \in [0, 1].$$

These two facts satisfy the requirements of Lemma 7 (see Appendix D), and hence there exists $P_{t_1} \in \mathcal{P}(R_1, R_2)$, $P_{t_2} \in \mathcal{P}(R_1, R_2)$, and $\theta \in [0, 1]$ for which

$$\begin{aligned} \theta H_1(P_{t_1}) + (1-\theta)H_1(P_{t_2}) &\leq D_1 \\ \theta H_2(P_{t_1}) + (1-\theta)H_2(P_{t_2}) &\leq D_2. \end{aligned}$$

Timeshareing³ between distributions P_{t_1} and P_{t_2} with probabilities θ and $(1-\theta)$, respectively, yields a distribution $P^* \in \mathcal{P}(R_1, R_2)$ which satisfies $H_1(P^*) \leq D_1$ and $H_2(P^*) \leq D_2$. This proves the theorem. \square

4.3 A stronger converse

For the CEO problem, we are able to obtain a stronger converse result as discussed in Section 3.3. We can obtain a similar result for the multiterminal source coding problem. Indeed, the converse result we just proved continues to hold even when \hat{Y}_i^n is allowed to be any probability measure on \mathcal{Y}_i^n , rather than a product distribution. To see this, consider once again the proof of Theorem 6. Suppose (R_1, R_2, D_1, D_2) is achievable for this more general setting and replace (8) with

$$\mathbf{X} = \begin{cases} (\mathbf{Y}_1, 1) & \text{with probability } t \\ (\mathbf{Y}_2, 2) & \text{with probability } 1-t, \end{cases}$$

where \mathbf{Y}_1 and \mathbf{Y}_2 are the discrete, memoryless super-sources with corresponding super-symbols Y_1^n and Y_2^n , respectively. Then for the CEO problem defined by $(\mathbf{X}, \mathbf{Y}_1, \mathbf{Y}_2)$:

$$\begin{aligned} &\frac{1}{n} \sum_{j=1}^n \mathbb{E}d(\mathbf{X}_j, \hat{\mathbf{X}}_j) \\ &= \frac{t}{n} \sum_{j=1}^n \mathbb{E} \log \left(\frac{1}{t \hat{\mathbf{Y}}_{1,j}(\mathbf{Y}_{1,j})} \right) + \frac{1-t}{n} \sum_{j=1}^n \mathbb{E} \log \left(\frac{1}{(1-t) \hat{\mathbf{Y}}_{2,j}(\mathbf{Y}_{2,j})} \right) \\ &= h_2(t) + \frac{t}{n} \sum_{j=1}^n \mathbb{E}d(\mathbf{Y}_{1,j}, \hat{\mathbf{Y}}_{1,j}) + \frac{1-t}{n} \sum_{j=1}^n \mathbb{E}d(\mathbf{Y}_{2,j}, \hat{\mathbf{Y}}_{2,j}) \\ &\leq h_2(t) + tD_1 + (1-t)D_2, \end{aligned}$$

where we used the extended definition of $d(\cdot, \cdot)$ given in Section 3.3. Now, we simply invoke the strengthened converse for the CEO problem to yield a

³The timesharing scheme can be embedded in the timesharing variable Q , increasing the cardinality of \mathcal{Q} by a factor of two.

distribution $P_t \in \mathcal{P}(R_1, R_2)$ satisfying the single-letter inequality

$$\begin{aligned} h_2(t) + tD_1 + (1-t)D_2 &\geq H(X|U_1^{(t)}, U_2^{(t)}, Q^{(t)}) \\ &= h_2(t) + tH(Y_1|U_1^{(t)}, U_2^{(t)}, Q^{(t)}) \\ &\quad + (1-t)H(Y_2|U_1^{(t)}, U_2^{(t)}, Q^{(t)}), \end{aligned}$$

and we proceed as before.

4.4 An Example: The Daily Double

The *Daily Double* is a single bet that links together wagers on the winners of two consecutive horse races. Winning the Daily Double is dependent on both wagers winning together. In general, the outcomes of two consecutive races can be correlated (e.g. due to track conditions), so a gambler can potentially use this information to maximize his expected winnings. Let \mathcal{Y}_1 and \mathcal{Y}_2 be the set of horses running in the first and second races respectively. If horses y_1 and y_2 win their respective races, then the payoff is $o(y_1, y_2)$ dollars for each dollar invested in outcome $(Y_1, Y_2) = (y_1, y_2)$.

There are two betting strategies one can follow:

1. The gambler can wager a fraction $b_1(y_1)$ of his wealth on horse y_1 winning the first race and parlay his winnings by betting a fraction $b_2(y_2)$ of his wealth on horse y_2 winning the second race. In this case, the gambler's wealth relative is $b_1(Y_1)b_2(Y_2)o(Y_1, Y_2)$ upon learning the outcome of the Daily Double. We refer to this betting strategy as the *product-wager*.
2. The gambler can wager a fraction $b(y_1, y_2)$ of his wealth on horses (y_1, y_2) winning the first and second races, respectively. In this case, the gambler's wealth relative is $b(Y_1, Y_2)o(Y_1, Y_2)$ upon learning the outcome of the Daily Double. We refer to this betting strategy as the *joint-wager*.

Clearly the joint-wager includes the product-wager as a special case. However, the product-wager requires less effort to place, so the question is: how do the two betting strategies compare?

To make things interesting, suppose the gamblers have access to rate-limited information about the first and second race outcomes at rates R_1, R_2 respectively. Further, assume that $R_1 \leq H(Y_1)$, $R_2 \leq H(Y_2)$, and $R_1 + R_2 \leq H(Y_1, Y_2)$. For (R_1, R_2) and $p(y_1, y_2)$ given, let $\mathcal{P}(R_1, R_2)$ denote the set of joint pmf's of the form

$$p(q, y_1, y_2, u_1, u_2) = p(q)p(y_1, y_2)p(u_1|y_1, q)p(u_2|y_1, q)$$

which satisfy

$$\begin{aligned} R_1 &\geq I(Y_1; U_1|U_2, Q) \\ R_2 &\geq I(Y_2; U_2|U_1, Q) \\ R_1 + R_2 &\geq I(Y_1, Y_2; U_1, U_2|Q) \end{aligned}$$

for alphabets $\mathcal{U}_1, \mathcal{U}_2, \mathcal{Q}$ satisfying $|\mathcal{U}_i| \leq |\mathcal{Y}_i|$ and $|\mathcal{Q}| \leq 5$.

Typically, the quality of a bet is measured by the associated doubling rate (cf. [18]). Theorem 6 implies that the optimal doubling rate for the product-wager is given by:

$$\begin{aligned} W_{\text{p-w}}^*(p(y_1, y_2)) &= \sum_{y_1, y_2} p(y_1, y_2) \log b_1^*(y_1) b_2^*(y_2) o(y_1, y_2) \\ &= \mathbb{E} \log o(Y_1, Y_2) - \inf_{p \in \mathcal{P}(R_1, R_2)} \{H(Y_1|U_1, U_2, Q) + H(Y_2|U_1, U_2, Q)\}. \end{aligned}$$

Likewise, Theorem 4 implies that the optimal doubling rate for the joint-wager is given by:

$$\begin{aligned} W_{\text{j-w}}^*(p(y_1, y_2)) &= \sum_{y_1, y_2} p(y_1, y_2) \log b^*(y_1, y_2) o(y_1, y_2) \\ &= \mathbb{E} \log o(Y_1, Y_2) + \min\{R_1 - H(Y_1|Y_2), R_2 - H(Y_2|Y_1), \\ &\quad R_1 + R_2 - H(Y_1, Y_2)\}. \end{aligned}$$

It is important to note that we do not require the side informations to be the same for each type of wager, rather, the side informations are only provided at the same rates. Thus, the gambler placing the joint-wager receives side information at rates (R_1, R_2) that maximizes his doubling rate, while the gambler placing the product-wager receives (potentially different) side information at rates (R_1, R_2) that maximizes his doubling rate. However, as we will see shortly, for any rates (R_1, R_2) , there always exists rate-limited side information which simultaneously allows each type of gambler to attain their maximum doubling rate.

By combining the expressions for $W_{\text{p-w}}^*(p(y_1, y_2))$ and $W_{\text{j-w}}^*(p(y_1, y_2))$, we find that the difference in doubling rates is given by:

$$\begin{aligned} \Delta(R_1, R_2) &= W_{\text{j-w}}^*(p(y_1, y_2)) - W_{\text{p-w}}^*(p(y_1, y_2)) \\ &= \min\left\{R_1 - H(Y_1|Y_2), R_2 - H(Y_2|Y_1), R_1 + R_2 - H(Y_1, Y_2)\right\} \\ &\quad + \inf_{p \in \mathcal{P}(R_1, R_2)} \{H(Y_1|U_1, U_2, Q) + H(Y_2|U_1, U_2, Q)\} \tag{9} \\ &= \inf_{p \in \mathcal{P}(R_1, R_2)} \min\left\{R_1 - I(Y_1; U_1|U_2, Q) + I(Y_1; Y_2) - I(Y_1; U_2, Q) + H(Y_2|U_1, U_2, Q), \right. \\ &\quad R_2 - I(Y_2; U_2|U_1, Q) + I(Y_2; Y_1) - I(Y_2; U_1, Q) + H(Y_1|U_1, U_2, Q), \\ &\quad \left. R_1 + R_2 - I(Y_1, Y_2; U_1, U_2|Q) + I(Y_1; Y_2|U_1, U_2, Q)\right\} \\ &= \inf_{p \in \mathcal{P}(R_1, R_2)} I(Y_1; Y_2|U_1, U_2, Q). \tag{10} \end{aligned}$$

The final equality (10) follows since

- $R_1 \geq I(Y_1; U_1|U_2, Q)$ and $R_2 \geq I(Y_2; U_2|U_1, Q)$ for any $p \in \mathcal{P}(R_1, R_2)$.

- $I(Y_2; Y_1) \geq I(Y_2; U_1, Q)$ and $I(Y_1; Y_2) \geq I(Y_1; U_2, Q)$ for any $p \in \mathcal{P}(R_1, R_2)$ by the data processing inequality.
- The infimum in (9) is attained by a $p \in \mathcal{P}(R_1, R_2)$ satisfying $R_1 + R_2 = I(Y_1, Y_2; U_1, U_2|Q)$. See Lemma 8 in Appendix E for details.
- By definition of conditional mutual information,

$$H(Y_i|U_1, U_2, Q) \geq I(Y_1; Y_2|U_1, U_2, Q)$$

for $i = 1, 2$.

Let $p^* \in \mathcal{P}(R_1, R_2)$ be the distribution that attains the infimum in (9) (such a p^* always exists), then (10) yields

$$\begin{aligned} & W_{\text{j-w}}^*(p(y_1, y_2)) - W_{\text{p-w}}^*(p(y_1, y_2)) \\ &= \sum_{u_1, u_2, q} p^*(u_1, u_2, q) \sum_{y_1, y_2} p^*(y_1, y_2|u_1, u_2, q) \log \frac{p^*(y_1, y_2|u_1, u_2, q)}{p^*(y_1|u_1, u_2, q)p^*(y_2|u_1, u_2, q)} \\ &= \mathbb{E}_{p^*} \log o(Y_1, Y_2)p^*(Y_1, Y_2|U_1, U_2, Q) \\ &\quad - \mathbb{E}_{p^*} \log o(Y_1, Y_2)p^*(Y_1|U_1, U_2, Q)p^*(Y_2|U_1, U_2, Q). \end{aligned}$$

Hence, we can interpret the auxiliary random variables corresponding to p^* as optimal rate-limited side informations for *both* betting strategies. Moreover, optimal bets for each strategy are given by

1. $b^*(y_1, y_2) = p^*(y_1, y_2|u_1, u_2, q)$ for the joint-wager, and
2. $b_1^*(y_1) = p^*(y_1|u_1, u_2, q)$, $b_2^*(y_2) = p^*(y_2|u_1, u_2, q)$ for the product-wager.

Since $\mathcal{P}(R_1, R_2) \subseteq \mathcal{P}(R'_1, R'_2)$ for $R_1 \leq R'_1$ and $R_2 \leq R'_2$, the function $\Delta(R_1, R_2)$ is nonincreasing in R_1 and R_2 . Thus, the benefits of using the joint-wager over the product-wager diminish in the amount of side-information available. It is also not difficult to show that $\Delta(R_1, R_2)$ is jointly convex in (R_1, R_2) .

Furthermore, for rate-pairs (R_1, R_2) and (R'_1, R'_2) satisfying $R_1 < R'_1$ and $R_2 < R'_2$, there exist corresponding optimal joint- and product-wagers $b^*(y_1, y_2)$ and $b_1^*(y_1)b_2^*(y_2)$, and $b^{*'}(y_1, y_2)$ and $b_1^{*'}(y_1)b_2^{*'}(y_2)$, respectively, satisfying

$$D\left(b^{*'}(y_1, y_2) \middle| \middle| b_1^{*'}(y_1)b_2^{*'}(y_2)\right) < D\left(b^*(y_1, y_2) \middle| \middle| b_1^*(y_1)b_2^*(y_2)\right). \quad (11)$$

So, roughly speaking, the joint-wager and product-wager look “more alike” as the amount of side information is increased. The proof of the strict inequality in (11) can be inferred from the proof of Lemma 8 in Appendix E.

To conclude this example, we note that $\Delta(R_1, R_2)$ enjoys a great deal of symmetry near the origin in the sense that side information from either encoder contributes approximately the same amount to the improvement of the product-wager. We state this formally as a theorem:

Theorem 7. Define $\rho_m(Y_1, Y_2)$ to be the Hirschfeld-Gebelein-Rényi maximal correlation between random variables Y_1 and Y_2 . Then, $\Delta(R_1, R_2) \geq I(Y_1; Y_2) - \rho_m^2(Y_1, Y_2) \cdot (R_1 + R_2)$. Moreover, this bound is tight as $(R_1, R_2) \rightarrow (0, 0)$.

Proof. If $R_2 = 0$, then it is readily verified that $\Delta(R_1, 0)$ can be expressed as follows:

$$\Delta(R_1, 0) = I(Y_1; Y_2) - \max_{\substack{p(u_1|y_1): I(U_1; Y_1) = R_1, \\ U_1 \rightarrow Y_1 \rightarrow Y_2, |\mathcal{U}_1| \leq |\mathcal{Y}_1| + 1}} I(U_1; Y_2).$$

By symmetry:

$$\Delta(0, R_2) = I(Y_1; Y_2) - \max_{\substack{p(u_2|y_2): I(U_2; Y_2) = R_2, \\ U_2 \rightarrow Y_2 \rightarrow Y_1, |\mathcal{U}_2| \leq |\mathcal{Y}_2| + 1}} I(U_2; Y_1).$$

Here, we can apply a result of Erkip [19, Theorem 10] to evaluate the gradient of $\Delta(R_1, R_2)$ at $(R_1, R_2) = (0, 0)$:

$$\left. \frac{\partial}{\partial R_1} \Delta(R_1, R_2) \right|_{(R_1, R_2) = (0, 0)} = \left. \frac{\partial}{\partial R_2} \Delta(R_1, R_2) \right|_{(R_1, R_2) = (0, 0)} = -\rho_m^2(Y_1, Y_2). \quad (12)$$

Taking this one step further, for $\nu_1, \nu_2 > 0$, we can evaluate the one-sided derivative:

$$\lim_{\lambda \downarrow 0} \frac{\Delta(\lambda\nu_1, \lambda\nu_2) - \Delta(0, 0)}{\lambda} = -\rho_m^2(Y_1, Y_2) \cdot (\nu_1 + \nu_2). \quad (13)$$

Since $\Delta(R_1, R_2)$ is convex, we obtain a lower bound given by the supporting hyperplane tangent to $\Delta(R_1, R_2)$ at $(R_1, R_2) = (0, 0)$:

$$\Delta(R_1, R_2) \geq I(Y_1; Y_2) - \rho_m^2(Y_1, Y_2) \cdot (R_1 + R_2).$$

By (12) and convexity of $\Delta(R_1, R_2)$, this bound is tight as $(R_1, R_2) \rightarrow (0, 0)$.

To see that (13) holds, it suffices to prove the lower bound

$$\lim_{\lambda \downarrow 0} \frac{\Delta(\lambda\nu_1, \lambda\nu_2) - \Delta(0, 0)}{\lambda} \geq -\rho_m^2(Y_1, Y_2) \cdot (\nu_1 + \nu_2)$$

since the reverse inequality follows from convexity of $\Delta(R_1, R_2)$ and the known derivatives (12) along the coordinate axis. To this end, fix $\lambda, \nu_1, \nu_2 > 0$ and

observe that

$$\begin{aligned} & \frac{\Delta(\lambda\nu_1, \lambda\nu_2) - \Delta(0, 0)}{\lambda} \\ &= \frac{1}{\lambda} \inf_{p \in \mathcal{P}(\lambda\nu_1, \lambda\nu_2)} \left\{ I(Y_1, Y_2; U_1, U_2|Q) - I(Y_1; U_1, U_2|Q) - I(Y_2; U_1, U_2|Q) \right\} \end{aligned} \quad (14)$$

$$= (\nu_1 + \nu_2) - \frac{1}{\lambda} \left(I_{p^*}(Y_1; U_1, U_2|Q) + I_{p^*}(Y_2; U_1, U_2|Q) \right) \quad (15)$$

$$\begin{aligned} &= (\nu_1 + \nu_2) - \frac{1}{\lambda} \left(I_{p^*}(Y_1; U_1|U_2, Q) + I_{p^*}(Y_1; U_2|Q) \right. \\ &\quad \left. + I_{p^*}(Y_2; U_2|U_1, Q) + I_{p^*}(Y_2; U_1|Q) \right) \\ &\geq (\nu_1 + \nu_2) - \frac{1}{\lambda} \left(\nu_1 + I_{p^*}(Y_1; U_2|Q) + \nu_2 + I_{p^*}(Y_2; U_1|Q) \right) \end{aligned} \quad (16)$$

$$= -\frac{1}{\lambda} \left(I_{p^*}(Y_1; U_2|Q) + I_{p^*}(Y_2; U_1|Q) \right) \quad (17)$$

$$\geq -\frac{1}{\lambda} \left(\max_{\substack{p(\tilde{u}_2|y_2): I(\tilde{U}_2; Y_2) \leq \lambda\nu_2, \\ \tilde{U}_2 \rightarrow Y_2 \rightarrow Y_1, |\tilde{U}_2| \leq |Y_2|+1}} I(\tilde{U}_2; Y_1) + \max_{\substack{p(\tilde{u}_1|y_1): I(\tilde{U}_1; Y_1) \leq \lambda\nu_1, \\ \tilde{U}_1 \rightarrow Y_1 \rightarrow Y_2, |\tilde{U}_1| \leq |Y_1|+1}} I(\tilde{U}_1; Y_2) \right) \quad (18)$$

$$\xrightarrow{\lambda \downarrow 0} \nu_1 \frac{\partial}{\partial R_1} \Delta(R_1, R_2) \Big|_{(R_1, R_2)=(0,0)} + \nu_2 \frac{\partial}{\partial R_2} \Delta(R_1, R_2) \Big|_{(R_1, R_2)=(0,0)} \quad (19)$$

$$= -\rho_m(Y_1, Y_2)^2 \cdot (\nu_1 + \nu_2). \quad (20)$$

In the above string of inequalities

- (14) is simply follows from the definition of $\Delta(R_1, R_2)$.
- Equality (15) follows since Lemma 8 guarantees that the infimum is attained in (14) for some $p^* \in \mathcal{P}(\lambda\nu_1, \lambda\nu_2)$ satisfying $I_{p^*}(Y_1, Y_2; U_1, U_2|Q) = \lambda(\nu_1 + \nu_2)$. Here, we write $I_{p^*}(Y_1, Y_2; U_1, U_2|Q)$ to denote the mutual information $I(Y_1, Y_2; U_1, U_2|Q)$ evaluated for the distribution p^* .
- The inequality (16) follows since $I(Y_1; U_1|U_2, Q) \leq \lambda\nu_1$ and $I(Y_2; U_2|U_1, Q) \leq \lambda\nu_2$ for any $p \in \mathcal{P}(\lambda\nu_1, \lambda\nu_2)$.
- (18) follows since maximizing each term separately in (17) over a set of distributions that includes $\mathcal{P}(\lambda\nu_1, \lambda\nu_2)$ yields a lower bound.
- The limit (19) follows from (12) and positive homogeneity of the one-sided derivatives.
- Finally, (20) follows from the known derivatives along the coordinate axis, given by (12).

□

4.5 An Application: List Decoding

In the previous example, we did not take advantage of the stronger converse result which we proved in Section 4.3. In this section, we give an application that requires this strengthened result.

Formally, a 2-list code (of blocklength n consists) of encoding functions:

$$g_i^{(n)} : \mathcal{Y}_i^n \rightarrow \{1, \dots, M_i^{(n)}\} \text{ for } i = 1, 2$$

and list decoding functions

$$\begin{aligned} L_1^{(n)} &: \{1, \dots, M_1^{(n)}\} \times \{1, \dots, M_2^{(n)}\} \rightarrow 2^{\mathcal{Y}_1^n} \\ L_2^{(n)} &: \{1, \dots, M_1^{(n)}\} \times \{1, \dots, M_2^{(n)}\} \rightarrow 2^{\mathcal{Y}_2^n}. \end{aligned}$$

A list decoding tuple $(R_1, R_2, \Delta_1, \Delta_2)$ is achievable if, for any $\epsilon > 0$, there exists a 2-list code of blocklength n satisfying the rate constraints

$$\begin{aligned} \frac{1}{n} \log M_1^{(n)} &\leq R_1 + \epsilon \\ \frac{1}{n} \log M_2^{(n)} &\leq R_2 + \epsilon, \end{aligned}$$

and the probability of list-decoding error constraints

$$\begin{aligned} \Pr \left[Y_1^n \notin L_1^{(n)} \left(g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n) \right) \right] &\leq \epsilon, \\ \Pr \left[Y_2^n \notin L_2^{(n)} \left(g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n) \right) \right] &\leq \epsilon. \end{aligned}$$

with list sizes

$$\begin{aligned} \frac{1}{n} \log |L_1^{(n)}| &\leq \Delta_1 + \epsilon \\ \frac{1}{n} \log |L_2^{(n)}| &\leq \Delta_2 + \epsilon. \end{aligned}$$

With a 2-list code so defined, the following theorem shows that the 2-list decoding problem and multiterminal source coding problem under logarithmic loss are equivalent (inasmuch as the achievable regions are identical):

Theorem 8. *The list decoding tuple $(R_1, R_2, \Delta_1, \Delta_2)$ is achievable if and only if*

$$\begin{aligned} R_1 &\geq I(U_1; Y_1 | U_2, Q) \\ R_2 &\geq I(U_2; Y_2 | U_1, Q) \\ R_1 + R_2 &\geq I(U_1, U_2; Y_1, Y_2 | Q) \\ \Delta_1 &\geq H(Y_1 | U_1, U_2, Q) \\ \Delta_2 &\geq H(Y_2 | U_1, U_2, Q). \end{aligned}$$

for some joint distribution

$$p(y_1, y_2, u_1, u_2, q) = p(y_1, y_2)p(u_1|y_1, q)p(u_2|y_2, q)p(q),$$

where $|\mathcal{U}_1| \leq |\mathcal{Y}_1|$, $|\mathcal{U}_2| \leq |\mathcal{Y}_2|$, and $|\mathcal{Q}| \leq 5$.

Remark 2. We note that a similar connection to list decoding can be made for other multiterminal scenarios, in particular the CEO problem.

To prove the theorem, we require a slightly modified version of [20, Lemma 1]:

Lemma 2. If the list decoding tuple $(R_1, R_2, \Delta_1, \Delta_2)$ is achieved by a sequence of 2-list codes $\{g_1^{(n)}, g_2^{(n)}, L_1^{(n)}, L_2^{(n)}\}_{n \rightarrow \infty}$, then

$$\begin{aligned} H(Y_1^n | g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n)) &\leq |L_1^{(n)}| + n\epsilon_n \\ H(Y_2^n | g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n)) &\leq |L_2^{(n)}| + n\epsilon_n, \end{aligned}$$

where $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$.

Proof. The proof is virtually identical to that of [20, Lemma 1], and is therefore omitted. \square

Proof of Theorem 8. First observe that the direct part is trivial. Indeed, for a joint distribution $p(y_1, y_2, u_1, u_2, q) = p(y_1, y_2)p(u_1|y_1, q)p(u_2|y_2, q)p(q)$, apply the Berger-Tung achievability scheme and take $L_i^{(n)}$ to be the set of y_i^n sequences which are jointly typical with the decoded quantizations (U_1^n, U_2^n) . This set has cardinality no larger than $2^{n(H(Y_i|U_1, U_2, Q) + \epsilon)}$, which proves achievability.

To see the converse, note that setting

$$\hat{Y}_i^n = \Pr \left[Y_i^n | g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n) \right]$$

achieves a logarithmic loss of $\frac{1}{n} H(Y_i^n | g_1^{(n)}(Y_1^n), g_2^{(n)}(Y_2^n))$ for source i in the setting where reproductions are not restricted to product distributions. Applying the strengthened converse of Theorem 6 together with Lemma 2 yields the desired result. \square

5 Relationship to the General Multiterminal Source Coding Problem

In this section, we relate our results for logarithmic loss to multiterminal source coding problems with arbitrary distortion measures and reproduction alphabets.

As before, we let $\{Y_{1,j}, Y_{2,j}\}_{j=1}^n$ be a sequence of n independent, identically distributed random variables with finite alphabets \mathcal{Y}_1 and \mathcal{Y}_2 , respectively, and joint pmf $p(y_1, y_2)$.

In this section, the reproduction alphabets $\check{\mathcal{Y}}_i$, $i = 1, 2$, are arbitrary. We also consider generic distortion measures:

$$\check{d}_i : \mathcal{Y}_i \times \check{\mathcal{Y}}_i \rightarrow \mathbb{R}^+ \text{ for } i = 1, 2,$$

where \mathbb{R}^+ denotes the set of nonnegative real numbers. The sequence distortion is then defined as follows:

$$\check{d}_i(y_i^n, \check{y}_i^n) = \frac{1}{n} \sum_{j=1}^n \check{d}_i(y_{i,j}, \check{y}_{i,j}).$$

We will continue to let $d(\cdot, \cdot)$ and $\hat{\mathcal{Y}}_1, \hat{\mathcal{Y}}_2$ denote the logarithmic loss distortion measure and the associated reproduction alphabets, respectively.

A rate distortion code (of blocklength n) consists of encoding functions:

$$\check{g}_i^{(n)} : \mathcal{Y}_i^n \rightarrow \{1, \dots, M_i^{(n)}\} \text{ for } i = 1, 2$$

and decoding functions

$$\check{\psi}_i^{(n)} : \{1, \dots, M_1^{(n)}\} \times \{1, \dots, M_2^{(n)}\} \rightarrow \check{\mathcal{Y}}_i^n \text{ for } i = 1, 2.$$

A rate distortion vector (R_1, R_2, D_1, D_2) is strict-sense achievable if there exists a blocklength n , encoding functions $\check{g}_1^{(n)}, \check{g}_2^{(n)}$ and a decoder $(\check{\psi}_1^{(n)}, \check{\psi}_2^{(n)})$ such that

$$R_i \geq \frac{1}{n} \log M_i^{(n)} \text{ for } i = 1, 2 \quad (21)$$

$$D_i \geq \mathbb{E} \check{d}_i(Y_i^n, \check{Y}_i^n) \text{ for } i = 1, 2. \quad (22)$$

Where

$$\check{Y}_i^n = \check{\psi}_i^{(n)}(\check{g}_1^{(n)}(Y_1^n), \check{g}_2^{(n)}(Y_2^n)) \text{ for } i = 1, 2.$$

For these functions, we define the quantity

$$\beta_i(\check{g}_1^{(n)}, \check{g}_2^{(n)}, \check{\psi}_1^{(n)}, \check{\psi}_2^{(n)}) := \frac{1}{n} \sum_{j=1}^n \mathbb{E} \log \left(\sum_{y_i \in \mathcal{Y}_i} 2^{-\check{d}_i(y_i, \check{Y}_{i,j})} \right) \text{ for } i = 1, 2. \quad (23)$$

Now, let $\beta_i(R_1, R_2, D_1, D_2)$ be the infimum of the $\beta_i(\check{g}_1^{(n)}, \check{g}_2^{(n)}, \check{\psi}_1^{(n)}, \check{\psi}_2^{(n)})$'s, where the infimum is taken over all codes that achieve the rate distortion vector (R_1, R_2, D_1, D_2) .

At this point it is instructive to pause and consider some examples.

Example 1 (Binary Sources and Hamming Distortion). For $i = 1, 2$, let $\check{\mathcal{Y}}_i = \mathcal{Y}_i = \{0, 1\}$ and let \check{d}_i be the α -scaled Hamming distortion measure:

$$\check{d}_i(y_i, \check{y}_i) = \begin{cases} 0 & \text{if } \check{y}_i = y_i, \\ \alpha & \text{if } \check{y}_i \neq y_i. \end{cases}$$

In this case,

$$\sum_{y_i \in \mathcal{Y}_i} 2^{-\check{d}_i(y_i, \check{Y}_{i,j})} = 2^0 + 2^{-\alpha}, \quad (24)$$

so $\beta_i(R_1, R_2, D_1, D_2) = \log(1 + 2^{-\alpha})$ for any (R_1, R_2, D_1, D_2) . This notion that $\beta_i(R_1, R_2, D_1, D_2)$ is a constant extends to all distortion measures for which the columns of the $|\mathcal{Y}_i| \times |\check{\mathcal{Y}}_i|$ distortion matrix are permutations of one another.

Example 2 (Binary Sources and Erasure Distortion). For $i = 1, 2$, let $\mathcal{Y}_i = \{0, 1\}$, $\check{\mathcal{Y}}_i = \{0, 1, e\}$ and let \check{d}_i be the standard erasure distortion measure:

$$\check{d}_i(y_i, \check{y}_i) = \begin{cases} 0 & \text{if } \check{y}_i = y_i \\ 1 & \text{if } \check{y}_i = e \\ \infty & \text{if } \check{y}_i \in \{0, 1\} \text{ and } \check{y}_i \neq y_i. \end{cases}$$

In this case,

$$\sum_{y_i \in \mathcal{Y}_i} 2^{-\check{d}_i(y_i, \check{Y}_{i,j})} = \begin{cases} 2^{-\infty} + 2^0 = 1 & \text{if } \check{Y}_{i,j} \in \{0, 1\} \\ 2^{-1} + 2^{-1} = 1 & \text{if } \check{Y}_{i,j} = e. \end{cases} \quad (25)$$

so $\beta_i(R_1, R_2, D_1, D_2) = 0$ for any (R_1, R_2, D_1, D_2) . This result can easily be extended to erasure distortion on larger alphabets by setting the penalty to $\log |\check{\mathcal{Y}}_i|$ when $\check{Y}_i = e$.

Theorem 9. Suppose (R_1, R_2, D_1, D_2) is strict-sense achievable for the general multiterminal source coding problem. Then

$$\left. \begin{aligned} R_1 &\geq I(U_1; Y_1 | U_2, Q) \\ R_2 &\geq I(U_2; Y_2 | U_1, Q) \\ R_1 + R_2 &\geq I(U_1, U_2; Y_1, Y_2 | Q) \\ D_1 &\geq H(Y_1 | U_1, U_2, Q) - \beta_1(R_1, R_2, D_1, D_2) \\ D_2 &\geq H(Y_2 | U_1, U_2, Q) - \beta_2(R_1, R_2, D_1, D_2) \end{aligned} \right\} \quad (26)$$

for some joint distribution $p(y_1, y_2)p(q)p(u_1|y_1, q)p(u_2|y_2, q)$ with $|\mathcal{U}_i| \leq |\mathcal{Y}_i|$ and $|\mathcal{Q}| \leq 5$.

Proof. Since (R_1, R_2, D_1, D_2) is strict-sense achievable, there exists a block-length n , encoding functions $\check{g}_1^{(n)}, \check{g}_2^{(n)}$ and a decoder $(\check{\psi}_1^{(n)}, \check{\psi}_2^{(n)})$ satisfying (21)-(22). Given these functions, the decoder can generate reproductions $\check{Y}_1^n, \check{Y}_2^n$ satisfying the average distortion constraints (22). From the reproduction \check{Y}_i^n , we construct the reproduction \hat{Y}_i^n as follows:

$$\hat{Y}_j(y_i) = \frac{2^{-\check{d}_i(y_i, \check{Y}_{i,j})}}{\sum_{y'_i \in \mathcal{Y}_i} 2^{-\check{d}_i(y'_i, \check{Y}_{i,j})}}.$$

Now, using the logarithmic loss distortion measure, observe that \hat{Y}_i^n satisfies

$$\begin{aligned}
\mathbb{E}d(Y_i^n, \hat{Y}_i^n) &= \frac{1}{n} \sum_{j=1}^n \mathbb{E} \log \left(2^{\check{d}_i(Y_{i,j}, \check{Y}_{i,j})} \right) + \frac{1}{n} \sum_{j=1}^n \mathbb{E} \log \left(\sum_{y'_i \in \mathcal{Y}_i} 2^{-\check{d}_i(y'_i, \check{Y}_{i,j})} \right) \\
&= \frac{1}{n} \sum_{j=1}^n \mathbb{E} \check{d}_i(Y_{i,j}, \check{Y}_{i,j}) + \beta_i \left(\check{g}_1^{(n)}, \check{g}_2^{(n)}, \check{\psi}_1^{(n)}, \check{\psi}_2^{(n)} \right) \\
&\leq D_i + \beta_i \left(\check{g}_1^{(n)}, \check{g}_2^{(n)}, \check{\psi}_1^{(n)}, \check{\psi}_2^{(n)} \right) \\
&:= \tilde{D}_i.
\end{aligned}$$

Thus, $(R_1, R_2, \tilde{D}_1, \tilde{D}_2)$ is achievable for the multiterminal source coding problem with the logarithmic loss distortion measure. Applying Theorem 6 and taking the infimum over all coding schemes that achieve (R_1, R_2, D_1, D_2) proves the theorem. \square

This outer bound is interesting because the region is defined over the same set of probability distributions that define the Berger-Tung inner bound. While the β_i 's can be difficult to compute in general, we have shown that they can be readily determined for many popular distortion measures. As an application, we now give a quantitative approximation of the rate distortion region for binary sources subject to Hamming distortion constraints. Before proceeding, we prove the following lemma.

Lemma 3. *Suppose $(R_1, R_2, \tilde{D}_1, \tilde{D}_2)$ is strict-sense achievable for the multiterminal source coding problem with binary sources and \check{d}_i equal to the α_i -scaled Hamming distortion measure, for $i = 1, 2$. Then the Berger-Tung achievability scheme can achieve a point (R_1, R_2, D_1, D_2) satisfying*

$$D_i - \tilde{D}_i \leq \left(\frac{\alpha_i}{2} - 1 \right) H_i + \log(1 + 2^{-\alpha_i})$$

for some $H_i \in [0, 1]$, $i = 1, 2$.

Proof. By Theorem 9, $(R_1, R_2, \tilde{D}_1, \tilde{D}_2)$ satisfy (26) for some joint distribution $p(y_1, y_2)p(q)p(u_1|y_1, q)p(u_2|y_2, q)$. For this distribution, define the reproduction functions

$$\check{Y}_i(U_1, U_2, Q) = \arg \max_{y_i} p(y_i|U_1, U_2, Q) \text{ for } i = 1, 2. \quad (27)$$

Then, observe that for $i = 1, 2$:

$$\begin{aligned}
\mathbb{E}\check{d}_i(Y_i, \check{Y}_i) &= \sum_{u_1, u_2, q} p(u_1, u_2, q) \left[\alpha_i \cdot \min_{y_i} p(y_i|u_1, u_2, q) + 0 \cdot \max_{y_i} p(y_i|u_1, u_2, q) \right] \\
&= \alpha_i \sum_{u_1, u_2, q} p(u_1, u_2, q) \cdot \min_{y_i} p(y_i|u_1, u_2, q) \\
&\leq \frac{\alpha_i}{2} \sum_{u_1, u_2, q} p(u_1, u_2, q) \cdot H(Y_i|U_1, U_2, Q = u_1, u_2, q) \\
&= \frac{\alpha_i}{2} H(Y_i|U_1, U_2, Q).
\end{aligned} \tag{28}$$

Where (28) follows from the fact that $2p \leq h_2(p)$ for $0 \leq p \leq 0.5$. Thus, $D_i = \frac{\alpha_i}{2} H(Y_i|U_1, U_2, Q)$ is achievable for rates (R_1, R_2) using the Berger-Tung achievability scheme. Combining this with the fact that $\tilde{D}_i \geq H(Y_i|U_1, U_2, Q) - \log(1 + 2^{-\alpha_i})$, we see that

$$D_i - \tilde{D}_i \leq \frac{\alpha_i}{2} H(Y_i|U_1, U_2, Q) - H(Y_i|U_1, U_2, Q) + \log(1 + 2^{-\alpha_i}).$$

□

Lemma 3 allows us to give a quantitative outer bound on the achievable rate distortion region in terms of the Berger-Tung inner bound.

Corollary 1. *Suppose $(R_1, R_2, \tilde{D}_1^{(1)}, \tilde{D}_2^{(1)})$ is strict-sense achievable for the multiterminal source coding problem with binary sources and \check{d}_i equal to the standard 1-scaled Hamming distortion measure, for $i = 1, 2$. Then the Berger-Tung achievability scheme can achieve a point $(R_1, R_2, D_1^{(1)}, D_2^{(1)})$, where*

$$D_i^{(1)} - \tilde{D}_i^{(1)} \leq \frac{1}{2} \log\left(\frac{5}{4}\right) < 0.161 \text{ for } i = 1, 2.$$

Proof. For rates (R_1, R_2) , note that distortions $(\tilde{D}_1, \tilde{D}_2)$ are strict-sense achievable for the α_i -scaled Hamming distortion measures if and only if distortions $(\tilde{D}_1^{(1)}, \tilde{D}_2^{(1)}) = (\frac{1}{\alpha_1} \tilde{D}_1, \frac{1}{\alpha_2} \tilde{D}_2)$ are strict-sense achievable for the 1-scaled Hamming distortion measure. Likewise, the point (R_1, R_2, D_1, D_2) is achieved by the Berger-Tung coding scheme for the α_i -scaled Hamming distortion measures if and only if $(R_1, R_2, \frac{1}{\alpha_1} D_1, \frac{1}{\alpha_2} D_2)$ is achieved by the Berger-Tung coding scheme for the 1-scaled Hamming distortion measure.

Thus, applying Lemma 3, we can use the Berger-Tung achievability scheme to achieve a point $(R_1, R_2, D_1^{(1)}, D_2^{(1)})$ satisfying

$$\begin{aligned}
D_i^{(1)} - \tilde{D}_i^{(1)} &= \frac{1}{\alpha_i} (D_i - \tilde{D}_i) \\
&\leq \frac{1}{\alpha_i} \left(\frac{\alpha_i}{2} - 1 \right) H_i + \frac{1}{\alpha_i} \log(1 + 2^{-\alpha_i}) \\
&= \left(\frac{1}{2} - \frac{1}{\alpha_i} \right) H_i + \frac{1}{\alpha_i} \log(1 + 2^{-\alpha_i})
\end{aligned} \tag{29}$$

for some $H_i \in [0, 1]$. We can optimize (29) over α_i to find the minimum gap for a given H_i . Maximizing over $H_i \in [0, 1]$ then gives the worst-case gap. Straightforward calculus yields the saddle-point:

$$\begin{aligned} & \max_{H_i \in [0, 1]} \inf_{\alpha_i > 0} \left\{ \left(\frac{1}{2} - \frac{1}{\alpha_i} \right) H_i + \frac{1}{\alpha_i} \log(1 + 2^{-\alpha_i}) \right\} \\ &= \inf_{\alpha_i > 0} \max_{H_i \in [0, 1]} \left\{ \left(\frac{1}{2} - \frac{1}{\alpha_i} \right) H_i + \frac{1}{\alpha_i} \log(1 + 2^{-\alpha_i}) \right\} \\ &= \frac{1}{2} \log \left(\frac{5}{4} \right) < 0.161, \end{aligned}$$

which is achieved for $\alpha_i = 2$ and any $H \in [0, 1]$. \square

Remark 3. *We note briefly that this estimate can potentially be improved if one knows more about the source distribution.*

6 Concluding Remarks

One immediate direction for further work would be to extend our results to more than two encoders. For the CEO problem, our results can be extended to an arbitrary number of encoders. This extension is proved in Appendix B.

On the other hand, generalizing the results for the two-encoder source coding problem with distortion constraints on Y_1 and Y_2 poses a significant challenge. The obvious point of difficulty in the proof is extending the interpolation argument to higher dimensions so that it yields a distribution with the desired properties. In fact, a “quick-fix” to the interpolation argument alone would not be sufficient since this would imply that the Berger-Tung inner bound is tight for more than two encoders. This is known to be false (even for the logarithmic loss distortion measure) since the Berger-Tung achievability scheme is not optimal for the lossless modulo-sum problem studied by Körner and Marton in [21].

Acknowledgement

The authors would like to thank Professors Suhas Diggavi and Aaron Wagner for the helpful discussions on this topic.

A Cardinality Bounds on Auxiliary Random Variables

In order to obtain tight cardinality bounds on the auxiliary random variables used throughout this paper, we refer to a recent result by Jana. In [17], the author carefully applies the Caratheodory-Fenchel-Eggleston theorem in order to obtain tight cardinality bounds on the auxiliary random variables in the

Berger-Tung inner bound. This result extends the results and techniques employed by Gu and Effros for the Wyner-Ahlsvede-Körner problem [22], and by Gu, Jana, and Effros for the Wyner-Ziv problem [23]. We now state Jana's result, appropriately modified for our purposes:

Consider an arbitrary joint distribution $p(v, y_1, \dots, y_m)$ with random variables V, Y_1, \dots, Y_m coming from alphabets $\mathcal{V}, \mathcal{Y}_1, \dots, \mathcal{Y}_m$ respectively.

Let $d_l : \mathcal{V} \times \hat{\mathcal{V}}_l \rightarrow \mathbb{R}$, $1 \leq l \leq L$ be arbitrary distortion measures defined for possibly different reproduction alphabets $\hat{\mathcal{V}}_l$.

Definition 6. Define \mathcal{A}^* to be the set of $(m + L)$ -vectors $(R_1, \dots, R_m, D_1, \dots, D_L)$ satisfying the following conditions:

1. auxiliary random variables U_1, \dots, U_m exist such that

$$\sum_{i \in \mathcal{I}} R_i \geq I(Y_{\mathcal{I}}; U_{\mathcal{I}} | U_{\mathcal{I}^c}), \text{ for all } \mathcal{I} \subseteq \{1, \dots, m\}, \text{ and}$$

2. mappings $\psi_l : \mathcal{U}_1 \times \dots \times \mathcal{U}_m \rightarrow \hat{\mathcal{V}}_l$, $1 \leq l \leq L$ exist such that

$$\mathbb{E} d_l(V, \psi_l(U_1, \dots, U_m)) \leq D_l$$

for some joint distribution

$$p(v, y_1, \dots, y_m) \prod_{j=1}^m p(u_j | y_j).$$

Lemma 4 (Lemma 2.2 from [17]). Every extreme point of \mathcal{A}^* corresponds to some choice of auxiliary variables U_1, \dots, U_m with alphabet sizes $|\mathcal{U}_j| \leq |\mathcal{Y}_j|$, $1 \leq j \leq m$.

In order to obtain the cardinality bounds for the CEO problem, we simply let $L = 1$, $V = X$, and $\hat{\mathcal{V}}_1 = \hat{\mathcal{X}}$. Defining

$$d_1(x, \hat{x}) = \log \left(\frac{1}{\hat{x}(x)} \right),$$

we see that $\overline{\mathcal{RD}}_{CEO}^* = \text{conv}(\mathcal{A}^*)$, where $\text{conv}(\mathcal{A}^*)$ denotes the convex hull of \mathcal{A}^* . Therefore, Lemma 4 implies that all extreme points of $\overline{\mathcal{RD}}_{CEO}^*$ are achieved with a choice of auxiliary random variables U_1, \dots, U_m with alphabet sizes $|\mathcal{U}_j| \leq |\mathcal{Y}_j|$, $1 \leq j \leq m$. By timesharing between extreme points, any point in $\overline{\mathcal{RD}}_{CEO}^*$ can be achieved for these alphabet sizes.

Obtaining the cardinality bounds for the multiterminal source coding problem proceeds in a similar fashion. In particular, let $L = m = 2$, $V = (Y_1, Y_2)$, and $\hat{\mathcal{V}}_j = \hat{\mathcal{Y}}_j$, $j = 1, 2$. Defining

$$d_j((y_1, y_2), \hat{y}_j) = \log \left(\frac{1}{\hat{y}_j(y_j)} \right) \text{ for } j = 1, 2,$$

we see that $\overline{\mathcal{RD}}^* = \text{conv}(\mathcal{A}^*)$. In this case, Lemma 4 implies that all extreme points of $\overline{\mathcal{RD}}^*$ are achieved with a choice of auxiliary random variables U_1, U_2 with alphabet sizes $|\mathcal{U}_j| \leq |\mathcal{Y}_j|$, $1 \leq j \leq 2$. By timesharing between extreme points, any point in $\overline{\mathcal{RD}}^*$ can be achieved for these alphabet sizes.

In order to obtain cardinality bounds on the timesharing variable Q , we can apply Caratheodory's theorem (cf. [24]). In particular, if $C \subset \mathbb{R}^n$ is compact, then any point in $\text{conv}(C)$ is a convex combination of at most $n + 1$ points of C . Taking C to be the closure of the set of extreme points of \mathcal{A}^* is sufficient for our purposes (boundedness of C can be dealt with by a standard truncation argument).

B Extension of CEO Results to m Encoders

In this appendix, we prove the generalization of Theorem 3 to m encoders, which essentially amounts to extending the argument in the proof of Theorem 3 to the general case. We begin by stating the m -encoder generalizations of Theorems 1 and 2, the proofs of which are trivial extensions of the proofs given for the two-encoder case and are therefore omitted.

Definition 7. Let $\mathcal{R}_{CEO,m}^i$ be the set of all (R_1, \dots, R_m, D) satisfying

$$\sum_{i \in \mathcal{I}} R_i \geq I(Y_{\mathcal{I}}; U_{\mathcal{I}} | U_{\mathcal{I}^c}, Q) \text{ for all } \mathcal{I} \subseteq \{1, \dots, m\}$$

$$D \geq H(X | U_1, \dots, U_m, Q).$$

for some joint distribution $p(q)p(x) \prod_{i=1}^m p(y_i|x)p(u_i|y_i, q)$.

Theorem 10. All rate distortion vectors $(R_1, \dots, R_m, D) \in \mathcal{R}_{CEO,m}^i$ are achievable.

Definition 8. Let $\mathcal{R}_{CEO,m}^o$ be the set of (R_1, \dots, R_m, D) satisfying

$$\sum_{i \in \mathcal{I}} R_i \geq \sum_{i \in \mathcal{I}} I(U_i; Y_i | X, Q) + H(X | U_{\mathcal{I}^c}, Q) - D \text{ for all } \mathcal{I} \subseteq \{1, \dots, m\} \quad (30)$$

$$D \geq H(X | U_1, \dots, U_m, Q). \quad (31)$$

for some joint distribution $p(q)p(x) \prod_{i=1}^m p(y_i|x)p(u_i|y_i, q)$.

Theorem 11. If (R_1, \dots, R_m, D) is strict-sense achievable, then $(R_1, \dots, R_m, D) \in \mathcal{R}_{CEO,m}^o$.

Given the definitions of $\mathcal{R}_{CEO,m}^i$ and $\mathcal{R}_{CEO,m}^o$, the generalization of Theorem 3 to m encoders is an immediate consequence of the following lemma:

Lemma 5. $\mathcal{R}_{CEO,m}^o \subseteq \mathcal{R}_{CEO,m}^i$.

Proof. Suppose $(R_1, \dots, R_m, D) \in \mathcal{R}_{CEO,m}^o$, then by definition there exists $p(q)$ and conditional distributions $\{p(u_i|y_i, q)\}_{i=1}^m$ so that (30) and (31) are satisfied. For the joint distribution corresponding to $p(q)$ and conditional distributions $p\{(u_i|y_i, q)\}_{i=1}^m$, define $\mathcal{P}_D \subset \mathbb{R}^m$ to be the polytope defined by the inequalities (30). Now, to show $(R_1, \dots, R_m, D) \in \mathcal{R}_{CEO,m}^i$, it suffices to show that each extreme point of \mathcal{P}_D is dominated by a point in $\mathcal{R}_{CEO,m}^i$ that achieves distortion at most D .

To this end, define the set function $f : 2^{[m]} \rightarrow \mathbb{R}$ as follows:

$$\begin{aligned} f(\mathcal{I}) &:= I(Y_{\mathcal{I}}; U_{\mathcal{I}}|U_{\mathcal{I}^c}, Q) - (D - H(X|U_1, \dots, U_m, Q)) \\ &= \sum_{i \in \mathcal{I}} I(U_i; Y_i|X, Q) + H(X|U_{\mathcal{I}^c}, Q) - D. \end{aligned}$$

It can be verified that the function f and the function $f^+(\mathcal{I}) = \max\{f(\mathcal{I}), 0\}$ are supermodular functions (see Appendix C). By construction, \mathcal{P}_D is equal to the set of (R_1, \dots, R_m) which satisfy:

$$\sum_{i \in \mathcal{I}} R_i \geq f^+(\mathcal{I}).$$

It follows by basic results in submodular optimization (see Appendix C) that, for a linear ordering $i_1 \prec i_2 \prec \dots \prec i_m$ of $\{1, \dots, m\}$, an extreme point of \mathcal{P}_D can be greedily computed as follows:

$$\tilde{R}_{i_j} = f^+(\{i_1, \dots, i_j\}) - f^+(\{i_1, \dots, i_{j-1}\}) \text{ for } j = 1, \dots, m.$$

Furthermore, all extreme points of \mathcal{P}_D can be enumerated by looking over all linear orderings $i_1 \prec i_2 \prec \dots \prec i_m$ of $\{1, \dots, m\}$. Each ordering of $\{1, \dots, m\}$ is analyzed in the same manner, hence we assume (for notational simplicity) that the ordering we consider is the natural ordering $i_j = j$.

Let j be the first index for which $\tilde{R}_j > 0$. Then, by construction,

$$\tilde{R}_k = I(U_k; Y_k|U_{k+1}, \dots, U_m, Q) \text{ for all } k > j.$$

Furthermore, we must have $f(\{1, \dots, j'\}) \leq 0$ for all $j' < j$. Thus, \tilde{R}_j can be expressed as

$$\begin{aligned} \tilde{R}_j &= \sum_{i=1}^j I(Y_i; U_i|X, Q) + H(X|U_{j+1}, \dots, U_m, Q) - D \\ &= I(Y_j; U_j|U_{j+1}, \dots, U_m, Q) + f(\{1, \dots, j-1\}) \\ &= (1 - \theta)I(Y_j; U_j|U_{j+1}, \dots, U_m, Q), \end{aligned}$$

where $\theta \in [0, 1)$ is defined as:

$$\begin{aligned} \theta &= \frac{-f(\{1, \dots, j-1\})}{I(Y_j; U_j|U_{j+1}, \dots, U_m, Q)} \\ &= \frac{D - H(X|U_1, \dots, U_m, Q) - I(U_1, \dots, U_{j-1}; Y_1, \dots, Y_{j-1}|U_j, \dots, U_m, Q)}{I(Y_j; U_j|U_{j+1}, \dots, U_m, Q)}. \end{aligned}$$

By the results of Theorem 10, the rates $(\tilde{R}_1, \dots, \tilde{R}_m)$ permit the following coding scheme: For a fraction $(1 - \theta)$ of the time, a codebook can be used that allows the decoder to recover U_j^n, \dots, U_m^n with high probability. The other fraction θ of the time, a codebook can be used that allows the decoder to recover U_{j+1}^n, \dots, U_m^n with high probability. As $n \rightarrow \infty$, this coding scheme can achieve distortion

$$\begin{aligned}
\tilde{D} &= (1 - \theta)H(X|U_j, \dots, U_m, Q) + \theta H(X|U_{j+1}, \dots, U_m, Q) \\
&= H(X|U_j, \dots, U_m, Q) + \theta I(X; U_j|U_{j+1}, \dots, U_m, Q) \\
&= H(X|U_j, \dots, U_m, Q) + \frac{I(X; U_j|U_{j+1}, \dots, U_m, Q)}{I(Y_j; U_j|U_{j+1}, \dots, U_m, Q)} \times \\
&\quad [D - H(X|U_1, \dots, U_m, Q) - I(U_1, \dots, U_{j-1}; Y_1, \dots, Y_{j-1}|U_j, \dots, U_m, Q)] \\
&\leq H(X|U_j, \dots, U_m, Q) + D - H(X|U_1, \dots, U_m, Q) \\
&\quad - I(U_1, \dots, U_{j-1}; Y_1, \dots, Y_{j-1}|U_j, \dots, U_m, Q) \tag{32}
\end{aligned}$$

$$\begin{aligned}
&= D + I(X; U_1, \dots, U_{j-1}|U_j, \dots, U_m, Q) \\
&\quad - I(U_1, \dots, U_{j-1}; Y_1, \dots, Y_{j-1}|U_j, \dots, U_m, Q) \\
&= D - I(U_1, \dots, U_{j-1}; Y_1, \dots, Y_{j-1}|X, U_j, \dots, U_m, Q) \\
&\leq D. \tag{33}
\end{aligned}$$

In the preceding string of inequalities (32) follows since U_j is conditionally independent of everything else given (Y_j, Q) , and (33) follows from the non-negativity of mutual information.

Therefore, for every extreme point $(\tilde{R}_1, \dots, \tilde{R}_m)$ of \mathcal{P}_D , the point $(\tilde{R}_1, \dots, \tilde{R}_m, D)$ lies in $\mathcal{R}_{CEO,m}^i$. This proves the lemma. \square

Finally, we remark that the results of Appendix A imply that it suffices to consider auxiliary random variables U_1, \dots, U_m with alphabet sizes $|\mathcal{U}_j| \leq |\mathcal{Y}_j|$, $1 \leq j \leq m$. The timesharing variable Q requires an alphabet size bounded by $|\mathcal{Q}| \leq m + 2$.

C Supermodular Functions

In this appendix, we review some basic results in submodular optimization that were used in Appendix B to prove Lemma 5. We tailor our statements toward supermodularity, since this is the property we require in Appendix B.

We begin by defining a supermodular function.

Definition 9. *Let $E = \{1, \dots, n\}$ be a finite set. A function $s : 2^E \rightarrow \mathbb{R}$ is supermodular if for all $S, T \subseteq E$*

$$s(S) + s(T) \leq s(S \cap T) + s(S \cup T). \tag{34}$$

One of the fundamental results in submodular optimization is that a greedy algorithm minimizes a linear function over a supermodular polyhedron. By

varying the linear function to be minimized, all extreme points of the supermodular polyhedron can be enumerated. In particular, define the supermodular polyhedron $\mathcal{P}(s) \subset \mathbb{R}^n$ to be the set of $x \in \mathbb{R}^n$ satisfying

$$\sum_{i \in T} x_i \geq s(T) \text{ for all } T \subseteq E.$$

The following theorem provides an algorithm that enumerates the extreme points of $\mathcal{P}(s)$.

Theorem 12 (See [25–27]). *For a linear ordering $e_1 \prec e_2 \prec \dots \prec e_n$ of the elements in E , Algorithm C.1 returns an extreme point v of $\mathcal{P}(s)$. Moreover, all extreme points of $\mathcal{P}(s)$ can be enumerated by considering all linear orderings of the elements of E .*

Algorithm C.1: GREEDY(s, E, \prec)

comment: Returns extreme point v of $\mathcal{P}(s)$ corresponding to the ordering \prec .

for $i = 1, \dots, n$

 Set $v_i = s(\{e_1, e_2, \dots, e_i\}) - s(\{e_1, e_2, \dots, e_{i-1}\})$

return (v)

Proof. See [25–27]. □

Theorem 12 is the key tool we employ to establish Lemma 5. In order to apply it, we require the following lemma.

Lemma 6. *For any joint distribution of the form $p(q)p(x) \prod_{i=1}^m p(y_i|x)p(u_i|y_i, q)$ and fixed $D \in \mathbb{R}$, define the set function $f : 2^{[m]} \rightarrow \mathbb{R}$ as:*

$$\begin{aligned} f(\mathcal{I}) &:= I(Y_{\mathcal{I}}; U_{\mathcal{I}}|U_{\mathcal{I}^c}, Q) - (D - H(X|U_1, \dots, U_m, Q)) \\ &= \sum_{i \in \mathcal{I}} I(U_i; Y_i|X, Q) + H(X|U_{\mathcal{I}^c}, Q) - D, \end{aligned} \quad (35)$$

and the corresponding non-negative set function $f^+ : 2^{[m]} \rightarrow \mathbb{R}$ as $f^+ = \max\{f, 0\}$. The functions f and f^+ are supermodular.

Proof. In order to verify that f is supermodular, it suffices to check that the function $f'(\mathcal{I}) = I(Y_{\mathcal{I}}; U_{\mathcal{I}}|U_{\mathcal{I}^c}, Q)$ is supermodular since the latter two terms in (35) are constant. To this end, consider sets $T, S \subseteq \{1, \dots, m\}$ and observe

that:

$$\begin{aligned}
f'(S) + f'(T) &= I(Y_S; U_S | U_{S^c}, Q) + I(Y_T; U_T | U_{T^c}, Q) \\
&= H(U_S | U_{S^c}, Q) - H(U_S | Y_S, Q) + H(U_T | U_{T^c}, Q) - H(U_T | Y_T, Q) \\
&= H(U_S | U_{S^c}, Q) + H(U_T | U_{T^c}, Q) \\
&\quad - H(U_{S \cup T} | Y_{S \cup T}, Q) - H(U_{S \cap T} | Y_{S \cap T}, Q) \tag{36}
\end{aligned}$$

$$\begin{aligned}
&= H(U_{S \setminus T} | U_{S^c}, Q) + H(U_{S \cap T} | U_{(S \cap T)^c}, Q) + H(U_T | U_{T^c}, Q) \\
&\quad - H(U_{S \cup T} | Y_{S \cup T}, Q) - H(U_{S \cap T} | Y_{S \cap T}, Q) \tag{37}
\end{aligned}$$

$$\begin{aligned}
&= H(U_{S \setminus T} | U_{S^c}, Q) + H(U_T | U_{T^c}, Q) - H(U_{S \cup T} | Y_{S \cup T}, Q) \\
&\quad + I(U_{S \cap T}; Y_{S \cap T} | U_{(S \cap T)^c}, Q) \\
&\leq H(U_{S \setminus T} | U_{(S \cup T)^c}, Q) + H(U_T | U_{T^c}, Q) - H(U_{S \cup T} | Y_{S \cup T}, Q) \\
&\quad + I(U_{S \cap T}; Y_{S \cap T} | U_{(S \cap T)^c}, Q) \tag{38} \\
&= I(U_{S \cup T}; Y_{S \cup T} | U_{(S \cup T)^c}, Q) + I(U_{S \cap T}; Y_{S \cap T} | U_{(S \cap T)^c}, Q) \\
&= f'(S \cap T) + f'(S \cup T).
\end{aligned}$$

The labeled steps above can be justified as follows:

- (36) follows since U_i is conditionally independent of everything else given (Y_i, Q) .
- (37) is simply the chain rule.
- (38) follows since conditioning reduces entropy.

Next, we show that $f^+ = \max\{f, 0\}$ is supermodular. Observe first that f is monotone increasing, i.e., if $S \subset T$, then $f(S) \leq f(T)$. Thus, fixing $S, T \subseteq \{1, \dots, m\}$, we can assume without loss of generality that

$$f(S \cap T) \leq f(S) \leq f(T) \leq f(S \cup T).$$

If $f(S \cap T) \geq 0$, then (34) is satisfied for $s = f^+$ by the supermodularity of f . On the other hand, if $f(S \cup T) \leq 0$, then (34) is a tautology for $s = f^+$. Therefore, it suffices to check the following three cases:

- Case 1: $f(S \cap T) \leq 0 \leq f(S) \leq f(T) \leq f(S \cup T)$. In this case, the supermodularity of f and the fact that $f^+ \geq f$ imply:

$$\begin{aligned}
f^+(S \cup T) + f^+(S \cap T) &\geq f(S \cup T) + f(S \cap T) \\
&\geq f(S) + f(T) = f^+(S) + f^+(T).
\end{aligned}$$

- Case 2: $f(S \cap T) \leq f(S) \leq 0 \leq f(T) \leq f(S \cup T)$. Since f is monotone increasing, we have:

$$f^+(S \cup T) + f^+(S \cap T) = f(S \cup T) + 0 \geq f(T) + 0 = f^+(S) + f^+(T).$$

- Case 3: $f(S \cap T) \leq f(S) \leq f(T) \leq 0 \leq f(S \cup T)$. By definition of f^+ :

$$f^+(S \cup T) + f^+(S \cap T) = f(S \cup T) + 0 \geq 0 + 0 = f^+(S) + f^+(T).$$

Hence, $f^+ = \max\{f, 0\}$ is supermodular. □

D Amplifying a Pointwise Convexity Constraint

Lemma 7. *Let $r_1, r_2 \in \mathbb{R}$ be given, and suppose $f_1 : K \rightarrow \mathbb{R}$ and $f_2 : K \rightarrow \mathbb{R}$ are continuous functions defined on a compact domain $K \subset \mathbb{R}^n$. If there exists a function $h : [0, 1] \rightarrow K$ satisfying*

$$t(f_1 \circ h)(t) + (1-t)(f_2 \circ h)(t) \leq tr_1 + (1-t)r_2 \quad \text{for all } t \in [0, 1], \quad (39)$$

then there exists $x_1^, x_2^* \in K$ and $t^* \in [0, 1]$ for which*

$$\begin{aligned} t^* f_1(x_1^*) + (1-t^*) f_1(x_2^*) &\leq r_1 \\ t^* f_2(x_1^*) + (1-t^*) f_2(x_2^*) &\leq r_2. \end{aligned}$$

Before we prove the lemma, we make a few remarks. At first glance, this lemma appears somewhat bizarre. Indeed, the set K need only be compact (e.g., connectedness is not required) and h can be an arbitrarily complicated function, as long as it satisfies (39). The strange nature of the lemma is echoed by the proof in that we merely prove the existence of the desired x_1^*, x_2^* and t^* ; no further information is obtained. Stripped to its core, the existence of the desired x_1^*, x_2^* and t^* essentially follows from the pigeonhole principle, which manifests itself in the sequential compactness of K .

Despite its strange nature, Lemma 7 is crucial in establishing the converse result for the multiterminal source coding problem under logarithmic loss. In this application, K is taken to be a closed subset of a finite-dimensional probability simplex and f_1, f_2 are conditional entropies evaluated for probability distributions in K .

Finally, we remark that the Lemma 7 can be generalized to a certain extent. For example, the function h need only be defined on a dense subset of $[0, 1]$ and the set K can be a more general sequentially compact space.

Proof of Lemma 7. Since f_1, f_2 are continuous⁴ and K is compact, there exists $M < \infty$ such that f_1 and f_2 are bounded from above and below by M and $-M$, respectively. Fix $\epsilon > 0$, and partition the interval $[0, 1]$ as $0 = t_1 < t_2 < \dots < t_m = 1$, such that $|t_{j+1} - t_j| < \frac{\epsilon}{M}$. For convenience define $x_{t_j} := h(t_j)$ when t_j is in the partition.

Now, for $i = 1, 2$ define piecewise-linear functions $g_1(t), g_2(t)$ on $[0, 1]$ by:

$$g_i(t) = \begin{cases} f_i(x_{t_j}) & \text{if } t_j \text{ is in the partition} \\ \theta f_i(x_{t_j}) + (1-\theta) f_i(x_{t_{j+1}}) & \text{if } t \text{ is in the interval } (t_j, t_{j+1}), \end{cases} \quad (40)$$

where $\theta \in (0, 1)$ is chosen so that $t = \theta t_j + (1-\theta)t_{j+1}$ when t is in the interval (t_j, t_{j+1}) .

⁴Although not required for our purposes, we can assume f_1 and f_2 are defined and continuous over all of \mathbb{R}^n . This is a consequence of the Tietze extension theorem.

With $g_1(t)$ and $g_2(t)$ defined in this manner, suppose $t = \theta t_j + (1 - \theta)t_{j+1}$ for some j and θ . Then some straightforward algebra yields:

$$\begin{aligned}
tg_1(t) + (1 - t)g_2(t) &= (\theta t_j + (1 - \theta)t_{j+1}) (\theta f_1(x_{t_j}) + (1 - \theta)f_1(x_{t_{j+1}})) \\
&\quad + (1 - \theta t_j - (1 - \theta)t_{j+1}) (\theta f_2(x_{t_j}) + (1 - \theta)f_2(x_{t_{j+1}})) \\
&= \theta^2 [t_j f_1(x_{t_j}) + (1 - t_j)f_2(x_{t_j})] \\
&\quad + (1 - \theta)^2 [t_{j+1} f_1(x_{t_{j+1}}) + (1 - t_{j+1})f_2(x_{t_{j+1}})] \\
&\quad + \theta(1 - \theta) [(1 - t_j)f_2(x_{t_{j+1}}) + (1 - t_{j+1})f_2(x_{t_j}) \\
&\quad\quad + t_{j+1} f_1(x_{t_j}) + t_j f_1(x_{t_{j+1}})] \\
&\leq \theta^2 [t_j f_1(x_{t_j}) + (1 - t_j)f_2(x_{t_j})] \\
&\quad + (1 - \theta)^2 [t_{j+1} f_1(x_{t_{j+1}}) + (1 - t_{j+1})f_2(x_{t_{j+1}})] \\
&\quad + \theta(1 - \theta) [(1 - t_{j+1})f_2(x_{t_{j+1}}) + (1 - t_j)f_2(x_{t_j}) \\
&\quad\quad + t_j f_1(x_{t_j}) + t_{j+1} f_1(x_{t_{j+1}})] + \epsilon \\
&\leq \theta^2 [t_j r_1 + (1 - t_j)r_2] \\
&\quad + (1 - \theta)^2 [t_{j+1} r_1 + (1 - t_{j+1})r_2] \\
&\quad + \theta(1 - \theta) [(1 - t_{j+1})r_2 + (1 - t_j)r_2 \\
&\quad\quad + t_j r_1 + t_{j+1} r_1] + \epsilon \\
&= (\theta t_j + (1 - \theta)t_{j+1})r_1 + (1 - \theta t_j - (1 - \theta)t_{j+1})r_2 + \epsilon \\
&= tr_1 + (1 - t)r_2 + \epsilon, \tag{41}
\end{aligned}$$

where the first inequality follows since $|t_{j+1} - t_j|$ is small, and the second inequality follows from the fact that (39) holds for each t_j in the partition. Notably, this implies that it is impossible to have

$$g_1(t) > r_1 + \epsilon \quad \text{and} \quad g_2(t) > r_2 + \epsilon$$

hold simultaneously for any $t \in [0, 1]$, else we would obtain a contradiction to (41). Also, since we included the endpoints $t_1 = 0$ and $t_m = 1$ in the partition, we have the following two inequalities:

$$g_1(1) \leq r_1, \quad \text{and} \quad g_2(0) \leq r_2.$$

Combining these observations with the fact that $g_1(t)$ and $g_2(t)$ are continuous, there must exist some $t^* \in [0, 1]$ for which

$$g_1(t^*) \leq r_1 + \epsilon, \quad \text{and} \quad g_2(t^*) \leq r_2 + \epsilon$$

simultaneously. An illustration of this is given in Figure 5, which is a mere variation on the classical intermediate value theorem.

Applying this result, we can find a sequence $\{x_1^{(n)}, x_2^{(n)}, t^{(n)}\}_{n=1}^{\infty}$ satisfying

$$\begin{aligned}
t^{(n)} f_1(x_1^{(n)}) + (1 - t^{(n)})f_1(x_2^{(n)}) &\leq r_1 + \frac{1}{n} \\
t^{(n)} f_2(x_1^{(n)}) + (1 - t^{(n)})f_2(x_2^{(n)}) &\leq r_2 + \frac{1}{n}
\end{aligned}$$

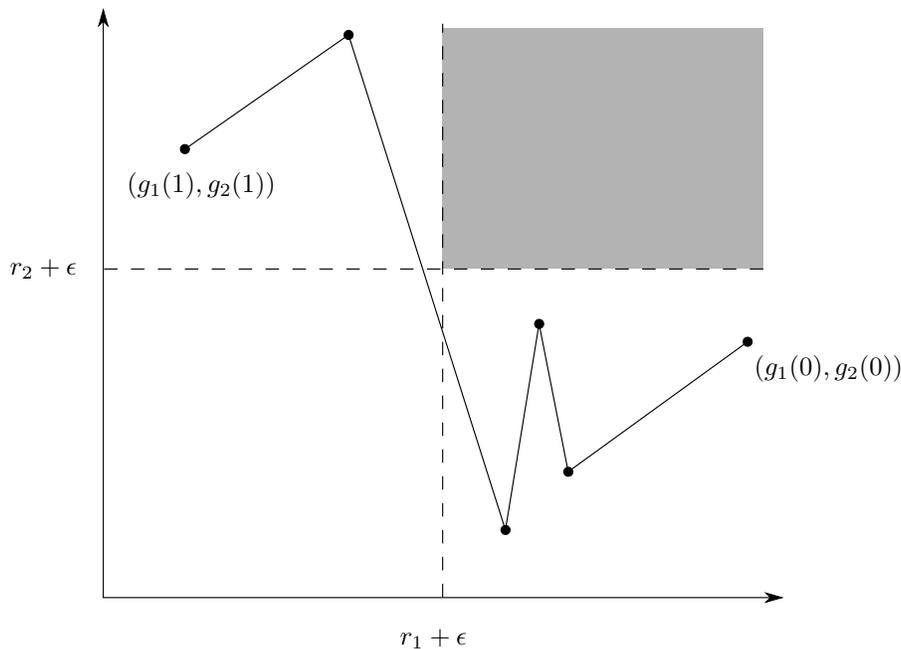


Figure 5: A parametric plot of the function $\varphi : t \mapsto (g_1(t), g_2(t))$. Since $\varphi(t)$ is continuous, starts with $g_2(0) \leq r_2 + \epsilon$, ends with $g_1(1) \leq r_1 + \epsilon$, and doesn't intersect the shaded area, $\varphi(t)$ must pass through the lower-left region.

for each $n \geq 1$. Since $K \times K \times [0, 1]$ is sequentially compact, there exists a convergent subsequence $\{n_j\}_{j=1}^{\infty}$ such that $(x_1^{(n_j)}, x_2^{(n_j)}, t^{(n_j)}) \rightarrow (x_1^*, x_2^*, t^*) \in K \times K \times [0, 1]$. The continuity of f_1 and f_2 then apply to yield the desired result. \square

E A Lemma for the Daily Double

For a given joint distribution $p(y_1, y_2)$ on the finite alphabet $\mathcal{Y}_1 \times \mathcal{Y}_2$, let $\mathcal{P}(R_1, R_2)$ denote the set of joint pmf's of the form

$$p(q, y_1, y_2, u_1, u_2) = p(q)p(y_1, y_2)p(u_1|y_1, q)p(u_2|y_1, q)$$

which satisfy

$$\begin{aligned} R_1 &\geq I(Y_1; U_1 | U_2, Q) \\ R_2 &\geq I(Y_2; U_2 | U_1, Q) \\ R_1 + R_2 &\geq I(Y_1, Y_2; U_1, U_2 | Q) \end{aligned}$$

for given finite alphabets $\mathcal{U}_1, \mathcal{U}_2, \mathcal{Q}$.

Lemma 8. For R_1, R_2 satisfying $R_1 \leq H(Y_1)$, $R_2 \leq H(Y_2)$, and $R_1 + R_2 \leq H(Y_1, Y_2)$, the infimum

$$\inf_{p \in \mathcal{P}(R_1, R_2)} \{H(Y_1|U_1, U_2, Q) + H(Y_2|U_1, U_2, Q)\}$$

is attained by some $p^* \in \mathcal{P}(R_1, R_2)$ which satisfies $R_1 + R_2 = I(Y_1, Y_2; U_1^*, U_2^*|Q^*)$, where U_1^*, U_2^*, Q^* correspond to the auxiliary random variables defined by p^* .

Proof. First, note that the infimum is always attained since $\mathcal{P}(R_1, R_2)$ is compact and the objective function is continuous on $\mathcal{P}(R_1, R_2)$. Therefore, let U_1^*, U_2^*, Q^* correspond to the auxiliary random variables which attain the infimum.

If $H(Y_1|U_1^*, U_2^*, Q^*) + H(Y_2|U_1^*, U_2^*, Q^*) = 0$, then we must have $I(Y_1, Y_2; U_1^*, U_2^*|Q^*) = H(Y_1, Y_2)$. Thus, $R_1 + R_2 = I(Y_1, Y_2; U_1^*, U_2^*|Q^*)$.

Next, consider the case where $H(Y_1|U_1^*, U_2^*, Q^*) + H(Y_2|U_1^*, U_2^*, Q^*) > 0$. Assume for sake of contradiction that $R_1 + R_2 > I(Y_1, Y_2; U_1^*, U_2^*|Q^*)$. For any $p \in \mathcal{P}(R_1, R_2)$:

$$I(Y_1; U_1|U_2, Q) + I(Y_2; U_2|U_1, Q) \leq I(Y_1, Y_2; U_1, U_2|Q).$$

Hence, at most one of the remaining rate constraints can be satisfied with equality. If none of the rate constraints are satisfied with equality, then define

$$(\tilde{U}_1, \tilde{U}_2) = \begin{cases} (U_1^*, U_2^*) & \text{with probability } 1 - \epsilon \\ (Y_1, Y_2) & \text{with probability } \epsilon. \end{cases}$$

For $\epsilon > 0$ sufficiently small, the distribution \tilde{p} corresponding to the auxiliary random variables $\tilde{U}_1, \tilde{U}_2, Q^*$ is still in $\mathcal{P}(R_1, R_2)$. However, \tilde{p} satisfies

$$H(Y_1|\tilde{U}_1, \tilde{U}_2, Q^*) + H(Y_2|\tilde{U}_1, \tilde{U}_2, Q^*) < H(Y_1|U_1^*, U_2^*, Q^*) + H(Y_2|U_1^*, U_2^*, Q^*),$$

which contradicts the optimality of p^* .

Therefore, assume without loss of generality that

$$\begin{aligned} R_1 &= I(Y_1; U_1^*|U_2^*, Q^*) \\ R_1 + R_2 &> I(Y_1, Y_2; U_1^*, U_2^*|Q^*). \end{aligned}$$

This implies that $R_2 > I(Y_2; U_2^*|Q^*)$. Now, define

$$\tilde{U}_2 = \begin{cases} U_2^* & \text{with probability } 1 - \epsilon \\ Y_2 & \text{with probability } \epsilon. \end{cases}$$

Note that for $\epsilon > 0$ sufficiently small:

$$\begin{aligned} I(Y_2; U_2^*|Q^*) &< I(Y_2; \tilde{U}_2|Q^*) < R_2 \\ I(Y_1, Y_2; U_1^*, U_2^*|Q^*) &< I(Y_1, Y_2; U_1^*, \tilde{U}_2|Q^*) < R_1 + R_2, \end{aligned}$$

and for any $\epsilon \in [0, 1]$:

$$\begin{aligned} R_1 &= I(Y_1; U_1^* | U_2^*, Q^*) \geq I(Y_1; U_1^* | \tilde{U}_2, Q^*) \\ H(Y_1 | U_1^*, U_2^*, Q^*) + H(Y_2 | U_1^*, U_2^*, Q^*) &\geq H(Y_1 | U_1^*, \tilde{U}_2, Q^*) + H(Y_2 | U_1^*, \tilde{U}_2, Q^*). \end{aligned} \tag{42}$$

Since $R_2 \leq H(Y_2)$, as ϵ is increased from 0 to 1, at least one of the following must occur:

1. $I(Y_2; \tilde{U}_2 | Q^*) = R_2$.
2. $I(Y_1, Y_2; U_1^*, \tilde{U}_2 | Q^*) = R_1 + R_2$.
3. $I(Y_1; U_1 | \tilde{U}_2, Q^*) < R_1$.

If either of events 1 or 2 occur first then the sum-rate constraint is met with equality (since they are equivalent in this case). If event 3 occurs first, then all rate constraints are satisfied with strict inequality and we can apply the above argument to contradict optimality of p^* . Since (42) shows that the objective is nonincreasing in ϵ , there must exist a $\tilde{p} \in \mathcal{P}(R_1, R_2)$ which attains the infimum and satisfies the sum-rate constraint with equality. \square

References

- [1] D. Slepian and J. Wolf, “Noiseless coding of correlated information sources,” *Information Theory, IEEE Transactions on*, vol. 19, pp. 471 – 480, Jul 1973.
- [2] R. Ahlswede and J. Körner, “Source coding with side information and a converse for degraded broadcast channels,” *Information Theory, IEEE Transactions on*, vol. 21, pp. 629 – 637, Nov 1975.
- [3] A. Wyner, “On source coding with side information at the decoder,” *Information Theory, IEEE Transactions on*, vol. 21, pp. 294 – 300, May 1975.
- [4] A. Wyner and J. Ziv, “The rate-distortion function for source coding with side information at the decoder,” *Information Theory, IEEE Transactions on*, vol. 22, pp. 1 – 10, Jan. 1976.
- [5] T. Berger and R. Yeung, “Multiterminal source encoding with one distortion criterion,” *Information Theory, IEEE Transactions on*, vol. 35, pp. 228 – 236, Mar 1989.
- [6] A. Wagner, S. Tavildar, and P. Viswanath, “Rate region of the quadratic gaussian two-encoder source-coding problem,” *Information Theory, IEEE Transactions on*, vol. 54, pp. 1938 – 1961, May 2008.
- [7] V. Prabhakaran, D. Tse, and K. Ramachandran, “Rate region of the quadratic gaussian ceo problem,” in *Information Theory, 2004. ISIT 2004. Proceedings. International Symposium on*, p. 119, june-2 july 2004.

- [8] Y. Oohama, “Gaussian multiterminal source coding,” *Information Theory, IEEE Transactions on*, vol. 43, pp. 1912–1923, nov 1997.
- [9] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. New York, NY, USA: Cambridge University Press, 2006.
- [10] T. Courtade and R. Wesel, “Multiterminal source coding with an entropy-based distortion measure,” in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pp. 2040–2044, Aug. 2011.
- [11] P. Harremoës and N. Tishby, “The information bottleneck revisited or how to choose a good distortion measure,” in *Information Theory, 2007. ISIT 2007. IEEE International Symposium on*, pp. 566–570, June 2007.
- [12] T. Andre, M. Antonini, M. Barlaud, and R. Gray, “Entropy-based distortion measure for image coding,” in *Image Processing, 2006 IEEE International Conference on*, pp. 1157–1160, Oct. 2006.
- [13] T. Berger, Z. Zhang, and H. Viswanathan, “The ceo problem [multiterminal source coding],” *Information Theory, IEEE Transactions on*, vol. 42, pp. 887–902, may 1996.
- [14] T. Berger, *Multiterminal Source Coding*. In G. Longo (Ed.), *The Information Theory Approach to Communications*. New York, NY, USA: Springer-Verlag, 1977.
- [15] S.-Y. Tung, *Multiterminal Source Coding*. PhD thesis, Cornell University, Ithaca, NY, 1978.
- [16] A. El Gamal and Y.-H. Kim, *Network Information Theory*. Cambridge University Press, 2012.
- [17] S. Jana, “Alphabet sizes of auxiliary random variables in canonical inner bounds,” in *Information Sciences and Systems, 2009. CISS 2009. 43rd Annual Conference on*, pp. 67–71, March 2009.
- [18] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, 2nd ed., 2006.
- [19] E. Erkip, *The Efficiency of Information in Investment*. PhD thesis, Stanford University, 1996.
- [20] Y.-H. Kim, A. Sutivong, and T. Cover, “State amplification,” *Information Theory, IEEE Transactions on*, vol. 54, pp. 1850–1859, may 2008.
- [21] J. Korner and K. Marton, “How to encode the modulo-two sum of binary sources (corresp.),” *Information Theory, IEEE Transactions on*, vol. 25, pp. 219–221, Mar 1979.

- [22] W. Gu and M. Effros, “On approximating the rate region for source coding with coded side information,” in *Information Theory Workshop, 2007. ITW '07. IEEE*, pp. 432 –435, Sept. 2007.
- [23] W. Gu, S. Jana, and M. Effros, “On approximating the rate regions for lossy source coding with coded and uncoded side information,” in *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, pp. 2162 – 2166, July 2008.
- [24] H. Witsenhausen, “Some aspects of convexity useful in information theory,” *Information Theory, IEEE Transactions on*, vol. 26, pp. 265 – 271, May 1980.
- [25] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. Berlin: Springer-Verlag, 2003.
- [26] S. Fujishige, *Submodular Functions and Optimization*. Berlin: Elsevier Science, 2nd ed., 2010.
- [27] S. McCormick, *Submodular Function Minimization*. In *Discrete Optimization*, K. Aardal, G. Nemhauser, and R. Weismantel, eds. *Handbooks in Operations Research and Management Science*, vol. 12. Elsevier, 2005.